TES.
BYRN.1

Department of Psychology,

Trinity College, Dublin.

THE CONTEXTUAL NATURE OF CONDITIONAL REASONING

by

Ruth M.J. Byrne, B.A.

ISPA
6092

Ref. 6092

Instituto Superior de Psicologia Aplicada
BIBLIOTECA

A thesis submitted in fulfilment of the requirements for the degree of Doctor of Philosophy, to the University of Dublin.

I hereby declare that this thesis is entirely my own work and that it has not been submitted as an exercise for a degree at any other university.

I also give my permission to the Library to lend or copy this thesis upon request.

Signed : Ruth Byrne

Date : 30/12/86

Acknowledgements

My heartfelt thanks go to my Da and Ma, Paul and Mary, for their encouragement at all times and for their support, both personal and financial. My thanks also to my family: Linda, Cathy, Noel, Tony, Paul, Steven, Jenny, and Molly. I am deeply indebted to my companion and best critic, Mark Keane, for encouraging, motivating and critically assessing this work in all it's stages, and for forcing me always to shape ideas into coherence.

I am grateful to Nick McDonald for supervising this thesis, and to everyone else who has read or otherwise shown interest in this work: Chris Simms, Ross Skelton, Howard Smith, Leo Baker, Paul Dowding, John Masterson, Alison Black, Angus Gellatly, Ken Gilhooly and Phil Johnson-Laird.

I am grateful to Trinity Trust for awarding me a postgraduate award for the years 1984-1986, and to the Medical Research Council, Applied Psychology Unit for providing facilities to produce this thesis. Thanks also to my sister, Cathy, for providing me with secretarial facilities and advice. I am indebted to the undergraduates who took part in the experiments.

Finally, special thanks to Lisa, for keeping my spirits up.

Summary

In this thesis, it is proposed that the products of human reasoning are a function of the context in which that reasoning occurs. Consequently, the effects of two kinds of contexts are addressed, in two series of experiments.

In the first series, the environmental context, in which conditional reasoning is embedded, was explored. In Experiment 1 subjects were given conditional arguments, accompanied by contextual information consisting of a second conditional. The antecedent in the second conditional was either an alternative condition or an additional condition. It was found that alternative conditions suppress the inferences traditionally considered fallacious, while additional conditions suppress the inferences sanctioned as logically valid. It was also found that when these alternative or additional conditions are conjoined or disjoined in the minor premises that accompany the conditional arguments, this suppression no longer occurs.

In Experiment 2, groups of subjects were again given conditional arguments that were either accompanied by extra contextual information or not. Further groups were given conditional arguments that were prefaced either by information about the relevance of contextual information, or by tasks requiring the retrieval of contextual information. It was found that neither recognising the relevance of contextual information nor retrieving that information in isolation, is enough to suppress the inferences. It was found once more, that

contextual information, which embodies both of these features, suppresses the inferences, replicating the primary results of Experiment 1.

In Experiment 3, groups of subjects were given conditional inducements. One group received inducements unaccompanied by any contextual information. Other groups received them accompanied by information about the situation in which the utterance was made, or by information about the duration of the utterance (either that it was long or short), or accompanied by information about both the situation and the duration. It was found that information about both the situation and that it was a long duration suppresses the invalid inferences, while information about the situation and that it was a short duration suppresses the valid inferences. Neither situational information, which was hypothesised to enable the retrieval of relevant information, nor durational information, which was hypothesised to enable the recognition of the relevance of other information, were singly sufficient to suppress the inferences. A model of inference-testing, based on the conjoined operation of the processes of recognising the relevance of other information and retrieving specific instances of the relevant information, is described.

A second way that the context in which conditional reasoning is embedded can affect reasoning was investigated. In Experiment 4, the frequency of inferences made from sequences of conditional premises was compared to the frequency of inferences made from those premises in isolation from each other. Subjects were given conditional premises of different forms. It was found that fewer inferences are made on sequences of premises than from those premises in isolation from each other, for certain forms of premises.

In Experiment 5, two features, identified as potentially responsible for this, were manipulated. Some subjects received the sequences of conditional premises as before, while others received the premises in isolation, again as before, both of the form for which differences were observed in Experiment 4. Other subjects received arguments constructed to ensure that either the joint presentation of the premises was preserved, or that the uncertainty of using an intermediate conclusion as a minor premise from which to make an inference was preserved, to establish the respective roles of these factors. It was found that the source of differences in the frequency of inferences between sequences of premises and those premises in isolation can be attributed to the joint representation of the premise information.

Experiment 6 determined the role of some of the features of the form of argument for which differences were observed. Groups of subjects were presented with two different forms of sequences of premise, and comparable premises in isolation. One of these new forms possessed one of the features, that of a negative in the first conditional, while the second sequence possessed two of the features, that

of a negative in the first conditional with that negative located at the consequent of the conditional. Neither possessed the third feature, of the atransitivity of the argument. It was found that neither of these sequences of premises differed from comparable premises in isolation, in the frequency of inferences made. This indicates that all three features are responsible for the joint representation of arguments of certain sequences leading to fewer inferences. A model is proposed that shows how the representation of sequences of premises differs from the representation of premises in isolation. Finally, the implications of the effects of context on conditional reasoning found in these experiments, for the four theories of reasoning described in the introductory review, and for cognitive psychology in general, are explored in the final chapter.

Table of Contents

List of figures

List of Tables

Appendix A

Appendix B

Appendix C

PROLOGUE: HUMAN REASONING

Conditional reasoning

Most people can think about hypothetical situations. The ability to do this, and to consider various possible consequences of our actions without having to perform those actions, is a labour-saving and often, a life-saving attribute. The capacity to think about the conditions under which an action will occur is equally invaluable, helping us to plot how to bring an action about, or how to prevent an action occurring. Planning future events or predicting eventual outcomes, and understanding the reasons for past events is made possible because of our ability to think about the hypothetical relation between certain conditions and consequences. This kind of hypothetical or conditional reasoning lies not only at the heart of all other kinds of reasoning but also, arguably, at the heart of thinking in general. As such it is the fundamental underpinning of our ability to behave in an intelligent and meaningful way.

Conditional reasoning, as it occurs in our planning and also in the communication of our plans; in our expectations about the future and our beliefs about the past; and in our consideration of the connections between things in the world, pervades all of our everyday experience. How do we accomplish such reasoning ? We are, by and large, reasonable creatures: granted both the peculiar circumstances of each situation we find ourselves in and the vagaries in the kinds of information we focus on. One approach to understanding our reasonableness is to ask whether or not it is "logical". The meaning of the word "logical" is frequently equated with "reasonable", in both vernacular and technical usage. However, "logical" has come to be more precisely defined to mean "reasoning in accordance with the rules of some particular logical system". That is, to be logical is to agree that the inferences deemed valid in a certain logical system are indeed valid, and those considered invalid are also indubitably so.

Propositional logic

For conditional reasoning, the standard against which inferential validity is judged is propositional logic. This is a logical system that treats sentential connectives such as AND, OR and IF. The logical analysis of conditional reasoning focuses on its embodiment in IF sentential frames. So for a sentence like "if we go the cinema then we will have a good time", the propositional calculus allows one to appraise whether or not it is valid to deduce one component of the conditional, for example the <u>consequent</u>: "we will have a good time", from knowing that the <u>antecedent</u>: "we go to the cinema" occurs. An inference is valid whenever the conclusion, (in this case, "we will have a good time"), will always be true whenever the premises, ("if we go to the cinema then we will have a good time" and "we go to the cinema") are true.

On a logical analysis, the truth of a conditional sentence depends solely on the truth or falsity of the components that are connected in it. An IF sentence in the propositional calculus is defined as material implication. That is, a conditional as a whole is true only when it s antecedent is false, or it's consequent is true. A truth table, tabulating all possible arrangements of truth values of the components, will illustrate this:

if we go to the cinema then we will have a good time

we go to the cinema	we have a good time
true	true
true	false
false	true
false	false

An implicative sentence is false only in the second situation above, where the antecedent is true but the consequent is false. From this truth-functional account of material implication, rules can be derived that sanction some inferences and outlaw others. These rules may be phrased in a natural deduction format, as in most introductory textbooks of propositional logic (e.g., Copi, 1982). The rules are stated in terms of what conclusion may be made from the premises.

The propositional calculus is entirely formal that is, it is a syntactic, content-independent system. So the validity of inferences is judged solely on the basis of the form they take, not on the basis of their meaning or of their truth. Any content can be put into IF sentences and no matter how nonsensical or how false the conclusion of an inference from such a sentence is, the inference will still be valid, provided that the form in which it is phrased is a valid one. Of the four symmetrical inferences arising from a conditional sentence, two are deemed valid and two are considered fallacious in the propositional calculus. From a premise like "if it is raining then she will get wet", and the minor premise:

(1) it is raining

a valid *modus ponens* inference is

therefore she will get wet

and from the minor premise

(2) she will not get wet

the valid *modus tollens* inference is

therefore it is not raining

From the denial of the antecedent

(3) it is not raining

an invalid conclusion is

therefore she will not get wet

and from the affirmation of the consequent

(4) she will get wet

it is invalid to conclude

therefore it is raining

Logical systems are prescriptive or normative. That is, they dictate which the inferences people *should* make. As we will see in the following chapters, the prescriptions of logics are too context-independent to hold in the many situations in which people find themselves. Many logicians (e.g., Haack, 1978; Mates, 1972) maintain that IF is never used in a material implicative sense in everyday usage. Capturing the meaning of conditionals by a logical system is a major concern of contemporary philosophy (e.g., Anderson and Belnap, 1975).

Human reasoning

Propositional logic, however, has been considered to be descriptive of human reasoning by some researchers (e.g., Henle, 1962). This view has led to the claim that people are not only "logical" in the sense of making conclusions that are in accord with the conclusions sanctioned by a logical system. It is claimed further that people reach these conclusions through the operation of mental rules of logic corresponding to the rules for material implication in the propositional calculus (e.g., Braine, 1978).

Other theorists, as we will see, have argued that people do not possess mental rules of logic (e.g., Griggs, 1983; Johnson-Laird, 1983). People frequently sanction inferences that are deemed fallacious in propositional logic. Their conclusions are influenced by a variety of factors, not merely by the syntactic form nor even just by the truth of each component. Consequently, some researchers have concluded that people do not possess logical rules, and further that they are not capable of deliberately making conclusions that are logical (e.g., Pollard, 1982). We will see that a growing body of research indicates that it is no longer feasible to hold that people have rules of logic in their head s, which they apply when they encounter a situation that requires a reasoned conclusion (e.g., Griggs and Cox, 1982). However, we also see that it has been argued that it is nonetheless possible to conceive of people being capable of logical reasoning. That is, people are capable of making inferences that are deemed valid by propositional logic, without relying on the operation of mental rules of logic (e.g., Johnson-Laird, 1982; 1983).

In order to assess the relative merits of these opposing views of human reasoning it is necessary to describe the primary theories developed from these views, and to survey the empirical data for which they attempt to account.

There are four main theoretical approaches to reasoning. The oldest and once dominant tradition is that of the "mental logic" approach. Challenging this are accounts of reasoning which emphasize it's reliance on the retrieval of memories of past experience from long-term memory on the one hand, and it's dependence on non-logical response biases on the other. A third approach compromises between mental logic and memory retrieval, proposing instead that reasoning is achieved through rules that are intermediate in generality between abstract syntactic rules and specific memories. Finally, the fourth theory, superseding a memory-based account, proposes that reasoning consists of manipulating the representations or models that people construct of the information about which they are reasoning. These four theoretical approaches will be outlined in turn. Their principal tenets will be described, illustrated by a resume of the main models arising from them. The empirical evidence that each theory accounts for will be recounted and evaluations of the strength of this evidence in supporting the theory will be made. Criticisms of each theory will then be entertained.

A number of considerations will hold in making these criticisms, arising from criteria which any theory of human reasoning must meet (Johnson-Laird, 1983; Rips, 1983a). One such criterion is that the mechanisms considered responsible for reasoning must be able to apply in all the situations in which people reason. They must be capable of applying to new situations, where a person does not simply have a stored response to the problem. Another criterion is that these mechanisms must be able to account for the effect that the situation, the context and the content of the problem, has on reasoning. These dual criteria of generality and specificity mutually constrain the success of mechanisms proposed at different levels of abstraction. A third criterion is that a theory of reasoning must be able to characterize the relation of reasoning to other abilities, such as problem-solving, comprehension, memory retrieval, and working memory constructions. A fourth criterion is that the mechanisms of reasoning must be parsimonious, as it seems likely that human cognition does not include a vast repetition of structures. These are four of the most important criteria that will be applied to evaluating the theories of reasoning that are reviewed in the next chapter. It will be assumed throughout this thesis that a general theory of human reasoning is possible, that is, that individual differences (e.g.,Johnson-Laird, 1983; Rips and Conrad, 1983; Sternberg, 1985) are not so pervasive as to render a theory of reasoning obsolete.

CHAPTER 1:

THEORIES OF HUMAN REASONING

Mental logic

The mental logic view of reasoning attributes correct reasoning to the operation of formal logical rules, derived from the rules of propositional logic. These rules are formal in the sense of being abstract, content-independent rules. Propositional logic is assumed to describe the process of reasoning entirely, by some researchers (e.g., Henle, 1962; Piaget and Inhelder, 1973). Others consider that some of the inferences sanctioned by the propositional calculus, such as modus tollens, are so demonstrably difficult that they cannot be accomplished simply by the operation of a single elementary rule. These latter theorists have attempted to develop logical systems descriptive of a 'mental logic' which consist solely of simple elementary rules. (Braine, 1978; Braine, Reiser and Rumain, 1984; Osherson, 1975, 1976; Rips, 1983b).

Natural deduction models, as mental logics are sometimes called, specify a repertoire of elementary deductive steps that take a reasoner from one inferential step to another. All of the elementary rules are intended to be intuitively obvious principles. They are expected to be both logically and psychologically valid. These depart from propositional logic in that some rules in the propositional calculus will not be in the mental logic system, because of this criterion of elementariness in mental logic systems. So while all the rules of the mental logics are ones that are sanctioned by propositional logic, the mental logic rules are not exhaustive of the propositional calculus ones.

Mental logic systems share the "natural deduction" format adopted in modern accounts of propositional logic. That is, the rules of the mental logic are phrased in terms of propositions which allow some components to be deduced rather than other components. This is preferred over other formats: it is never proposed that mental logic is represented in terms of truth tables, or in terms of axioms or primitives from which people must generate rules. It is assumed that the logical rules are embodied in people's heads along with instructions for their use. The rules operate in a uniform way on standard translations. For example, in present implementations of mental logic (e.g., Braine, 1978; Rips, 1983b), the rules apply only to the term "if" and to abstract tokens such as p's and q's.

Reasoning errors are defined relative to the inferences that are sanctioned by the propositional calculus. Because all of the mental logic rules are ones which are valid in propositional logic, there is no such thing as a rule which embodies a logically incorrect principle. But it is possible for a sentence to be mapped to an inappropriate rule. That is, a person may initially misunderstand a sentence and so apply a rule to it which is not the rule that should be applied to that sentence. In this way, errors are attributed to components other than the operation of the formal rules.

The theory maintains a strict division between the representation of information, that is, understanding what is to be reasoned about, and the operations that process it, that is, "pure reasoning". On the mental logic account, the representation of information is achieved by a comprehension component, which translates meaningful information into abstract syntactic code (Falmagne, 1975). While the inferential mechanisms of mental logic are isomorphic to a normative theory, the encoding or comprehension component clearly is not. The competence model specified by the natural deduction rules must be incorporated within a performance model (Kahn and Rips, 1983).

The burden for this approach is to identify the factors responsible for errors in the comprehension component and how they affect specific inferences (Rumain, Connell and Braine, 1983; Staudenmayer, 1975). Errors are attributed to a variety of sources within this interpretative component, such as the incorrect input of the material, (e.g.,misreading the premises); or the incorrect translation of the material, (such as mapping a sentence to an inappropriate piece of abstract code). The interpretative process is more inductive than deductive. There is an intrusion of habits that are characteristic of everyday practical reasoning (Braine, 1978) and ordinary comprehension (Braine and Rumain, 1983) into formal reasoning. Practical reasoning refers to the decisions that people make in everyday situations, which may concern actions rather than verbal utterances or text. Formal reasoning, it is claimed, requires compartmentalising information and discovering the minimum commitments of utterances. This is said to be alien to practical reasoning.

Let us turn now to two of the main instantiations or models derived according to these theoretical principles.

Mental logic models

The first model to be outlined is a recent one proposed by Rips (1983a, 1983b, Rips and Conrad, 1983). The basic assumption made is that people reason through the construction of mental proofs or derivations of conclusions, which are constructed step by elementary step. Rips has provided an implementation of a natural deduction system (ANDS), that is specialised for propositional connectives. The core of the natural deduction system is the inference rules. These are embodied in a set of computational "routines", which are productions systems, or abstract schemas. An example of such a structure is:

Conditions:
(1) Assertion tree contains proposition X = If p, q
(2) X has not been used by the modus ponens schemas
(3) Assertion tree contains proposition p

Action:
(1) Add q to assertion tree

The schemas dictate which deductions are possible, add the deduced propositions to the memory trees and keep the procedure moving towards a solution.

The mental logic rules take the form of abstract schemas and produce hierarchical memory structures. They amount to specific ways to manipulate propositions in memory. There are two types of inference schemas which operate in conjunction with two interconnected memory trees. "Forward" moving inference rules insert propositions into an "assertion" tree. This tree holds the premises, and any immediately derivable propositions, such as simplifications of the premises. It also holds suppositions, but these are kept in separate nodes from assertions, to indicate their suppositional status. Everything that is deduced from a supposition is inserted into the supposition node. This is a useful way of exploring the consequences of propositions whose truth value is unknown. "Backward" moving inference rules insert propositions into these supposition nodes and also into a "subgoal" tree. The subgoal tree ensures that cognisance is taken of the conclusion from the start, to avoid drawing irrelevant subconclusions. As subgoal chaining is possible, multiple attempts at an argument can be made. Each routine or schema is checked for its applicatability to an argument in a fixed order.

It is claimed that this system can be extended easily to apply to causality. Its innovations compared with other natural deduction systems (e.g., Braine, 1978; Osherson, 1975) rest in its use of backward rules, which allow the creation of subgoals, and the use of conclusions in constructing a proof.

One of the primary problems for ANDS is that the verbal input to the system must match the logical constants, such as IF, exactly. That is, ANDS accepts only formal notation (e.g., If p, q) as input. A theory of human reasoning needs some provision for handling natural language sentences. Rips (1983b) suggests that this can be achieved either by adapting the deductive rules so that they apply directly to natural language, or by parsing natural language into the more formal system in which the rules are currently specified. This, as we will see however, is the proverbial can of worms.

The second natural deduction system we will examine is that proposed by Braine (1978; Braine, Reiser and Rumain, 1984). On this account, a complete theory of reasoning requires several components. Firstly, a comprehension mechanism, to understand natural language in terms of the abstract schemas that embody the inferential rules, is required. A second necessary component is a heuristic reasoning "program" containing the strategies that will put together a chain of inferences and that will select the schemas that are to be applied at each point in an argument's derivation. A third necessary provision is a "fallback" component, in case the system cannot produce a conclusion by the reasoning program. The fourth, and core, component, (and also the only component specified in any detail by natural deduction theorists) is a logical system, which specifies the repertoire of inferential steps available, in terms of the abstract schemas. There is also a need for some definition of the amount of working memory available (Braine, 1978).

The schemas are structures phrased with premises above a line and conclusions below this line. For example, the schema for double negation is:

$$\frac{\text{False (negative P)}}{\text{P}}$$

Braine's inference-line theory (1978) maintains that the logical function of IF is to state inference rules. IF effectively gives a warranty to conclude the consequent of a sentence given its antecedent (even though a sentence itself provides no justification for that warranty). Hence, there is no need for a modus ponens rule in the 1978 version as this inference is embodied in the function of IF. However, in the 1984 version of the model this inference is distinctly specified. Modus tollens requires both a mental logic rule and some further computations (because of its backward directionality). As inference rules have no application when their antecedent is false, there is no reason to assign them in such cases.

The rules outlined in Braine's system are claimed to constitute a valid, elementary, complete and universal set. The reasoning program that Braine outlines involves these "direct" reasoning schemas and also some "indirect" reasoning schemas. The reasoning program attempts to find the shortest route to a conclusion, and so direct reasoning schemas will be applied prior to indirect ones, in a fixed order. The indirect reasoning program amounts to several strategies such as reductio ad absurdum (assuming that what one wants to prove true is false , and showing that on this assumption one reaches a contradiction), and lemma production (creating suppositions and making deductions from these suppositions).

The first step of the reasoning process is, as in all natural deduction models, to decode the information into the abstract schema of the system. The process of this translation is as yet unspecified. Errors of reasoning are attributed to the misconstrual of the premise information, and also to processing difficulties, such as lapses of attention, or memory limitations.

Experimental evidence

The empirical findings cited as evidence for the psychological reality of mental logic can be categorised into those that bear on the existence of the logical rules and those that bear on the existence of the ancillary comprehension component. The former includes evidence that relates to the various schemas and processing components proposed, while the latter involves investigations into the effects of content and more recently, context.

Inference schemas

Braine's evidence for schemas is two-fold. Their elementariness is testified to by the infrequency with which people make errors on them. Their use in mental derivations is testified to by the technique (Osherson, 1975) of predicting the difficulty of an argument by the sum of the difficulties of its component inferential steps. Of the three measures of problem difficulty that Braine describes, in practise only the first, that of subjective ratings, is used. The other measures prove to be too problematic to be useful: reaction times are too long and confounded by reading differences, and reasoning errors are too few on these problems. Abstract content is used to minimise comprehension errors. The subject's task is typically to evaluate an argument as true or false, or as indeterminate, and then to rate its difficulty on a 9 point scale. These ratings are expected to correlate with the difficulty of the argument (as predicted either by the difficulty of each step or by the number of steps, regardless of their nature).

Braine provides no evidence for the proposed reasoning program. Instead, he argues that if a reasoning program is to use the repertoire of schemas proposed, to solve all the problems subjects do, to predict that direct reasoning is easier than indirect reasoning and also to predict that the difficulty of direct reasoning is a function of the length of the chain of inferences to be made, then it will be similar to the program suggested. Braine also argues that the two-part structure of direct and indirect reasoning is supported by the observation that direct reasoning appears to develop earlier than indirect reasoning and involves fewer errors than indirect reasoning.

Rips provides different kinds of evidence in support of his proposed repertoire of schemas. He provides four categories of evidence for his model. The first is that the features of his model are characteristic of human reasoning. People seem to use conclusions as the goal of their reasoning, they use suppositions, simplify the premises by paraphrasing and work backwards from the conclusion. They can also use a reductio ad absurdum strategy and can do logical doubletakes, (that is abandon unsuccessful strategies and initiate new ones). All of these features are captured by ANDS. However, as they were all observed in the protocol data that were used to shape the theory, Rips allows that they cannot be counted as evidence, but rather are indicative of the compatibility of ANDS with human reasoning. A second indicator of ANDS compatibility with human reasoning is the judged similarity of ANDS reasoning with that of human subjects, as rated by two judges (Rips, 1983b).

Evidence for the proposed processing structures of ANDs hierarchical tree structures in memory, and of the placement of suppositions in separate nodes, is found by Marcus (1982). She showed that people can recall nonsuppositions better than suppositions. The strongest evidence that Rips provides for the repertoire of schemas is from mathematically fitting the model to subject's decisions about whether the conclusions of reasoning arguments are necessarily true or are not necessarily true. In these problems Rips used materials concerning the location of people in states (if Betty is in Denver...), and about the operation of machines. He makes the strong assumption that no misinterpretations of the premises occur. The model predicts how often people will be correct through asserting what rules will be needed, and through utilising the crucial assumption of the nonavailability of some of these rules on some proportion of trials (due either to retrieval difficulties, a failure to recognise the rules as applicable or a failure to apply them properly). This ensures that the model fits the data reasonably well. Further evidence (Rips and Conrad, 1983) suggests that this model with these assumptions, can account for individual differences that have been found on reasoning arguments.

Comprehension component

Experiments on the comprehension component are more diverse. They concentrate on the factors that may influence the interpretation of conditionals. Evidence has been found of the effects of linguistic variables on reasoning, such as the content of the propositions (e.g., Fillenbaum, 1975, 1976; Marcus and Rips, 1979; Rips and Marcus, 1977; Staudenmayer, 1975). The semantic relation that holds between the propositions, such as class inclusion or causality, also effects the frequency of different inferences (Marcus and Rips, 1979; Staudenmayer, 1975). The connective used, such as IF, or CAUSE is also an influential factor (Staudenmayer, 1975).

Differences in response patterns due to non-linguistic variables, such as factual or empirical knowledge, contextual presuppositions, and the pragmatics of conversation have also been found (Markovits, 1984, 1985; Rumain, Connell and Braine, 1983; Staudenmayer, 1975). A third category of influential variables that have been identified (Staudenmayer, 1975) are response biases, strategies, emotional involvement, social attitudes and the interpretation of the nature of the task per se.

Errors may be considered to indicate either features of the performance system (Cohen, 1981), or actual misinterpretations (Henle, 1978). A recent departure is to view them as systematic distortions (Grice, 1975). On Grice's view, taken up by Braine (Braine and Rumain, 1983; Rumain, Connell and Braine, 1983), comprehension processes create illusionary distortions around the underlying rule to which the premise information is to be mapped. Grice specified various conversational postulates that speakers and hearers abide by, such as giving each other all the relevant information for the current purposes of the exchange, and giving information that is clear, true and relevant to the topic of conversation.

Since people will thus assume that they have been given all the relevant information, some logically unnecessary inferences may be made. People may appear to make the inferences that follow from an equivalence rule (all four of the inferences outlined) rather than just the two that follow from an implicative rule. It is argued that people in fact, reason according to the correct implicative rule. That is, people consider only the two valid inferences to be <u>necessary</u> and the other fallacious inferences are merely "invited" (Geis and Zwicky, 1971) by the conversational comprehension processes. The fallacies thus arise through the logic of conversation (Grice, 1975), rather than through the logical properties of IF.

These invited inferences can be suppressed when people are prevented from their usual conversational understanding of the sentences. Explicit instructions suppress the fallacies (Rumain, Connell and Braine, 1983, Experiment 1) and implicit information about other conditions will also do so (Rumain, Connell and Braine, 1983, Experiment 2). Rumain, Connell and Braine (1983) argue that the lexical entry for IF is not defective, that is, it is not the equivalence interpretation. Instead, the inference rule marks the modus ponens inference, at least, as necessarily valid and marks the fallacies as invited. Markovits (1984, 1985) has further shown that implicit information about extra conditions creates variations in performance that vary systematically with a subject's initial, spontaneous performance on a reasoning argument. He attributes this to a competence system that allows us the ability to create hypothetical extensions.

Criticisms

Evaluation of evidence

There are some problems with the evidence proffered in support of the mental logic schemas. There is considerable doubt about what the difficulty ratings, which are the primary measure of the psychological reality of Braine's abstract schemas, actually measure (Gellatly, in press). The correlation between the difficulty ratings and the difficulty of an argument as predicted in the mental logic accounts, are sometimes low. Furthermore, the subjective difficulty of arguments is also consistent with many other models, not only with other schema models but also with entirely different non-rule based models (Gellatly, in press).

A major problem with using protocol data from subject's "thinking aloud" while solving reasoning arguments is in deciding what should be ignored, or attributed to momentary aberrations, for example one of Rip's (1983b) subjects makes a logically erroneous inference, yet this is not taken as indicating that the subject possesses an abstract schema corresponding to that erroneous inference. Conversely, there are difficulties in ascertaining what one should infer as implicit in a subject's reasoning, when explicit mention of, say a certain inference, is omitted from their protocol.

There are also difficulties in the investigations of content effects in reasoning. Differences were revealed through inferring the truth tables that people must have in order to make the responses that they do. This method imputes the variability to the subject rather than to the materials the subject is reasoning with. So it tries to find a single truth table that will account for the subject's responses to a variety of problems. The difficulty with this however, is that many subject's responses simply do not correspond consistently to any one pattern. Even when subjects do appear to conform to a particular truth table, this is no evidence that they are in fact using this truth table to reach their answer (Evans, 1982).

<u>Theoretical criticisms</u>

There are a great many intractable problems for the mental logic approach. A laudably high premium is placed on mental operations that cover as much ground as possible with the minimum cognitive overload (Rips, 1983a). This generality is desirable as it is plausible that cognitive mechanisms are parsimonious. However, the mental logic position approaches a level of elegance that is taut and overly restrictive. Hence, the application of the abstract schemas is severely limited, unless heavily buffered by considerable ad hoc, ancillary assumptions about a comprehension component, despite recent attempts to relate the abstract schemas to knowledge in memory (Sperber and Wilson, 1986). There is no good reason to suppose that economy is to be attained in only this way, that is, by positing inference mechanisms at the most abstract level possible. Nor is there any reason to suppose that economy will, in fact, be achieved in this way. Mental logicians have constructed systems corresponding to simple propositional connectives, and most suggest that it will be necessary to expand their systems to encompass more complicated inferences, such as causal relations, or relations between items whose properties must be taken into account. So the quantity of abstract schemas will proliferate.

A second problem, already mentioned, is that the enterprise is maintained only by keeping to a strict division between comprehension or interpretation and "pure" reasoning. That is, the operation of the reasoning mechanisms can account for human reasoning only through attributing all erroneous responses to a comprehension component.

is results in the necessity for two separate mechanisms to explain rors and correct responses. Correct responses are attributed to the ›eration of the abstract schemas. Incorrect responses are attributed › the comprehension component. The abstract schemas never lead a ubject into error, nor does the comprehension component ever ndependently lead a subject to make a correct response. This is :learly unparsimonious.

A third, and related, problem is the resultant emphasis on the comprehension component as a source of errors, such that it has become a catch-all for all vagaries. No account of the comprehension component has been provided. It has proved theoretically more tractable to tackle the rule-based schema component, in a computational manner. However, it is possible that when efforts are made to parse natural language into the abstract schemas, or instantiate the abstract schemas in some natural language system, the schemas will be found to be inadequate for accounting for all the inferences to be made. For example, the schema corresponding to material implication can produce adequate responses for abstract conditionals. However, because it is entirely truth-functional, it cannot provide adequate responses for the kinds of conditionals that are used daily, such as "if you come to the party, will you bring your wife ?" (Johnson-Laird, 1983). It is not the case that this conditional is understood in terms of the truth of each of its components. Nor is it the case that from knowing the antecedent holds ("you come to the party"), a person should infer the consequent ("will you bring your wife?"). In other words, as conditionals serve a variety of illocutionary purposes, their force can rarely be captured simply by their truth conditions, and the inferences that people make from the variety of conditionals cannot be captured by a material implication schema. It is also probable that causal relations (e.g., Downing, Sternberg and Ross, 1985; O'Brien, Costa and Overton, 1986; Schustack and Sternberg, 1981) will prove difficult to capture. The abstract or syntactic nature of the mental logic rules or schemas is ultimately incompatible with the findings of the effect of content on reasoning, which will be described in the next section.

While proponents of mental logic have always agreed that the influence of content and contextual factors should be taken into account (particularly by anyone attempting to refute the mental logic viewpoint), they do not attempt to specify how these factors are to be modelled in conjunction with their schema systems. The analogy is often drawn by defenders of rationality (e.g., Cohen, 1981) between the mental logic system and a linguistic grammar system: both are competence systems and both must be embodied in performance systems. However, firstly, the idea of a competence system of reasoning, analogous to linguistic competence is not a complete analogy. For example, (Kahneman, 1981) in speaking of linguistic competence it is not meant that a speaker who can produce correct grammatical sentences about hockey may not be able to produce grammatical sentences about, say, planets. In speaking of a mental logic competence system, the invocation of a competence - performance distinction is often meant precisely to account for such content differences: a reasoner who can produce correct logical inferences about hockey may not be able to produce correct logical inferences about planets. More damagingly, once one assumes that people have a competence system, there is no way of testing the model, as it will always be protected by the performance model. That is, once the competence system becomes embedded within a performance system, there is no way to invalidate the former. A final criticism of mental logic systems is that no reason is ever given as to why the mental logic rules or schemas must only be ones that are sanctioned by formal logical systems, such as the propositional calculus. The motivation to include only rules sanctioned by logic in the mental logic, is based on a priori considerations rather than on empirical observation.

Despite these many criticisms, the mental logic approach had the virtue of demonstrating that it is possible to theorise about reasoning processes in a precise manner. However, empirical research brought on the realisation that while to be "logical", as mentioned earlier, means to reach conclusions which a logical system would condone, there is no a priori reason to suppose that to reach logical conclusions one needs to use, as mental mechanisms, the rules that the logical system uses. Instead, it is entirely feasible that one can reason logically without mentally utilising a logical rule system. There are many other mental operations which will allow people to make logically correct inferences. It is to some proposals along these lines that we now turn.

Memory-based reasoning

Evidence concerning the effects of the material that a person is reasoning about on their performance casts serious doubt on the psychological validity of abstract mental logic inference rules or schemas. These findings, which are outlined below, led to the formulation of hypotheses that reasoning consists of the retrieval of particular kinds of information from long-term memory.

Experimental evidence

The evidence for the effect of content on reasoning comes from experiments using the selection task. This task involves the presentation of four cards. In its abstract guise, these cards display, for example, an A, a D, a 4 and a 7. Subjects are given a regulation, such as "if there is an A on one side of the card then there must be a 4 on the other side" and their task is to indicate those cards that must be turned over in order to test whether or not the regulation is true or false. The logically correct response to this task is to turn over the cards displaying the A (and so equivalent to the modus ponens inference) and the 7 (equivalent to the modus tollens inference). However, only about 10% of subjects succeed in making these card choices (see Evans, 1982; Pollard, 1982, for reviews of the abstract version of the selection task).

In contrast to this, it has been found that with a regulation, and cards, which have realistic referents, such as "if an envelope is sealed then it must have a 40 lira stamp on it", the frequency of logically correct responses increases dramatically (Bracewell and Hidi, 1974; Gilhooly and Falconer, 1974; Pollard, 1981; Johnson-Laird, Legrenzi and Sonino Legrenzi, 1972; Van Duyne, 1974, 1976; Wason and Shapiro, 1971). The realism of the materials does not, in itself, facilitate logical performance, however. This was discovered when a number of attempts to replicate these findings failed. These attempted replications involved not only the same materials as the original experiments, but also departed from these to use realistic but arbitrarily connected regulations, such as "if I eat haddock, then I drink gin" (Manketelow and Evans, 1979; Griggs and Cox, 1982; Reich and Ruth, 1982). It was thus suggested that the effects of realistic materials are related to the past experience that the subject has had with the referents of the regulation and cards (Johnson-Laird, 1983; Manktelow and Evans, 1979). Materials relating to subject's past experience reliably effect performance (Cox and Griggs, 1982; D'Andrade, in Rumelhart and Norman, 1981; Griggs and Cox, 1982, 1983; Pollard and Evans, 1980).

It was suggested that realistic materials do not simply affect the translation of a sentence to abstract code, but instead manipulate the availability of specific memories. These memories may be of counterexamples to the regulation, such as an envelope that is sealed but has only a 20 lira stamp on it. Remembering this instance will influence the subjects decision about which cards to turn over. Several such memory-based accounts have emerged.

Retrieval accounts

One of the earliest accounts is that people simply remember the correct answer to reasoning arguments (Manktelow and Evans, 1979). If the 'associative strength' between the antecedent and the negation of the consequent of a conditional is greater than the associative strength between the antecedent and the consequent of the conditional, then this will lead to the former being retrieved (Pollard and Evans, 1981). This view was superseded by an account in terms of the availability of a counterexample to the rule (Pollard, 1982). Tversky and Kahneman's (1973) original availability hypothesis held that there must be more counterexample instances than confirming instances in order for the former to be more available than the latter. Pollard's account modifies this to suggest instead that counterexample instances must be simply more available. This allows for the content and context of a problem to cue counterexamples from memory. So in some situations, even though there may be more confirming than disconfirming instances, the latter will be recalled or will be more available. This has been elaborated into a memory-cueing hypothesis. Griggs (1983, 1984, Griggs and Cox, 1982) has advanced the position that the content of a problem cues familiar relevant information in memory, including prominent counterexamples to a regulation. This direct cueing is considered to be a short-circuiting mechanism, having no rational substrate, but being able to lead to decisions that are optimal in real life situations (Griggs, 1983; Pollard, 1982).

These accounts have had to be extended to encompass the findings that the frequency of correct responses increases even when subjects are given realistic materials, for which they could have had no direct experience, such as "if a purchase exceeds 30 dollars then the receipt must be signed" (Cox and Griggs, 1982; D'Andrade, 1981; Griggs and Cox, 1983; Cheng and Holyoak, 1985, Cheng, Holyoak, Nisbett and Oliver, 1986). These increases have been found not only for materials that may be within the vicarious experience of subjects, but even for materials which are inconsistent with past experience. Consequently it has been suggested that memory cueing of counterexamples does not involve a simple rote retrieval of certain pieces of information, but instead requires some generalisation of that information to have also occurred (Manktelow, 1981). This has been extended to suggestions of a reasoning-by-analogy strategy, where analogous situations are utilised to derive counterexamples (Griggs, 1983).

Criticisms

Evaluation of evidence

Content undoubtedly effects reasoning performance in systematic ways. Much of these effects are too entrenched to be explained away with reference to comprehension difficulties. Instead, the evidence suggests strongly that information in memory has some direct effect on the reasoning process itself. However, it is not clear what the role of memory is, relative to the role of instructions to look for falsifying instances in testing the regulation. Pollard (1981) has emphasised the importance of role-playing in order to make counterexamples available in a particular context. Cheng, Holyoak, Nisbett and Oliver (1986) suggest that information as to what constitutes violations, as well as simply orienting people to look for such violations, is necessary. Griggs (1983) has suggested that reasoning from a regulation, as people do in the realistic tasks is different from reasoning about the truth or falsity of a regulation. On the other hand, there are claims (Griggs 1984; Chrostowski and Griggs, 1985) that memory by itself is both sufficient and necessary to improve performance, and while instructions to look for violating instances may affect facilitation, they are not necessary.

Secondly, the attempted replications of content effects are often messy, using different methodologies (Pollard, 1981) or different procedures (Bracewell and Hidi, 1974) from the original studies. Furthermore, although there is an intuitive appeal of the materials used there has been a lack of any systematicity in the manipulations of the properties of those materials. The most trenchant criticisms of this approach however, are theoretical.

Theoretical criticisms

The memory-cueing hypothesis remains ill-specified. The exact conditions under which memories are cued has received no treatment. Consequently, it has received no real tests. It has not been related to any specification of long-term memory organisation. The processes of retrieval, or specific retrieval strategies have not been outlined. The basic memory-cueing hypothesis cannot explain how people respond to reasoning arguments about topics they have never considered before, as people cannot simply have stored all of the conclusions to all of the arguments they can reason with (Rips, 1983a). Modifications of the theory in the light of further evidence, such as that people do not need to have direct personal experience of the referents of the regulation for their performance to improve, have also been ill-specified. For example, there are no descriptions of the processes of reasoning by analogy, such as analogical mapping or analogical retrieval (e.g., Keane, 1985, in press).

Recent studies have shown that counterexamples can be generated (Hoch and Tschirgi, 1983) rather than just retrieved. It has also been claimed that simple access to such counterexamples may not be enough, as it may be necessary to be aware of the pertinence of the retrieved counterexamples (Markovits, 1984). There is nothing in a retrieved counterexample that sanctions or outlaws a conclusion to an argument. So there must be some mechanism for evaluating the significance of retrieved counterexamples (Rips, 1983a). In other words, it is not enough to simply retrieve counterexamples, and the memory-based account is seriously deficient in this respect.

Two issues arising from this approach have been tackled in alternative approaches to reasoning. One centers on the allegation that reasoning and memory have become too intertwined in these studies, and their relative effects must be disentangled. This has led to advocations of the merits of focusing on reasoning with abstract materials. The remainder of this section includes an outline of the findings of reasoning performance with abstract materials. The second issue, which has been mentioned already, is that subject's performance has been found to be facilitated by materials for which they have no direct experience. The theory to be outlined in the next section attempts to account for this phenomenon.

Biases

It was considered that investigations of reasoning performance on abstract tasks would provide more direct measures of people's reasoning abilities, undistorted by the materials being reasoned about (Evans, 1983a). Reasoning performance with abstract materials is characterised by response patterns that have been attributed to "biases". Two main kinds of biases have been observed, one a pattern of responses referred to as a "negative conclusion bias" (Roberge, 1971, 1974, 1978; Evans 1972a 1977, Pollard and Evans, 1980). Here, subject's responses can be parsimoniously described as a tendency to endorse more negative conclusions than affirmative conclusions. This has been explained as the result of a 'caution heuristic', since it is more likely that negative conclusions (e.g., 'not A') will be right, rather than affirmative conclusions (e.g., A), because the former are less specific than the latter.

The second bias is termed a "matching bias". Subjects responses refer only to an item that has been explicitly mentioned in a regulation. This bias has been found on a variety of reasoning tasks (Evans, 1972b, Evans and Newstead, 1977; Evans and Lynch, 1973; Manktelow and Evans, 1979). However, it is found only on IF p, q regulations, not on Only If p, q regulations (Evans and Newstead, 1977, Evans, 1977) nor is it found on disjunctives (Van Duyne, 1974). This suggests that matching bias depends on the syntactic formulation of the regulation.

Two contrasting theories have arisen to account for these biases. One attributes them to attentional processes (Evans, 1983b). For example, it is claimed that matching bias arises because people selectively attend to the card items named in a regulation. For example, for the regulation "if A is on one side of the card then 4 is on the other", the cards "A" and "4" match the items in the regulation. People do not attend to the cards not mentioned in the regulation. Although "B" is the implicit negation of A (that is, it is an instance of "not A") and 7 is the implicit negation of 4, people are biased towards attending only to items explicitly mentioned in the regulation, so they do not choose these.

The second theory attributes "biases" to the availability or salience of the items in the stimulus presentation, or to the availability of the items themselves (Evans, 1982, 1984). That is, B does not seem to be about A at all. In support of this view, there is evidence (Evans, 1984) that when items are explicitly negated (for example "not A" as the negation of "A", rather than "B" as the negation of "A") matching bias disappears. Hence, the explanation for biases receives a treatment similar to the explanation for facilitation by realistic materials: some items are more available than others.

Criticisms. The identification of biases in performance in reasoning arguments is fraught with difficulties. One is that, although the descriptions of the response patterns is the most parsimonious possible, to label these descriptions as "biases" is of dubious value. "Biases" is a cognitively vacuous term which gives no explanation of the underlying processes responsible for the effects (Van Duyne, 1973). The second problem is that these biases have no ecological validity, as they are not found on realistic materials. Consequently, their import is questionable. Thirdly, the explanations in terms of attention or availability, proposed to account for these biases are ill-specified. Hence, it is difficult to distinguish competing explanations, and hard to test them.

Despite the criticisms offered here, both the memory-based reasoning approach, and the non-logical biases approach have served useful functions. They have redressed the balance after the mental logic theories emphasis on reasoning with recourse to abstract logical rules. The memory-based and non-logical biases accounts herald a move away from the idea that errors are attributable only to a comprehension component or interpretative stage prior to reasoning proper. They have led to serious consideration being given to the role of past experience and heuristics in reasoning.

Pragmatic schemas

An attempt at a compromise between the two positions, of a rule-based mental logic on the one hand, and memory-based retrieval on the other, has been made (Cheng and Holyoak, 1985; Cheng, Holyoak, Nisbett and Oliver, 1986). This derives from earlier suggestions concerning the role of domain-specific schemas in reasoning. In response to the necessity to explain how people can reason easily with realistic materials for which they have no direct experience, it was proposed that there are schemas which organise information in such a way that people can reason flexibly from them (Wason, 1983; Rumelhart and Norman, 1981). It was suggested that specific classes of events are stored in organised structures in long-term memory. The elicitation of such structures allows the elements within them to be manipulated (Wason, 1983). It was also proposed that schemas are composed of specialised bits of procedural knowledge, so reasoning knowledge is embedded in a relatively inaccessible procedural format. Understanding a problem and reasoning about it are thus almost the same thing (Rumelhart and Norman, 1981).

Taking this general approach several steps further, Cheng and Holyoak (1985) propose that there exist schemas whose core consists of generalised rules of inference. They argue that while specific experiences may play a role in reasoning, experience alone cannot explain the phenomenon of, for example, there being no differences in performance on familiar and slightly less familiar problems. The proposed schemas are intermediate in generality between memories of specific past experience and abstract rules. Understanding a problem on this approach consists of mapping statements onto a particular set of context-sensitive rules that are attached to a particular schema. These reasoning rules are based on our pragmatic interpretation of a situation. They are generalised sets of rules that are defined in relation to classes of goals. So a sentence like "if she's drinking beer then she must be over 18" is understood with reference to a generalised permission schema such as "if the action is to be taken then the precondition must be met". If the semantic aspects of a situation suggest that one is dealing with a permission type situation, then all the rules about permissions in general can be called upon. Purpose is a major cue in the evocation of schemas. Once evoked, the rules of the schema are all available.

The schemas are the products of induction from reasoning experience with classes of goal-related situations, such as permissions, obligations, covariations and causations. The schemas are pragmatic because they inculcate useful inferences rather than those that are sanctioned by propositional logic as being valid inferences. Instead, the rules are useful heuristics. Each schema's core is succinctly summarised by four conditional rules that are isomorphic to the valid (modus ponens and modus tollens) inferences and to the invalid (denial of the antecedent and affirmation of the consequent) inferences in the propositional calculus that were outlined earlier. However, each rule may or may not have a modal in it, that is, it may or may have a deontic or permission term, such as "must" or "may". It is claimed that the presence of modals in the rules will lead to some inferences being made and others not being made. For example, the presence in a schema of the rule "if the precondition is met then the action may be taken" ensures that people will not make the inference from "the precondition is met" to "the action is taken".

Logical errors are explained in terms of the ease of mapping concrete situations to pragmatic schemas, as well as by the degree to which the evoked schemas generate inferences that conform to standard logic. If the rules in a schema that correspond to the inferences deemed fallacious by propositional logic, have modals in them, then those inferences will not be sanctioned by the pragmatic schema either.

Cheng and Holyoak speculate that specific memories, abstract schemas, and pragmatic schemas may all be different ways in which people can reason. They claim however that pragmatic schemas are a high priority or optimal level of reasoning. So abstract regulations such as "if there is an A on one side of the card then there is a 4 on the other side" will be interpreted in terms of some pragmatic reasoning schema. Otherwise it will be understood by a person's knowledge of formal logic or by nonlogical strategies such as matching bias.

Experimental evidence

The best evidence provided for the pragmatic reasoning schemas approach is a demonstration that improved performance is found consistently on reasoning problems that can be categorised as permissions and as obligations, such as "if a person is drinking beer then the person must be over 18" (Cheng and Holyoak, 1985). This is attributed to the presence, in the permission pragmatic schema, of rules which ensure that the pragmatic schema sanctions the same inferences as does propositional logic (because of the presence of modals). The observation that performance on sentences interpreted as permissions is characterised by more logically correct responses than sentences interpreted by a causal schema (Cheng, Holyoak, Nisbett and Oliver, 1986) is also deemed supportive, as causal pragmatic schemas are said to sanction all four inferences.

Further supporting evidence comes from the finding that the provision of a rationale for seemingly arbitrary rules improves performance. This indicates that the rationale aids mapping to the correct pragmatic schema. Training on pragmatic reasoning schemas also improves performance. Some support is derived from the finding that training on both abstract rules and specific examples together improves performance, but training on either one alone does not. Furthermore, a whole logic course does not effect any improvement. So it is taken that these are not the optimal levels of reasoning. Finally, some ancillary support is derived for the nature of the permission pragmatic schemas outlined, from the systematicity of the translations from IF sentences to Only if sentences, in that the order of components and the introduction of modals in these translations is as predicted by the pragmatic reasoning schema theory.

Criticisms

Evaluation of evidence

While performance improves given a general statement, such as "if the action is to be taken then the precondition must be fulfilled", relative to performance on the abstract cards materials, this is not evidence that there is a schema corresponding to this general statement. If it was, than facilitation by a sentence, such as "if she is drinking beer then she must be over 18" would indicate that there is a specific schema corresponding to that sentence. By this kind of hypothesis, any materials that a person can reason with are indicative that a schema corresponding to those materials exists. This is both implausible and unparsimonious.

Secondly, as Cheng and Holyoak point out, the evidence, that the provision of a rationale improves performance, is confounded by the fact that the rationale is not content-free. The rationale may change the nature of the relevant experience. Even without the rationale, performance on the problems is high, which suggests that while a rationale may improve performance it is not necessary for good performance.

Theoretical criticisms

The main criticisms of this account centre on the arbitrariness of the level of generality that is assumed. It is assumed that the level of generality intermediate between abstract schemas and specific memories that is optimal for reasoning is the level of "actions" and "preconditions". But intermediate rules could reside at other levels of intermediacy, such as "desirable actions", or "desirable interpersonal actions". Cheng and Holyoak provide no motivation for the level of intermediacy that they choose. Their theory is unrelated to any specification of memory organisation, and so it is not possible to evaluate the general level of pragmatic schemas in relation to long-term memory structures. A subsequent problem is the generality of the application of a schema. Its intended application to all permission situations (and not just action-based ones) appears to conflict with the context-specificity of the rules attached to the schema (that is, only action-based ones). No clues as to the organisational relations between the schemas are given. For example, a permission and an obligation appear to differ only in the direction of the sentences denoting them. So a permission sentence may be "if the action is to be taken then the precondition must be met" while an obligation may be "if the precondition is met then the action must be taken". How their differences in meaning may be distinguished is not clear on a pragmatic schema account.

More damagingly there are difficulties in attempting to ascertain how people could distinguish which schema to map a sentence to. There is no specification of what constitutes an effective cue to a particular schema. So although the purpose of the sentence is intended to ensure that it is mapped to the correct pragmatic schema, it is not clear how the purpose is to be identified.

Furthermore, the role of modals is misleading. For a permission schema, for example, the rules include modals such as `if the action is to be taken then the precondition _must_ be met' and `if the precondition is met then the action _may_ be taken'. This modal is deontic, that is it refers to what is allowed and disallowed. It is claimed that it is the presence of these modals that makes people refrain from making certain inferences. However, the only way such modals would make people refrain from inferences is if they were epistemic, that is, modals referring to possibility and necessity. An epistemic modal would introduce the uncertainty necessary to ensure that people refrain from an inference. For example, from the regulation "if the precondition is fulfilled then the action may, that is, _might_ be taken", people will not infer that the action is taken, when they know that the precondition is fulfilled. A deontic modal will not ensure this. Finally, conditionals are typically organised in a condition - consequence sequence. That is, the antecedent contains the condition, and the consequent holds the consequence. In pragmatic schemas, the organisation is in the opposite direction. The significance of this chosen order is opaque.

Pragmatic reasoning schemas have nonetheless served to emphasise that while it is necessary to take account of people's world knowledge in reasoning, it is also essential to provide some mechanism for evaluating the significance of that world knowledge to the task at hand.

Reasoning procedures

Models

The final theory to be outlined provides a way to conceive both of the role of knowledge in reasoning, and of the operation of inferential mechanisms for evaluating the import of this knowledge. It heralds a move towards general reasoning procedures that can make use of more basic cognitive processes. Here we find an emphasis on the construction of representations of the premise information that draw on the knowledge the reasoner has in memory. This theory emphasises general procedures for combining the represented information into integrated models.

A precursor of this approach is the transitive-chain theory (Sternberg, 1985; Guyote and Sternberg, 1981). In this model, it is assumed that an internal representation is used in encoding premise information. This consists of pairs of informational components represented in some symbolic form. The structure of the symbolic representation allows the specification of a performance algorithm for combining the information to yield a new representation. The information is integrated via the formation of transitive chains that rearrange the components in the original representation. These transitive chains are formed by two simple inferential heuristics. The source of errors in reasoning on this early account is a failure to combine all of the encoded representations, because of limitations on the amount of processing capacity available. Preferences for working with simpler representations, (such as symmetrical relations and relations with no negatives in them), and a general tendency to use simplifying heuristics are also considered responsible for errors.

That people's understanding of, and reasoning with information can be modelled through general procedures that make use of their knowledge has been more fully explored in the mental models theory of reasoning (Johnson-Laird, 1983). In this account general procedures for constructing representations of the events described by the premises are used. A single representative sample from the set of models that satisfies an assertion is constructed by some procedures. Other procedures test whether there are alternative constructions or models that are also consistent with the premises but that lead to different conclusions. In attempting to construct such alternative models people abide by the principle that a conclusion is valid when there is no way for the premises to be true while the conclusion is false. This formulation provides the machinery for modelling both a person's understanding of the meaning of events and their ability to make reasonable inferences from what they have represented.

Mental models are based on the meanings of the premises and on implicit inferences from general knowledge (Johnson-Laird, 1983). After a basic propositional representation of the information, an optional stage of representation is available, of using this superficial representation as a partial basis for constructing a mental model. The model's structure is analogous to the state of affairs described. The procedural semantics used to relate language to mental models consists of several procedures which construct mental models on the basis of the meanings of expressions. These include procedures that construct a new mental model when an assertion makes no reference to entities in the current model. There is a procedure which, when one entity in an assertion is represented in the current model, adds the other entities, properties or relations in an appropriate way. There are also procedures to integrate two or more separate models if an assertion interrelates them, and to verify, when all the entities in an assertion are represented, that all the asserted properties or relations predicated of them are also represented. There is also a procedure to add a property or relation ascribed in an assertion to the model in an appropriate way.

Implicit and explicit inferences are distinguished by whether there is a deliberate search for alternative models of the premises. If a systematic and exhaustive search reveals no such alternatives then the conclusion is valid. However, limited working memory constrains people's performance, as they lack guidelines for systematic searches for counterexamples, and do not have definite principles for deriving a conclusion.

On this account, what underlies the meaning of a conditional is the ability to envisage states of affairs that may or may not correspond to reality : these may be possible states of affairs (ones that may or may not hold) or hypothetical states (ones that are possible or imaginary, or are assumed for the sake of argument) or they may be imaginary states (ones that are physically or logically impossible). IF is a verbal cue to consider such situations. One can conceive of a single uniform meaning of IF, out of which its logical vagaries arise, because of the role played by the propositions IF connects (Johnson-Laird, in press). The logical vagaries emerge from the background knowledge used in interpreting the meaning of the antecedent and the consequent, and in evaluating the relation between them. People use the beliefs provoked by their initial interpretation of the conditional to construct a mental model of a scenario in which the antecedent is realised, and they then interpret the consequent in terms of that model. Thus it is not simply the bald truth or falsity of the antecedent and consequent that determines how people will reason with a conditional, but rather their meaning as embodied in the representation of them.

The model constructed of the conditional will suggest a putative conclusion. This may be suggested by the operation of heuristics. Potential heuristics are those derived from the principles of conversational exchange, such as that conclusions should contain at least as much information as the premises on which they are based, and must be stated parsimoniously. The suggested conclusion will be tested by attempting to construct a model of the same state of affairs, but that is inconsistent with the conclusion. There are two recursive procedures that cope with constructing the right model. The first specifies that when the assertion is true in the current model, an attempt should be made to modify the model, in a way that is consistent with the previous assertions, but that will make the conclusion false. If such a modified model is not possible then the conclusion is a valid one. The second procedure specifies that when the assertion is false in the current model, then an attempt to modify the model should be made, to keep it consistent with all the previous assertions, but rendering the current assertion true as well. If this is not possible, then the assèrtion is inconsistent with the model.

Experimental evidence

The mental model theory is compatible with the evidence offered so far, in support of other positions. It can account for the evidence offered by the mental logicians, concerning the differential recall of suppositions in an argument, and also for the subjective ratings of the difficulties of arguments, in terms of the number and kind of models to be constructed. It can also encompass the effects of content and of context on reasoning, as these factors will affect the construction of initial models and especially, of alternative models. For example, the content-specificity of reasoning, indicated by the results of the selection task, are accounted for in terms of the ease of construction of alternative models. If one has information that can be easily used to modify a model constructed of the premise information, then those inferences considered fallacious in logic will also come out as fallacious by one's model modification. Furthermore, the theory of mental models also incorporates the findings of improved reasoning with realistic materials for which one has no direct personal experience, and for materials intermediate in generality, (such as "if the precondition is met then the action must be taken"), because models of these materials can still be constructed, and supplemented with general knowledge.

There is also evidence (Wason and Johnson-Laird, 1972, ch. 7) to indicate that IF is interpreted in a non-truth functional way. That is, the way people understand a conditional is not wholly dependent on the truth of it's component parts. So conditionals like "if you come to the party, bring your wife" are not understood by the truth of the constituents. As noted earlier, IF serves a variety of illocutionary functions and a theory of reasoning needs to incorporate the semantics and pragmatics of sentences, rather than simply the truth of their parts. Further evidence for mental models theory comes from the finding that the cognitive load of a task can lead to a failure to search for alternative models (Oakhill and Johnson-Laird, 1985; Wason and Johnson-Laird, 1972). It is also supportive that people apparently do not make trivial conclusions, (such as deducing from "p and q" that "p and p and q"). Their conclusions are usually informative ones.

This approach has the advantage that it accepts all that has been attributed to a comprehension component of reasoning in earlier research as being crucial to reasoning. It claims, however, that there is no need to posit the operation of some reasoning process after comprehension. One of the most interesting and innovative steps of this theory is the proposal that the procedure responsible for conditional reasoning is one which tests whether the conclusion follows invariably.

Criticisms

Mental models have been criticised for failing to be truly simulatory or dynamic (Rips, 1983a). That is, there has been some confusion over the actual representational claims of the theory. The nature of the representations has been criticised as being unclear (Rips, 1983a). It has also been claimed that evidence in support of mental models is not evidence against the mental logic abstract schemas (Sperber and Wilson, 1986) and that the two may co-exist. As has been demonstrated however, there are many arguments against abstract schemas. Arguments about the non-mutual exclusivity of theories are always possible but rarely prove to be useful or informative. The theory that general procedures, which call on basic cognitive processes, are responsible for reasoning can economically account for the data, as has been shown. Since the procedures consist of other processes, they are themselves an economical architecture, in the sense of avoiding repetition of structures.

Overview

The main thrust of this thesis is that it is necessary to model how people understand in order to model everyday reasoning. So it is necessary to know how people construct representations or build models, especially within the constraints imposed by the limitations of working memory. It is also important to ascertain how people use their past knowledge to modify the models they construct, and how they combine and connect information in a principled manner in order to examine various possibilities and outcomes.

The main empirical thrust is that to achieve this, we need to know how people reason in everyday situations. From knowing how people combine their knowledge with information from the situation, in their construction of representations, we will begin to understand everyday reasoning. The focus of this thesis is on how people reason with realistic materials. It will attempt to systematically evaluate the effect of <u>context</u> on reasoning in two ways. Firstly, the effect of the external situation or <u>environmental</u> <u>context</u> will be explored, in a series of experiments reported in chapters 2 and 3. Emerging from this study will come a model of the procedure responsible for testing representations or models by constructing other models of the premises (described in Chapter 4). The second way in which the effect of context on everyday reasoning will be explored is by examining the effect of what else a person is thinking about at the time, or the <u>cognitive context</u>. The second series of experiments (reported in Chapter 5) will focus on this. Arising from this investigation is a model of how people build a representation in the first place, by using heuristics born out of working-memory limitations (outlined in Chapter 6).

One of the advantages of studying the contexts in which people reason is that it can serve to further decide on the relative merits of the four theories outlined here. This will be discussed in Chapter 7. An implicit thrust of much of the work to follow is that it is important to understand everyday reasoning in terms of what is known about how people understand language, how they retrieve information, and how they build representations despite memory limitations. This will be returned to in Chapter 7.

CHAPTER 2

THE SUPPRESSION OF ALL INFERENCES

Introduction

In the next three experiments the effects of contextual information on the inferences subjects make will be explored. In making inferences in naturalistic settings, people's world knowledge plays an intrinsic role in the inferences they make. It will be argued that it is necessary to take account of such influences on inferential behaviour. It is suggested that subject's attempt to understand and reason with the impoverished materials usually provided in experimental investigations of reasoning, in the same way that they would understand and reason in naturalistic settings. This, arguably, ensures that their performance in experiments is as uncharacteristic of their natural performance as the experimental setting is uncharacteristic of a naturalistic setting.

It will also be shown that it is possible to experimentally manipulate the world knowledge that people take into account. In this way we can approach a more ecologically-valid study of reasoning. Attempts will also be made here to establish what processes mediate this everyday inferential behaviour. It will be argued that the ways in which people understand conversation, by assuming that the information they have been given is complete, and is true, and by attempting to relate it to information that they have in long term memory, underlies their inferential mechanisms.

Experiment 1

A useful route to understanding how people do something is to examine how they refrain from doing it. When people are given a realistic conditional sentence such as "if it is raining then she will get wet", they typically make all of the symmetrical inferences that can be made. That is,

(1) given: "it is raining", they make the modus ponens inference to "she will get wet',

(2) given: "she did not get wet", they make the modus tollens inference to "it is not raining",

(3) given the denial of the antecedent: "it is not raining", they conclude "she did not get wet" and

(4) given the affirmation of the consequent: "she will get wet", they conclude " its raining".

What makes people refrain from making these inferences ?

Conditional contextual information

Some recent research (e.g.,Rumain, Connell and Braine, 1983) has been concerned with how presenting conditionals in certain contexts can prevent people from making some of these inferences. For instance, one can accompany the first conditional premise "if it is raining then she will get wet", with the extra conditional information "if it is snowing then she will get wet".

Then when people are given

(1) the denial of the antecedent "it is not raining", they will no longer infer "she did not get wet", presumably because they realise

that she may get wet from the snow.

similarly, given

(2) the affirmation of the consequent "she got wet" people will not infer "it is raining", again because it may have been snowing.

So the alternative possibility of snow suppresses these inferences.

The two inferences suppressed are exactly those that are invalid by the rules of propositional logic, for material implication. The presence of snow as an alternative antecedent capable of making a person wet has no effect on the other two inferences. So given

(3) "it is raining" people will still make the modus ponens inference to "she got wet", as the possibility of snow is not pertinent to this inference.

similarly, given

(4) "she did not get wet" people still make the modus tollens inference to "it is not raining", because snow is also ruled out. These latter two inferences are the two that are valid in the propositional calculus.

The present experiment addresses the possible effects of contextual information on these valid inferences. On a mental model account of reasoning, these valid inferences should also be suppressable. Once the model of the premise information can be modified in a way that suggests the conclusion could be false while the premises are true, then the inference suggested by the initial model will no longer hold. What would modify the model of "if it is raining then she will get wet, it is raining" in such a way that people will no longer conclude "she will get wet" ? As we have seen, an alternative antecedent has no effect on them. What might have an effect is to accompany the initial conditional "if it is raining then she will get wet" with "if she goes out for a walk then she will get wet". It is plausible that, given

(1) "it is raining", people will no longer make the modus ponens conclusion "she will get wet", since they may realise that she could be indoors.

It also seems likely that given

(2) "she did not get wet" people will no longer reach the modus tollens conclusion "it is not raining", again because she may not be outdoors.

Since the <u>additional</u> antecedent of being outdoors may not hold, these inferences will be suppressed.

The potential undermining of the additional presupposition, that she is outdoors, should have no effect on the invalid inferences dealt with earlier. So given

(3) "it is not raining", people are still likely to conclude "she did not get wet", regardless of any information about being outdoors

and given

(4) "she got wet", it is plausible that they will still infer "it is raining", as the presupposition of being outdoors does seem to have been met.

The primary aim of the experiment reported here is to demonstrate that given contextual information, people will refrain from making the modus ponens and modus tollens inferences. Information about additional conditions or presuppositions like being outdoors, will suppress these valid inferences. This experiment will also attempt to replicate earlier findings that contextual information can lead to the suppression of the so-called fallacies. Information about alternative conditions like snow, has been found to suppress these inferences. So it will be shown that contextual information can suppress both the valid and the invalid inferences. In this experiment, some subjects will receive information about alternatives while others will receive information about additional conditions (see the conditions displayed in Appendix A, Table 1). It is hypothesised that this contextual information will differentially effect the inferences in the ways described above. Thus this experiment will show that all of the inferences can be suppressed by the appropriate contextual information.

Contextually extended minor premises

I have outlined how contextual information is expected to effect people's behaviour with conditionals (for example "if it is raining then she will get wet, if it is snowing then she will get wet") accompanied by a categorical premise (for example, "it is not raining").

How would people behave if they were given conditional premises accompanied by different minor premises ? Instead of giving subjects a simple categorical premise (like "it is not raining"), one could give a disjunction of the antecedents in the minor premise. So one could give subjects:
"If it is raining then she will get wet, if it is snowing then she will get wet, it is not raining or it is not snowing" (see Appendix A, Table 2).

Disjunction of alternatives

It seems most likely , even though the disjunctive premise is less informative than the categorical premise "it is not raining", that people will make the denial of the antecedent conclusion "she did not get wet" from this disjunctive minor premise. So it is expected, on the grounds of intuitive plausibility that this inference will no longer be suppressed.
But recall that, according to a logical analysis, they should still consider alternative possibilities. That is, people should take into account that there may be a further alternative to rain and snow, such as sleet. To be logically correct they should still refrain from

the denial of the antecedent inference to "she did not get wet". similarly, given the above conditionals and the affirmation of the consequent: "she got wet" would people consider "it is raining or it is snowing" to follow ? On a logical analysis they should not, as again there may be an alternative to rain and snow such as sleet, but intuitively it seems entirely plausible, and so it is expected, that this inference will be made.

Information about alternative conditions deters people from the fallacious forms of inference, as past research has already shown. People may refrain from the fallacies because of a more abstract logical insight or an awareness that there may always be alternative conditions to be taken into account. If this latter is the case then the mere reference to other alternative conditions should provoke some awareness that further alternatives may be pertinent. That is, if people refrain from making inferences in the categorical minor premise condition because they have attained some insight into the pertinence of alternatives, then that awareness should also occur even when the categorical premise is replaced by a disjunction. Hence, people should still refrain from making the fallacious inference to

(1) "she did not get wet" given the denial of the antecedent: "it is not raining or it is not snowing" as a minor premise.

similarly given

(2) the affirmation of the consequent: "she got wet", they should refrain from endorsing the fallacious conclusion "it is raining or it

is snowing"
if they are aware of the pertinence of alternatives in general. This is not what we expect on the grounds of intuitive plausibility however.

Of course, when given disjunctive minor premises of extra contextual alternatives like rain and snow, no effect on the modus ponens and modus tollens inferences would be expected. Given
(3) "if it is raining then she will get wet, if it is snowing then she will get wet" and the minor premise "it is raining or it is snowing" people would still be expected to make the inference to "she got wet".
Given
(4) "she did not get wet" people will also still be expected to endorse the modus tollens inference to "it is not raining or it is not snowing", even though this disjunction is again less informative than a categorical premise of "it is not raining".

<u>Conjunction of additional conditions</u>

In contrast, for the modus ponens and modus tollens inferences, a similar pattern of continued suppression is expected on a logical analysis, when the amalgamation of information in the minor premise is a <u>conjunction</u> of the additional conditions (see Appendix A, Table 3). Given
(1) "if it is raining then she will get wet" and "if she goes out for a walk then she will get wet" and the minor premise "it is raining

and she went out for a walk" it seems intuitively likely, and so it is predicted that people will still make the modus ponens inference that "she will get wet".

But should not they take account of other additional conditions? For example, the person may be wearing protective rain-gear in which case they still would not get wet. So again if people are aware that additional conditions are generally important then they will refrain from making the modus ponens inference to "she got wet" even when they are given the conjunction of additional conditions. It is not clear what would ensue for the conjunction of additional conditions for the modus tollens inference. So given
(2) "if it is raining then she will get wet, if she goes out for a walk then she will get wet, she did not get wet", it is not intuitively clear whether or not people will endorse the conclusion "it is not raining and she did not go for a walk".
This conclusion does not seem as intuitively plausible as the ones outlined so far, so no specific prediction is made.

For the denial of the antecedent and affirmation of the consequent inferences, the provision of additional conditions should not inhibit the inferences. So the conjunction of additional conditions should also presumably have no effect. That is, given
(3) the denial of the antecedent: "it is not raining and she does not go out for a walk", along with the conditional premises above, the inference to "she did not get wet" seems to follow just as much as

from "it is not raining" alone.

Also, given

(4) the affirmation of the consequent: "she got wet" the inference to "it is raining and she goes out for a walk" is as unhindered as the inference to "it is raining".

So to recap, the extra contextual information can be amalgamated in the minor premise. For alternatives like snow, the appropriate amalgamation is through disjunction (rain or snow). It seems plausible that given a disjunction of alternative conditions, people will no longer refrain from making the denial of the antecedent and affirmation of the consequent inferences. However, if one assumes that people gain some logical insight, then reference to alternative conditions may provoke such insight. If this latter is the case, then people should refrain from the denial of the antecedent and the affirmation of the consequent inferences even given disjunctive minor premises.

For an additional condition such as going out for a walk, the appropriate amalgamation is conjunction (raining and going out for a walk). The same situation to the disjunction of alternatives holds. It seems intuitively likely, and so it is predicted that people will make the modus ponens inference when given the conjunction of rain and out for a walk. But if there is a logical insight into the pertinence of presuppositions then the mention of additional conditions should ensure that people still refrain from the modus ponens inference. (The situation concerning modus tollens is less clear, and so no prediction is made).

The second aim of this experiment then, is to see how people perform when they are given the extra contextual conditional information, (either of alternatives or additional conditions) in an amalgamated form (either disjoined or conjoined) in the minor premise. So subjects receiving information about alternative conditions will not only receive this information along with a categorical premise, they will also receive it with the alternatives disjoined in the minor premise. Meanwhile, subjects receiving information about additional conditions will receive categorical premises first, and then they will receive these additional conditions conjoined in the minor premises.

Conjunction of alternatives and disjunction of additionals

It is also of peripheral interest to observe how people behave with the conjunction of alternatives (see Appendix A, Table 3) and the disjunction of additionals (see Appendix A, Table 2).
So subjects receiving alternatives will also receive these alternatives conjoined (e.g., if it is raining then she will get wet, if it is snowing then she will get wet, it is raining and it is snowing). Those receiving additional conditions will also receive these additional conditions disjoined in the minor premises (if it is raining then she will get wet, if she goes out for a walk then she will get wet, it is raining or she goes out for a walk). No specific

expectations exist as to how subjects will perform in these conditions.

Extended minor premises alone

I have described how people might be expected to behave with categorical premises and with conjunctive and disjunctive premises, when they accompany conditional premises. What would happen if people were given extra contextual information, but this time it was not presented explicitly in an extra conditional premise? (see Appendix A, Table 4). How would people perform if reference to the extra conditions was "made in passing" in the minor premise alone ?

Disjunction of alternatives

That is, given "if it is raining then she will get wet, it is not raining or it is not snowing" would people conclude "she did not get wet"? Or would the mere reference to an alternative to rain, such as snow, provoke some insight into the pertinence of alternatives in general, thus leading people to refrain from this denial of the antecedent inference? Presenting information in this way (without the extra conditional premise) enables one to address the same question asked in the previous section, but this time with a less explicit presentation of information.

Conjunction of additionals

The same question may be asked of information about additional conditions (see Appendix A, Table 5): given just "if it is raining then she will get wet, it is raining and she goes out for a walk" will people conclude that "she will get wet"? Or will the mention of an additional condition provoke the insight that these additions may not hold, and so lead to the suppression of this modus ponens inference ? As before, it seems more intuitively plausible that people will make these inferences and so this is predicted, despite the "logical insight" hypothesis to the contrary.
Again it is of peripheral interest to see how people behave with the conjunction of alternatives and the disjunction of additionals, when these are presented in the minor premises, without any extra conditional. And again no specific predictions are made for this.

To examine these possibilities in this experiment some subjects received no conditional contextual information at all. These subjects were given the single conditional premise on its own, accompanied by a categorical minor premise (e.g.,"if it is raining then she will get wet, it is raining"). These served as standard control conditions to the groups who received contextual conditional information (either about alternative or additional conditions), when these latter also received only a categorical minor premise (see Appendix A, Table 6). Furthermore, these control subjects also received minor premises with an extra antecedent (either an alternative or an additional condition) conjoined to the given antecedent (for example,"if it is raining then she will get wet, it is raining and it is snowing" or "if it is raining then she will get wet, it is raining and she goes out for a walk"). similarly, these subjects received disjoined alternatives ("if it is raining then she will get wet, it is raining or it is snowing") or disjoined additional conditions ("if it is raining then she will get wet, it is raining or she goes out for a walk"). These manipulations will be outlined in greater detail in the next section.

Method

Materials and Design

All subjects received items from three content domains drawn from mundane temporal events. One of these domains has been used in illustrating the purposes of the experiment, and concerns rain as a condition for getting wet. The other two concerned meeting one's family as a condition for going to a play, and finishing an essay as a condition for working late in a library (see Appendix B, Table 1). These materials were generated on the grounds of intuitive plausibility. An alternative antecedent was considered to be a candidate if it was capable of leading to the consequent by itself (for example "if it is snowing then she will get wet"), instead of the initial antecedent ("if it is raining.."). The alternative antecedent was always from the same general superordinate category as the initial condition (for example, snow and rain are both from the same category of "weather"). An additional antecedent was considered if it too could lead to the consequent by itself (for example "if she goes out for a walk then she will get wet"), but must also hold in conjunction with the initial antecedent ("if it is raining.."). The additional antecedent was chosen from a different superordinate category than the initial antecedent (for example, rain and going out for a walk are from different categories of, say "weather" and "exposure" respectively).

The presence of contextual conditional information was varied between subjects. Half of the subjects received extra contextual conditional information. So the first group of subjects received an extra alternative conditional, either snow, meeting friends or reading textbooks respectively. The second group received conditionals with additional conditions, such as going out for a walk, having enough money for a play, or the library remaining open respectively. The other two groups of subjects received only a single conditional premise with no extra contextual conditional information at all.

The presentation of minor premises was varied within subject groups. Each subject received a blocked presentation of the conditionals with each type of minor premise (see Appendix B, Table 2). Subjects first received conditional(s) with categorical premises, then with conjoined premises and then with disjoined premises. The order of the presentation of the types of minor premise was fixed across subjects. All subjects received the appropriate conditional(s) with a minor categorical premise to invite the four symmetrical inferences. So for the premise "if it is raining then she will get wet", each got the minor categorical premises of

(1) it is raining (to invite the modus ponens inference)

(2) it is not raining (to invite the denial of the antecedent inference),

(3) she will get wet (to invite the affirmation of the consequent inference) and

(4) she will not get wet (to invite the modus tollens inference).

(The examples used will continue to be in the "if it is raining then she will get wet" domain, but are intended to apply to all three of the content domains).

After these 12 items (four inferences for each of three content domains) were completed, all subjects received the appropriate conditional(s) with a conjoined minor premise. The first group of subjects received "if it is raining then she will get wet, if it is snowing then she will get wet", and, for example, "it is raining and it is snowing". The second group received "if it is raining then she will get wet, if she goes out for a walk then she will get wet", and, for example, "it is raining and she goes out for a walk". The third group received "if it is raining then she will get wet", and for example, "it is raining and it is snowing", and the fourth group received "if it is raining then she will get wet", and for example, "it is raining and she goes out for a walk". For a conditional like: "if it is raining then she will get wet" each subject received the minor premises of

(1) "it is raining and it is snowing" (or "it is raining and she goes out for a walk") to invite the modus ponens inference, and the conclusion choices of (a) "she will get wet" (b) "she will not get wet" (c) "she may or may not get wet".

They also received

(2) "it is not raining and it is not snowing" (or "it is not raining and she does not go out for a walk"), to invite the denial of the antecedent inference, with the same choice of conclusions as for (1).

They received the minor premise affirming the consequent

(3) "she will get wet", and the conclusion choices of (a) "it is raining and it is snowing" (or "it is raining and she goes out for a walk") (b) "it is not raining and it is not snowing" (or "it is not raining and she does not go out for a walk") (c) "it may or may not be raining and it may or may not be snowing" (or "it may or may not be raining and she may or may not go out for a walk").

They also received

(4) "she will not get wet", to invite the modus tollens inference, with the same conclusion choices as for (3).

After completing these 12 items subjects received a further 12 with the information disjoined in the minor premise. Inference type was also a within-subject factor, with each subject receiving premises to invite the modus ponens and tollens inferences and the denial of the antecedent and affirmation of the consequent fallacies (in random order) for each minor premise type.

For example, subjects in the condition receiving alternative conditions were given arguments with alternative antecedents, such as:

If it is raining then she will get wet
If it is snowing then she will get wet
It is raining
Therefore:
(a) she will get wet
(b) she will not get wet
(c) she may or may not get wet

They received 12 items like this, for the four inferences instantiated in the three content domains. Then they received another 12 items like these:

If it is raining then she will get wet
If it is snowing then she will get wet
It is raining and it is snowing
Therefore:
(a) she will get wet
(b) she will not get wet
(c) she may or may not get wet

And then they received a final 12 items of:

If it is raining then she will get wet
If it is snowing then she will get wet
It is raining or it is snowing
Therefore:
(a) she will get wet
(b) she will not get wet
(c) she may or may not get wet

A subject in the group receiving the explicit presentation of additional antecedents was given almost exactly the same materials, but that for every instance of an alternative ("it is snowing") they were given an additional condition ("she goes out for a walk"). A subject in the group receiving the implicit presentation of alternatives received almost the same materials as the explicit presentation of alternatives group but for the absence of the second conditional. So they received 12 items of:

If it is raining then she will get wet
It is raining
Therefore:
(a) she will get wet
(b) she will not get wet
(c) she may or may not get wet

for the four inferences instantiated in the three content domains. In this first part of their task they acted as a standard control group. They then received 12 items of:

If it is raining then she will get wet
It is raining and it is snowing
Therefore:
(a) she will get wet
(b) she will not get wet
(c) she may or may not get wet

And then they received a final 12 items of:

If it is raining then she will get wet
It is raining or it is snowing
Therefore:
(a) she will get wet
(b) she will not get wet
(c) she may or may not get wet

A subject in the additional conditions, implicit presentation group received virtually the same materials but that they were given additional conditions ("it is raining and she goes out for a walk") rather than alternatives.

To summarise, the kind of extra information given (either alternative or additional conditions) was a between-subject factor, that is, different groups of subjects received different kinds of antecedents. So too was the method of presentation of the extra contextual information (either explicitly presented as an extra contextual conditional, or more implicitly, mentioned only in the minor premises). The nature of the minor premise, either categorical, conjoined or disjoined was a within-subject factor, that is, all subjects received all three types. All subjects received each of these kinds of minor premise in the order above. The inference type, either modus ponens, modus tollens, denial of the antecedent or affirmation of the consequent was also a within-subject factor, so all subjects received all four types of inference. All subjects received each of the four inferences in random order in each of the three minor premise conditions, that is, for a total of 12 trials. They received these twelve trials instantiated in three content domains, summing to a total of 36 trials.

The subject's task was one which required them to select the conclusion from a choice of three conclusions. Given the major conditional premise(s) "if it is raining then she will get wet (if it is snowing/ she goes out for a walk) then she will get wet" and the minor premise "it is raining (it is raining and/or it is snowing/she goes out for a walk", they chose one conclusion from a choice of three

(a) she will get wet

(b) she will not get wet

(c) she may or may not get wet,

by ticking their choice. As well as the dependent measure of the choice of conclusion selected, there was an ancillary measure of response certainty. The ancillary measure was taken by accompanying each inference item with a "certainty scale". Subjects were instructed to circle a number on the certainty scale to indicate how certain they were that the conclusion they had chosen did follow from the premises. This was a simple line with the numbers one to five marked on it, where 1 was to indicate uncertainty and five was to indicate certainty.

Procedure

Subjects were tested in groups of between five and 10 people. They received the 36 items in a single booklet with specific instructions printed on the front page. These instructions explained the task with reference to an example. This example employed only a single major premise, a categorical minor premise, and a content not used in the experiment.

The example, which was used for all subjects in all conditions, was:

If it is sunny then she will go for a swim
it is sunny
Therefore:
(a) she will go for a swim
(b) she will not go for a swim
(c) she may or may not go for a swim.

Subjects were asked to assume that the premises were true and to "choose whichever of the conclusions - (a), (b) or (c) - you think follows from the sentences". They were also instructed to "circle, on the certainty scale, how certain you are that the conclusion you have chosen follows from the sentences". Finally, subjects were asked to read each item carefully and to work from beginning to end at their own pace, without changing any responses or skipping any items.

Subjects

The 32 subjects, from Trinity College, Dublin, who participated in this experiment were randomly assigned to the four groups (n = 8). None had received tuition in logic. Four subjects were replaced because of their failure to complete the items. These replacements were made during the course of the experiment, prior to any data analysis.

Results and Discussion

The data (both response choice and certainty judgements) were subjected to separate analyses of variance, each of which involved the factors of (1) type of information (either alternative or additional antecedents (2) presentation of information (either explicitly, in an extra conditional premise, or implicitly in the minor premises alone), (3) minor premise form (either categorical, conjunctive, or disjunctive, and (4) inference type (either modus ponens, tollens, denial of the antecedent or affirmation of the consequent). There were repeated measures on the last two factors.

Response choice data

The response choice data will be presented first (see the percentage of these scores in Appendix A, Table 7). Recall that subjects chose a conclusion from a choice of three conclusions.
The response choices of the appropriate invited inferences were subjected to analysis. Only 0.3% of responses overall were of the inappropriate inference (for example, from "if it is raining then she will get wet, it is raining" to "she will not get wet"). So the conclusion choices were primarily binary, between the appropriate invited inference, and the indeterminate choice ("she may or may not get wet"). Hence the responses to this latter are a mirror-image of the response choices of the appropriate invited inferences. Of the large variety of significant effects (see the results of the analysis in Appendix C, Table 1), the four-way interaction of information type, presentation type, minor premise form and inference type is the most informative [$F(4.81, 134.67) = 3.66$, $p = 0.0044$, using Greenhouse - Geisser adjusted degrees of freedom, as the variance - covariance matrice assumptions were not met]. This interaction signifies the variety of differences created by contextual information, each of which will be outlined in turn. Conservative post hoc tests (Tukey's Honestly Significant Difference (HSD), Winer, 1971) were conducted on this interaction to test for the various expected differences outlined earlier (see the expectations outlined in Appendix A, Table 6).

<u>Conditional contextual information</u>. First, the valid modus ponens and modus tollens inferences were suppressed by contextual information as predicted.

(1) When people were given extra additional conditions such as "if it is raining then she will get wet, if she goes out for a walk then she will get wet" and were told "it is raining", they made the modus ponens inference to "she will get wet" only 37.5% of the time (see Appendix A, Table 7, row 3, column 1) significantly less than when they have not been given any extra contextual information, just the premises "if it is raining then she will get wet, it is raining" (96%, $p < 0.01$, see row 4, column 1).

(2) Given the extra additional condition premise as above and the minor premise "she will not get wet" they made the modus tollens inference to "it is not raining" 33% of the time, significantly less than when given no extra information (92%, $p < 0.01$).

(3) Being given contextual information about additional conditions had no effect on the denial of the antecedent inference from "it is not raining" to "she will not get wet" (62.5% with extra additional conditions and 46% without).

(4) Nor did it have any significant effect on the affirmation of the consequent inference from "she will get wet" to "it is raining" (54% with additional condition information and 71% without).

Giving extra information about alternative conditions <u>did</u> effect these latter inferences however. This replicates earlier findings (Braine, Connell and Rumain, 1983; Markovits, 1984, 1985).

(1) When subjects are given "if it is raining then she will get wet, if it is snowing then she will get wet" and the minor premise "it is not raining", they made the denial of the antecedent inference to "she will not get wet" only 4% of the time (see row 1, column 3), significantly less often than when no extra alternative condition is given (46%, see row 2, column 3, $p < 0.05$).

(2) Given the above conditionals and the minor premise "she will get wet" subjects made the affirmation of the consequent inference to "it is raining" only 12.5% of the time, significantly less than when they were given no extra alternative (58%, $p < 0.05$).

(3) Contextual information about alternative conditions had no effect on the modus ponens inference from "it is raining" to "she will get wet" (96% with the alternative, 100% without).

(4) Nor did it have any effect on the modus tollens inference from "she will not get wet" to "it is not raining" (96% with an alternative, 75% without).

These data suggest that contextual information can suppress both the formally fallacious inferences and the formally valid inferences. (see summary in Appendix A, Table 6). Although the frequency with which the fallacies were endorsed in the control conditions is less than the frequency with which the valid inferences are endorsed, the magnitude of the suppression of both kinds of inferences is equivalent. Both the valid and invalid inferences are decreased by approximately 50%. So while the initial rates at which these inferences are made may differ as a function of the initial representation of the premise information, and as a function of the heuristics which suggest putative conclusions, they are nonetheless suppressed at the same rates.

Contextually extended minor premises. Secondly, when people were given amended minor premises, how do they treat the inferences ? The disjunction of alternatives will be treated first.

(1) When given "it is not raining or it is not snowing" as a minor premise, after having been given the two (initial and extra alternative) conditional premises, people made the denial of the antecedent inference to "she will not get wet" (62.5%, see row 9 column 3) significantly more than when given only "it is not raining" (4%, see row 1, column 3, $p < 0.01$).

(2) When given the affirmation of the consequent "she will get wet" people chose the conclusion "it is raining or snowing" (75%) significantly more than they choose "it is raining" (12.5%, $p < 0.01$).

So people did not behave as though they had some insight into the general principle that alternatives were relevant.

(3) Alternatives do not effect the modus ponens and modus tollens inferences, as we saw in the previous section. So it is not surprising that there were no differences in the frequency of modus ponens conclusions of "she will get wet" when people were given "it is raining or it is snowing" (100%) and when they were given just "it is raining" (96%).

(4) There were also no differences, when they were given "she will not get wet", in modus tollens choices of "it is not raining or it is

not snowing" (96%) and "it is not raining" (96%, see summary in Table 7).

What of the conjunction of additional conditions?

(1) People made the modus ponens inference to "she will get wet" (100%, row 5, column 1) from "it is raining and she goes for a walk" (after the two conditionals have been given) significantly more than from just "it is raining" (37.5%, row 3, column 1, $p < 0.01$).

So people do not seem to have insight that other presuppositions may be pertinent but appear to only take note of those that were mentioned.

(2) For the modus tollens inference from "she will not get wet" there is no difference in the frequency of conclusions to "it is not raining and she does not go out for a walk" (33%) and "it is not raining" (33%). This inference was the one noted earlier as having no intuitively clear response, so perhaps this is not surprising.

(3) Additional conditions have no effect on the denial of the antecedent and affirmation of the consequent inferences, as we saw in the previous section. So it remains unsurprising that there were no significant differences in the frequency of denial of the antecedent conclusions to "she will not get wet" from "it is not raining and she does not go for a walk" (92%), and from just "it is not raining" (62.5%).

(4) There were also no significant differences between the frequency of inferences from the affirmation of the consequent of "she will get wet" to "it is raining and she goes out for a walk" (71%) and to just

"it is raining" (54%, see summary in Appendix A, Table 6).

Recall that, as a peripheral concern, subjects were also given minor premises with the alternative conditions conjoined, and ones with the additional conditions disjoined. The disjunction of additional conditions will be dealt with first.

(1) Given "it is raining or she goes out for a walk" (along with the two appropriate conditionals) there is no significant difference in the frequency of modus ponens inferences to "she will get wet" (62.5%, row 11, column 1) than given just "it is raining" (37.5%, row 3, column 1).

(2) Given "she will not get wet" there is no significant difference in frequency of modus tollens inferences to "it is not raining or she does not go out for a walk" (67 %) than to "it is not raining" (33%). Additional conditions have no effect on the denial of the antecedent and the affirmation of the consequent inferences. So we can still be unsurprised that there were no differences between these inferences accompanied by a categorical minor premise and accompanied by a disjunction of the additional conditions. That is,

(3) there were no differences between the frequency of the denial of the antecedent inferences from "it is not raining or she does not go out for a walk" (83%) to "she will not get wet", than when given just "it is not raining" (62.5%).

(4) similarly, given "she will get wet", there were no differences in the frequency of affirmation of the consequent inferences to "it is raining or she goes out for a walk" (50%) and to "it is raining"

(54%).

So subjects did not treat the disjunction of additional conditions in the same way that they treated the conjunction of additional conditions. The disjunction of additional conditions leaves the modus ponens and tollens inferences suppressed (see summary in Appendix A, Table 6).

What of the conjunction of alternatives ? Here there <u>are</u> differences for some of the inferences.

(1) There were more denial of the antecedent inferences to "she will not get wet" given "it is not raining and it is not snowing" (79%) (see row 5, column 3), than given just "it is not raining" (04%) (see row 1, column 3, $p < 0.01$).

(2) There were also more affirmation of the consequent inferences from " she will get wet" to (the seemingly anomalous) "it is raining and it is snowing" (50%) than to just "it is raining" (12.5%, $p < 0.05$).

Alternatives have no effect on the modus ponens and modus tollens inferences as we have seen already, and there remain no differences between the frequency of these inferences from a simple categorical premise and from a minor premise involving the conjunction of alternatives.

(3) So given "it is raining and it is snowing" there is no difference in the frequency of modus ponens inferences to "she will get wet" (100%) than given just "it is raining" (96%)

(4) similarly, given "she will not get wet", there were no

differences in the frequency of modus tollens inferences to "it is not raining and it is not snowing" (87.5%) than to "it is not raining" (96%).

Even when alternatives are conjoined they are treated similarly to when they are disjoined, as if they exhaust all the alternatives people have considered (see summary in Appendix A, Table 6).

Extended minor premises alone. Finally, how do people cope when just given the extra alternative or additional condition in the minor premise, conjoined or disjoined with the initial condition, (without having first been given this extra condition in an extra conditional premise) ? The short answer is that they perform almost exactly as they have been observed to perform in the previous section, when they are given both the extra conditional and the conjunction or disjunction of the conditions in the minor premise.

The conjunction of additional conditions will be reported first. (1) Given "if it is raining then she will get wet, it is raining and she goes out for a walk" people made the modus ponens inference to "she will get wet" (100%, row 8, column 1) just as they do given "if it is raining then she will get wet, if she goes out for a walk then she will get wet, it is raining and she goes out for a walk" (100%, row 7, column 1). The response frequency to the two conditionals ("if it is raining then she will get wet, if she goes out for a walk then she will get wet") accompanied by a conjoined minor premise ("it is raining and she goes out for a walk"), and to the initial conditional ("if it is raining then she will get wet") accompanied by a conjoined minor premise ("it is raining and she goes out for a walk") is not significantly different, not just for the modus ponens inference,

(2) but for the modus tollens inference (67% versus 33%),

(3) the denial of the antecedent inference (83% versus 92%) and

(4) the affirmation of the consequent inference (79% versus 71%).

For the disjoined minor premises of alternatives, such as "if it is raining then she will get wet, it is not raining or it is not snowing",

(1) subjects made the denial of the antecedent inference to "she will not get wet" just as much (58%, row 10, column 3) as when given "if it is raining then she will get wet, if it is snowing then she will get wet, it is not raining or it is not snowing" (62.5%, row 9, column 3). This lack of significant differences in pattern again applied to

(2) the modus tollens inference (62.5% versus 96%) and

(3) the affirmation of the consequent inference (50% versus 75%) as well.

(4)There was however a significant difference in the modus ponens inferences. From "if it is raining then she will get wet, it is raining or it is snowing", the frequency of modus ponens inferences to "she will get wet" is only 37.5%, which is significantly less than from the premises "if it is raining then she will get wet, if it is snowing then she will get wet, it is raining or it is snowing" (100%,

$p < 0.01$).
One can speculate that this pattern is exactly the one that would be expected if the disjunction of these alternatives was interpreted exclusively in the former instance and inclusively in the latter, but such an isolated difference in a pattern of similarity is fundamentally uninteresting.

Recall that for completeness, the minor premises of conjoined alternative conditions and disjoined additional conditions were included. The pattern of responses to these inferences was again equivalent to the pattern for them when both conditional premises were presented (see Appendix A, Table 7, to compare the frequency of inferences in rows 5 and 6, and in rows 11 and 12).

Certainty scale data

The certainty scale data were subjected to a similar analysis of variance to that already outlined (Table 8, in Appendix A displays the percentages of certainty judgements, where 100% is equivalent to endorsing a certainty rating of 5 for each inference). These scores are of the certainty ratings made regardless of the response choice a subject made. They reflect how certain a subject was that the response they chose did follow, regardless of whether the response they chose was of the appropriate, inappropriate or indeterminate conclusion. It would not be meaningful to segregate certainty scores according to whether subjects made the appropriate inference or not, as in that case one could not differentiate between, say, a low score that indicates that a subject was uncertain in their choices of the appropriate invited inferences, and a low score indicating that a subject had made no choices of the appropriate invited inferences and so their certainty judgements were not used in the analysis. The effects from the analysis are presented in Appendix c, Table 2. The only significant effects in this analysis were interactions between the presentation of information and the type of minor premise [F (1.28, 35.72) = 5.18, p = 0.0217, by Greenhouse - Geisser degrees of freedom], and between minor premise type and inference type [F(4.46, 124.78) = 3.83, p = 0.0042]. There was also a main effect of inference type [F (2.15, 60.08) = 7.26, p = 0.0012].

Post hoc tests (Tukey's HSD) on the first of these interactions showed that subjects were more certain that the conclusion they had chosen followed from the premises, given premises like "if it is raining then she will get wet, it is raining" then given premises like "if it is raining then she will get wet, it is raining or it is snowing" (see Table 8, part ii). That is, subjects were more certain when using categorical minor premises (93%) than using disjunctive minor premises (87%), but only when no extra contextual conditional had been presented. This applies to all of the four inference types ($p < 0.05$). As there is less information in a disjunction than in an assertion (Johnson-Laird, 1983, ch. 2), this seems to indicate that subjects were more certain the more information they had. This also seems to account for the second interaction (see Table 8, part iii). Post hoc tests on this latter revealed that subjects were more certain that the modus ponens conclusions they had chosen followed from conjoined minor premises (97%) than from disjoined minor premises (90%), regardless of whether an extra contextual conditional has been given and regardless of whether the extra condition was an alternative or additional one ($p < 0.05$).

Conclusions

This experiment has shown that appropriate extra, contextual conditional information can lead people to refrain from making not only the fallacies but also the valid inferences. The appropriate contextual information to suppress the modus ponens and modus tollens inferences is additional conditions, while the appropriate contextual information to suppress the denial of the antecedent and affirmation of the consequent inferences is alternative conditions. The implications that this has for the logical validity of inferences and for those theoretical positions which maintain that people reason through the operation of mental rules of inference, will be discussed first. When it has been shown that these positions are severely undermined by these results, alternative positions will be entertained. Suggestions about how people are using the contextual information will then be discussed.

First, it will be argued that these findings suggest that no inference can be considered necessary in an absolute sense, but only valid relative to the context within which it is made. So in some contexts all four inferences are valid, in others the traditionally valid inferences (the modus ponens and modus tollens inferences) are invalid inferences, while in still other contexts, the traditionally invalid (denial of the antecedent and affirmation of the consequent) inferences are valid inferences.

It was shown that when the extra contextual information is presented in a conjoined or disjoined fashion, the suppression of the inferences no longer occurs. This latter holds true even when the extra contextual information is not presented in an extra conditional but is merely mentioned without prior introduction in the minor premise. These findings indicate that people are not grasping an abstract principle that other information is pertinent, such as underlies the motivation for the material implication interpretation of a conditional within propositional logic. Rather they seem to be adjusting their interpretation of the information to accommodate only the information that is given, no more and no less. So they do not become prompted by the presence of an alternative like snow, to realise that alternatives in general are important. The instance provided does not initiate a general appreciation that the possibility of alternatives is pertinent. Instead, they amend the initial premise "if it is raining then she will get wet" to "if it is raining or it is snowing then she will get wet", when they are given the extra conditional information "if it is snowing then she will get wet". From this amended premise the denial of the antecedent and affirmation of the consequent inferences do not follow.

We can see from this fact that the logical validity and invalidity of the four inferences rests on an interesting asymmetry. The denial of the antecedent and affirmation of the consequent inferences are invalid, according to the rules of the propositional calculus, and people should recognise that alternative conditions, like snow, might obtain. From a logical perspective, to understand "if it is raining then she will get wet" as a material implication rather than as an equivalence, where getting wet also entails that it is raining, people must reason from an amended premise. That is, they must respond on the basis of a premise like "if it is raining or it is snowing then she will get wet" rather than from the "if it is raining then she will get wet" that has been given. On the other hand, the modus ponens and modus tollens inferences are valid precisely because people are <u>not</u> to consider such amended state of affairs as "if it is raining and she goes out for a walk then she will get wet". Instead they are to <u>ignore</u> other conditions and are to reason from the given "if it is raining then she will get wet" premise.

The lobsidedness of this state of validity results in it being easier to get the modus ponens and modus tollens inferences logically correct because one need only stick to what is given, whereas to get the denial of the antecedent or the affirmation of the consequent inferences logically correct one must amend the given premise with other extra information.

What do these findings suggest about the correctness or otherwise of taking other information into account ? From exploring this, we will see that mental logicians are caught on the horns of a dilemma by the results of this experiment. On one hand, one can say that if people are expected to ignore contextual information in order to get the modus ponens and modus tollens inferences right, then a parsimonious mechanism of inference-making should ensure that they also ignore contextual information in making the denial of the antecedent and affirmation of the consequent inferences. On the other hand, if one must take contextual information into account to refrain from the denial of the antecedent and the affirmation of the consequent, then the same constraint of parsimony should ensure that contextual information is also taken into account for the modus ponens and modus tollens inferences.

So while it could be argued that the inference "she will get wet" from the amended premises "if it is raining and she goes out for a walk then she will get wet, it is raining" is not a modus ponens inference, one would then be committed also to the position that the inference of "she will not get wet" from the amended premises "if it is raining or it is snowing then she will get wet, it is not raining" is not a denial of the antecedent inference. Similarly, it cannot be argued that the premise "if she goes out for a walk then she will get wet" causes subjects to doubt the truth of the initial condition "if it is raining then she will get wet", without also arguing that the premise "if it is snowing then she will get wet" also causes such doubt. That is, any qualms about the logical status of a conditional modified by additional antecedents must be paralleled by qualms about the logical status of a conditional modified by alternative antecedents.

However, recall that the effects of contextual information about alternatives, such as snow, in suppressing the denial of the antecedent and affirmation of the consequent inferences, has been hailed, in earlier studies (Braine and Rumain, 1983; Markovits, 1984, 1985; Rumain, Connell and Braine, 1983) as indicating that people have the "logically correct" representation of material implication in their heads. In this representation, modus ponens and modus tollens are <u>necessary</u> inferences while the denial of the antecedent and affirmation of the consequent inferences are merely "invited" by conversational comprehension processes. But if contextual information can also suppress the modus ponens and modus tollens inferences, then by this same analysis, these latter inferences are also merely "invited" ones and are not necessary.

As there is no feasible way for the natural logic position to escape from the horns of this dilemma, these experimental findings prove embarrassing for any position that maintains that people reason by applying internally represented logical rules. So these findings suggest that it may be fruitful to attempt to explain reasoning by some other, and more context-sensitive, means. How do people reason? How do they take account of the information they have been given and, relating it to what they know, draw meaningful connections between the pieces of information ? The processes by which they achieve this will now be discussed, and investigated in the next two experiments.

CHAPTER 3

THE SOURCE OF THE SUPPRESSION EFFECT

Introduction

It has been found in the previous experiment that contextual information can lead to the suppression of all conditional inferences, both the valid ones and the invalid ones. In the initial discussion of the aims of the first experiment a description of the way in which this contextual information can affect inference-making, was offered. This concentrated on the way this information alters the status of the antecedent in the target conditional. However what remains to be established is <u>how</u> contextual information leads people to refrain from making inferences. What processes are brought to bear upon the contextual information in order that it may be used, and ultimately that it may suppress the inferences ?

Experiment 2

This question is most easily answered if one considers what people would have to do if they were <u>not</u> given the extra contextual information in order to perform in the same way as they do, when they are given this information. The findings of two rival research traditions will be used in attempting to derive an answer to this. Two distinct operations that people must perform if they are to spontaneously refrain from making inferences have been identified. As we will see, one of these is that people must become aware that other conditions are pertinent to the inference under consideration. The second operation is that people must have some previous experience of another condition that suppresses the invited inference, in order to refrain from making that inference. That is, people must be able to retrieve counterexamples to the inference in order to refrain from making it. These two operations will now be discussed in turn.

Relevance of other information

The more recent of the two traditions investigates the suppression of the denial of the antecedent and affirmation of the consequent inferences through the provision of extra contextual alternative conditions. This has been alluded to earlier (Markovits, 1984, 1985; O'Brien and Overton, 1980, 1982; Rumain, Connell and Braine, 1983). It has led to an emphasis on a general appreciation of the relevance of alternatives, as the necessary process in suppressing the denial of the antecedent and affirmation of the consequent inferences. A failure to appreciate the pertinence of alternatives is attributed to conversational comprehension processes (Braine and Rumain, 1983; Fillenbaum, 1975; Geis and Zwicky, 1971; Grice, 1975; Levinson, 1983; Markovits, 1984, 1985; Rumain, Connell and Braine, 1983). These involve the observance of default assumptions such as the "maxim of quantity", described by Grice (1975). This maxim proposes that speakers give their hearers as much information as they need for the current purposes of the exchange. By this maxim people assume that they need not go beyond the information that they have been given.

Hence, presenting extra contextual information indicates to people that the extra information is pertinent. The presentation of conditional contextual information, about alternative or additional conditions, embodies this property of indicating the relevance of extra information, by virtue of its presentation. Giving subjects "if it is raining then she will get wet, if it is snowing then she will get wet" indicates to them that "snowing" is relevant because it has been presented. Perhaps it is this property alone, of indicating the relevance of another condition, that is crucial? If this is the case then reference to actual instances is superfluous. However, as we have seen in the first experiment, the mere reference to alternative or additional conditions does not appear to initiate a general insight that such alternatives or additional conditions are relevant. Instead the effect appeared to be firmly tied to the specific instances given

In the next experiment an attempt will be made to establish whether the awareness of the relevance or pertinence of other conditions can, by itself, suppress the appropriate inference. That is, will indicating the relevance of other conditions, without reference to any specific instances of such other conditions, suppress the inferences ? If people are only told to bear in mind that there may be alternatives to rain that could also lead to getting wet, will they still infer that "she will not get wet" when they are given the denial of the antecedent, "it is not raining"? Similarly, if they are told that there may be additional conditions as well as rain that must also hold for a person to get wet, will they then continue to made the modus ponens inference from "it is raining" to "she will get wet" ? According to the mental logic position, it seems that indicating the relevance of other information should be enough to suppress the inferences. This possibility will be addressed in the next experiment.

Retrieval of specific counterexamples

In stark contrast to this emphasis on a general awareness of the relevance of other conditions, an older research tradition has identified the retrieval of specific counterexamples as the main operation capable of suppressing inferences. The findings supporting this view have come from a radically different paradigm to that used here (the selection task). In the selection task the end result of such counterexamples is not to suppress the inferences but to provoke explicit testing of them (see Evans, 1982; Griggs, 1983; Wason, 1983 for reviews).

Recall that in the selection task a subject is presented with four cards, on which may be printed the instances "drinking beer", "drinking coke", "over 19" and "under 19" (e.g., Griggs and Cox, 1983). The subject's task is to imagine that these instances represent four people sitting in a bar. They are asked to indicate which of these people need to be checked in order to see whether they are violating the rule: "If a person is drinking beer then s/he must be over 19". With content like this people usually have no trouble identifying the logically correct choices of checking (1) the person drinking beer and (2) the person under 19 (about 70% of correct decisions are made). This is an easy task compared to one where the cards refer only to letters and numbers (about 10% correct choices made). This superior performance in the former task has been attributed to the subject's prior experience with a counterexample to the rule, such as a person under 19 drinking beer.

Translated across paradigms to the syllogism tasks dealt with here, such a counterexample is similar to having an alternative consequence (such as "if a person is drinking beer then they must be over 19, if a person is drinking beer then they must pretend to be over 19"). Reasoning from an amended premise such as "if a person is drinking beer then they must be over 19 or must pretend to be over 19" allows people to refrain from making the inference from "drinking" to "over 19", (because the person may be only pretending to be over 19), so subjects <u>test</u> this inference instead of assuming that it is true. They also refrain from making the inference from "the person is under 19" to "the person is drinking coke", and test it as well, (again because of the alternative consequent).

Although this series of experiments will concentrate on alternative and additional conditions rather than on alternative or additional consequents, the principle derivable from the selection task research tradition still applies. That is, if people are able to retrieve a specific alternative or additional condition, then this mere retrieval will lead them to refrain from making the appropriate inference. This retrieval does not entail any general awareness of the relevance of alternatives or additional conditions. By this account, the presentation of an extra alternative or additional condition obviates the subject's need to retrieve such an instance. Perhaps it is this property of the conditional contextual information that suppresses the inferences, and any explicit indication of the relevance of alternative or additional conditions is superfluous ? In this case, incidental retrieval of alternatives or additional conditions should suppress the inferences, regardless of any appreciation of the relevance of such extra conditions.

In the next experiment an attempt will be made to establish whether the retrieval of counterexamples is enough to suppress the inferences. If subjects are given a task, such as to name several alternatives to rain that can also lead to a person getting wet; and they are subsequently given an inference task with the premise "if it is raining then she will get wet", the earlier incidental retrieval of alternatives may lead them to refrain from inferring "she will not get wet", from the denial of the antecedent "it is not raining". Similarly if subjects are asked to name things that must hold in addition to it being raining for a person to get wet, this retrieval may suppress their subsequent modus ponens inferences from the premises "it is raining" to "she will get wet". Note that this retrieval of instances will not have been directed towards making any inferences, but will have been in a separate task prior to the inference task. Furthermore, subjects will not have been informed of the connection between the tasks. According to the memory-based accounts of reasoning, this retrieval should be enough to suppress the inferences.

What if both the recognition of the relevance of other conditions and the retrieval of specific instances are vital processes in refraining from making inferences ? It is expected here that both of these are important in modifying the model one has constructed of the premise information. If this is the case, then neither of the experimental manipulations will have any effect on the frequency of inferences made. In this case inferences will be suppressed only when both of the processes are facilitated. This is the case when conditional contextual information is provided, as both properties are present in this manipulation. The next experiment will include not only conditions where the operation of one or other of these properties is facilitated. It will also include conditions where both are facilitated (through the provision of extra conditional contextual information). A further condition will consist of neither the facilitation of the recognition of relevance nor the retrieval of specific instances. This will use materials similar to those used in the control condition of the last experiment. For a summary of the experimental conditions, see Appendix A, Table 9. It is hypothesised that only the condition where both recognising the relevance of other information and retrieving that information is facilitated will lead to the suppression of the inferences since both of these are present in the condition where suppression has been observed.

Method

Instituto Superior de Psicologia Aplicada
BIBLIOTECA

Materials and Design

The same three content domains employed in Experiment 1 were used in this experiment, plus a fourth content domain (see Appendix B, Table 1). This latter concerned finding some leftover paint as a condition for redecorating a room, with the alternative condition of being able to spare some money; and the additional condition of being able to spare some time. Subjects received the four content domains in random order. Each subject received the four inferences (modus ponens and tollens, the affirmation of the consequent and denial of the antecedent) for each of the four content domains. So they received 16 (four contents by four inferences) items in a block. These inferences were presented in entirely random order for each content domain.

The type of conditional information presented was manipulated within subjects as well. Each subject received one block of 16 items with alternative condition information, and one block with additional condition information (see Appendix A, Table 9). The order of presentation of conditionals was counterbalanced, so that half of the subjects received the information relevant to alternatives first and additional conditions second, while the other half received the additional conditions first and the alternatives second. This allowed investigation of any effects of one kind of material upon the other. Altogether subjects received 32 items to complete.

The between-subject factors were the experimental manipulations of (1) indicating the relevance of other conditions and thus obviating the need for subjects to infer their relevance; and (2) ensuring the retrieval of instances of other conditions, thus obviating the subjects need to infer that they must retrieve them. One group of subjects received items like the two conditionals in Experiment 1, with the extra, conditional contextual information given. So these subjects received items like "if she finds some leftover paint then she will redecorate her room, if she can spare some money then she will redecorate her room" and it is expected that given the denial of the antecedent, for example, "she does not find some leftover paint" the inference to "she will not redecorate her room" will be suppressed. Given "if she finds some leftover paint then she will redecorate her room, if she can spare some time then she will redecorate her room" it is expected that, for example, the modus ponens inference from "she finds some leftover paint" to "she will redecorate her room" will be suppressed. In this "relevance and retrieval" condition it is assumed that both properties, of indicating the relevance of other conditions, and of obviating the need to retrieve instances of the other conditions are present.

Subjects in a second condition received items like those in the control groups of Experiment 1. That is, they received no extra contextual information, but only a basic item such as "if she finds some leftover paint then she will redecorate her room, she finds some leftover paint". This is assumed to have neither of the properties under investigation. For subjects in the remaining two groups, it was attempted to experimentally manipulate the presence and absence of the two properties, to ensure that only one or the other but not both were present.

In the first of these to be described, it was attempted to ensure that the relevance of other conditions was appreciated, without presenting specific instances of such conditions. Subjects in this third group were informed that they would be asked questions about a sentence such as "if it is raining then she will get wet". In the alternatives block they were informed that they should bear in mind that there may be alternatives to rain that can also lead to getting wet. They then received the four items, of modus ponens and tollens inferences, and the denial of the antecedent and affirmation of the consequent inferences, for this content domain. After this they received another content domain, with a similar reminder about the relevance of alternatives to the condition of that content domain, and then the four inference items were presented to be answered. In the additional conditions block exactly the same format was adopted, except that this time the reminder informed subjects that there may be other things in addition to the condition of that content that must also hold for the consequent to occur. So the subjects in this "relevance" condition received general instructions that other conditions may be relevant, but no specific instances were given (see Procedure for further details).

The second of the experimental manipulations attempted to ensure that specific instances are retrieved without indicating the relevance of those instances. Subjects in this fourth, and final group, the "retrieval" condition, were given retrieval tasks prior to each of the inferential tasks. So in the alternatives block, subjects received a premise such as "if it is raining then she will get wet" and were asked to list two things, other than rain, that could also lead to getting wet. They were then asked to answer the four inferential items of the modus ponens and tollens inferences and the denial of the antecedent and affirmation of the consequent inferences, for this content domain. After they had done this they were given a premise from another content domain. They were given the retrieval task for this content, followed by the inference task (see Appendix A, Table 9).

Each item was accompanied by a certainty scale, identical to those used in Experiment 1. The main dependent measure was the response choice of each subject, as measured by their choice of one of three conclusions: (a) she will redecorate her room, (b) she will not redecorate her room, (c) she may or may not redecorate her room. An ancillary measure was the certainty judgements.

Procedure

Subjects were tested in groups of three to five. The instructions were almost identical to those provided in Experiment 1. The general instructions were modified slightly for the two experimental conditions. The instructions for the relevance condition included a general reminder that "it is important that you bear in mind that they (the premises) may not give all the relevant information". For the retrieval condition the general instructions were also modified to include the prelude "You will be given items that involve several kinds of contents and you will be asked to give extra information about each content". These subjects were told that "this is to provide a check on differences between people in the knowledge they have about the contents of the items", so that they would not infer that the two parts of the task (the retrieval task and the inference items) were related.

For the relevance and the retrieval conditions there was a general indicator prior to each contextual information block establishing what kind of information was being referred to, either alternatives or additional conditions. So for the relevance condition, prior to the alternatives block, this read "In this section it is important to note that there may be relevant alternatives, that is, there may be other things instead of, for example, it being sunny that can also lead to going for a swim " (referring to the example used in the general instructions). Prior to the additional conditions block, this indicator read: "In this section it is important to note that there may be relevant additions, that is, there may be other things as well as, for example, it being sunny that must hold for going for a swim". For the retrieval condition, each block was preceded by similar indicators. Prior to the alternatives block, there was: "In this section you will be asked to name alternatives, that is, other things instead of, for example it being sunny that can also lead to going for a swim". The additional conditions block was preceded by a similar indicator referring to additions.

Finally, before each specific content domain block, subjects in the two experimental conditions received the more specific instructions about their tasks. That is, in the relevance condition, each content block of four inferences was preceded by the general reminder that other (either alternative or additional conditions as appropriate to the information block) conditions to, for example, paint could lead her to redecorate her room. Similarly, for the retrieval condition, each content block was preceded by a retrieval task.

Subjects

The 32 undergraduates from Trinity College, Dublin, who participated in this experiment, were randomly assigned to one of four groups (n = 8). Four subjects were replaced during the course of the experiment, prior to any data analysis. Two were replaced due to having received tuition in logic, and two through failing to finish the items in the booklet.

Results and Discussion

Response choice data

Presentation order. Before we can discuss the results of the experimental manipulations, it is necessary to consider whether or not the order in which subjects received the alternative and additional conditions information affected their inferential behaviour. Since the experiment used a block design, with some subjects receiving alternative information before additional information and other subjects receiving additional information before alternative information, the first thing to check is whether the order of presentation has an effect. If order has no effect then we expect that the invalid inferences will be suppressed by alternative conditions and the valid inferences will be suppressed by additional conditions, not only when these are presented in the first block, but also when they are presented in the second block (see summary of expectations in Appendix A, Table 9). To do this, the response choices (presented in Appendix A, Table 10) were subjected to a five-way analysis of variance. The response choices that were analysed were again the choice of the appropriate invited inference. Overall choices of the inappropriate conclusion was only 0.4%, and so the choices of the appropriate inference and of the indeterminate conclusion mirror each other. This involved the between-subjects variables of (a) relevance (b) retrieval and (c) order of presentation; and the within-subject factors of (d) information type and (e) inference type (see Appendix C, Table 3).

It was noted that there were some significant effects of order. That is, whether people received alternatives first or additional conditions first does seem to effect the frequency of inferences made. This is indicated by the interactions of order of presentation with the retrieval condition; its interaction with the retrieval condition relevance condition and the type of inferences; and its interaction with the relevance condition, the information type and the type of inferences. Consequently, post hoc tests (Tukey's HSD) were performed on the non-significant five-way interaction (see Winer, 1971), to establish exactly how the order of presentation confounds the results.

We may note firstly, that the findings of the first experiment regarding the suppression of the modus ponens and modus tollens inferences were replicated here, when the additional conditions information was presented prior to the alternative conditions information (see Appendix A, Table 10, rows 5 and 8, columns 1-4).

(1) So given "if she finds some leftover paint then she will redecorate her room, if she can spare some time then she will redecorate her room" and "she finds some paint", only 37.5% of modus ponens inferences to "she will redecorate her room" were made, significantly less than when just given "if she finds some leftover paint then she will redecorate her room" and "she finds some leftover paint" (100%, $p < 0.01$).

(2) Given the two conditionals above and "she will not redecorate her room", there were less modus tollens inferences to "she does not find some leftover paint" (37.5%) than from just the first conditional and "she will not redecorate her room" (75%, $p < 0.05$).

(3) No differences were observed for the denial of the antecedent "she does not find some leftover paint" to "she will not redecorate her room", given both conditionals (62.5%) and given just the first

conditional (69%).

(4) There were also no differences for the affirmation of the consequent inference from "she will redecorate her room" to "she finds some leftover paint", given the two conditionals (56%) and given just the first of them (69%).

The findings for the suppression of the denial of the antecedent and affirmation of the consequent inferences by extra alternative conditions also replicates the findings of Experiment 1 (see rows 1 and 4, columns 1 to 4, Appendix A, Table 10).

(1) There were significantly less inferences from the denial of the antecedent "she does not find some leftover paint" to "she will not redecorate her room" given the two conditionals "if she finds some leftover paint then she will redecorate her room, if she can spare some money then she will redecorate her room" (12.5%) than from just the first of these conditionals (56%, $p < 0.01$).

(2) For the affirmation of the consequent "she will redecorate her room" there were fewer inferences to "she finds some leftover paint" (12.5%) from the two conditionals above than from just the first conditional (44%). Although this difference does not reach significance by Tukey's HSD, it is significant by the less conservative Newman - Keuls test ($p < 0.05$).

(3) There were no differences for the modus ponens or tollens inferences given information about alternatives. So with both conditions and "she finds some leftover paint" there were the same frequency of modus ponens inferences as with only the first conditional (100% vs

94%).

(4) Similarly there is no difference in the modus tollens inference from two conditionals (94%) than from the first conditional only (75%).

So the pattern of primary findings from Experiment 1 were replicated in Experiment 2 when the extra information was presented <u>first</u>. That is, when people received the alternatives block first the invalid inferences were suppressed and when they received the additional conditions block first then the valid inferences were suppressed. What happens when the information is presented in the second block ?

The case with additional conditions will be dealt with first. For this kind of information the suppression of valid inferences no longer occurs (see rows 13 and 16, columns 1 to 4).

(1) The modus ponens inference from "she finds some leftover paint" to "she will redecorate her room" is made with the same frequency from the two conditional premises "if she finds some leftover paint then she will redecorate her room, if she can spare some time then she will redecorate her room" (87.5%) as from just the first conditional (87.5%).

(2) The modus tollens inference from "she will not redecorate her room" to " she does not find some leftover paint" is made with the same frequency given two conditionals (69%) as just given the first conditional (75%).

(3) For additional conditions information there were no differences for the denial of the antecedent (31% with two conditionals, 50% with one)

(4) nor any for the affirmation of the consequent inferences (37.5% with two conditionals, 31% with only one).

So preceding the additional information block by the alternative information block confounds the results seriously, with the result that the suppression of the modus ponens and modus tollens inferences no longer occurs. (see Appendix A, Table 9, for a summary of these results).

How does the suppression of the denial of the antecedent and affirmation of the consequent inferences by alternatives information fare, when that information is in the second block, after the additional conditions block ? (see Table 10, rows 9 and 12, columns 1 to 4).

(1) For the two conditionals "if she finds some leftover paint then she will redecorate her room, if she can spare some money then she will redecorate her room" and the denial of the antecedent "she does not find some leftover paint" the frequency of inferences to "she will not redecorate her room" was only 19%, significantly less than given only the first conditional (62.5%, $p < 0.01$).

(2) Given the affirmation of the consequent "she will redecorate her room", the frequency of the inferences to "she finds some leftover paint" was only 12.5%, significantly less than when only one condition was given (69%, $p < 0.01$).

So the denial of the antecedent and affirmation of the consequent inferences were still suppressed, even when the alternative block was presented after the additionals block. This is in contrast to what happens to the modus ponens and modus tollens inferences in this alternatives block. We expect no difference in the frequency of these inferences to ensue when given conditional alternatives information.

(3) Yet while there were no differences for the modus ponens inferences between being given the two conditionals above (81%) and being given only one (94%), there were differences in the frequency of modus tollens.

(4) Given the two conditionals above, ("if she finds some leftover paint then she will redecorate her room, if she can spare some money then she will redecorate her room") and "she will not redecorate her room" there were less modus tollens inferences to "she does not find some leftover paint" (31%) than given just the first conditional (75%, $p < 0.01$).

This is inexplicable unless one assumes some kind of interference from the preceding additional information block.

So overall, presenting the alternative and additional information in a within-subject design has led to some confounding interference. Having to consider alternatives after additional conditions, or vice ve rsa, has led subjects to behave in ways that they do not behave when they have to consider one or other type of information unpreceded by the other type. The most plausible reason for this is that subjects remember their response to the previous trials with the other kind of information, and this interferes with the response they make to the similar kind of information when it is encountered. To avoid confounding the results of the experimental manipulations of relevance and retrieval, the blocks presented <u>second</u> will be omitted from further consideration. In this way the data to be considered will be from those cells where subjects were dealing only with alternatives or with additional conditions, unpreceded by any other information. (The relevant response choices are now those in Table 10, rows 1 to 8 only).

<u>Relevance and Retrieval</u>. The success of the experimental manipulations of relevance and retrieval at individually suppressing the inferences, will be assessed relative to the revised four-way analysis of variance. This analysed the between-subject factors of relevance, retrieval and information type, and the within-subject factor of inference type. The effects from this analysis are presented in Appendix C, Table 4. Since there was a significant interaction between all four terms, post hoc tests (Duncan's multiple range test) were performed on this.

Can inferences be suppressed by facilitating the retrieval of specific instances of other conditions, or by indicating the relevance of other conditions alone ? The findings for facilitating the retrieval of other conditions will be presented first. Take the case where people were given a retrieval task, of naming alternatives to, for example finding leftover paint, that can also lead to redecorating a room (see Appendix A, Table 10, row 3).

(1) When they were subsequently given a conditional premise "if she finds some leftover paint then she will redecorate her room" and the denial of the antecedent "she does not find some leftover paint", they concluded "she will not redecorate her room" (69%), with the same frequency as they did from these premises when they were not preceded by a retrieval task (56%, row 4), and significantly more than they do when they have been given the extra conditional information (12.5%, row 1, $p < 0.01$).

(2) When given the retrieval task as above and the affirmation of the consequent "she will redecorate her room" people concluded "she finds the leftover paint" (56%) with the same frequency as when not given a retrieval task (44%), and significantly more then when given the two conditionals (12.5%, $p < 0.05$).

So getting people to name alternatives prior to making inferences did not suppress the denial of the antecedent and affirmation of the consequent inferences. No differences were expected for the modus ponens or tollens inferences and none were found.

(3) Preceding the modus ponens inference by a retrieval of alternatives task resulted in the same frequency of inferences (94%) with just a single conditional (94%) and with the two conditionals (100%).

(4) similarly, the retrieval of alternatives task prior to the modus tollens inference produced the same frequency of inferences (56%) as for just a single conditional (75%) and for both conditionals (94%).

How does a task requiring people to name <u>additional</u> conditions fare in suppressing the modus ponens and modus tollens inferences ? (see Appendix A, Table 10, row 7).

(1) Given the retrieval task, and the premises "if she finds some leftover paint then she will redecorate her room, she finds some leftover paint", subjects made the modus ponens inference to "she will redecorate her room" (87.5%) with a similar frequency to when not given the retrieval task (100%), and significantly more than when given the extra conditional information (37.5%, $p < 0.05$).

(2) Given the same retrieval task as above, and the modus tollens premise "she will not redecorate her room", people concluded "she does not find the leftover paint" with a frequency (44%) that is not significantly different from when they were not given the retrieval task (75%). This was not different from when given the two conditionals (37.5%).

So while the retrieval task prior to the inference task does not suppress the modus ponens inference, as indicated by the significant differences between the retrieval and the extra conditional conditions, it appears to suppress the modus tollens inference, as indicated by the lack of significant differences between the retrieval and the extra conditional conditions.

(3) The retrieval of additional conditions task was not expected to suppress the denial of the antecedent or the affirmation of the consequent inferences, and it did not. There were no differences between the denial of the antecedent inference after the retrieval

task (44%) and without a retrieval task (69%) or with the two conditional premises (62.5%).

(4) Nor was there any significant difference between the affirmation of the consequent inference after a retrieval task (37.5%) and without the retrieval task (69%), nor with the two conditionals (56%).

To summarise the success of the incidental retrieval of counterexample instances in suppressing inferences: retrieving alternatives does not succeed in suppressing the invalid inferences; retrieving additional conditions does not succeed in suppressing the valid modus ponens inference, but does appear to suppress the valid modus tollens inference (see the summary in Appendix A, Table 9).

How successful is facilitating a general appreciation or awareness of the relevance of alternatives or additions, in suppressing the inferences ? (see Appendix A, Table 10, row 6).

(1) When given the reminder "remember there may be things in addition to finding leftover paint that must also hold for a person to redecorate their room", and then given "if she finds some leftover paint then she will redecorate her room, she finds some leftover paint", subjects made the modus ponens inference to "she will redecorate her room" (81%) with the same frequency as when they had no general reminder, but just the single conditional premises (100%) and significantly more than when given two conditions (37.5%, $p < 0.05$).

(2) Similarly, given this general reminder, and the inferential premises, people made the modus tollens inference from "she will not redecorate her room" to "she does not find the leftover paint" (87.5%) with the same frequency as when they were not given any reminder (75%), and significantly more than when they have been given two conditionals (37.5%, $p < 0.05$).

(3) The reminder about additional conditions was not expected to suppress the denial of the antecedent or the affirmation of the consequent inferences. However given the additional condition reminder, the denial of the antecedent "she does not find some leftover paint" led to fewer conclusions of " she will not redecorate her room" (19%) than without any general reminders (69%) or than with the two conditionals (62.5%, $p < 0.05$).

(4) For the affirmation of the consequent "she will redecorate her room" there were also significantly fewer conclusions to "she finds some leftover paint" (19%) than without such a general reminder (69%, $p < 0.05$) and this was also almost significantly different from when the two conditionals were given (56%).

So instructions about the relevance of additional conditions failed to suppress the modus ponens and modus tollens inferences, but suppressed the denial of the antecedent and affirmation of the consequent inference instead. The most obvious and plausible explanation for this finding is that the term "additions" in the reminder was not specific enough to distinguish between additional conditions and alternative conditions, and that people were in fact treating the term additions as if it meant the same as alternatives. This ad hoc speculation is unfortunately circular however. The finding may be illuminated, by examining how people perform when they are given a reminder about alternatives being relevant (see Appendix A, Table 10, row 2).

(1) Given the reminder "there may be alternatives to finding leftover paint that could also lead to redecorating" and the premises "if she finds some leftover paint then she will redecorate her room, she does not find some leftover paint", subjects concluded "she does not redecorate her room" (69%) with the same frequency as when given these premises without a general reminder (56%) and significantly more than when given the two conditionals (12.5%, $p < 0.01$). So the general reminder did not succeed in suppressing the denial of the antecedent inference.

(2) Given the affirmation of the consequent "she will redecorate her room", subjects conclude "she finds some leftover paint" with the same frequency (44%) as when they were not given a reminder (44%) and apparently more often (although not significantly so) than when given the two conditionals (12.5%).

So it does not seem that the general reminder about alternatives had much success in suppressing the invalid inferences.

(3) The general reminder about alternatives was not expected to suppress the modus ponens and tollens inferences. This holds for the modus ponens inferences, where there was no significant difference in the frequency of inferences made after the reminder (62.5%) and

without the reminder (94%) and when given two conditionals (100%). (4) Although there was no difference, for the modus tollens inference, in the frequency of inferences when given the reminder, (44%) than without the reminder (75%), this was significantly less than when given two conditionals (94%, $p < 0.01$). So it seems that the modus tollens inference was unexpectedly suppressed by a reminder about alternatives.

This parallels the curious suppression of the denial of the antecedent and affirmation of the consequent inferences by information about additional conditions which we saw earlier. Clearly it is plausible to attribute this to obscurities in the terms (alternatives and additions) used, while acknowledging that this ad hoc attribution is somewhat unsatisfactory.

Certainty judgements

The certainty scale data were subjected to a similar analysis of variance as reported for the response choices. This analysis was also conducted only on the responses to the blocks of items presented first, to prevent any confounding of the results. The scores are presented in Appendix A, Table 11, and the effects from the analysis are in Appendix C, Table 5. The only significant effects in this analysis were (1) an interaction of the relevance and retrieval factors [$F\ (1, 24) = 6.30$, $p < 0.0192$], (2) an interaction of retrieval, information type and inferences [$F\ (1.97, 47.37) = 3.47$, $p = 0.397$, using Greenhouse-Geisser adjusted degrees of freedom as the variance-covariance matrice assumptions were not met]. The third interaction was of all four factors, relevance, retrieval, information type and inferences [$F(1.97, 47.37) = 3.68$, $p = 0.0331$, using Greenhouse-Geisser degrees of freedom].

Post hoc tests (Duncan's multiple range tests) were conducted on the latter. These reveal that when subjects were given the general reminder of the relevance of alternatives,

(1) they were less certain in making the modus ponens inference (66%) than when given the extra conditional information (80%, $p < 0.01$).

(2) The same applies to the modus tollens inference (67% versus 80 %, $p < 0.05$).

(3) There were no differences between subject's certainty about their conclusions of the denial of the antecedent (72% versus 73%) and

(4) affirmation of the consequent inferences (72% versus 69%) when they were given the alternatives reminder than when they were given the extra conditional information.

Given the reminder of the relevance of additional conditions

(1) subjects were less certain about making the denial of the antecedent inferences (59%) than when they were given the two conditionals (79%, $p < 0.01$), and

(2) they were less certain in the affirmation of the consequent inferences (56%) than when they were given the two conditionals (78%, $p < 0.01$).

(3) Also given only the additional conditions reminder, people were

less certain in the modus tollens inference (66%) than given the two conditional premises (80%, $p < 0.01$).
(4) There were no differences for the modus ponens inferences when subjects were given the additional conditions reminder.

By and large, this means that the general reminder made people less certain on the inferences that were to be made, that is, those inferences that were not expected to be suppressed (the modus ponens and modus tollens inferences by the alternatives reminder; and the denial of the antecedent and affirmation of the consequent inferences by the additional reminder).

Given the retrieval task, for the retrieval of alternatives:
(1) people were less certain of the modus ponens inferences (68%) than given the two conditionals (80%, $p < 0.01$).
(2) They were also less certain given the retrieval task prior to the modus tollens inference (67%) than when given the two conditionals (80%, $p < 0.05$).
(3) There were no differences in people's certainty about the denial of the antecedent (68% versus 73%) and
(4) affirmation of the consequent inferences (74% versus 69%).
This also seems to indicate that people were less certain in the conclusions that they made when the inferences were not suppressed.

There were no such differences for the incidental retrieval of additional conditions on any of the inferences (compare rows 5 and 7, columns 1 to 4, in Appendix A, Table 11).

To summarise the results of Experiment 2: The order of presentation of information about alternative and additional conditions can confound the effects that this extra information has on the inferences. When alternatives were presented first we saw again that they suppress

the invalid inferences. When additional conditions were presented first we also found again that this suppresses the valid inferences.

We can conclude that the incidental retrieval of counterexamples was not successful in suppressing any of the inferences. That is, retrieving alternatives does not succeed in suppressing the denial of the antecedent and affirmation of the consequent inferences, and retrieving additional conditions does not succeed in suppressing the modus ponens inferences although it does suppress the modus tollens inference.

However, reminding people of the relevance of other conditions gives unsystematic suppression of the inferences. Reminding of the relevance of alternatives does not suppress the denial of the antecedent or the affirmation of the consequent inferences, nor the modus ponens inference, but it does suppress the modus tollens inference. Reminding of the relevance of additional conditions does not suppress the modus ponens and modus tollens inferences but it does suppress the denial of the antecedent and affirmation of the consequent inferences.

It is difficult to conclude anything from these latter results about the role of the appreciation of the relevance of other conditions in suppressing the inferences. The results are inconclusive, notwithstanding the suggestion that it was the experimental manipulation that was unreliable, because subjects did not fully understand the terms (alternatives to, and in addition to) used in the experimental instructions. It is apparent however, that the incidental retrieval of counterexamples is not enough to lead to the suppression of any of the inferences (with the exception that retrieving additional conditions does suppress the modus tollens inference).

Conclusion

Experiment 2 has replicated the suppression effects of the first experiment, showing again that it is reliably possible to suppress all inferences. The finding that indicating the relevance of other information gives rise only to unsystematic suppression of inferences is important. It is especially significant in the light of the mental logician's claims that being aware of the relevance of other conditions should by itself lead to the suppression of the inferences. The experiment has also established that the retrieval of counterexample instances by itself is not enough to suppress the inferences. In the light of the memory-based accounts of reasoning, where it is claimed that the retrieval of such instances will suffice to suppress the inferences, it is useful to have established this.

The next experiment will continue to examine the respective roles of an appreciation of the relevance of alternative or additional conditions, and the retrieval of specific instances of these conditions. The procedure in the next experiment has been modified to avoid the pitfalls of the previous experiment. After this examination, it should be possible to suggest how recognising the relevance of other conditions, and retrieving specific counterexample instances operate, and what the relation between these two factors is. Once we have established how they operate, it will be possible to understand why they are not effective in isolation from each other. Firstly however, the next experiment will establish whether or not they can ever be effective in isolation.

Experiment 3

How can we establish, in a more reliable manner, whether or not the appreciation of the relevance of other conditions by itself or the retrieval of counterexample instances by itself, is the crucial property of the extra contextual information that suppresses the inferences ? It may be possible to manipulate the presence and absence of these properties in an ecologically-valid way by availing of the conversational conventions that people abide by when understanding utterances.

To elaborate further, given the information that, for example, a gunman said to a bank-clerk "if you shout then I will shoot you", people tend to make all of the inferences (Fillenbaum, 1975, 1976, 1978; Geis and Zwicky, 1971).

(1) They make the modus ponens inference from "the bank-clerk shouted" to "the gunman shot him" and

(2) the modus tollens inference from "the gunman did not shoot the bank-clerk" to "the bank-clerk had not shouted".

(3) They also make the logically invalid but highly reasonable inference from the denial of the antecedent "the bank-clerk did not shout" to "the gunman did not shoot him" and

(4) the fallacious inference from the affirmation of the consequent "the gunman shot the bank-clerk" to "the bank-clerk had shouted".

For a threat like this what would constitute an alternative condition to shouting that could also lead to getting shot ? What additional conditions must also hold as well as shouting in order to be shot ? It is not at all easy to see what these other conditions might be, given this threat in isolation from any context. A speaker asserting a condition under which s/he will carry out a promise or threat establishes that condition within the conversational exchange. The conditions for the consequent of an inducement are very often not pre-established ones. For example, snowing and raining may be considered pre-established physical conditions for getting wet, and the company of family or friends pre-established social conditions for going to a play. But for the threat "if you shout I will shoot you", "shouting" appears to be peculiar to the time and place of the utterance. Without having information about the kind of situation that the utterance is made in, and with no information about the duration against which the utterance is to be a understood, it is difficult to conceive of the nature of other conditions.

Retrieval and situational information

If it is known what kind of situation the utterance is made in, it seems plausible that it will be easier to think of what kind of alternatives or additional conditions are important. For example, imagine knowing that the utterance is made during a bank-siege. General knowledge about bank-sieges delimits the range of other possible conditions that must be considered. So it is hypothesised that situational information helps the retrieval of specific instances. An alternative to shouting that could also lead to getting shot, in a bank-siege might be trying to escape or pressing a panic button.

This suggests that given:

(1) 'During the siege, the gunman said to the bank-clerk "if you shout I will shoot you" and given also: "the bank-clerk did not shout", the denial of the antecedent inference to "the gunman did not shoot the bank-clerk" may be suppressed, because people may think about alternatives such as the bank-clerk trying to escape.

(2) similarly, given the affirmation of the consequent that "the gunman shot the bank-clerk", people may think that this could be either because the bank-clerk shouted or because s/he tried to escape and so the inference to "the bank-clerk had shouted" would be suppressed. Thinking of such alternative conditions would not suppress the modus ponens and modus tollens inferences, of course.

These valid inferences could be suppressed by other information. People could surmise that additional conditions must hold, as well as shouting, for the gunman to shoot the bank-clerk. For example, the gunman must be a murdering sort, or the pressure of the situation must be such that shooting the bank-clerk seems a viable thing to do.
(3) Given that "the bank-clerk shouted", people might think about such additional conditions as the siege going smoothly enough or the gunman already making his getaway at the time of the shouting, so that the modus ponens inference to "the gunman shot him" will not be made.
(4) Similarly they will refrain from making the modus tollens inference from "the gunman did not shoot the bank-clerk" to "the bank-clerk did not shout", if they have thought of such additional conditions as the gunman being safely on his way at the time, or not wishing to jeopardise his getaway. Obviously the denial of the antecedent and affirmation of the consequent inferences would remain unaffected by these additional conditions.

So it is hypothesised that information about the kind of situation during which an utterance is made delimits the range of conditions that are pertinent. This should help people in trying to retrieve specific alternatives or additional conditions instances. This situational information can be conveyed simply by prefacing the report of the utterance by the situational information (e.g., "during the siege..") It is assumed that situational information does not indicate to people that other information is relevant. Situational information only helps a person to focus on a particular range of alternative or additional conditions that are pertinent to the particular situation specified, and to retrieve these.

But what if the retrieval of specific instances is only effective when a person is already aware of the relevance of such instances, as seems to be strongly indicated by the findings of the previous experiment? In this case we need to provide people not only with information about the situation, to help them retrieve specific instances, but also with information that will help them appreciate the relevance of other conditions. What kind of information will do this ? Again we may use the conversational conventions that govern utterances, to decide.

<u>Relevance and durational information</u>

Hearers assume that they have been given enough information by their speaker "for the <u>current purposes</u> of the exchange" (Grice, 1975, p. 45, emphasis mine). So it seems likely that if people know that they must assess the utterance outside of the normal duration of the utterance, this may prod them into thinking that other conditions are relevant. It is hypothesised that <u>durational</u> information will help people appreciate the <u>relevance</u> of other conditions. Examples will clarify this.

If people are told that 'During the two week siege, the gunman said to the bank-clerk "if you shout then I will shoot you"' and are then given the denial of the antecedent "the bank-clerk did not shout", the long duration against which the utterance is to be assessed will make people think that alternatives, such as trying to escape, are relevant. So the intuitively plausible hypothesis is that information that the utterance occurred during a long duration will help people think of alternatives. The length of the situation will suggest to people all of the other things a person may do during a long period of time, in such a situation. As they have also been given information about the situation, that is, that it is a siege, they should be able to retrieve specific instances of alternatives (like escaping). Consequently:

(1) the inference to "the gunman did not shoot him" will be suppressed.

(2) So too will the inference from the affirmation of the consequent "the gunman shot the bank-clerk" to "the bank-clerk shouted".

To summarise, giving people information that the utterance is to be assessed relative to a long duration (such as two weeks) and giving them information about the kind of situation that the utterance is made in (e.g., a siege) should both indicate that alternative conditions are relevant, and should facilitate the retrieval of instances of such alternatives (see Appendix A, Table 12). Information indicating the relevance of alternatives, and helping the retrieval of an alternative, such as the long duration and situational information described, is not expected to have any effect on the modus ponens and modus tollens inferences however. It is assumed that long durational information does not make people think of additional conditions, but only of alternative conditions.

What does effect the modus ponens and modus tollens inferences, as we have seen repeatedly, is additional conditions. How can we indicate the relevance of additional conditions within this scheme? It seems that this can also be achieved through indicating the length of the duration against which the utterance is to be assessed, but this time, that it is of a very <u>short</u> duration. An uncharacteristically short duration may suggest that additional conditions are relevant, that is that the presuppositions of the utterance may be undermined. Because the situation is shorter than usual, the typical assumptions may not hold, and so this will call these presuppositions into questions. The intuitively plausible hypothesis is that <u>short</u> durational information will help people appreciate the relevance of <u>additional</u> conditions. Again, examples will clarify this.

Imagine being given information about a short duration such as 15 minutes, as well as information that the utterance is made during a siege (that is, given 'during the 15 minute siege...'). This should indicate that additional conditions are relevant and should also facilitate the retrieval of specific instances of such additional conditions. An additional condition might be that the duration is so short that the siege must have gone smoothly, so the gunman would not endanger his getaway by shooting anyone. Hence,

(1) given 'during the 15 minute siege, the gunman said to the bank-

clerk "if you shout then I will shoot you"' and given "the bank-clerk shouted", the modus ponens inference to "the gunman shot the bank-clerk" should be suppressed.

(2) Similarly, given "the gunman did not shoot the bank-clerk", the modus tollens inference to "the bank-clerk did not shout" should also be suppressed.

Needless to say, we would not expect this indication of the relevance of additional conditions, and aid to retrieving specific instances of such additional conditions to have any effect on the denial of the antecedent or affirmation of the consequent inferences. It is assumed that short durational information will not indicate the relevance of alternatives.

To recap, giving people information about the duration for which an utterance is to be assessed should indicate the relevance of other conditions. Information that the duration is a long one, such as two weeks, should indicate the pertinence of alternatives. It will not indicate the relevance of additional conditions, because it is hypothesised that long durational information ensures people think of alternatives as being relevant. Indicating that the duration is a short one, such as 15 minutes, should indicate the pertinence of additional conditions. Again it is assumed that short duration information will not indicate the relevance of alternatives. Given information about the situation in which the utterance is made, such as that it is during a siege, should delimit the range of other conditions for the consequence of the utterance and so it should help the retrieval of specific instances of alternatives or additional conditions. Information about the duration and the situation together should then suppress the inferences, because the relevance of alternatives or additional conditions is indicated, and specific instances of them can be retrieved.

What if only one of these, indicating the relevance of other conditions or facilitating the retrieval of specific instances is necessary, and the other is superfluous ? If the retrieval of counterexamples is all that is necessary to suppress the inferences then information about the situation alone should be enough to have a suppression effect. Recall however, that in the previous experiment, the retrieval of instances does not seem to be enough in itself to suppress the inferences. So it is expected here that retrieval, that is, situational information, will not suppress the inferences.

Of more interest in the next experiment is whether or not indicating the relevance of other conditions is all that is needed to suppress the inferences. If it is, then information about the duration should be all that is necessary to create a suppression effect. Information about a <u>long</u> duration should suppress the invalid denial of the antecedent and affirmation of the consequent inferences, as it should prompt awareness of the relevance of alternatives. Information about a <u>short</u> duration should suppress the valid modus ponens and tollens inferences, as it should prompt awareness of the relevance of additional conditions (see Appendix A, Table 12). It is expected that recognising the relevance of other information, that is, durational information, will not be enough to suppress the inferences. So it is expected that long durational information by itself will not suppress the invalid inferences, and that short durational information by itself will not suppress the valid inferences.

The next experiment addresses these issues by comparing the frequency of inferences made in several conditions. One group of subjects were given just the basic information (e.g., the gunman said to the bank-clerk "if you shout, I will shoot you"). It is predicted that the inferences will not be suppressed in this condition. The other groups of subjects received this basic information prefaced by some extra, contextual information. In two of these groups, subjects received information to foster both an awareness of the relevance of other conditions and the retrieval of specific instances together and so it is expected that this will suppress the inferences. They received information about both the duration and the situation. One group of these subjects received long duration and situation information ('during the two week siege the gunman said to the bank-clerk "if you shout I will shout you"), and it is expected that the invalid inferences will be suppressed by this; the other receive short duration and situation information ('during the 15 minute siege..'), and it is expected that the valid inferences will be suppressed by this.

In the remaining three groups, subjects received information to foster either the awareness of the relevance of other conditions, or the retrieval of specific instances, but not both, so it is predicted that the inferences will not be suppressed for these groups. So one group of subjects received only information about the situation ('during the siege..'); another received only information about a long duration ('during the next two weeks..'), and the last receive only information about a short duration ('during the next 15 minutes...').

Method

Materials and design

Subjects were given items reporting promises, threats and bribes made by one actor to another. These were drawn from four everyday content domains (see Appendix A, Table 3). Subjects received the four inferences (modus ponens and tollens and the denial of the antecedent and affirmation of the consequent inferences) for each of these four content domains. Each subject received 16 items to complete.

The presentation of situational and durational information was varied between subjects. Three of the groups of subjects received information about the situation while the other three groups did not. Of these sets of three groups, one group received no information about the duration, another received information that it was a long duration, while a third received information that it was a short duration. So there were six conditions.

Procedure

The instructions were again almost identical to those of the first and second experiment, both in presentation and in content. The example used was an inducement. It was of a different content from those used in the experiment, and contained no information about the situation or the duration. It read:

> The teacher said to the pupil:
> if you misbehave, I will give you extra homework.
> The pupil misbehaved.
> Did the teacher give the pupil extra homework ?
> (a) yes
> (b) no
> (c) maybe
> (d) can't say.

Note that the response choice format was slightly different from the previous experiments. This seemed more in keeping with the materials. The dependent measure was response choices. Again each item was accompanied by a certainty scale, and certainty judgements were used as an ancillary dependent measure.

Subjects

The 48 undergraduates, of Trinity College, Dublin, who participated in this experiment were randomly assigned to one of six groups (n = 8). Two were replaced because they had studied logic. These replacements were made prior to any data analysis, during the experiment.

Results and Discussion

The response choice data and certainty judgements were subjected to separate analyses of variance. The between-subjects factors were situation information (either given or not) and durational information (either not given, long or short duration information). The within subject factor was inference type. The response choice data will be discussed first.

Response choice data

The response choice frequencies are presented in Appendix A, Table 13 and the significant effects from the analysis are in Appendix C, Table 6. Once again, the appropriate invited conclusions were analysed. Since choices of the inappropriate inferences were only 2% overall, the appropriate and indeterminate ("can't" and "may be") responses are mirror images of each other. There was a significant three-way interaction between the situational, durational and inference factors [F (6, 126) = 2.22, $p < 0.0454$]. Post hoc tests (Duncan's multiple range tests) were performed on this.

Does provision of information about the situation and that it is of a long duration, (which we have surmised indicates the relevance of alternative conditions and helps the retrieval of specific instances of alternatives), suppress the denial of the antecedent and affirmation of the consequent inferences ? (see Appendix A, Table 13, row 3)

(1) Given 'during the two week siege, the gunman said to the bank-clerk "if you shout, I will shoot you", the bank-clerk did not shout' people agreed that the denial of the antecedent inference to "the gunman did not shoot him" follows with a frequency of 19%. This was significantly less than when given no extra information, just: 'the gunman said to the bank-clerk "if you shout I will shoot you", the bank-clerk did not shout' (47%, row 1, $p < 0.05$). This was also significantly less than when they were situational information only, in: 'during the siege the gunman said to the bank-clerk "if you shout I will shoot you", the bank-clerk did not shout' (44%, $p < 0.05$, row 2). It was significantly less than when they were just given long durational information in: 'during the next two weeks, the gunman said to the bank-clerk "if you shout I'll shoot you", the bank-clerk did not shout' (66%, $p < 0.01$, row 4, column 3). Thus for the denial of the antecedent, information about the situation and a long duration suppressed the inference as expected. Furthermore, neither the long durational information nor the situation information was enough by itself to suppress the inferences.

(2) What of the affirmation of the consequent inference from "the gunman shot the bank-clerk" to "the bank-clerk shouted" ? This conclusion was agreed with at the same rate when it was prefaced with 'during the two week siege..' (31%) as when it was not prefaced with any extra information (37.5%). This was also no different from the frequency when it was prefaced by 'during the siege..' (28%), although it was significantly less than when prefaced by 'during the

next two weeks..' (59%, $p < 0.05$). So for the affirmation of the consequent inference durational and situational information together did not suppress the inference. Note however that the initial frequency of these inferences was low: the frequency of inferences when prefaced by both kinds of information, by situational information and by no information at all was low. It was only prefacing by durational information alone that failed to suppress the inference.

The modus ponens and tollens inferences were not expected to be effected by alternatives, and this expectation holds.

(3) The frequency of the modus ponens inference was the same without any information prefacing (56%) as with both kinds of information (56%), as with just situational information (72%) and as with just durational information (50%).

(4) However, the frequency of modus tollens inferences did differ when no extra information was given (34%) than when both kinds were given (59%, $p < 0.05$), when just situational information was given (66%, $p < 0.05$) but not when just durational information was given (50%).

To summarise so far: information about the situation and a long duration succeeds in suppressing the denial of the antecedent inference. Situation and long duration information on their own do not. Information about the situation and long duration may suppress the affirmation of the consequent inference, but these inferences seem to be equally suppressed with none of this information. It was only durational information that fails to suppress this inference. The

modus ponens and tollens inferences were unaffected by long duration and situation information, either jointly or separately, as we
However, modus tollens were made infrequently when no information was given.
expect. So these results indicate that situational and long durational information can suppress the invalid inferences (notwithstanding global difficulties with the affirmation of the consequent inference) and does not effect the valid inferences (see Appendix A, Table 12). This is, by and large, as was predicted.

What is the effect of presenting information about the situation extended over a short duration ? It was surmised that this would indicate the relevance of additional conditions and allow their retrieval, and consequently, would suppress the valid modus ponens and tollens inferences.

(1) The frequency of modus ponens inferences made when prefaced by 'during the 15 minute siege..' (41%, row 5) was not significantly different from the frequency when they were not prefaced (56%, row 1). This was significantly different from the frequency when prefaced only by situational information (72%, $p < 0.05$, row 2), but not significantly different from when prefaced by durational information (31%, row 6). There are similar irregularities here in the suppression of the modus ponens inference as we saw earlier in the suppression of the affirmation of the consequent inference.

(2) The suppression of the modus tollens inference was almost as variable. The frequency with which it was made when prefaced by 'during the 15 minute siege' (25%) was not significantly different from it was frequency when it was not prefaced (34%). Note that these

were both low enough to suggest that the inference was suppressed without any extra information. The inferences were significantly less frequent than when prefaced by just situational information (66%, $p < 0.01$), and were also significantly less frequent than when prefaced by just durational information (53%, $p < 0.05$). This is compatible with the suggestion that the provision of no extra information suppressed this inference initially.

The denial of the antecedent and affirmation of the consequent inferences were not expected to be suppressed by additional conditions.

(3) The affirmation of the consequent was not affected, there being no difference between it's being prefaced by both kinds of extra information (34%) and its not being prefaced at all (37.5%). Nor did this differ from it being prefaced just by situational information (28%) or by short durational information alone (44%).

(4) There were also no significant differences between the frequency of the denial of the antecedent inferences when prefaced by both kinds of information (25%) than when not prefaced at all (47%), or when prefaced just by situational information (44%). Surprisingly however, there was a significant difference between the frequency of these inferences when prefaced by both kinds of information (25%) as when prefaced by only durational information (50%, $p < 0.05$). It does not seem feasible to infer from this isolated finding that situational and short durational information leads people to think of alternatives as well as additional conditions, in the face of the

otherwise regular pattern.

So to summarise the effects of information about the situation and a short duration: the modus ponens inference appears suppressed by situational and durational information, and by durational information alone. Situational information on its own does not seem to suppress it. The modus tollens inferences appears to be suppressed by both kinds of information and by neither. But situational information on its own fails to suppress this inference and short durational information on its own also fails to suppress this inference. The denial of the antecedent and affirmation of the consequent inferences should be unaffected by this information. While the affirmation of the consequent was unaffected, the denial of the antecedent was made <u>more</u> when prefaced by short durational information (see Table 12).

To recap overall then: the inferences prefaced by both situational and durational information, are suppressed compared to inferences without any prefacing. The denial of the antecedent and affirmation of the consequent inferences are reduced when prefaced by situation and long duration information compared to when they are not prefaced at all. The modus ponens and modus tollens inferences are reduced when prefaced by situational and short durational information compared to when they are not prefaced at all. However, only the difference for the denial of the antecedent reaches statistical significance (see Table 12).

Comparison of the effects of both kinds of information with the effects of situational information alone suggests that the latter is not enough to suppress the inferences. There are significantly more inferences made in the latter case, for all but the affirmation of the consequent inference. Comparison of the effects of both kinds of information with the effects of durational information alone suggests that durational information is not enough to suppress the inferences either. There are significantly more inferences made in the latter case for all but the modus ponens inference.

Thus there are some disappointing irregularities in the suppression of the inferences, which could plausibly be attributed to the materials (e.g., '15 minute's may not be short enough, and reference to the actors by their role, such as 'the bank-clerk', may provide too much contextual information about the situation in the basic condition). Despite this, there is a regularity in the comparisons of the effectiveness of either of the types of contextual information in isolation from each other to the joint effectiveness of both types of contextual information. This suggests that both an appreciation of the relevance of other information, <u>and</u> the retrieval of specific instances of this information are necessary to suppress the inferences.

<u>Certainty data</u>

The certainty judgements were subjected to a similar analysis of variance as the analysis above. These judgement scores are presented in Appendix A, Table 14, and the effects of the analysis are displayed in Appendix C, Table 7). The only significant effect of this analysis was a main effect of information about the situation [F (1, 42) = 3.9, p = 0.055]. This indicates that subjects were more certain when they had not been given information about the situation (85%) than when they have been given such information (78%). In isolation, this result, although intriguing, is not particularly informative.

Conclusion

These three experiments have shown that contextual information can suppress all inferences. They have indicated that this suppression cannot be attributed to the recognition of the relevance of other information by itself, nor to the retrieval of counterexample instances by itself. Instead it has been shown that it is when the two of these operate in conjunction with each other that inferences become suppressed. A model of how people refrain from making the inferences will encapsulate the processes of appreciating the relevance of other conditions and of retrieving that information. Such a model will now be outlined.

CHAPTER 4:

A MODEL OF INFERENCE SUPPRESSION

"For Watt's concern, deep as it appeared, was not after all with what the figure was, in reality, but with what the figure appeared to be, in reality." Samuel Beckett [(1963) Watt. London: Calder. p. 226.]

A model of reasoning

The model of how people use knowledge to modify their models of the premise information will be outlined. It is an extension of the idea that people represent the premise information in a way that reflects their understanding of the connections between the things referred to in the premises. People build a model corresponding to the state of affairs referred to in the premises. Within this model, a potential conclusion "falls out".

This suggested conclusion is true within the representation or model that has been built of the premise information. People make inferences to reach conclusions that are <u>true</u> or <u>plausible</u> (Bucci, 1978; Pollard, 1982) rather than that are <u>valid</u>. That is, given true premises, people reach a conclusion that is true on the basis of those premises. However, people do have an understanding (be it explicit or tacit) of the essential semantic principle underlying making valid inferences, which is that a conclusion is valid if, and only if, there is no way that the premises could be true while the conclusion is false (Johnson-Laird, 1983). To made a valid inference, a person must test to see if the premises could be true while the conclusion is false. If this state of affairs could be the case, then the conclusion is invalid and must be refrained from.

However, as we have seen, by and large, people only refrain from making inferences when the relevance of other information is recognised and the retrieval of that information is possible. In the absence of either or both of these properties people continue to make the inferences. So the test of whether the conclusion could be false while the premises are true is omitted unless the context or content indicates that other information is relevant and helps the retrieval of that information.

The overarching procedure for testing inferences suggested here is:

> IF there is an appreciation of the relevance of some specifiable other information to the inference at hand
>
> AND there is a retrieval of some instance or token corresponding to that information
>
> THEN the inference will be suppressed.

It will be instructive to look more closely at the operation of the two procedures that are called by this overarching procedure (see Figure 1).

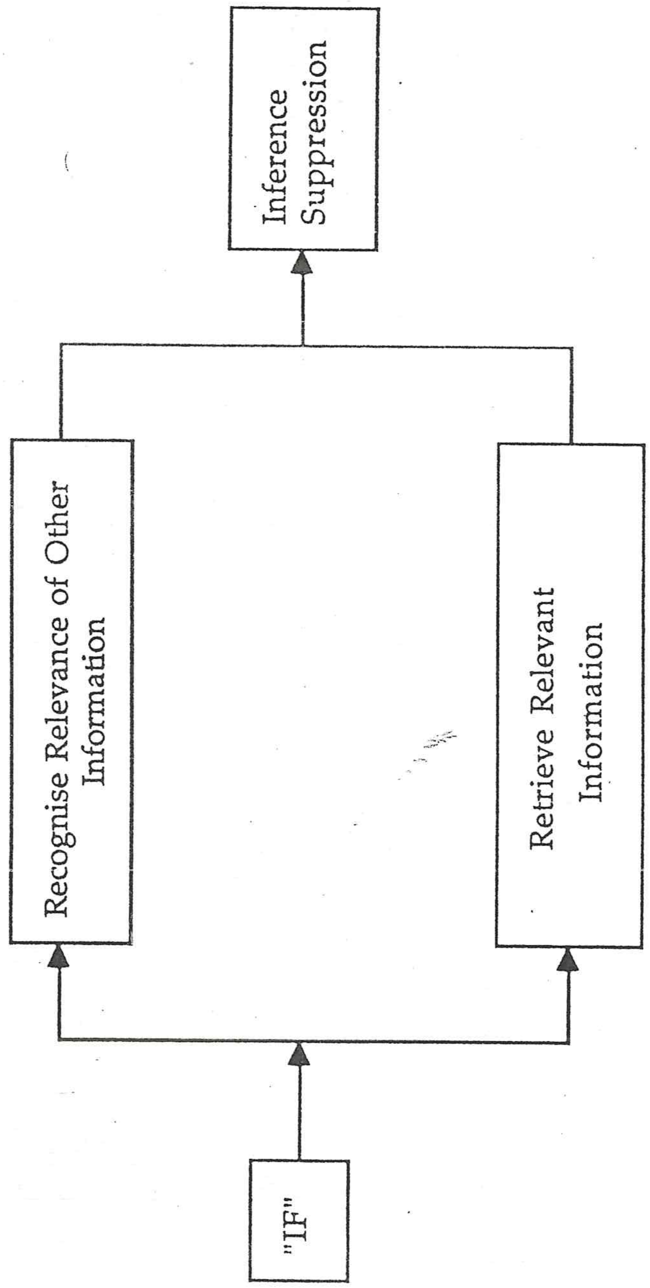

Figure 1 An overview of a model of inference suppression

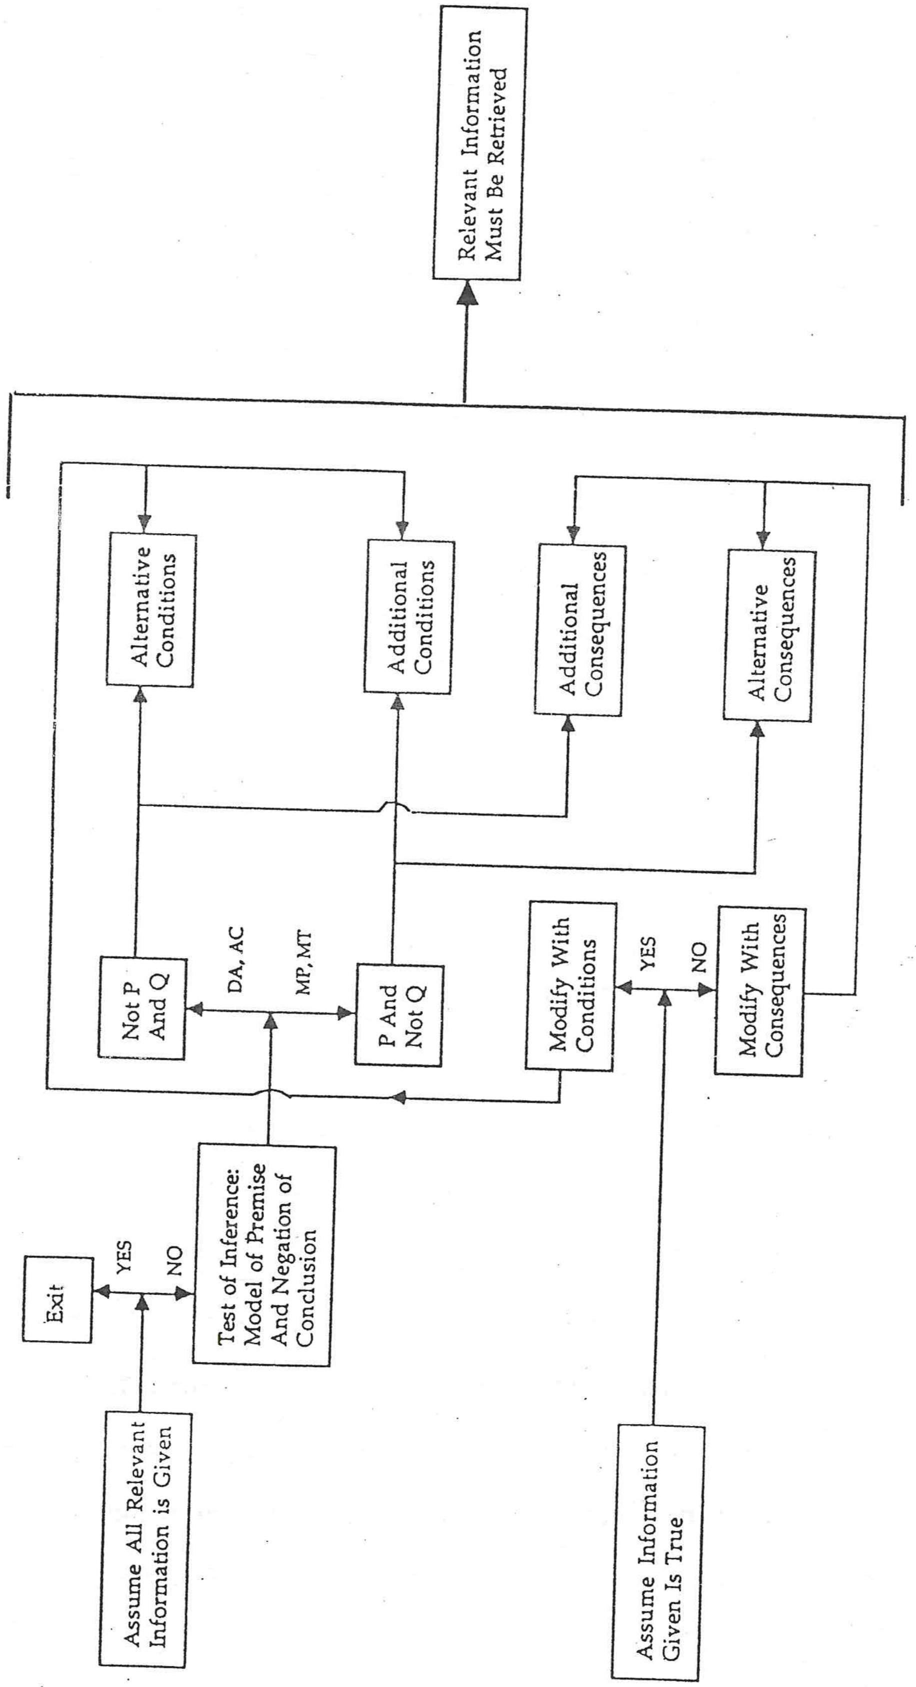

Figure 2 The subprocedures of the general procedure of recognising the relevance of other information

Recognising the relevance of other information

The aim of this procedure is to establish the kind or class of information that is relevant. It consists of two subprocedures which can be thought of as the operations of default assumptions (see Figure 2).

Relevance

Assumption holds. The first of these is the default assumption that all the information relevant to making the inference has been given (Grice, 1975). If this holds, then the procedure ends before it even really begins.

Assumption is violated. However, if the default assumption is violated, then there must be a test of whether the premises can be true while the conclusion is false. The form or the structure of this test will depend on the form of the premises and the conclusion to be tested.

Take the modus ponens inference, from "if this works we can take the evening off" and "this works", to "we can take the evening off". The test is whether the premise (this works) can be true while the conclusion is false (we do not take the evening off). For the modus tollens inference, from "we do not take the evening off" to "this does not work", the test is to see whether the premise is true (we do not take the evening off) while the conclusion is false (this works). It is apparent that these two tests consist of the same state of affairs -this working and us not taking the evening off- simply stated in a different order.

For the denial of the antecedent inference from "this does not work" to "we do not take the evening off", the test is whether the premise (this does not work) could be true while the conclusion is false (we take the evening off). Similarly, the affirmation of the consequent, from "we take the evening off" to "this works", also has the test of a true premise (we take the evening off) and a false conclusion (this does not work). So we can see that these two have the same test -this not working and us taking the evening off -just stated in a different order.

The result of the violation of the default assumption of relevance is the establishment of the <u>form</u> of the test of an inference -it will be either of the form "this works and we do not take the evening off" or of the form "this does not work and we take the evening off".

<u>Truth</u>

<u>Assumption holds</u>. The second subprocedure of the relevance procedure can be thought of as a default assumption of truth. It is again a conversational convention that hearers assume that their speakers have given them true information (Grice, 1975). The assumption that the information given ("if this works then we can take the evening off"), is true can either hold or be violated. Let us think

about what it means to assume that an IF premise is true.

The meaning of the word "IF" can be characterised as a verbal cue, which invites the supposition of the state of affairs expressed in the condition ("this works") and asserts that under that condition, the state of affairs referred to as a consequence of that condition ("we can take the evening off") will hold or be true (Johnson-Laird, 1983; Rips and Marcus, 1977). This is essentially the core of a mental model view of conditional reasoning. A model of the state of affairs in the condition will embrace or incorporate a model of the state of affairs in the consequent (Johnson-Laird, in press).

As we have seen from the previous subprocedure, which specifies the form of the test of the conclusion, this test is an attempt to construct either (a) a model of the state of affairs referred to in the condition, where the consequent does not hold, or to (b) build a model where the state of affairs referred to in the condition does not hold, but the consequent does hold. The default assumption of truth narrows down further what kind of information must be added to the model to help to make this test.

If the default assumption of truth holds, the initial premise "if this works then we can take the evening off" is taken to be true. Reasoners will attempt to construct a test model, (of either "this working and us not taking the evening off", or of "this not working and us taking the evening off", as appropriate) that does <u>not</u>

contradict the model already constructed. That is, they attempt to construct a model which is compatible with the model of the initial premises. How can one construct a model of "this working and us not taking the evening off", or of "this not working and us taking the evening off", that is compatible with the model already constructed of "this working and us taking the evening off" ? In precisely the way that has been suggested throughout this thesis: other conditions can modify the model of "this working and us taking the evening off", without contradicting or denying the truth of it. So when the assumption of truth holds, conditions will be used to modify the model.

To sum up so far: when (a) the default assumption that all relevant information has been given, is violated, this constrains the test to be carried out to be a particular form. And when (b) the assumption of truth holds, this yields the result that the model to be built must be compatible with the initial model of "if this works then we will take the evening off". In order to be compatible it must be modified by other conditions.

Combination. What will the combined results of these two constraints be ? For tests of the form "this works and we do not take the evening off", yielded by the relevance subprocedure (tests of the modus ponens and modus tollens inferences), for models to be modified by conditions (yielded by the truth subprocedure), combining these two results will mean that additional conditions are the general class of information that must be used to modify the model. The <u>conditions</u> that can modify a model of this <u>form</u> must be <u>additional</u> ones.

To jump ahead, for clarity's sake, and provide some specific examples: given "if this works then we will take the evening of" and "this works" the conclusion "we will take the evening off" is suggested. Given that the default assumption, that all relevant information has been given, is violated, the conclusion must be tested, and the form of this test is of "this working and us not taking the evening off". Given that the default assumption that the initial premise is true holds, the general class of information that is relevant to this test is additional conditions. A specific instance of an additional condition (but remember that people would not have retrieved these yet and the example here is only for illustrative purposes) would be, say, "we can prove that it works". So even though this works, since we have to be able to prove that it works we might not take the evening off.

Now consider the case where the form of the test is "this does not work and we take the evening off", yielded by the relevance subprocedure (tests of the denial of the antecedent and affirmation of the consequent inferences). Given that the assumption of truth holds, it is still other <u>conditions</u> that are the relevant class of information. However, for this form of test it is <u>alternative</u> conditions. For tests of this <u>form</u> to be modified by <u>conditions</u> it is <u>alternative</u> conditions that are required.

So given "if this works then we will take the evening off" and "this does not work", the conclusion "we do not take the evening off" is suggested. But if the assumption of relevance is violated then the conclusion will be tested and the form of the test will be "this not working and us taking the evening off". As the assumption of truth holds, the class of information that is relevant is other alternative conditions. A specific example of an alternative condition might be "we get thoroughly fed up with it". So even though this does not work, we might get fed up with it and so take the evening off. Once again, it must be emphasised that these specific instances are for clarification at this juncture, and it is not suggested that people have as yet retrieved anything. Rather the two subprocedures of the relevance procedure have delimited what class of information is relevant.

Assumption is violated. What if the assumption of truth is violated ? In that case the reasoner can construct a model that is incompatible, or that contradicts the initial model of the premises, as the premise is not taken to be true. What will a model of say, "this works and we do not take the evening off", or "this does not work and we take the evening off", look like if it does not have to be compatible with the initial model of "if this works then we will take the evening off" ? In this case, it will be a model modified by other consequences. When the assumption of truth is violated, models can be modified by consequences.

<u>Combination</u>. Assume then, that the default assumption, that all the relevant information has been given, is violated, and the form of the test is: "this works and we do not take the evening off", (tests of the modus ponens and tollens inferences). Assume also that the assumption of truth is violated, so <u>consequences</u> will modify the model. Combining these two factors, it is <u>alternative</u> consequences that will be the class of relevant information which the reasoner must bring to bear on the model. A specific example of an alternative consequence might be that "we spend all evening trying to get the rest of it to work too". So "this works and we don't take the evening off, instead we spend all evening working".

Imagine, instead that the form of the test is: "this does not work and we take the evening off". Then imagine that the assumption of truth is violated. In this case, the class of information that is relevant is <u>additional</u> consequences. An example of this might be that "we realise we can get one of the experts to work it out tomorrow", so that "this doesn't work but we take the evening off because we know we can leave it until tomorrow".

So far we have seen the operation of the default assumption of relevance, and the default assumption of truth, and how the results of each combine. The end results of these satisfies the aim of the procedure of appreciating the relevance of other conditions. That is, it establishes the class or kind of information that is relevant: either alternative or additional conditions, or alternative or additional consequences.

<u>Retrieval of specific counterexample instances</u>

Let us turn now to the other part of the overarching procedure for the suppression of the inferences (see Figure 3).

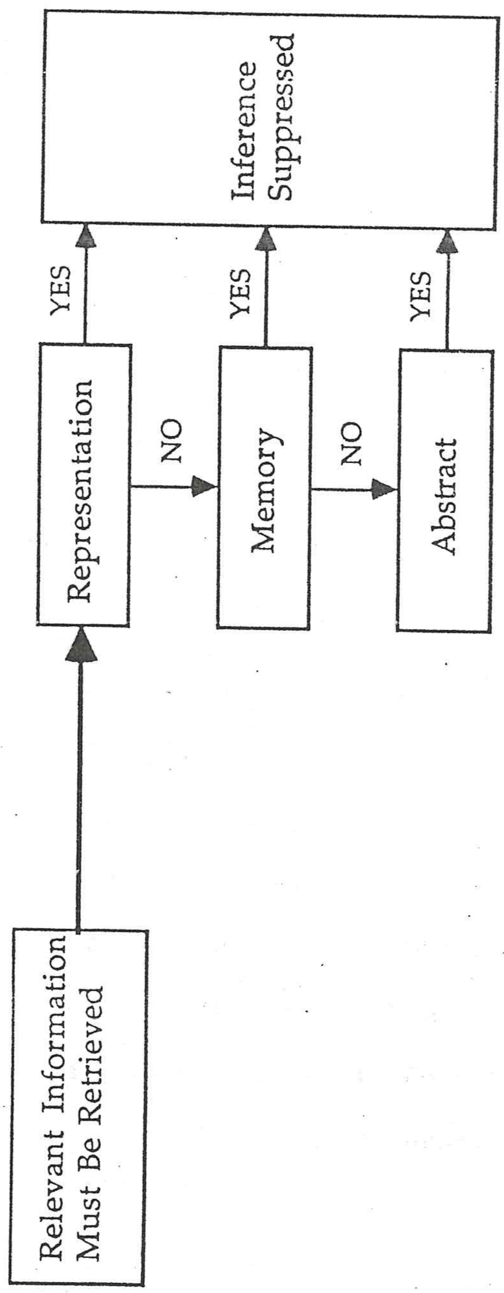

Figure 3 The places-to-search for the general procedure of retrieving specific instances.

The aim of the retrieval of specific instances procedure is to provide a specific instance of the class of information that has been identified as relevant. The retrieval procedure is made up of three subprocedures, each of which is an alternative to the one preceding it (that is, if the first subprocedure works, the others are not tried; if the first does not work then the second is tried and so on. If all the subprocedures are tried and fail then the procedure is not satisfied. The three subprocedures of the retrieval procedure can be thought of as places-to-search in an effort to find specific instances.

Representations

The first place to search is the representation that has been built up of the premise information. This is to check whether an instance of the class of information that has been identified as relevant has already been given. For example, given "if this works then we will take the evening off, if we get fed up with this then we will take the evening off" these premises will be represented in an initial model. Given also, for example, "this does not work", the conclusion "we do not take the evening off" is suggested. If the assumption that all relevant information has been given is violated then the form of the test will be: "this does not work and we take the evening off". Given that the assumption of truth holds, the relevant information is other conditions. Combining these two, the class of information that is relevant is alternative conditions. The first place to search for the relevant information is in the representation of the initial information. Here we find an instance of an alternative condition: getting fed up with it, so the retrieval procedure is satisfied by the information in the first place searched. Consequently, the inference becomes suppressed.

It is possible that the extra information that has been identified as relevant is not provided, and so is not represented in the model of the premises. In this case, the procedure moves on to the next place to search, which is long term memory.

Memory

The following discussion of memory organisation relies on an assumption of a particular memory organisation. This is of how memory is organised for the purpose of conditional reasoning. However, the retrieval sub-procedure is entirely independent of this assumed organisation. Since no direct tests of this assumed memory organisation are presented here, the following discussion must be viewed as an aside.

It will be assumed that conditions and consequences are organised by superordinate structures. So "rain" might be organised by the superordinate structure of "weather", "this works" by the superordinate structure of, say, "motivation to continue". It will be assumed that alternative conditions or consequences are organised by one superordinate structure. For example, "rain" and "snow" are alternative conditions and they are both organised by one superordinate structure, "weather"; "this working", and "us getting fed up with it" are alternative conditions organised by the structure "motivation to continue". It will be assumed that additional conditions and consequences are organised by separate superordinate structures, (that is, not by the same one). Thus, "rain" and "going out for a walk" are additional conditions: "rain" is organised by the superordinate condition of weather, while "going out for a walk" is organised by the superordinate structure of, say, "exposure". Similarly, "this working" is organised by the structure "motivation", while "proving it works" is organised by the structure of say, "scientific respectability".

The main consequence of these assumptions for the present work is that from such an organisation it is predicted that it will be easier to retrieve alternative conditions than additional conditions. That is, given the condition "rain", one can access an alternative to "rain" via rain's superordinate structure of "weather", and searching within that structure for an alternative (e.g., "snow"). To retrieve an additional condition, accessing rain's superordinate structure will be useless. Consequently a person must first identify what other superordinate structure is relevant (for example, "exposure"). Only when they access such a structure, can a person search for an instance (such as "going out for a walk"). This involves altogether more work than retrieving alternatives.

But to return to the model. If the search through memory is successful, an instance of the class of other conditions identified as relevant is found. This satisfies the retrieval procedure of the overarching procedure, and so the inference is suppressed.

Abstract tokens

It may be however that no relevant instance can be retrieved (either because there is none in memory or because none can be found). In this case, the procedure makes its third and final search. It would not be too fanciful to call this a search through an abstract space, since at this point, the reasoner attempts to represent an instance of the class identified as relevant with an abstract token, rather than by a specific instance (Johnson-Laird, 1983). When a person knows that a particular class of information is relevant but is unable to find an instance of it s/he can allow a variable to stand for the instance. This abstract token is incorporated into their model of the premises, instead of a definite instance. Since this is quite ingenious, it is assumed that most people do not do it, either because they give up before this, or because it never occurs to them. But if they do this, it will satisfy the retrieval procedure. Combined with the relevance procedure in the overarching procedure, this will suppress the inference.

To summarise the model: the procedure by which people refrain from making inferences operates after they have constructed a model of the premises and have reached a potential conclusion. It is a procedure consisting of the appreciation of the relevance of other conditions and the retrieval of specific instances, the combination of which will lead to the suppression of the inference.

The relevance procedure consists of two subprocedures. Firstly, there is a default assumption of relevance, which, if violated, specifies the form of the test to be conducted on the model. Secondly, there is a default assumption of truth. If this holds then other conditions may modify the model. If it is violated, then other consequences may modify the model. These two subprocedures combine to yield the class of information that is relevant: either alternative or additional conditions or alternative or additional consequences.

The resultant class of relevant information yielded by this procedure must combine with the results of the second procedure, that of retrieving specific instances. This consists of three sequential subprocedures, of searching for specific instances in particular places. Failure to find an instance in one ensures an attempt to search in the next, success ensures termination of the search. The three places-to-search are (a) in the representation or model itself, (b) in memory, and (c) in an abstract space. If one of these yields an instance (or a pseudo-instance in the last case) then the two procedures of the overarching procedure will have been satisfied. In this case, the inference will be suppressed.

Comparison of the model to others

The model which has been described requires that both the recognition of the relevance of other information and the retrieval of specific instances of that information must occur for inferences to be suppressed. This requirement means that it accounts for, or "postdicts" the data better than rival models.

Recall that these two operations have been identified in rival research traditions as the pertinent factors in inference suppression. However, the data of Experiment 2 and more especially, of Experiment 3, have indicated that neither of the operations of relevance or retrieval are enough by themselves to suppress the inferences. Hence a model encapsulating and superseding the findings of both traditions is appropriate. The suggested model addresses some of the persistent problems that have prevailed in reasoning research.

Firstly, the present experimental findings and modelling efforts shed new light on the use of conditionals in inducements (Fillenbaum, 1975, 1976, 1978; Geis and Zwicky, 1971). It is known that the pragmatics of conversation -such as that speakers give their hearers both enough and true information- govern such context-sensitive utterances as inducements (Grice, 1975; Levinson, 1983). By incorporating these conventions into the processes of reasoning, this model shows how the use of conditional inducements is firmly rooted to the time and place of the utterance. The results of Experiment 3 have shown how such utterances when 'set in context' elicit radically different patterns of inference than when presented in isolation.

Rival accounts

Memory-based accounts

Secondly, the model casts doubt on purely memory-based accounts of people's inferential behaviour on the selection task. It was found (in Experiment 2 and 3) that the retrieval of counterexamples by itself is not enough to suppress inferences. Consequently, accounts which emphasise retrieval alone (e.g.,Griggs, 1983; Pollard, 1982) tell only half a tale. While information in long-term memory has an active role in reasoning, it is a role which only comes to the fore when prodded by some recognition of the relevance of the information. In emphasising the joint operation of the two procedures of relevance and retrieval, the model presented here encompasses the experimental findings that there are interactive effects on inferential behaviour, of subject's memory of counterexamples and of instructions to think of instances that "violate the regulation" (Griggs, 1984).

Mental logic accounts

Thirdly, this model has implications for accounts of the suppression of the denial of the antecedent and affirmation of the consequent inferences, which appeal to the detachment of conversationally-invited inferences (Braine and Rumain, 1983; Rumain, Connell and Braine, 1983). This view is untenable in the face of the suppression of all of the inferences. This overall suppression conflicts with the tenets of the rule-based view, as has been argued earlier. Instead, the suppression of all inferences points to the need for a sufficiently rich machinery to capture the influence of the semantics and pragmatics of the situation, on inference-making.

Pragmatic schemas

The suppression of all inferences also undermines the pragmatic schema approach (Cheng and Holyoak, 1985; Cheng, Nisbett, Holyoak and Oliver, 1986). Attributing the suppression of inferences to the presence of modals in rules (e.g., if a person is over 19 then they *may* be drinking beer) will not work. On the one hand, this commits one to the view that these modals are somehow omitted from the rule when the context changes to a situation where the inference is not suppressed. Or else, one is committed to the view that when the context changes, the premise is mapped to a different modal-less rule. This is clearly cumbersome.

Inference suppression

Finally, the model provides suggestions about how people refrain from making inferences. Other models have not so far, addressed this issue. This is primarily because rule-based views envisage inference making as the operation of a rule and the failure to make an inference has simply been attributed to the absence of a rule (e.g, Osherson, 1975; Rips and Conrad, 1983). Similarly, memory-based views also explain the failure of a subject to make an inference by attributing it to the absence of instances of past experience in memory. This model shows instead that the machinery for making inferences, in building models of the state of affairs referred to in the premise, carries within it the wherewithal to test the inferences made.

Seeing the model in operation

The motivation for the construction of the model has been to account for reasoning in everyday situations. It would be appropriate to see its operation in situations even closer to common-sense reasoning than those dealt with so far. One such everyday type of conditional reasoning occurs when planning (Fillenbaum, 1975; Marcus and Rips, 1979; Staudenmayer, 1975). In planning, for example, future courses of action, one must often think of a set of contingencies that are interconnected. Each action is a condition for, or a consequence of, another action. An example of this is planning that:
"If I work at this harder then it will become clearer. If it becomes clearer then I will take a break"
Often such planning sequences are made more complicated since they embody alternatives, such as:
"If I work at this harder then it will become clearer, If it does not become clearer then I will take a break"
or:
"If I work at this harder then it will not become clearer, if it becomes clearer then I will take a break"

Note that in these characteristic planning sequences there is not one conditional in isolation, but several interconnected ones (two in the above examples). Here we have a naturalistic situation, where a condition, such as "working harder", is embedded in a context where other conditions and consequences related to it are also "active" or are being thought about. That is, in such sequences, people are naturally engaged in thinking about interconnections between conditions and consequences.

The model or representation of a planning sequence will incorporate such interconnections as can be made. Note also that the relevance of these interconnections to the inferences suggested in the model will be obvious to the planner. The retrieval of the specific instances of relevance to the inference being made has thus, already been accomplished and so the specific instances are already represented in the model. Some of the interconnections may constitute an alternative or additional condition or alternative or additional consequence. In this case, the inference will be suppressed. For example, from the sequence, "if I work harder then this will not become clearer, if it becomes clearer then I will take a break", could a connection be built between "working harder" and "taking a break" ? If one could, and it constituted an alternative consequence from "working harder", other than "this not becoming clearer", what would be the effects of this change in the representation of the sequence? By the model proposed earlier, such an alternative consequence would be recognised as being relevant to the suggested inference, and would already be retrieved, and so it would lead to the suppression of certain inferences. Before examining how this may be so, let us move on to how people behave with planning sequences.

CHAPTER 5:

SUPPRESSION IN PLANNING SEQUENCES

Introduction

To understand everyday reasoning with planning sequences we will have to examine the inferences that people make from these sequences. The same four inferences that we have looked at so far, modus ponens and tollens, the denial of the antecedent and affirmation of the consequent inferences, are invited within sequences of conditional premises. For example, from the extended argument that has the conditional premises: "If she goes to the family gathering then she will enjoy herself, if she enjoys herself then she will go to bed tired" and the minor premise "she went to the family gathering", a chain of modus ponens inferences may be made. These will lead to the conclusions that " she enjoyed herself" and that "she went to bed tired".

The same inferences can be made from sequences of conditional premises as can be made from single conditional premises in isolation (see Appendix A, Table 15). That is, these modus ponens inferences are just the same ones that would be made from a comparable single conditional premise argument, or standard argument. So from "if she goes to the family gathering then she will enjoy herself" and the minor premise "she went to the family gathering", the modus ponens inference to "she enjoyed herself" is invited. And from "if she enjoys herself then she will go to bed" in isolation, with just the minor premise "she enjoyed herself", the modus ponens inference to "she went to bed tired" may be made.

This similarity in inferences from extended and standard arguments has lead to what I will call the assumed generalisability of the standard argument. This is the assumption that the same frequency of inferences will be made on extended arguments as are made on standard arguments. Such an assumption takes no account of the representational differences which, as we noted briefly at the end of Chapter 4, may obtain between these two kinds of arguments. Before discussing such potential differences let us see how this assumption has affected previous research. The widespread acceptance of this assumption is indicated in three ways. Firstly, this assumption is explicitly stated within some of the foremost theories of reasoning. Secondly, there is a profusion of research using standard arguments and thirdly, there is a paucity of research on extended arguments. Each of these will be described in turn.

In natural deduction models of reasoning (Osherson, 1975, 1976; Rips, 1983a, 1983b; Braine, 1978, Braine, Reiser and Rumain, 1984) it is assumed that reasoning with extended arguments consists of a simple iterative extension of reasoning with standard arguments. What are iterated are the inference schemas or rules. It is assumed that these rules, which are applied to the available starting information in standard arguments, are applied in a repetitive fashion both to the starting information of an extended argument and also to the products of each previous step within an extended argument. On some accounts (e.g., Braine, Reiser and Rumain, 1984) the procedures that seek these inference schemas may also be applied iteratively. This iteration has also been used to account for responses to standard arguments that are made up of more than one connective (such as "if she does not both go to the family gathering and drink a lot then she will enjoy herself"). These necessitate "multiple operations", it is claimed in natural deduction models rather than "single operations" (e.g., Osherson, 1975; Rips, 1983b; Rips and Conrad, 1983; Braine, Reiser and Rumain, 1984).

Most research on reasoning employs standard arguments (for reviews, see Wason and Johnson-Laird, 1972; Evans, 1982). The experimental findings from these are assumed to generalise directly to everyday reasoning. So for example, the ecological validity of the effects of realistic content on inference making is frequently taken for granted. This is despite such content effects being explored solely using standard arguments and related single premise paradigms.

The final testimony to the widespread acceptance of the assumption, that reasoning with standard arguments is representative of reasoning with more extended types of argument, is the distinct paucity of research on the latter. While extended arguments have been used to investigate memory for arguments (Marcus, 1982) and as tasks analogous to the categorical syllogism (Guyote and Sternberg, 1981, Experiment 5), their use as tools to investigate the inferences people make has been neglected.

The first thing to establish then is whether or not the assumption that people make the same frequency of inferences on standard and extended arguments is warranted. It has been suggested already (in Chapter 4) that there is at least one feature of extended arguments (which is also present in everyday reasoning) that is not characteristic of standard arguments. This is that the person reasons from an extended representation with extended arguments. That is, there may be several actions or events represented in the model one constructs of the premises of an extended argument, unlike the mere two actions or events embodied in the model constructed from a standard argument. Interconnecting links may be forged between some of these extra actions or events in an extended argument representation. We have seen in the last chapter that when people consider extra information, then the inferences they make can be affected. Hence, the possibility of some inferences becoming suppressed is greater from the rich representation constructed for extended arguments, than from the comparatively impoverished model constructed of a standard argument. Before attempting to distinguish exactly how fewer inferences could come to be made from extended arguments, the first step must be simply to establish whether or not there are any differences between the frequency of inferences made on extended arguments and on comparable pairs of standard arguments. The first experiment here on everyday planning sequences addresses this question.

Experiment 4

This experiment was designed to test whether or not there are differences in the frequency of inferences that are made from standard arguments and from extended arguments. Several different forms of arguments were used (for examples of each of these, see Appendix A, Table 15). The simplest form of extended argument used was a transitive one. This had affirmative conditional premises, such as:

"If she goes to the family gathering then she will enjoy herself, if she enjoys herself then she will go to bed tired".

Another form of extended argument used had a negative component in the first conditional. This resulted in the sequence:

"If she goes to the family gathering then she will not enjoy herself, if she enjoys herself then she will go to bed tired".

The third form of extended argument also had a negative component but this time in the second conditional:

"if she goes to the family gathering then she will enjoy herself, if she does not enjoy herself then she will go to bed tired"

Note that the components negated in these last two extended arguments are the middle components, that is, the component that is common to both conditionals ("she enjoys herself"). Because of this, these sequences are atransitive. Arguments of different forms like these may lead to different interconnections being made between the actions or events referred to in the premises. The interconnections may constitute alternative or additional conditions, or alternative or additional consequences. So this could lead to some inferences being suppressed by some forms of arguments and not by others.

Transitive extended arguments

For the transitive extended arguments accompanying the conditional premises with

(1) the minor premises of "she went to the family gathering" invites a chain of modus ponens inferences to "she enjoyed herself" and "she went to bed tired",

(2) while accompanying them with the minor premise "she did not go to the family gathering" invites a chain of denial of the antecedent inferences to "she did not enjoy herself" and to "she did not go to bed tired".

(3) Accompaniment by the minor premise "she went to bed tired" invites an affirmation of the consequent chain of inferences to "she enjoyed herself" and "she went to the family gathering",

(4) while accompaniment by "she did not go to bed tired" invites the modus tollens chain of inferences to "she did not enjoy herself" and "she did not go to the family gathering".

However, in contrast to the first of these descriptions (that the minor premise "she went to the family gathering" invites a chain of modus ponens inferences), some theorists have claimed that the inference from "she went to the family gathering" to "she went to bed tired" is reached by a single modus ponens inference. This is claimed to be a product of a transitive inference schema or rule (e.g., Braine, 1978, Braine and Rumain, 1983). With this transitive rule, people may infer "if she goes to the family gathering then she will go to bed tired" from the initial two conditional premises, and then make the single modus ponens inference from "she went to the family gathering" to "she went to bed tired". However Braine and Rumain (1983) have reported on the lack of consistency in the results of several studies about this transitive schema (e.g., Mason, Bramble and Must, 1975; Roberge, 1970). So it is not clear that people do make a direct inference from "she went to the family gathering" to "she went to bed tired", without going through the modus ponens steps that include "she enjoyed herself". Furthermore, even if people did reason according to such a transitive principle instead of according to a modus ponens chain, it is not clear what way they would reason for the other inferences. For example, it is not specified whether the minor premise "she did not go to bed tired" (described here as inviting a chain of modus tollens inferences) would lead directly to the conclusion "she did not go to the family gathering". As there is no single schema for modus tollens in Braine's system, his system predicts that this argument would require more than two elementary steps. So the idea that people make a series of inferences in these transitive extended arguments seems to be a consistent, and as yet unfalsified hypothesis, to adopt.

A complete transitive extended argument involves conditional premises such as "if she goes to the family gathering then she will enjoy herself, if she enjoys herself then she will go to bed tired" accompanied by the minor premises of either

(1) "she went to the family gathering" or

(2) "she did not go to the family gathering" or

(3) "she went to bed tired" or

(4) "she did not go to bed tired".

Atransitive extended arguments

The inclusion of atransitive extended arguments was motivated by three concerns. Firstly the usage of atransitive sequences in every-day reasoning seems to be common, particularly for the expression of conditional alternatives. Secondly, they provide an opportunity to observe chains made up of _different_ inferences. That is, in both of the atransitive forms of extended arguments used here, it is not the case that, say, a modus ponens inference is followed by another modus ponens inference. Instead each inference is followed by an inference that is different from itself. Examples will clarify this (see also Appendix A, Table 15). Take firstly, the extended sequence:

"if she goes to the family gathering then she will enjoy herself, if she does not enjoy herself then she will go to bed tired"

(1) Accompanied by "she went to the family gathering", the modus ponens inference to "she enjoyed herself" is invited, followed by the denial of the antecedent inference to "she did not go to bed tired".

(2) Accompanied by the denial of the antecedent "she did not go to the family gathering", the inference to "she did not enjoy herself" is followed by the modus ponens inference to "she went to bed tired".

(3) When the sequence is accompanied by "she went to bed tired", the affirmation of the consequent inference to "she did not enjoy herself" is followed by the modus tollens inference to "she did not go to the family gathering".

(4) Accompaniment by "she did not go to bed tired" invites the modus tollens inference to "she enjoyed herself", followed by the affirmation of the consequent inference to "she went to the family gathering".

For the sequence:

"if she goes to the family gathering then she will not enjoy herself, if she enjoys herself then she will go to bed tired"

the same chains of inferences, modus ponens followed by the denial of the antecedent, and vice ve rsa; and modus tollens followed by the affirmation of the consequent, and vice ve rsa, will follow. For this sequence, the conclusions will, of course, be different, as inspection of the application of these inferences to this sequence will reveal (see also Appendix A, Table 15).

A third reason for manipulating the presence of negatives in the argument sequences is the evidence (using standard arguments) that different frequencies of inferences are made, as a function of the presence and location of a negative in a conditional sentence (e.g., Evans, 1972a, 1977; Roberge, 1971, 1974, 1978; Pollard and Evans, 1980; Adams, 1980; Griggs and Cox, 1983). So it is of interest to examine whether or not there are differences between extended arguments and standard arguments as a function of negativity.

A complete atransitive argument will include the conditional premises of:

"if she goes to the family gathering then she will enjoy herself, if she does not enjoy herself then she will go to bed tired" <u>or</u> "if she goes to the family gathering then she will not enjoy herself, if she enjoys herself then she will go to bed tired" each accompanied by the minor premises:

(1) "she went to the family gathering" or

(2) "she did not go to the family gathering" or

(3) "she went to bed tired" or

(4) "she did not go to bed tired".

<u>Standard argument counterpart pairs</u>

The standard argument <u>counterparts</u> to these extended arguments were comparable, separate standard argument pairs (see Appendix A, Table 15). The counterparts to the transitive extended arguments were the standard argument pairs:

(1) "if she goes to the family gathering then she will enjoy herself" accompanied by the minor premises of either

(a) "she went to the family gathering" or

(b) "she did not go to the family gathering" or

(c) "she enjoyed herself" or

(d) "she did not enjoy herself" and

(2) "if she enjoys herself then she will go to bed tired" accompanied by either

(a) "she enjoyed herself" or

(b) "she did not enjoy herself" or

(c) "she went to bed tired" or

(d) "she did not go to bed tired".

The counterparts to the atransitive extended argument of: "if she goes to the family gathering then she will enjoy herself, if she does not enjoy herself then she will go to bed tired" are the standard arguments of:

(1) "if she goes to the family gathering then she will enjoy herself" accompanied by the appropriate minor premises, and

(2) "if she does not enjoy herself then she will go to bed tired", accompanied by the appropriate minor premises.

Similarly the counterparts to the atransitive extended argument "if she goes to the family gathering then she will not enjoy herself, if she enjoys herself then she will go to bed tired" were the standard arguments of

(1) "if she goes to the family gathering then she will not enjoy herself", accompanied by the appropriate minor premises, and

(2) "if she enjoys herself then she will go to bed tired" accompanied by the appropriate minor premises. The frequency of inferences made from extended arguments will be compared to the frequency made from comparable standard arguments, to ascertain whether or not there are differences between them.

Method

Design

A factorial design with four factors was employed. The first between-subjects factor was the argument length. Three groups of subjects received extended (two conditional premises) arguments and the other three groups received standard (single conditional premise) arguments.

The second between-subject factor was the form of the conditional premises. One group of subjects received extended arguments with affirmative premises while another group received standard arguments with affirmative premises. Extended arguments and standard arguments of this first form involved all affirmative conditional premises. A third group received extended arguments with a negative at the antecedent of the second conditional, while a fourth group received comparable standard arguments. For extended arguments and standard arguments of the second form, half of the conditional premises were affirmative and half had a negated antecedent ("if she goes to the family gathering then she will enjoy herself, if she does not enjoy herself then she will go to bed tired"). For the remaining two groups of subjects, one group received extended arguments with a negative located at the consequent of the first conditional, while the other group received comparable standard argument pairs. For extended arguments and standard arguments of the third form, half of the conditional premises had a negated consequent and half were affirmative ("If she goes to the family gathering then she will not enjoy herself, if she enjoys herself then she will go to bed tired").

There were two within-subject factors. These concerned two salient characteristics of the four inferences investigated. The first within-subject factor was the direction of the inference, either forwards, moving from the first component of the conditional to the second (the modus ponens and denial of the antecedent inferences) or backwards, moving form the second component to the first (the affirmation of the consequent and modus tollens inferences). The second concerned the form of the minor premise. This may be either of the same form as the component it refers to in the conditional premise (the modus ponens and affirmation of the consequent inferences) or of a different form from that component (the denial of the antecedent and modus tollens inferences)

Materials

Extended arguments. In the extended arguments the minor premise may either assert or deny one of three components [either the first (going to the family gathering), middle (enjoying herself), or last (going to bed tired)]. This gives rise to six distinct argument forms. For each of these forms, two inferences may be made (one inference about each of the other two components). Our interest lies in the sequence of inferences people make from the first component to the other two components, or from the last component to the other two components. So two of the argument forms, those affirming and denying the middle components (enjoying herself) are not of interest here. Forms of this sort were included as filler items only and responses to them were discarded prior to any data analysis. Each of the remaining four forms gives rise to one inference based on the given minor premise, and one based on the intermediate conclusion reached by this first inference (see Appendix A, Table 15).

Subjects received two of each of the six forms. So they received, in random order, 12 items that required (potentially) 24 inferences to be made. Twelve pairs of contents were generated for the 24 premises. The contents involved realistic mundane events, similar to those in the example used to illustrate the rationale of this experiment. They involved the actions of a single actor (identified as "she" for female subjects and as "he" for male subjects). The nature of the actions and events was such that they could be either transitive or plausibly negated at either of the middle terms. The actions were arranged in a temporally-sequential order. For examples of the contents used see Appendix B, Table 4.

Standard arguments. As the standard argument pairs were designed to be comparable in all ways to their counterpart extended arguments, their construction involved (1) the separation of the two conditional premises and (2) the accompaniment of these premises individually by the appropriate minor premises. Take for example, the counterpart pair of standard arguments to the extended argument of: "if she goes to the family gathering then she will enjoy herself, if she enjoys herself then she will go to bed tired, she went to the family gathering". These involved the presentation of the standard argument:

(1)"if she goes to the family gathering then she will enjoy herself, she went to the family gathering"

and after the subjects response to this argument had been made, they were given the second standard argument:

(2) "if she enjoys herself then she will go to bed tired, she enjoyed herself"

Procedure

Subjects were tested individually and while the test materials were presented in both print and oral form, their responses were oral. These responses were recorded on a tape recorder that ran uninterrupted throughout the session.

The instructions to subjects in all conditions used the same affirmative premise(s) modus ponens example, (with the appropriate number of premises). So subjects in the extended arguments condition were given "if she gets up tired then she will go into town, if she goes into town then she will spend all her money". They were instructed to imagine that they had been told the information in the premises about someone. Then they were to imagine that "later on" they heard the information in the minor premise "she got up tired". They were asked to say what they considered to follow from this information, and were given examples of several responses. These examples were: "you might think it follows that she went into town, or that she did not go into town, or that you can't say for definite whether she went into town or not, and you might think it follows that she spent all her money, or that she did not spend all her money, or that you can't say for sure whether she spent all her money or not". For each extended argument, if a subject made only one inference, the prompt "anything else ?" was given. Instructions to the subjects in the standard argument conditions consisted of the first conditional of this example in a standard argument with the same minor premise. The examples of the responses for these subjects concerned only that conditional.

Subjects

The 60 undergraduates of Trinity College, Dublin, who participated in this experiment were randomly assigned to one of six conditions (n = 10). Six subjects were replaced because they had received tuition in logic. These replacements were made prior to any data analysis.

Scoring

Responses were given a score of 1 if they were definite endorsements of both of the invited inferences. Let us look more closely at these criteria. Definite conclusions are those that do not include such qualifiers as, for example, possibly, probably, or maybe. For extended arguments these endorsements were of both of the inferences in the sequence of inferences. In standard arguments these endorsements were of the inferences in both pairs of the counterpart standard arguments.

For example, in an extended argument such as "if she goes to the family gathering then she will enjoy herself, if she enjoys herself then she will go to bed tired, she went to the family gathering", the responses "she enjoyed herself" and "she went to bed tired" were altogether given a score of 1. Endorsements of "she enjoyed herself" without endorsing "she went to bed tired", or vice vearsa were scored 0. similarly in the corresponding pair of standard arguments, endorsements of both of the comparable inferences were given a score of 1, while endorsements of only one or other of the inferences on their own was scored 0 (see Appendix A, Table 15). This stringency is necessary to avoid artifactual constraints.

For example consider the hypothetical case where a first inference is endorsed infrequently on both standard arguments and extended arguments. For extended arguments this would impose a constraint that there would also be few endorsements of the second inferences. This is because the second inference depends on the prior conclusion reached by the first inference. No such constraint exists for standard arguments. This inequality could lead to differences being observed that are of little intrinsic interest (1).

Results and Discussion

The scored responses were subjected to an analysis of variance

---

(1) I am grateful to J.St.B.T. Evans for pointing out that such difficulties could exist.

involving the four factors of argument length, argument form, inference direction and inference form, with repeated measures on the last two factors. The response scores are presented in Appendix A, Table 16, part i. When subjects did not make both inferences they made just one or the other of them. In Appendix A, Table 16, part ii displays the responses where subjects made only the first inference in a pair of inferences. Part iii of this Table presents the frequency of responses of only the second inference in a pair of inferences. When subjects did not definitely endorse an inference, they made it with some probabilistic qualifier, such as "possibly". Few responses were of the inappropriate inference, and so the probabilistic responses are a mirror image of the definite endorsements. The only significant effect of this analysis was an interaction of argument length and form [$F\ (2,\ 54) = 3.53$, $p = 0.0361$]. The other effects of the analysis are displayed in Appendix C, Table 8.

Post hoc tests (Tukey's HSD test) were conducted on this interaction to discover its source. These revealed that there are fewer inferences made on extended arguments of the form "if she goes to the family gathering then she will not enjoy herself, if she enjoys herself then she will go to bed tired" (31%) than on their counterpart standard arguments of the same form (65%, $p < 0.01$, see Appendix A, Table 16, rows 5 and 6). However, the frequency of inferences endorsed on transitive extended arguments (65%) did not differ significantly from their counterpart standard arguments (46%). Nor did the frequency of inferences on atransitive extended arguments of the form "if she goes to the family gathering then she will enjoy herself, if she does not enjoy herself then she will go to bed tired" (55%) differ from their counterpart standard arguments of the same form (61%).

A discussion of the implications of this finding, that extended arguments of certain forms differ from their comparable standard arguments, for reasoning research in general, will be reserved until we have established what causes this difference. Why is it that only extended arguments of the form "if she goes to the family gathering then she will not enjoy herself, if she enjoys herself then she will go to bed tired" differ from their counterpart standard arguments ? As the extended arguments and standard arguments do not differ in the form of the inferences to be made, the observed differences cannot be attributed to, for example, "negative conclusion biases" (e.g., Evans, 1972a, 1977; Pollard and Evans, 1980). Is it the case that interconnections between going to the family gathering, enjoying herself and going to bed tired are built up in the model of the premises of this form, such that these suppress some of the inferences as we surmised earlier ?

Extended arguments of all forms differ from their counterpart standard arguments along at least two identifiable dimensions. One is that the information in extended arguments is represented in a joint representation. The second is that the first inference a person makes from an extended argument yields an intermediate conclusion which must then be used as the basis for a second inference. In contrast to this, in standard arguments each inference is based on the given minor premise. These two dimensions will be outlined in detail shortly, as the aim of the next experiment is to ascertain which of these dimensions, in interacting with the form of the arguments, is responsible for the differences observed.

Experiment 5

In this next experiment our interest is in two dimensions characteristic of ordinary everyday reasoning, which are captured in extended arguments but not in standard arguments. Either one of them, or both, may be responsible for the differences found between certain extended arguments and their standard argument counterparts. An attempt will be made in the next experiment to systematically manipulate the presence and absence of these two dimensions. First, an outline of each of the dimensions, followed by a brief account of the way in which its presence may be manipulated, will be given.

Interconnections

Extended arguments differ from standard arguments in the amount of information to be assimilated into a unified representation. In an extended argument, the actions and events referred to in both of the conditional premises will be related in one representation. In this single representation, relations between the two conditional premises will be established. So interconnecting relations between going to the family gathering, enjoying herself, and going to bed tired may be made. In contrast to this, in the comparable pair of standard arguments, the actions and events of each conditional premise will be represented separately. Thus representing relations between the two conditional premises will not be possible. For example, interconnections between going to the family gathering and going to bed tired cannot be made.

It may be that, in the course of establishing relations between the two conditional premises (in an extended argument), novel relations are forged between the end terms (see, for example, Guyote and Sternberg, 1981; Johnson-Laird, 1983, on the formation of such relations for categorical syllogisms). For example, for the sequence "if she goes to the family gathering she will not enjoy herself, if she enjoys herself then she will go to bed tired" the invited connection between the end terms would relate "going to the family gathering" to "not going to bed tired" (because from "going to the family gathering", a modus ponens inference to "not enjoying herself" may be made, followed by a denial of the antecedent inference to "not going to bed tired"). What if, for some reason, a different relation, say from "going to the family gathering" to "going to bed tired", was made instead ? If the novel relation established in the unified representation runs contrary to that indicated by the inferential chain for the argument, the inferences may be considered unwarranted. That is, forging novel, contrary interconnections may suppress some inferences.

It is hypothesised that it is this dimension that is responsible for the observed differences between extended and standard arguments. One way to establish whether this is so, is to hold all else constant and remove this dimension. If the differences disappear when this dimension is eliminated, then this will indicate that this dimension is responsible for the observed differences. To eliminate the opportunity of representing the information in both conditional premises in a unified representation, the conditional premises can be presented separately (see Appendix A, Table 17).

Uncertainty in the use of intermediate conclusions

There is however, a second dimension characteristic of extended arguments which may be responsible for the differences observed, contrary to our expectations. In standard arguments, subjects are given a minor premise whose truth they are instructed to assume, on which to base their inferences. Thus from the pair of standard arguments:
(1)"If she goes to the family gathering then she will enjoy herself, she went to the family gathering" and
(2)"if she enjoys herself then she will go to bed tired, she enjoyed herself"
the minor premises ("she went to the family gathering" and "she enjoyed herself" respectively) are both given.

This is the case only for the first inference to be made in an extended argument. For example, for the extended argument:
"if she goes to the family gathering then she will enjoy herself, if

she enjoys herself then she will go to bed tired, she went to the family gathering" the minor premise "she went to the family gathering" is given and leads to the conclusion "she enjoyed herself". However the second inference (to "she went to bed tired" in this instance) must be based on the conclusion from the first inference ("she enjoyed herself" above), rather than on a given minor premise. So the intermediate conclusion must be used as a minor premise to make the second inference.

Subjects may be uncertain however about using these self-derived conclusions, as no external confirmation of their truth has been given. While the effects of uncertainty on reasoning have begun to receive attention (e.g., Rich, 1983), the use of prior conclusions has not, despite both being ubiquitous in everyday reasoning (e.g., Collins, 1978). However it is possible that uncertainty in the use of prior conclusions interacts with the argument form to lead to fewer inferences being made on extended arguments of the "if she goes to the family gathering then she will not enjoy herself, if she enjoys herself then she will go to be tired" form, than on comparable standard arguments.

So although it is hypothesised that this is not the crucial dimension, it is important to establish whether it plays a role in the observed differences. In order to eliminate this uncertainty, some external confirmation of the truth of the subject-derived intermediate conclusion in extended arguments can be given (see Appendix

A, Table 17). If the differences disappear when this is done, then this dimension is responsible for the differences.

Method

Design

The factorial design employed in this experiment had two between-subject factors and two within-subject factors. The between-subject factors concerned the experimental manipulations of the presence and absence of (1) the opportunity to represent both conditional premises in a unified representation, and (2) the uncertainty in using an intermediate conclusion, due to the lack of external confirmation of its truth. The form of all arguments was "if she works hard then she will not leave the library before dinnertime, if she leaves the library before dinnertime then she will spend the evening in the buttery", since this is the form for which differences have been observed between extended arguments and standard arguments in Experiment 4.

One group of subjects received extended arguments of this form, just as in Experiment 4 (see Appendix A, Table 17). In this extended arguments condition both of the dimensions, of being uncertain because of having to use an intermediate conclusion as a minor premise, and of being able to represent the two conditional premises in one model, are present. Subjects in a second condition received comparable pairs of standard arguments. In this standard arguments condition neither of these dimensions are present.

A further two conditions were constructed in which only one or the other but not both of these dimensions were present. These conditions were termed the uncertainty elimination condition and the interconnection elimination condition. In the uncertainty elimination condition, the third group of subjects received the same materials as those used in the extended argument condition, except that after these subjects had made their first inference, they were informed that their response was correct. For the interconnections elimination condition, the fourth group of subjects received materials that were almost identical to those of the standard argument condition. These subjects received the same materials for the first standard argument in each pair. The difference was that after they had made their first inference in each task, the experimenter presented the next appropriate conditional premise without an accompanying minor premise. The subject was then informed "you concluded from the last sentences that "she enjoyed herself", so what follows from that for this sentence ?". In this way subjects had to use their self-derived intermediate conclusion but were not given the opportunity to represent the premises together.

The within-subjects factors were, as in Experiment 4, the direction of the inference and the similarity of the form of the minor premise to the component it referred to in the conditional premises.

Materials

Six of the content domains from Experiment 4 were employed (see Appendix B, Table 4). Each content domain was used for each of the four inferential forms. So there were 24 arguments with a potential of 48 inferences to be made. These were presented in random order.

Procedure

Subjects were tested individually, with the items presented both in print and orally as before. Their responses were oral and were recorded by tape recorder as in Experiment 4. The instructions to subjects were also identical to those of Experiment 4. For subjects in the extended argument condition and the standard argument condition the procedure was also identical to that of Experiment 4. For subjects in the uncertainty elimination condition and interconnections elimination condition there were some slight procedural modifications. The first of these required the experimenter to inform subjects that their intermediate conclusion was correct. The second modification, for the interconnection elimination condition, involved the experimenter restating the subjects intermediate conclusion so that the subject could apply it to the second presented conditional.

Subjects

The 32 undergraduates from Trinity College, Dublin who participated in this experiment were randomly assigned to one of four conditions (n= 8). Two subjects were subsequently eliminated, one as a non-native English speaker, and one because experimental error had destroyed most of his responses. These eliminations were made prior to any data analysis. None of the subjects had received tuition in logic.

Scoring

The responses were scored as in the analysis of Experiment 4: definite endorsements of both of the invited inferences were given a score of 1, all other responses were scored 0.

Results and Discussion

The scored responses were analysed through a four-way analysis of variance. This analysis involved the factors of interconnection elimination and uncertainty elimination, as well as the factors of inference direction and inference form, outlined earlier. The analysis involved repeated measures on the last two factors and was weighted for unequal sample sizes. The response scores are displayed in Appendix A, Table 18, part i. The effects from this analysis are shown in Appendix C, Table 9. When subjects did not make both inferences, they made either the first one only (see Appendix A, Table 18, part ii), or they made the second inference alone (see Appendix A, Table 18, part iii). The remainder of their responses were primarily probabilistic, and so follow the same pattern as the definite responses.

Interconnections or uncertainty ?

Tukey HSD tests were conducted on the significant interaction of the two between-subject factors [F (1, 26) = 5.16, p = 0.0316]. These showed that more inferences are made in the interconnection elimination condition (73%) than in the extended argument condition (37.5%, $p < 0.05$). The frequency of inferences endorsed in the uncertainty elimination condition (51%) however, does not differ from those in the extended argument condition (37.5%). This shows that the possibility of constructing interconnections between the two conditional premises is what leads to fewer inferences being made from extended arguments, and not the uncertainty of using an intermediate conclusion.

The analysis also revealed a main effect of inference direction, signifying that that fewer inferences were made in the forward (modus ponens and denial of the antecedent inferences) direction (43%) than in the backward (affirmation of the consequent and modus tollens) direction [57%, F (1, 26) = 14.2, p = 0.0009]. The directionality of an inference has been identified as influential in previous research (e.g., Braine, 1978). A main effect of inference form was also found, indicating that fewer inferences from minor premises that are the same as the component referred to in the conditional premise (modus ponens and affirmation of the consequent) were made (44.5%) than from minor premises that are different (denial of the antecedent and modus tollens [56%, F (1, 26) = 12.76, p = 0.0014].

Extended arguments and standard arguments

It is surprising to note in this analysis however, that the frequency of inferences from the extended argument condition (37.5%) and from the standard argument condition (40%) does not differ. Is this a genuine failure to replicate the findings of Experiment 4, that fewer inferences are made on extended arguments of the "if she goes to the family gathering then she will not enjoy herself, if she enjoys herself then she will go to bed tired" form, then on comparable pairs of standard arguments ? To answer this, further statistical comparisons must be made.

Preliminary inspection of Table 18, in Appendix A, suggests that the same frequency of inferences is found in the extended arguments condition in this experiment as was found in the Experiment 4. However, it seems that inferences from the standard arguments in this experiment were also infrequent. Which is the representative case, a high frequency of inferences on these standard arguments, as in the Experiment 4 (see Table 16, row 2), or a low frequency from these, as in the present experiment (see Table 18, row 4) ? Fortunately this can be resolved.

Recall that subjects in the interconnections elimination condition received exactly the same materials as those in the standard argument condition for making their first inference. So there should be no differences between the frequency of first inferences in these conditions. That is, if the infrequency of inferences on the standard arguments is the representative case, then they should be equally infrequent on the interconnections elimination condition, when _first_ inferences are compared.

A similar analysis of variance to that conducted above was carried out on the responses made in the present experiment. Here however, the dependent measure was the definite endorsements of the invited _first_ inferences. That is, first inferences, that were the first of both inferences endorsed, in a pair of inferences were given a score of 1. So too were first inferences that were the only inference in a pair to be endorsed. Table 19, in Appendix A, displays the frequency of responses to the first invited inferences for the various conditions of this experiment. The effects of this analysis are shown in Appendix C, Table 10.

Tukey HSD tests were made on the interaction of the between-subjects factors [F (1, 26) = 5.16, p = 0.0316]. It was found that the frequency of first inferences made in the standard argument condition of Experiment 5 is significantly less (61%) than those made on the interconnections elimination condition (84.5%, $p < 0.05$). So this shows that the frequency of inferences in the standard argument condition of Experiment 5 is uncharacteristically low. On a best-out-of-three rationale, the results for standard arguments in Experiment 5 are the odd ones out. So we can conclude that it is not the case that there has been a genuine failure to replicate.

Interconnections

Thus we can conclude that it is the unified representation which leads to fewer inferences being endorsed on the extended argument of the form "if she goes to the family gathering then she will not enjoy herself, if she enjoys herself then she will go to bed tired" than on the counterpart standard arguments. It is not due to the uncertainty in the use of the intermediate conclusion.

We can now attempt to ascertain what it is in the representation of the extended arguments of this form that leads to fewer inferences being endorsed from it than from comparable standard arguments. And we can also address the issue of why this this representational difference has no effect for extended arguments of the other forms explored in Experiment 4. There are several unique aspects of the "if she goes to the family gathering then she will not enjoy herself, if she enjoys herself then she will go to bed tired" form, that distinguish it from other forms. In the next experiment the role of these unique features in affecting the frequency of inferences, will be explored, in an attempt to understand why fewer inferences are made on this form.

Experiment 6

There are at least three aspects of the extended argument of the form "if she goes to the family gathering then she will not enjoy herself, if she enjoys herself then she will go home tired" which differentiate it from either of the other two forms investigated so far. That is, there are features of arguments of this form which are not shared by arguments of the form "if she goes to the family gathering then she will enjoy herself, if she does not enjoy herself then she will go to bed tired", nor by arguments of the form "if she goes to the family gathering then she will enjoy herself, if she enjoys herself then she will go to bed tired".

One of these aspects is the presence of a negative in the first conditional premise. The second is the location of the negative at the consequent of the conditional. A third distinguishing feature is the atransitivity of the argument. What are the respective roles of these three features ? It is obvious from the results of Experiment 4 that the latter, atransitivity, by itself is not sufficient to lead to the differences observed. If it was, then similar differences would also be expected between extended arguments and standard arguments of the form "if she goes to the family gathering then she will enjoy herself, if she does not enjoy herself then she will go to bed tired". Recall that no such differences were found.

How can we examine whether the other two aspects are responsible for the differences found ? One can construct extended arguments to systematically isolate these factors. To ensure the presence of the first feature only, of there being a negative in the first conditional, an extended argument can be constructed which has a negative in the first conditional that is not located at the conse-

Instituto Superior de Psicologia Aplicada
BIBLIOTECA

quent, but at the antecedent instead. Such an argument will be transitive. An example of this argument form is "if she does not go to the family gathering then she will enjoy herself, if she enjoys herself then she will go to bed tired". Similarly, one can ensure the presence of both the first and the second features, of there being a negative in the first conditional and of that negative being located at the consequent of the conditional, but with the argument still remaining transitive. An example of such an argument form is: "if she goes to the family gathering then she will not enjoy herself, if she does not enjoy herself then she will go to bed tired". By comparing these extended arguments to their counterpart standard argument pairs, it will be possible to establish the role played by each of the three aspects unique to the form "if she goes to the family gathering then she will not enjoy herself, if she enjoys herself then she will go to bed tired", in leading to the effects found.

If the presence of a negative in the first conditional is enough to lead to the differences observed on this form, then it would be expected that each of the new forms will differ from their counterpart standard argument pairs, as both of them are in possession of this feature. If the location of the negative at the consequent of the first conditional is an essential property however, then differences between extended arguments and standard arguments would be expected on the "if she goes to the family gathering then she will not enjoy herself, if she does not enjoy herself then she will go to bed tired" form. At the same time, no differences would be expected between extended and standard arguments of the "if she does not go the family gathering then she will enjoy herself, if she enjoys herself then she will go to bed tired" form. If all three features, of the atransitivity of the argument, the negative in the first premise and the negative located at the consequent of the conditional, are all necessary for the differences observed to occur, then we would expect no differences between either of the extended arguments compared to their comparable counterpart standard arguments, as both are transitive. Since each of these three hypotheses are equally tenable, no expectations that the predictions from one will be met rather than the predictions of another, will be held.

Method

Design

The factorial design of this experiment had two between-subjects factors and two within-subjects factors. The first between-subjects factor was the argument length. Two groups of subjects received extended arguments and the other two groups received standard arguments. The second between-subject factor was the argument form, with one group of subjects receiving extended arguments of the form:
"if she does not go to the family gathering then she will enjoy herself, if she enjoys herself then she will go to bed tired",
while a second group of subjects received comparable, standard arguments of this form. A third group received arguments of the form:
"if she goes to the family gathering then she will not enjoy herself,

if she does not enjoy herself then she will go to bed tired",
while the fourth group were given comparable standard arguments of this form. The two within-subject factors, as in the preceding two experiments, were the direction of the inference, (either forwards or backwards), and the similarity of the minor premise to the component referred to in the conditional premise, (either the same form as the component referred to in the conditional, or a different form).

Materials

Four of the contents used in Experiment 5 were employed (see Appendix B, Table 4). These four contents were used for each of the four inferences. The resultant 16 arguments, with a potential of 32 inferences to be made, were presented in random order to each subject.

Procedure

The procedure was essentially the same as in the preceding two experiments except that the materials were presented in print form only, in A4 size booklets. The printed instructions on the cover of the booklet were a written version of the oral instructions given in Experiments 1 and 2. Subjects were instructed to record their responses in written form in the booklet.

It was necessary to have an analogous prompt to the "anything else?" query that was used in the preceding experiments, in cases where only one inference was made on the extended arguments. To this end, the presence of numbered spaces in which subjects were to record their responses, was included. Illustrations of this will clarify it. In the instructions the example responses for the extended arguments were numbered, as for instance

"you might think it follows that:

(1) she will go into town

(2) she will spend all her money"

Then for each item in the booklet, there were numbered blank spaces after each argument, for example:

If she goes to the family gathering then she will not enjoy herself

if she does not enjoy herself then she will go to bed tired

she went to the family gathering

What do you think follows from this ?

(1)

(2)

The subjects task was to write in these spaces what they considered to follow. Subjects were tested in groups of three to five.

Subjects

The 32 undergraduates from Trinity College, Dublin who took part in this experiment were randomly assigned to one of four conditions (n = 8). Two subjects were eliminated due to failure to comply with the requirements of the tasks. These eliminations were made prior to any data analysis. None of the subjects had received tuition in logic.

Scoring

Once again, the definite endorsements of both of the invited inferences were given a score of 1 while all other responses were scored 0.

Results and Discussion

A four-way analysis of variance with repeated measures on the third and fourth factors, weighted for unequal sample sizes, was conducted on the responses. The response scores are presented in Appendix A, Table 20, part i. Subjects responses, when they were not of both inferences, were of either the first alone (see Appendix A, Table 20, part ii) or they were of the second inference only (see Appendix A, Table 20, part iii). Otherwise, their responses were probabilistic, and so have the same pattern as the definite endorsements. The effects of the analysis are in Appendix C, Table 11. The analysis revealed only significant main effects of the direction of the inference [F (1, 26) = 4.31, p = 0.0479] and the similarity of the minor premise to the component referred to in the conditional premise [F (1, 26) = 25.53, p = 0.0000]. There was also an interaction of these variables [F (1, 26) = 4.51, p = 0.0433]. Tukey HSD tests on this latter interaction shows that there are more modus ponens inferences made (77%) overall than any of the other inferences ($p < 0.01$, for each comparison).

There were no differences between the extended arguments of different forms and their standard argument counterparts. For the form "if she does not go to the family gathering then she will enjoy herself, if she enjoys herself then she will go to bed tired", no significant differences were found between the extended arguments (60%) and the standard arguments (47%). This indicates that the presence of a negative in the first conditional is not enough to create differences between extended arguments and standard arguments. Furthermore no significant differences were found for the form "if she goes to the family gathering then she will not enjoy herself, if she does not enjoy herself then she will go to bed tired", between extended arguments (67%) and standard arguments (45.5%). This indicates that the location of the negative at the consequent of the first conditional is not enough to create differences between extended arguments and standard arguments either.

So the results of this experiment indicate that neither the presence of a negative in the first conditional, nor the location of that negative at the consequent of the first conditional is sufficient to lead to the differences observed between extended arguments of the form "if she goes to the family gathering then she will not enjoy herself, if she enjoys herself then she will go to bed tired" and their counterpart standard arguments. It is evident then, that all three properties -of the presence of a negative in the first premise, the location of the negative at the consequent of the first conditional, and the atransitivity of the argument- are necessary to create these differences.

Conclusion

This series of experiments has shown that extended arguments of certain forms ("if she goes to the family gathering then she will not enjoy herself, if she enjoys herself then she will go to bed tired") lead subjects to make fewer inferences than they make on comparable standard arguments. The source of this difference was located, in Experiment 5, in the representation of the premise information in extended arguments into unified representations, while the representation of standard arguments information is in two separate models. The properties of this form of extended argument, of the presence and location of the negative in the premises, and the atransitivity of the argument as a whole, affect this unified representation, to result in fewer inferences.

What we need now is a characterisation of the representations or models that people construct of extended arguments, that will reveal how these three properties lead to fewer inferences being made from certain extended arguments than from their counterpart standard arguments. A model of the way interconnections may be built up in the representations people construct to understand arguments, consistent with the results of the preceding experiments, will be presented in the next Chapter.

CHAPTER 6:

INTERCONNECTIONS

Modelling representational interconnections

Why are few inferences made on extended arguments of the form "if she goes to the family gathering then she will not enjoy herself, if she enjoys herself then she will go to bed tired" ? The experimental evidence has indicated that the source of the observed differences between extended arguments of this form and comparable standard argument pairs lies in the unified representation constructed of the information in extended arguments. The events or actions referred to in both of the conditional premises of an extended argument all may be represented into a unified representation. It has been suggested notionally that in these unified representations interconnections may be established between the events or actions. Furthermore, novel interconnections may be forged between the first and last mentioned events (that is, between "going to the family gathering" and "going to bed tired"). This interconnecting relation may be <u>contrary</u> to the relation that can be deduced by the chain of inferences from this argument.

A concrete example of a contrary relation will clarify its nature, before we discuss how such a relation could be constructed. From the premises: "if she goes to the family gathering then she will not enjoy herself, if she enjoys herself then she will go to bed tired", a representation or model will be constructed of the relation between "going to the family gathering" and "not enjoying herself". The connection of "not enjoying herself" leading to "going to bed tired" will also be represented within this initial representation. It is possible then that the relation of "going to the family gathering", itself leading to "going to bed tired", may be established, for reasons that will become apparent in the following sections. This relation runs contrary to the inferential chain from "going to the family gathering" to "not going to bed tired" (via the modus ponens inference from "going to the family gathering" to "not enjoying herself", and the denial of the antecedent inference from "not enjoying herself" to "not going to bed tired").

If a contrary relation like this is established, then the chain of inferences will be suppressed, since such a relation would constitute an alternative consequence to that from which the inference is suggested. But precisely how would such contrary interconnections come to be established within the representation or model constructed of the extended argument ? In the following sections a detailed model of the procedures underlying the construction of the representations is elaborated.

Representational assumptions

This model makes some simple assumptions about the representation of information, and about the procedures that operate on these representations. It is not intended to provide an account of people's understanding of conditionals. Instead, its aim is to explain how people create connections between the information they represent. How do people represent a negative event, such as "she did not go to the family gathering" ? Is "not going to the family gathering" something that one actively does, or is it to be understood only in relation to other things that one did instead ? More precisely, how do people represent a negative event as being the condition for something else to happen, or the consequence of something else happening ? Many previous theories of the representation of relations in which there are negative terms have relied on diagrammatic isomorphs, such as Euler circles (e.g. Adams, 1980). Aside from formal difficulties with these (c.f. Johnson-Laird, 1983) they have the disadvantage of being both abstract and ill-specified. Since these two characteristics are arguably antithetical to human reasoning, as we have also seen in preceding Chapters of this thesis, alternative representational descriptions have been offered.

One such alternative is the direct representation of the negated term, as in the componential isomorphs offered by Guyote and Sternberg (1981). However, the assumption that people represent a negative event (e.g., "not going to the family gathering") simply as that event negated, or as a non-event, is transparently problematic. Instead it could be argued that negative events are represented in terms of other events (e.g., "going to a film instead"). The problem with this latter suggestion however, is that the number of events that qualify as being members of the class of "not going to the family gathering" is so large as to be indeterminate (e.g., Staudenmayer and Bourne, 1978). Furthermore, unless actual alternative instances to "going to the family gathering" have been specified, they may not be considered relevant (e.g., Evans, 1983a). A third possibility is to represent relations that involve a negative term through partitioning the terms mentioned from each other (Johnson-Laird, 1983). So the relation between, for example, "not going to the family gathering and "enjoying herself" will be represented by these events segregated from each other:

She goes to the family
gathering

She enjoys herself

In this way, one is not committed to conceiving of non_events, nor to actual alternative instances. Instead one simply conceives of the two mentioned events as being mutually exclusive. This representational proposal will be adopted here.

Representational construction procedures

Granted that people use this kind of representation of negatives, how do they go on to build contrary interconnections ? It is proposed that there are two representation construction procedures that deal with forging interconnections between the events mentioned in the arguments (see Figure 4).

Figure 4 A model of the representation construction procedures

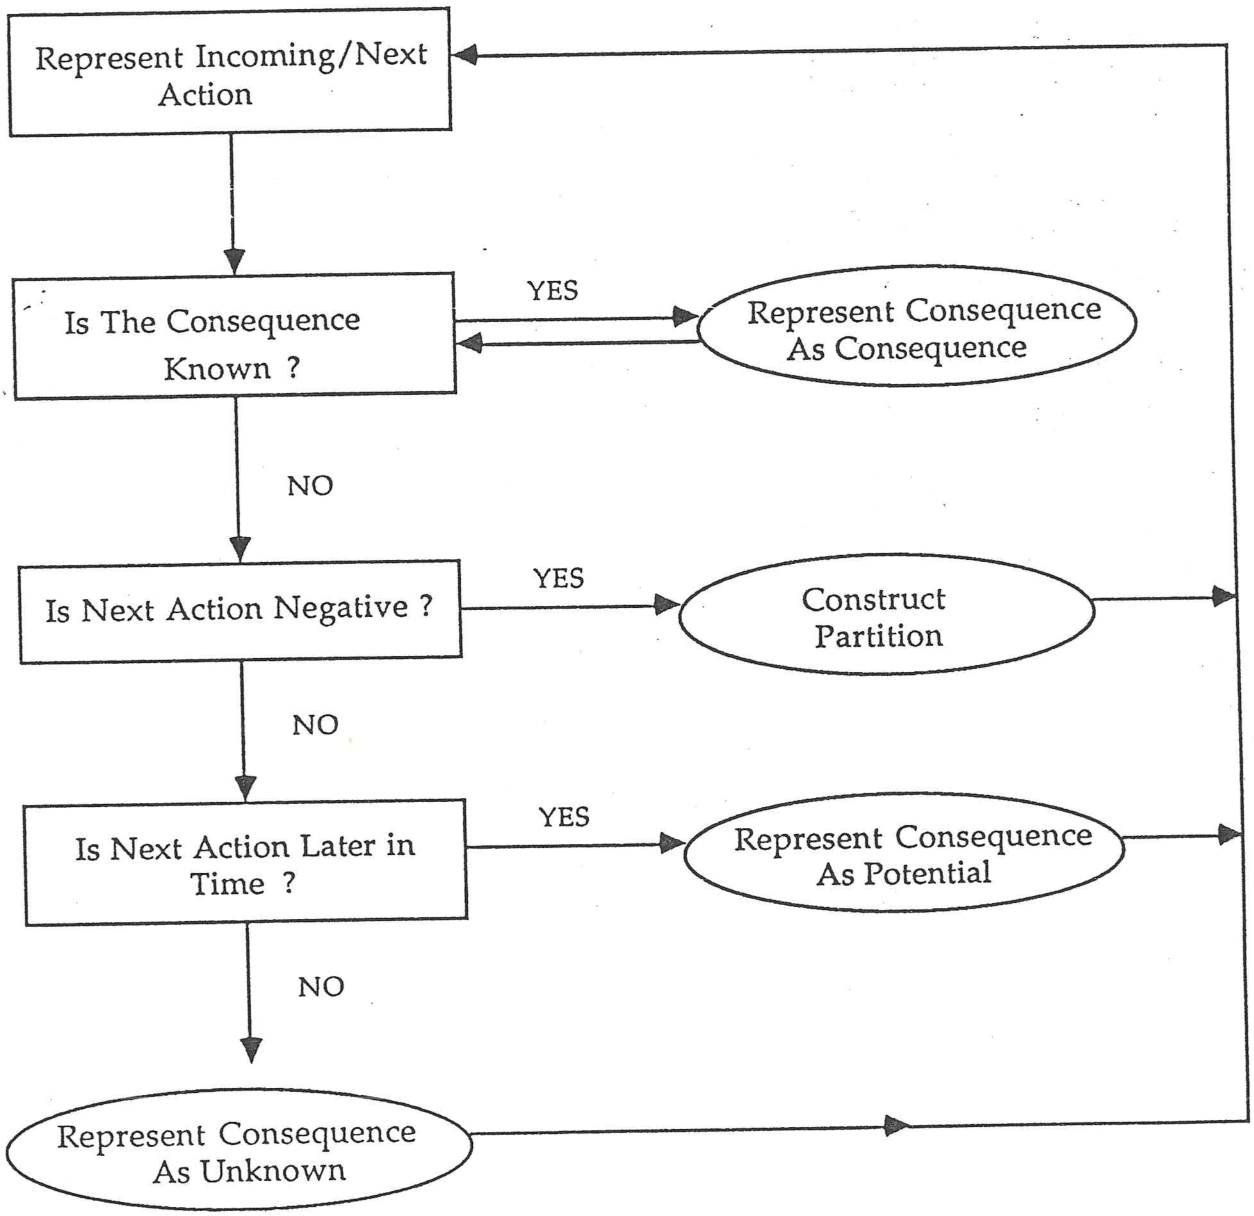

These representational construction procedures are consistent with the more general procedures advocated for the construction of mental models (Johnson-Laird, 1983, ch. 11). Their existence may be attributed to working memory restrictions which give rise to heuristics that will process incoming information into a model in a swift and efficient manner. These heuristics may not always lead to strictly accurate models, but overall, will ensure that information becomes encoded into the representation in a sensible way. As they operate in tandem, they will be described together.

The first representational construction procedure emerges from the directional nature of temporal information (2). It seems plausible that when people are attempting to take in a large amount of information, which is sequentially arranged by its temporal order of occurrence, they will be concerned with establishing what happens next, or with what follows from what. They will want to be clear about what comes next at each juncture. It is proposed that this first representational construction procedure processes temporally_sequential information by attempting to establish what <u>follows</u> from, or what are the <u>consequences</u> of each mentioned action or event.

---

(2) The influence of temporality on comprehension and reasoning has received support for planning in general (McDermott, 1982) and for conditional sentences of different forms in particular (e.g., Evans and Beck, 1981; Evans and Newstead, 1977; Rips and Marcus, 1977.

For example, in processing the information in the premise: "If she goes to the family gathering then she will enjoy herself" the first representational construction procedure examines each incoming mentioned action to ascertain whether or not its consequence is known. That is, the first action, going to the family gathering, is represented and the first representational construction procedure checks whether or not it's consequence is known. This may be represented graphically as:

going to the family gathering ?

This procedure inspects the forthcoming information for the consequence of "going to the family gathering". As it <u>is</u> known, the consequence is represented appropriately:

going to the family gathering --- enjoying herself

and the first representational construction procedure is applied again:

going to the family gathering --- enjoying herself --- ?

When an action examined by the first representational construction procedure has no explicitly given consequence, the second representational construction procedure comes into operation. As mentioned, it seems plausible that people try to ascertain what follows from what in a series of information. When they have not been given explicit information about what comes next in a sequence, it is likely that they will try to establish what <u>might</u> follow. The second representational construction procedure attempts to establish what the consequence of a given action <u>could</u> possibly be. In this search for a potential candidate, the operation of the second representational construction procedure is limited to the information that has been given. We have seen in the first series of experiments, reported in Chapters 2 and 3, that people frequently restrict the information they consider when they are reasoning to that which they have been given. So the second representational construction procedure examines all the <u>other</u> information that has been given, both that which has already been processed, and that which is still forthcoming, for a potential candidate consequence.

Potential candidates will be actions or events which satisfy the two constraints of this procedure. The first of these constraints is that to be a potential candidate consequence, an action or event must occur later in time than the action of which it is a potential consequence. Secondly, to be a potential candidate consequence, an action must not have been specifically partitioned from the action of which it is a potential consequence. Is there a potential consequence for "enjoying herself" in the example above ? The first action "going to the family gathering" violates the temporal constraint that candidate consequences must occur later in time. So the consequence of "enjoying herself" remains unknown in the example above.

Applying the procedures to the proposed representation

What happens when the representational assumption about negative events is combined with the representational construction procedures ? The combined effects are made clear in the following two examples. For the premise:
"If she goes to the family gathering then she will not enjoy herself" the assumption of the partitioned representation of negative events ensures that this will be represented as:

going to the family
gathering

---

enjoying herself

The attempt by the first representational construction procedure to ascertain the consequences of the <u>mentioned</u> actions renders this as:

going to the family
gathering --- ?

---

enjoying herself

The operation of the second representational construction procedure, in attempting to ascertain potential candidate consequences, will be unsuccessful. "Enjoying herself" cannot be a potential consequence as it is specifically partitioned from "going to the family gathering". Note that the first representational construction procedure is not applied to "enjoying herself", as it is applied only to events explicitly mentioned, not to ones that were presented as negative events.

Let us examine now the representation of the premise:
"If she does not go to the family gathering then she will enjoy herself"
This will also result in an initial representation as:

going to the family
gathering

---

enjoying herself

Representing negative events as partitioned from each other leads to this similarity in the initial representation of the forms "if she goes to the family gathering then she will not enjoy herself" and "if she does not go to the family gathering then she will enjoy herself".

However the similarity in initial representation is quickly lost when the representation construction procedures go to work. The attempt by the first representational construction procedure, to ascertain the consequences of the <u>mentioned</u> actions will lead this time to:

going to the family
gathering

———————————

enjoying herself --- ?

(and once more the second representational construction procedure will be unsuccessful). So the operation of the first representational construction procedure in examining only the <u>mentioned</u> actions to establish their consequences, ensures that different consequences are recognised as being unknown, for these two forms. The representational differences between these different forms does not have any effect on the inferences made on standard arguments of course, and so remain undetected.

It is in extended arguments that the way in which negative events are represented, and the operation of the representational construction procedures, begin to effect the frequency of inferences endorsed. Take the premises:

"if she goes to the family gathering then she will not enjoy herself,
if she enjoys herself then she will go to bed tired"

The negative representation assumption, combined with the operation of the first representational construction procedure, as in the examples above, will yield:

going to the family gathering --- ?

---

enjoying herself --- going to bed tired --- ?

The operation of the second representational construction procedure, in locating a candidate consequence of the first action, "going to the family gathering" will yield the last mentioned action, "going to bed tired" as a potential candidate. "Going to bed tired" has not been specifically partitioned from "going to the family gathering" and it is in the appropriate temporal order, so it satisfies the two constraints of the second representational procedure. So this possibility becomes represented as:

going to the family gathering --- ? going to bed tired ?

---

not enjoying herself --- going to bed tired --- ?

So here we see the operation of the representational construction procedures leading to the establishment of novel contrary relations.

<u>Contrary relations and inference suppression</u>

The contrary relation is <u>contrary</u> because it contradicts the assertion of both "going to the family gathering" and "not going to bed tired". This assertion would come about through chains of inferences of either the modus ponens inference followed by the denial of the antecedent inference, or of the modus tollens inference followed by the affirmation of the consequent inference. This relation also runs contrary to the joint assertion of "not going to the family gathering" and "going to bed tired". This assertion would result from the inferential chains of the denial of the antecedent and the modus ponens inference, or of the affirmation of the consequent followed by the modus tollens inference.

How does this contrary relation suppress the inferences ? The answer is : in exactly the same way as we have seen in Chapter 4. After a model of the information in the premises has been constructed, putative conclusions "fall out". If there is other information available, which the person realises is relevant, and can retrieve, then those conclusions become suppressed. So in the above model, the modus ponens inference from "she went to the family gathering" to "she did not enjoy herself" is suggested. However, the alternative consequence, that "she went to bed tired" is also represented. So this inference becomes dubious. What is more, in this chain, the denial of the antecedent inference, from "she did not enjoy herself" leads to "she did not go to bed tired". But this <u>chain</u> from "going to the family gathering" to "not enjoying herself", to "not going to bed tired" has had one of its links weakened by the alternative consequence. So while people may make one or other of the inferences in this chain, they will not make <u>both</u>. Making both of these inferences, and so completing this chain of inferences would contradict the interconnections built into the model. It is in this way that interconnections between the other actions that one thinks about in a planning sequence may lead to inferences being suppressed.

It should also be noted that non_contrary relations can be established between the first and last mentioned actions of the conditional premises in an extended argument. For example, for the extended argument:

"If she goes to the family gathering then she will not enjoy herself,
if she does not enjoy herself then she will go to bed tired"
construction procedure will yield the representation:

going to the family
gathering --- ?

---

enjoying herself

---

going to bed tired --- ?

The second representational construction procedure may propose "going to bed tired" as a potential candidate consequence of "going to the family gathering", as these have not been specifically partitioned from each other, and "going to bed tired" is later in time than "going to the family gathering". This will lead to the representation:

going to the family gathering --- ? going to bed tired ?

---

enjoying herself

---

going to bed tired --- ?

This relation between "going to the family gathering" and "going to bed tired" is not a contrary relation. It is the exact same relation that will result from the inferential chains in these <u>transitive</u> premises, as working through them will reveal. Hence, none of the inferences will be suppressed by the establishment of this relation.

<u>Consistency with the evidence</u>

Five forms of extended arguments have been compared to their counterpart standard arguments in the course of the preceding experiments. For one of these forms, "if she goes to the family gathering then she will not enjoy herself, if she enjoys herself then she will go to bed tired", the frequency of inferences made from the extended arguments was significantly less than the frequency of inferences made from the comparable, counterpart standard arguments. For the other four forms, there were no significant differences in the frequency of inferences endorsed on both extended arguments and on standard arguments. Does the proposed model predict such results for these forms ?

The representation of two of the forms has already been outlined in illustrating the model:

(1) For the first extended argument outlined above, "if she goes to the family gathering then she will not enjoy herself, if she enjoys herself then she will go to bed tired", which was examined in both Experiment 4 and Experiment 5, it was found that significantly fewer inferences are endorsed on the extended arguments of this form than on the counterpart standard arguments. As is evident from the outline above, the model predicts exactly this result.

(2) The second extended argument dealt with above: "if she goes to the family gathering then she will not enjoy herself, if she does not enjoy herself then she will go to bed tired" was found, in Experiment 6, to yield the same frequency of inferences as its counterpart standard arguments. As the outline above reveals, the model also predicts this result. It was demonstrated that despite the presence of a negative in the first conditional and the location of that negative at the consequent, the transitivity of the argument ensures no contrary relations are formed, and so no differences are expected. What of the other three forms used in the experiments ?

(3) For the extended argument of the form "if she does not go to the family gathering then she will enjoy herself, if she enjoys herself then she will go to bed early" the representation will be:

going to the family
gathering

---

enjoying herself --- going to bed tired --- ?

So no contrary relations are established. The negated <u>antecedent</u> leads to the first representational construction procedure <u>not</u> being applied to look for a consequence of "going to the family gathering". The model leads to the prediction that no differences between the extended arguments and standard arguments of this form are expected, and this is the finding of Experiment 6 for this form.

(4) For the extended argument of the form "if she goes to the family gathering then she will enjoy herself, if she does not enjoy herself then she will go to bed tired", the representation will be:

going to the family gathering --- enjoying herself --- ?

---

going to bed tired --- ?

and so no contrary relations are established for this form. The location of the negative in the second conditional ensures that the consequence of "going to the family gathering" is known from the outset. So no differences between this form and its counterpart standard arguments are expected, and none are found in Experiment 4.

(5) Finally, extended arguments of the form "if she goes to the family gathering then she will enjoy herself, if she enjoys herself then she will go to bed tired " will be represented as:

going to the family gathering ---	enjoying herself ---	going to bed tired --- ?

Again no relations, contrary or otherwise, will be established and so no differences are expected. The finding of Experiment 4 for this form was that there are no significant differences between extended arguments and standard arguments. Hence the proposed model is consistent with the results of all five forms explored in the experiments here.

Conclusion

This second series of experiments has shown that the frequency of inferences made in the standard arguments used in past research is not always generalisable to the frequency of inferences made in the extended arguments of everyday reasoning. This is because the model built of the premises of extended arguments may lead to relations being established between all of the actions and events. Such interconnections are impossible to achieve in the separate representations of a pair of counterpart standard arguments. These interconnections facilitate the operation of the processes modelled in Chapter 4. That is, they facilitate the process of recognising the relevance of other information and retrieving specific instances of that information. So the other things that a person may be thinking about, in, for example, planning, can lead to their model of what they are reasoning about being an extended one. This is one instance where inferences may become suppressed in a naturalistic way.

CHAPTER 7:

THE IMPLICATIONS OF THE CONTEXTUAL NATURE OF CONDITIONAL REASONING.

Discussion

The thrust of the experimental work reported here has been to demonstrate that to understand everyday conditional reasoning one must discover how the inferences a person makes are constrained or are warranted by everything else the person takes into account. Human thinking rarely occurs in a vacuum but is more often embedded firmly in a specific context. This context may be provided externally by the situation a person is in, or internally by the other things a person is thinking about. The recurring emphasis of these experiments has been that the other information a person considers in conditional reasoning is affected by the situation and the other things they are thinking about.

It has been demonstrated that the pattern of inferences people make can be manipulated by altering the context in which those inferences are made. This manipulation can be achieved explicitly, by presenting the specific aspects of the context that are relevant, such as alternative and additional conditions. Alternatively, in a more implicit manner, general characteristics of the context can be sketched, for example the kind of situation and its duration. An even more subtle manipulation than this manipulation of the environmental context, can be achieved through simulating the internal or cognitive context, within which a person's reasoning is naturally embedded, in, for example, planning events.

It has been shown that the inferences people make vary systematically according to the specific type of context in which they are embedded. The experiments in Chapters 2 and 3 have shown that altering the environmental context of the antecedent of a conditional alters the inferences made. So accompanying the conditional with other conditionals or with general information that provided additional antecedents to the initial conditional suppresses the valid modus ponens and tollens inferences and warrants the invalid denial of the antecedent and affirmation of the consequent inferences. In contrast, placing the same conditional in other contexts, reverses this pattern. suppressing the invalid inferences and warranting the valid ones. Furthermore, the experiments in Chapter 5 showed that alterations of the cognitive context can also affect the inferences made. Manipulating how much information a person represents in their mental representation or model of the situation described by the premises, also alters the inferences made. Certain interconnections built up in the model can suppress the inferences in a complete chain of inferences.

These findings, that both the cognitive and environmental context affect the inferences people make, have empirical and theoretical consequences.

Empirical consequences

The main empirical consequence concerns the use of standard arguments as a primary research tool in studying reasoning. In everyday situations people rarely reason from a single premise in isolation. As I have argued, to understand everyday reasoning, it is necessary to understand the complex, interconnected models people build of the information. Such models cannot be reliably captured through the use of standard arguments. Indeed, the models resulting from such standard arguments are impoverished in comparison. So the use of standard arguments as an investigative paradigm is of dubious value, producing results which, at worse may be artifactual, and at best, limited.

Theoretical consequences

The theoretical consequences concern different characterisations of the processes people use in reasoning. The experimental evidence suggests these processes are sensitive to context.

Implications for theories of reasoning.

One needs an account of reasoning processes which can capture the influence of environmental context and cognitive context within a rich conceptual model.

Mental logic. Mental logicians (e.g., Braine, Reiser and Rumain, 1984; Rips, 1983b) attempt to characterise people's semantic and pragmatic knowledge as a mere precursory influence along the way to an ultimately syntactic representation. By doing so, they emphasise syntax as the underlying basis of reasoning. From the evidence presented here however, syntax is arguably less important than semantic or pragmatic knowledge (although it is at present the easiest to model computationally).

As has been noted already, mental logicians are caught on one or other of the horns of a dilemma, by the demonstration here of the influence of context on both valid and invalid inferences in Experiment 1. They must either admit that valid inferences are merely invited inferences, or else admit instead that invalid inferences are necessary. The demonstration in Experiment 3, that simply by prefacing conditionals with situational and durational information can also suppress these inferences, is even more problematic for the mental logic view. Here the suppression effect cannot be explained away by claiming that reasoning from an amended rule (such as "if it is raining and she goes out for a walk then she will get wet") is governed by different principles than from reasoning is a straightforward rule (such as "if it is raining then she will get wet").

The position of the mental logicians is also challenged by the overall results of Experiments 2 and 3. These indicate that simply understanding the relevance of other information does not suppress the inferences. Instead, specific instances of relevant information must be retrieved from long-term memory. The mental logic view cannot stretch to encompass this evidence easily.

Finally, the natural logic models cannot account for the effects of extended arguments on inference making, demonstrated in Chapter 5. There is no clear reason, in the mental logic account, why inference rules should not be applied iteratively as a straightforward extension of their individual application. So differences between extended and standard arguments are inexplicable.

Memory-based accounts. The memory-based accounts of reasoning (Griggs, 1984; Pollard, 1982) fare equally badly. While they are compatible with the suppression of the inferences found in Experiment 1, the findings of Experiment 2 and 3 indicate that retrieval is not responsible by itself for this suppression.

Furthermore, there is nothing in the memory-based reasoning position to lead to an expectation that extended arguments warrant fewer inferences. That is, extended arguments and standard arguments will both lead to the recall of the same kinds of information, so no differences between them would be expected.

Pragmatic schemas. The pragmatic schemas theory (Cheng and Holyoak, 1985) also cannot account for the experimental findings. As mentioned earlier these pragmatic schemas prove to be cumbersome and unwieldy. For example, a sentence like 'the gunman said to the bank-clerk "if you shout then I'll shoot you"' will be mapped to a pragmatic schema embodying such rules as "if you do the action then I will bring about the outcome" and "if you do not do the action then I will not bring about the outcome". Is this sentence mapped to the same schema when it is prefaced by "during the two week siege..." ? Apparently not. To account for the results here the pragmatic reasoning theory must hold that prefacing this sentence causes it to be mapped to a different schema. This latter schema will have modals in the rules, such as "if you do not do the action then I may not bring about the outcome". When the sentence is prefaced by "during the 15 minute siege..." it must be mapped to a different schema from either of these two. This latter will consist of rules with modals positioned in, say the "if you shout then I may shoot you" rule, rather than the rule above. In other words, there must be considerable repetition of the basic schema, with different schemas differing only in the presence, and positioning of modals in the rules. This is clearly uneconomical, as such rules swiftly proliferate.

A second problem is that the pragmatic schemas approach cannot account for the suppression of the inferences in extended arguments. As each of the premises will be mapped to an appropriate schema, it should not make a difference whether or not premises are presented together in an extended argument or are presented separately in a standard argument. Amalgamations of pragmatic schemas, or rules in pragmatic schemas, would not alter the inferences warranted from the rules.

Models of everyday reasoning

One can however, capture the semantic and pragmatic knowledge people bring to bear on their inference-making through the machinery of the theory of the construction of mental representations of information (e.g., Johnson-Laird, 1983). Two models were proposed earlier which constitute specific instantiations of aspects of such a theory.

The first model is an instantiation of a semantic procedure for testing inferences. This procedure requires only an understanding that a conclusion is valid if there are no counterexamples to it. It relies upon a person having constructed a model of the information given to them. They can then systematically modify that representation in order to consider related possibilities.

As modelled here, the processes in this procedure are general ones called from more basic processes such as a person's conversational comprehension processes, and their retrieval processes. So it is a model which suggests how people's basic cognitive processes may be configured and reconfigured to deal with particular situations. In this way, conversational understanding and retrieval come together to make a powerful reasoning process.

The second model proposed here specifies some of the procedures used to construct representations in the first place. The procedures specified are born out of the need to cope with temporally arranged incoming information with a limited capacity working-memory. These limitations result in the use of "quick and dirty" heuristics to deal with incoming information in a uniform manner. These heuristics result in novel interconnections being constructed within the model of the information that the person is building.

Combining these two models leads to a distinctive view of reasoning. This can be summarised as follows :

* People construct models of the information given (Johnson-Laird,

1983). They may use their general knowledge to do this. Inferential conclusions will be suggested by the model of the information that people have constructed. The validity of a conclusion depends on there being no counterexamples to it.

* In the construction of a model, connections are made between the pieces of information about which people are reasoning.

* Some of the interconnections that are formed may constitute alternative or additional conditions for one of the consequences, or alternative or additional consequences following from one of the conditions. Such alternative or additional conditions or consequences may arise because the other information people have been given, or have retrieved, constitutes such a conditional or consequence. This is what happens when a person is given, for example, premises with alternative antecedents "if it is raining then she will get wet, if it is snowing then she will get wet". Or it may arise from the nature of the heuristics responsible for creating the interconnections in certain kinds of information, such as temporal information. These heuristics may lead to interconnections between pieces of information which are "contrary" and constitute, say, alternative consequences.

* If such a counterexample link has been made in the model, people are aware of its relevance to the invited inference. The presence of this information in the model ensures that the information is, or has been, recognised as being relevant.

* They can not only appreciate the general class of information that

is relevant, but they have already "retrieved" an instance of the appropriate class of information. They may already have retrieved the information from long-term memory, or they may have been given it.

* As their model has been coherently modified to fit alternative possibilities, the initial invited inference will be suppressed.

Further fields

Together the two models proposed outline in a detailed and precise way how people build models of the information given, and how they make inferences from that information. Both models have been shown in preceding Chapters to be consistent with both the spirit and the letter of the data. Do they appear as strong from a broader perspective ? That is, have these modelling efforts any import for cognitive psychology generally ?

Implications for cognitive processes

These models have been based on general cognitive processes of retrieval from long-term memory and of language understanding, and on principles of working-memory limitations. In these models, comprehension, retrieval and working-memory restrictions have been assigned radically different roles from those assigned previously in discussions of reasoning. A brief look at the relation of reasoning to each of these processes in turn will illustrate this.

Reasoning and language understanding

Conversational implicature has been characterised here as crucial to the way people reason. It effects whether the model people have built up of the information they have been given will be modified or not. In the past, comprehension processes have been regarded merely as embodied in a component operating prior to reasoning proper (e.g., Braine and Rumain, 1983). Frequently such processes have been considered to lead to erroneous representations of the information. It has already been noted in this discussion that attempting to partition semantic understanding from syntactic representation, and attributing more import to the latter, has neither theoretical nor empirical support. Furthermore, it is not considered here that people misunderstand the information given. Instead, they understand it according to the pragmatics of the situation. And in accordance with this, they will or will not attempt to modify the model they have constructed. Hence, language comprehension is an intrinsic part of the reasoning process.

Reasoning and memory

Reasoner's use of their own knowledge in reasoning has also been accorded a high priority in the suggested models of reasoning processes. What a person has in memory can make or break an inference. Unless people are able to use an abstract token to stand for an instance -and it is not plausible that this is a usual mode of reasoning- then they are reliant upon what is in memory and upon being able to retrieve it. Retrieval processes are thus a crucial part of reasoning. This differs from the mental logician's view that memory affects only comprehension and not reasoning proper. However, it also differs from the memory-based reasoning position, that the existence and retrieval of instances from memory is solely responsible for reasoning prowess. Instead retrieval has been awarded the less potent status of being firmly tied to an awareness of the relevance of particular kinds of information.

A further innovation is the proposal that memory organisation must be amenable to easy access in order that the inferences should become suppressed. The suggestion made earlier, that information in memory is organised within superordinate categories, is not new (e.g., Rosch, 1978; Schank, 1982). However, that this organisation differs for different categories of information, alternatives and additional information, is novel, as is the consequent entailment that the retrieval of these kinds of information may be achieved with differential ease. Thus different inferences can be suppressed with different degrees of success.

Reasoning and working-memory

Errors in reasoning have frequently been attributed to limitations of working-memory capacity (e.g., Braine, Reiser and Rumain, 1984; Guyote and Sternberg, 1981; Johnson-Laird and Steedman, 1978; Rips, 1983b). Often these limitations have been assumed to lead to omissions or failures to represent the information or to consider some of the possibilities. Here, however a different entailment of limited working-memory has been explored. This is that restrictions imposed by limited capacity lead to the development of procedures for constructing representations, which are heuristic rather than algorithmic. These procedures will sometimes lead to the construction of models that embody novel interconnections between pieces of information. These novel interconnections may be contrary to the connections that can be traced through inferential steps. Consequently some inferences are suppressed as a result of the procedures for the representation of information that arises from working-memory limitations, and so reasoning "errors" may occur.

Cognitive processes

The proposal that reasoning is derived from configurations of comprehension processes, retrieval processes and working-memory processes is distinct from previous studies. Previously comprehension processes and working-memory limitations were regarded as extraneous to, or interferences in "true reasoning". However, I have argued that these processes are intrinsic parts of reasoning. This moves away from the fragmentation of previous research towards a more integrated synthesis of different research areas. Reasoning marshals together various sub-processes, but is itself a process that can be called up in the pursuit of further aims. For example, the way in which we have conceived of reasoning here allows us to view it as a subprocedure of problem solving. Let us look at the relation between reasoning and- what it has become fashionable to call it instead- problem solving.

<u>Reasoning and problem solving</u>

The relation between reasoning and problem solving has always been a matter of speculation. Early students of problem solving (e.g., Maier, 1931) considered that they were studying reasoning. Recent students of reasoning have suggested that they are studying problem solving (Evans, 1982). These claims are not based merely on superficial similarities between tasks given to subjects in investigating these areas. Instead, they rest on the resemblance of the processes that underlie each. Within the framework adopted here these similarities may be sketched thus. When either reasoning or solving problems people will build models, where heuristics guide their representation of the information with which they must deal. With the greater information available when given a problem to solve, more interconnections can be built up. With a greater number of interconnections established within the model, more inferences will be suggested. The procedure for testing inferences will be implemented frequently and will result in various inferences being suppressed. Those that are not suppressed may become the input to subsequent representations of the problem. In this way, we can conceive of problem solving as a more "molar" activity than reasoning. That is, the "molecular" processes of reasoning are applied repeatedly in problem solving. Just as reasoning processes call on language understanding, retrieval and working-memory processes, so too problem solving calls on reasoning processes.

Conclusion

Rationality

Does the combination of processes from language understanding, memory retrieval and working-memory heuristics, within the procedures presented here, ensure rationality ? Or does this mixture buy abundant flexibility at the expense of reasonability ? The view of rationality that results from the theory advanced here is that people have the power to build models to understand the information they are dealing with. They can make connections between this information that were not previously explicit. They can bring other information to bear on this information, either from memory or from the world, or from both. They can modify their initial model to take this new information into account. These abilities ensure that people do have the power to reach reasonable conclusions. Whether they do or not depends on the context in which they are reasoning.

APPENDICES

APPENDIX A:

EXPERIMENTAL DATA

Appendix A

Table 1

The pattern of suppression of inferences expected (on the grounds of plausibility) by the provision of contextual information

major premise	"If it is raining then she will get wet"			
	MP	MT	DA	AC
minor premise	"raining"	"not wet"	"not raining"	"wet"
conclusion	"wet"	"not raining"	"not wet"	"raining"
Alternative antecedent e.g, snow	not suppressed	not suppressed	suppressed	suppressed
Additional antecedent e.g., out for a walk	suppressed	suppressed	not suppressed	not suppressed

Appendix A

Table 2

Patterns of suppression (plausibly) expected for the disjunction of alternative and additional conditions in the minor premise accompanying two conditionals

Disjunction of alternatives

major premises: "If it is raining then she will get wet if it is snowing then she will get wet"

	MP	MT	DA	AC
minor premise	"raining or snowing"	"not wet"	"not raining or not snowing"	"wet"
conclusion	"wet"	"not raining or not snowing"	"not wet"	"raining or snowing"
expected nature of response	not suppressed	not suppressed	not suppressed	not suppressed

Disjunction of additionals

major premises: "If it is raining then she will get wet if she goes out for a walk then she will get wet"

	MP	MT	DA	AC
minor premise	"raining or out for walk"	"not wet"	"not raining or not out for walk"	"wet"
conclusion	"wet"	"not raining or not out for walk"	"not wet"	"raining or out for walk"

No expectations as to nature of response

Table 3

Patterns of suppression plausibly expected for the conjunction of additionals and alternatives in the minor premises accompanying two conditionals.

Conjunction of additionals

major premises: "If it is raining then she will get wet if she goes out for a walk then she will get wet"

	MP	MT	DA	AC
minor premise	"raining and out for walk"	"not wet"	"not raining and not out for walk"	"wet"
conclusion	"wet"	"not raining and not out for walk"	"not wet"	"raining and out for walk"
expected nature of response	not suppressed	no expectations as to nature of response	not suppressed	not suppressed

Conjunction of alternatives

major premises: "If it is raining then she will get wet if it is snowing then she will get wet"

	MP	MT	DA	AC
minor premise	"raining and snowing"	"not wet"	"not raining and not snowing"	"wet"
conclusion	"wet"	"not raining and not snowing"	"not wet"	"raining and snowing"

no expectations as to nature of response

Appendix A

Table 4

Patterns of suppression expected for the disjunction of alternative and additional conditions in the minor premises alone, when no extra conditional is given

Disjunction of alternatives

major premises "If it is raining then she will get wet"

	MP	MT	DA	AC
minor premise	"raining or snowing"	"not wet"	"not raining or not snowing"	"wet"
conclusion	"wet"	"not raining or not snowing"	"not wet"	"raining or snowing"
expected nature of response	not suppressed	not suppressed	not suppressed	not suppressed

Disjunction of additionals

major premises "If it is raining then she will get wet"

	MP	MT	DA	AC
minor premise	"raining or out for walk"	"not wet"	"not raining or not out for walk"	"wet"
conclusion	"wet"	"not raining or not out for walk"	"not wet"	"raining or out for walk"

no expectations as to nature of response

Appendix A

Table 5

Patterns of suppression expected for the conjunction of additionals and alternatives in the minor premises, when no extra conditional information is given.

Conjunction of additionals

major premises	"If it is raining then she will get wet"			
	MP	MT	DA	AC
minor premise	"raining and out for walk"	"not wet"	"not raining and not out for walk"	"wet"
conclusion	"wet"	"not raining and not out for walk"	"not wet"	"raining and out for walk"
expected nature of response	not suppressed	no expect-ations as to nature of response	not suppressed	not suppressed

Conjunction of alternatives

major premises	"If it is raining then she will get wet"			
	MP	MT	DA	AC
minor premise	"raining and snowing"	"not wet"	"not raining and not snowing"	"wet"
conclusion	"wet"	"not raining and not snowing"	"not wet"	"raining and snowing"

no expectations as to nature of response

Appendix A

Table 6

The expected patterns of suppression and non-suppression under the conditions of Experiment 1.

Information type	Presentation	Minor premise and inference			
		Categorical			
		MP	MT	DA	AC
Alternatives	explicit	n+	n+	S+	S+
	implicit	n+	n+	n+	n+
Additionals	explicit	S+	S+	n+	n+
	implicit	n+	n+	n+	n+
		Conjoined			
Alternatives	explicit	n~	n~	n~	n~
	implicit	n~	n~	n~	n~
Additionals	explicit	n+	n~	n+	n+
	implicit	n+	n~	n+	n+
		Disjoined			
Alternatives	explicit	n+	n+	n+	n+
	implicit	n	n+	n+	n+
Additionals	explicit	n~	n~	n~	n~
	implicit	n~	n~	n~	n~

Key : S = suppressed
n = not suppressed
\+ = expectation met by experimental findings
~ = no expectation

Appendix A

Table 7

The percentages of inferences endorsed under the conditions of Experiment 1.

Information type	Presentation	Minor premise and inference			
		Categorical			
		MP	MT	DA	AC
Alternatives	explicit	96	96	04	12.5
	implicit	100	75	46	58
Additionals	explicit	37.5	33	62.5	54
	implicit	96	92	46	71
		Conjoined			
Alternatives	explicit	100	87.5	79	50
	implicit	71	67	50	42
Additionals	explicit	100	33	92	71
	implicit	100	67	83	79
		Disjoined			
Alternatives	explicit	100	96	62.5	75
	implicit	37.5	62.5	58	50
Additionals	explicit	62.5	67	83	50
	implicit	58	92	75	79

Each percentage is based on the responses of 8 subjects to three items.

Appendix A

Table 8

The judgements of certainty made under the conditions of Experiment 1

part i : The overall percentage of certainty judgements made under the conditions of Experiment 1.

Information type	Presentation	Minor premise and inference			
		Categorical			
		MP	MT	DA	AC
Alternatives	explicit	99	97.5	93	92
	implicit	100	93	92	92
Additionals	explicit	87	79	83	82
	implicit	97.5	92.5	91	92
		Conjoined			
Alternatives	explicit	100	99	93	86
	implicit	93	79	91	84
Additionals	explicit	100	87	95	87.5
	implicit	94	82.5	94	83
		Disjoined			
Alternatives	explicit	99	94	93	99
	implicit	88	86	92.5	86
Additionals	explicit	87	91	91	87
	implicit	86	84	90	81

Each percentage is based on the choice of 8 subjects of a rating between 1 and 5, to three items.

Table 8 continued

part ii : the percentage of certainty judgements relevant to the interactions of the minor premise form and information presentation.

	Categorical	Conjoined	Disjoined
explicit	89	93.5	93
implicit	93	88	87

Each percentage is based on the ratings of 8 subjects to 24 items.

part iii : the percentage of certainty judgements relevant to the interaction of minor premise and inference type

	MP	MT	DA	AC
Categorical	96	91	90	88.5
Conjoined	97	87	93	85
Disjoined	90	88	92	88

Each percentage is based on the ratings of 32 subjects of 3 items

Appendix A

Table 9

The conditions of Experiment 2

	MP	MT	DA	AC
	Block A : Alternatives			
Two conditionals (alternatives)	n+~	n+	S+~	S+~
Indicated relevance of alternatives	n+	n	n+	n+
Named alternatives	n+	n+	n+	n+
Single conditional	n+~	n+~	n+~	n+~
	Block B : Additionals			
Two conditionals (additionals)	S+	S+	n+~	n+~
Indicated relevance of additionals	n+	n+	n	n
Named additionals	n+	n	n+	n+
Single conditional	n+~	n+~	n+~	n+~

Half of the subjects received block A prior to block B, while the other half received block B prior to block A.

Key n = nonsuppressd
S = suppressed
+ = expectation met when subject received item first.
~ = expectation met when subject received item second, (applicable only to the two conditionals condition and single conditional condition).

Table 10

The percentage of inferences endorsed under the conditions of Experiment 2

Information type	operation	MP	MT	DA	AC
		First block			
Alternatives	relevance and retrieval	100	94	12.5	12.5
	relevance	62.5	44	69	44
	retrieval	94	56	69	56
	neither	94	75	56	44
Additionals	relevance and retrieval	37.5	37.5	62.5	56
	relevance	81	87.5	19	19
	retrieval	87.5	44	44	37.5
	neither	100	75	69	69

Table 10 continued

		Second block			
Alternatives	relevance and retrieval	81	31	19	12.5
	relevance	94	87.5	19	19
	retrieval	87.5	56	19	25
	neither	94	75	62.5	69
Additionals	relevance and retrieval	87.5	69	31	37.5
	relevance	56	25	56	62.5
	retrieval	81	62.5	81	62
	neither	87.5	75	50	31

Each percentage is based on the responses of four subjects to four items. Subjects who received the alternatives block first were given the additionals block second, and vice vearsa.

Table 11

The percentage of certainty judgements in Experiment 2

Information type	operation	MP	MT	DA	AC
Alternatives	relevance and retrieval	80	80	73	69
	relevance	66	67	72	72
	retrieval	68	67	68	74
	neither	74	73	77	72
Additionals	relevance and retrieval	78	80	79	78
	relevance	74	66	59	56
	retrieval	78	70	77	72
	neither	77	72	73	69

Each percentage is based on the ratings of 8 subjects on a scale of 1 to 5, for three items.

Table 12

The conditions of Experiment 3

	MP	MT	DA	AC
no extra information (the gunman said..)	n	n	n+	n
situational (during the siege..)	n+	n+	n+	n
situational and long durational (during the two week siege..)	n+	n+	S+	S+
long durational (during the next 2 weeks..)	n+	n+	n+	n+
situational and short durational (during the 15 minute siege..)	S+	S+	n	n+
short durational (during the next 2 weeks..)	n	n+	n+	n+

Key : n = non-suppressed
S = suppressed
\+ = expectation met

Appendix A

Table 13

The percentage of inferences endorsed under the conditions of Experiment 3

Information type	Information given	MP	MT	DA	AC
	None	56	34	47	37.5
	Situational	72	66	44	28
Alternatives	Durational and situation	56	59	19	31
	Durational	50	50	66	59
Additionals	Durational and situation	41	25	25	34
	Durational	31	53	50	44

Each percentage is based on the responses of eight subjects to four items.

Appendix A

Table 13

The percentage of inferences endorsed under the conditions of Experiment 3

Information type	Information given	MP	MT	DA	AC
	None	56	34	47	37.5
	Situational	72	66	44	28
Alternatives	Durational and situation	56	59	19	31
	Durational	50	50	66	59
Additionals	Durational and situation	41	25	25	34
	Durational	31	53	50	44

Each percentage is based on the responses of eight subjects to four items.

Table 14

The percentage of certainty responses under the conditions of Experiment 3

Information type	Information given	MP	MT	DA	AC	Overall no situational	Overall situational
	None	91	87	79	82.5	85	
	Situational	76	79	76	67.5		75
Alternatives	Durational and situation	82	81	78	81		80.5
	Durational	79	83	84	77.5	81	
Additionals	Durational and situation	74	77.5	74	85		78
	Durational	89	88	84	93	88.5	
						85	78

Each percentage is based on the ratings of 8 subjects on scales of 1 to 5, for 4 items.

Appendix A

Table 15

The inferences made in extended and standard arguments.

	Forwards MP so MP	DA so DA	Backwards MT so M T	AC so AC
(1) Transitive arguments:				
Extended if p then q if q then r	p so q so r	not p so not q so not r	not r so not q so not p	r so q so p
	(1)	(1)	(1)	(1)
Standard if p then q	p so q	not p so not q	not r so not q	r so q
if q then r	q so r	not q so not r	not q so not p	q so p
	(1)	(1)	(1)	(1)

Table 15 continued.

	MP so DA	DA so MP	MT so AC	AC so MT
(2) Atransitive arguments :				
Extended if p then q if not q then r	p so q so not r	not p so not q so r	not r so not q so p	r so q so not p
	(1)	(1)	(1)	(1)
Standard if p then q	p so q	not p so not q	not r so not q	r so q
if not q then r	q so not r	not p so r	not r so p	r so not p
	(1)	(1)	(1)	(1)
(3) Extended if p then not q if q then r	p so not q so not r	not p so q so r	not r so not q so p	r so q so not p
	(1)	(1)	(1)	(1)
Standard if p then not q	p so not q	not p so q	not r so not q	r so q
if q then r	not q so not r	q so r	not q so p	q so not p
	(1)	(1)	(1)	(1)

Note : The scoring system used in Experiments 4, 5 and 6 is indicated = a score of 1 was given to the two inferences marked.

Table 16

The percentage of inferences made in Experiment 4

Part i

The percentage of definite endorsements of both of the appropriate inferences for the arguments used in Experiment 4.

	Transitive arguments				
	Forwards		Backwards		
	MP so MP	DA so DA	AC so Ac	MT so MT	Overall
Extended if p then q if q then r	85	55	60	50	65
Standard if p then q if q then r	50	35	50	50	46

Table 16, part i, continued.

	Atransitive arguments				
	Forwards		Backwards		
	MP so DA	DA so MP	AC so MT	MT so AC	Overall
Extended if p then q if not q then r	70	45	60	45	55
Standard if p then q if not q then r	50	80	60	55	61
Extended if p then not q if q then r	15	30	45	35	31
Standard if p then not q if q then r	60	60	70	70	65

Each percentage is based on the responses of 10 subjects to two items.

Table 16 continued.

Part ii: The frequency of definite endorsements of the first inference only

Transitive arguments

	Forwards MP so MP	Forwards DA so DA	Backwards AC so Ac	Backwards MT so MT	Overall
Extended if p then q if q then r	15	20	25	0	15
Standard if p then q if q then r	10	30	5	0	11

Atransitive arguments

	Forwards MP so DA	Forwards DA so MP	Backwards AC so MT	Backwards MT so AC	Overall
Extended if p then q if not q then r	10	25	15	5	14
Standard if p then q if not q then r	5	15	5	10	9
Extended if p then not q if q then r	70	0	10	35	24
Standard if p then not q if q then r	30	10	20	5	16

Each percentage is based on the response of 10 subjects to two items.

Table 16 continued.

Part iii

The frequency of definite endorsements of the second inference alone

Transitive arguments

	Forwards MP so MP	Forwards DA so DA	Backwards AC so Ac	Backwards MT so MT	Overall
Extended if p then q if q then r	0	20	25	10	14
Standard if p then q if q then r	15	5	15	10	11

Atransitive arguments

	Forwards MP so DA	Forwards DA so MP	Backwards AC so MT	Backwards MT so AC	Overall
Extended if p then q if not q then r	15	5	15	15	12.5
Standard if p then q if not q then r	5	20	10	10	11
Extended if p then not q if q then r	5	0	10	0	4
Standard if p then not q if q then r	10	5	20	30	16

Each percentage is based on the responses of 10 subjects to two items.

Appendix A

Table 17

The conditions and expectations of Experiment 5

	Forward MP so DA	Forward DA so MP	Backwards AC so MT	Backwards MT so AC	Overall
Extended argument					s+
Certainty elimination					s+
representation elimination					n+
Standard argument					n

Key: s = suppressed
n = non-suppressed
\+ = expectation met in experiment

Table 18

The percentage of inferences made in Experiment 5

Part i

The percentage of definite endorsements of both appropriate inferences under the conditions of Experiment 5

	Forward		Backwards		Overall
	MP so DA	DA so MP	AC so MT	MT so AC	
Extended argument	26	36	38	50	37.5 (7)
Uncertainty elimination	46	54	48	56	51 (8)
representation elimination	67	65	67	92	73 (8)
Standard argument	21	31	43	64	40 (7)

					component
	40		49		44.5
		46.5		65.5	56
direction	43		58		

Percentages are based on the responses of eight (or seven as indicated) subjects to six items.

Table 18 continued

Part ii

The percentage of definite endorsements of the first inferences under the conditions of Experiment 5

	Forward MP so DA	DA so MP	Backwards AC so MT	MT so AC	Overall
Extended argument	52	5	12	17	21.5 (7)
Uncertainty elimination	35	6	8	31	20 (8)
representation elimination	25	2	2	19	48 (8)
Standard argument	50	7	26	12	24 (7)

Percentages are based on the responses of eight (or seven as indicated) subjects to six items.

Table 18 continued.

Part iii

The percentage of definite endorsements of the second inferences under the conditions of Experiment 5

	Forward MP so DA	DA so MP	Backwards AC so MT	MT so AC	Overall
Extended argument	0	9.5	2	2	3 (7)
Uncertainty elimination	0	2	2	0	1 (8)
representation elimination	0	6	0	2	2 (8)
Standard argument	7	59	7	24	7.5 (7)

Percentages are based on the responses of eight (or seven as indicated) to six items.

Table 19

The percentage of first inferences made in the conditions of Experiment 5.

	MP	DA	AC	MT	Overall
Extended argument	86	36	62	59.5	61 (7)
Uncertainty elimination	90	54	67	81	73 (8)
representation elimination	92	67	94	85	84.5(8)
Standard argument	74	33	69	67	61 (7)

Percentages are based on the responses of eight subjects (or seven as indicated) to six items.

Appendix A

Table 20

The percentage of inferences made in Experiment 6

Part i

The percentage of definite endorsements of both appropriate inferences for the arguments of different lengths and forms in Experiment 6

	MP so MP	DA so DA	AC so AC	MT so MT	Overall
Extended if p then not q if not q then r	93	50	71	54	67 (7)
Standard if p then not q if not q then r	61	29	46	46	45.5(8)
Extended If not p then q if q then r	87.5	56	47	50	60 (7)
Standard If not p then q if q then r	66	31	50	41	47 (8)
	77	41.5	53.5	48	

Each percentage is based on the responses of eight (or seven as indicated) subjects to four items.

Appendix A

Table 20

Part ii

The percentage of definite endorsements of the first inferences alone for the arguments in Experiment 6

	MP so MP	DA so DA	AC so AC	MT so MT	Overall	
Extended if p then not q if not q then r	4	4	11	14	8	(7)
Standard if p then not q if not q then r	0	6	0	0	4	(8)
Extended If not p then q if q then r	4	7	0	4	1.5	(7)
Standard If not p then q if q then r	3	6	3	0	3	(8)

Each percentage is based on the responses of eight (or seven as indicated) subjects to four items.

Appendix A

Table 20 continued.

Part iii

The percentage of definite endorsements of the second inferences only in Experiment 6

	MP so MP	DA so DA	AC so AC	MT so MT	Overall	
Extended if p then not q if not q then r	4	4	4	7	5	(7)
Standard if p then not q if not q then r	6	9	16	12.5	11	(8)
Extended If not p then q if q then r	0	0	7	14	5	(7)
Standard If not p then q if q then r	9	16	19	12.5	14	(8)

Each percentage is based on the responses of eight (or seven as indicated) subjects to four items.

APPENDIX B:

EXPERIMENTAL MATERIALS

Appendix B

Table 1

Examples of the materials used in Experiments 1 and 2

Basic conditional content	(1) if it is raining then she will get wet (2) if she meets her family then she will go to a play (3) if she has an essay to finish then she will study late in the library *(4) if she finds some leftover paint then she will re-decorate her room
Alternative	(1) if it is snowing then she will get wet (2) if she meets her friend then she will go to a play (3) if she has some textbooks to read then she will study late in the library *(4) if she can spare some money then she will re-decorate her room
Additional	(1) if she goes out for a walk then she will get wet (2) if she has enough money then she will go to a play (3) if the library stays open then she will study late in the library *(4) if she can spare some time then she will re-decorate her room

* = used in Experiment 2 only

Note : The basic conditional content was accompanied by an alternative or an additional antecedent conditional, but never by both at the same time.

Table 2

Examples of minor premises used in Experiment 1

Categorical

MP	it is raining, therefore (a) she will get wet (b) she will not get wet (c) she may or may not get wet
DA	it is not raining, therefore (a) she will not get wet (b) she will get wet (c) she may or may not get wet
MT	she will not get wet, therefore (a) it is not raining (b) it is raining (c) it may or may not be raining
AC	she will get wet, therefore (a) it is raining (b) it is not raining (c) it may or may not be raining

Conjunction of alternatives

MP	it is raining and it is snowing, therefore (a) she will get wet (b) she will not get wet (c) she may or may not get wet
DA	it is not raining and it is not snowing, therefore (a) she will not get wet (b) she will get wet (c) she may or may not get wet
MT	she will not get wet, therefore (a) it is not raining and it is not snowing (b) it is raining and it is snowing

(c) it may or may not be raining and
it may or may not be snowing

AC she will get wet, therefore
(a) it is raining and it is snowing
(b) it is not raining and it is not snowing
(c) it may or may not be raining and
it may or may not be snowing

Conjunction
of additionals

MP it is raining and she goes out for a walk, therefore
(a) she will get wet
(b) she will not get wet
(c) she may or may not get wet

DA it is not raining and she does not go out for a walk,
therefore
(a) she will not get wet
(b) she will get wet
(c) she may or may not get wet

MT she will not get wet, therefore
(a) it is not raining and she does not go out
for a walk
(b) it is raining and she goes out for a walk
(c) it may or may not be raining and
she may or may not go out for a walk

AC she will get wet, therefore
(a) it is raining and it is snowing
(b) it is not raining and it is not snowing
(c) it may or may not be raining and
it may or may not be snowing

Disjunction
of alternatives

MP it is raining or it is snowing, therefore
(a) she will get wet
(b) she will not get wet
(c) she may or may not get wet

DA it is not raining or it is not snowing,

therefore
(a) she will not get wet
(b) she will get wet
(c) she may or may not get wet

MT she will not get wet, therefore
(a) it is not raining or it is not snowing
(b) it is raining or it is snowing
(c) it may or may not be raining and it may or may not be snowing

AC she will get wet, therefore
(a) it is raining or it is snowing
(b) it is not raining or it is not snowing
(c) it may or may not be raining and it may or may not be snowing

Disjunction
of additionals

MP it is raining or she goes out for a walk, therefore
(a) she will get wet
(b) she will not get wet
(c) she may or may not get wet

DA it is not raining or she does not go out for a walk, therefore
(a) she will not get wet
(b) she will get wet
(c) she may or may not get wet

MT she will not get wet, therefore
(a) it is not raining or she does not go out for a walk
(b) it is raining or she goes out for a walk
(c) it may or may not be raining and she may or may not go out for a walk

AC she will get wet, therefore
(a) it is raining or it is snowing
(b) it is not raining or it is not snowing
(c) it may or may not be raining and it may or may not be snowing

Key : MP = modus ponens
MT = modus tollens

DA = denial of the antecedent
AC = affirmation of the consequent

Table 3

Examples of the materials used in Experiment 3

basic content	(1) the gunman said to the bank-clerk : "if you move, I will shoot you" (2) the tour-guide said to the tourist : "if you go shopping, I will lose you" (3) the mother said to her son : "if you go for a swim with your sister I will buy you an ice-cream" (4) the policeman said to the student protester: "if you enter the building I will arrest you"
situational prefaces	(1) during the siege (2) during the city-center tour (3) during the seaside trip (4) during the student protest
long durational prefaces	(1 to 4) during the next two weeks
short durational prefaces	(1 to 4) during the next 15 minutes
situational and long durational prefaces	(1) during the two-week siege (2) during the two-week city-center tour (3) during the two-week seaside trip (4) during the two-week student protest
situational and short durational prefaces	(1) during the 15 minute siege (2) during the 15 minute city-center tour (3) during the 15 minute seaside trip (4) during the 15 minute student protest

Table 4

Examples of the materials used in Experiments 4, 5 and 6.

(1) Extended argument	If she rings her friend then she will go to visit her, if she goes to visit her then she will listen to some music
counterpart standard arguments	If she rings her friend then she will go to visit her. If she visits her friend then she will listen to some music.
(2) Extended argument	If she gets up early then she will not drive to work, if she drives to work then she will go by the scenic route
counterpart standard arguments	If she gets up early then she will not drive to work. If she drives to work then she will go by the scenic route
(3) Extended argument	If she works hard then she will leave work early, if she does not leave work early then she will eat out
counterpart standard arguments	If she works hard then she will leave work early. If she does not leave work early then she will eat out
(4) Extended argument	If she studies hard this afternoon then she will not leave the library before dinnertime if she does not leave the library before dinnertime, then she will spend the evening in the buttery
counterpart standard arguments	If she studies hard this afternoon then she will not leave the library before dinnertime. If she does not leave the library before dinnertime, then she will spend the evening in the buttery

Table 4 continued

(5) Extended argument	If she does not meet her friends on Friday then she will go down the country if she goes down the country then she will have a pleasant weekend
counterpart standard arguments	If she does not meet her friends on Friday then she will go down the country. If she goes down the country then she will have a pleasant weekend

APPENDIX C:

STATISTICAL EFFECTS

Table 1

The effects of the four-way analysis of variance conducted on the responses to Experiment 1

Variable	df	F	p
A	1, 28	.55	.47
B	1, 28	.12	.74
AB	1, 28	5.16	.03
C	1.97, 55	7.61	.001
AC	1.97, 55	1.19	.31
BC	1.97, 55	17.60	.0000
ABC	1.97, 55	2.96	.06
D	1.94, 54	6.30	.004
AD	1.94, 54	6.20	.004
BD	1.94, 54	1.05	.36
ABD	1.94, 54	4.70	.01
CD	4.8, 134.7	11.30	.0000
ACD	4.8, 134.7	4.60	.0008
BCD	4.8, 134.7	2.10	.08
ABCD	4.8, 134.7	3.70	.004

Key : A = type of information, either alternatives or additionals
B = presence of information, either present or absent
C = kind of minor premise, either categorical, conjoined, or disjoined.
D = type of inference, either modus ponens, tollens, affirmation of the consequent or denial of the antecedent.

Table 2

The effects found in the four-way analysis of variance conducted on the certainty judgements made in Experiment 1

Variable	df	F	p
A	1, 28	2.25	.15
B	1, 28	.70	.41
AB	1, 28	1.70	.20
C	1.3, 35.7	.38	.59
AC	1.3, 35.7	1.42	.25
BC	1.3, 35.7	5.20	.02
ABC	1.3, 35.7	.71	.44
D	2.2, 60	7.30	.001
AD	2.2, 60	.61	.56
BD	2.2, 60	.70	.51
ABD	2.2, 60	.70	.51
CD	4.5, 124.8	3.83	.004
ACD	4.5, 124.8	1.40	.24
BCD	4.5, 124.8	1.21	.31
ABCD	4.5, 124.8	1.31	.30

Key : A = type of information, either alternatives or additionals
B = presence of information, either present or absent
C = kind of minor premise, either categorical, conjoined, or disjoined.
D = type of inference, either modus ponens, tollens, affirmation of the consequent or denial of the antecedent.

Table 3

The effects found in the five-way analysis of variance conducted on all the responses in Experiment 2

Variable	df	F	p
A	1, 24	6.76	.02
B	1, 24	1.60	.22
C	1, 24	.83	.37
AB	1, 24	.32	.58
AC	1, 24	.04	.83
BC	1, 24	4.44	.0456
ABC	1, 24	.71	.41
D	1, 24	.08	.78
AD	1, 24	.08	.78
BD	1, 24	3.24	.08
CD	1, 24	.69	.41
ABD	1, 24	.02	.89
ACD	1, 24	.08	.78
BDC	1, 24	.17	.68
ABDC	1, 24	.17	.68
E	2.1, 50	18.95	.0000
AE	2.1, 50	.54	.60
BE	2.1, 50	.23	.80
CE	2.1, 50	.56	.58
ABE	2.1, 50	.27	.78
ACE	2.1, 50	.32	.74
BCE	2.1, 50	2.28	.11
ABCE	2.1, 50	9.67	.0002
DE	2.2, 54.5	9.41	.0002
ADE	2.2, 54.5	3.80	.03
BED	2.2, 54.5	6.14	.003
CED	2.2, 54.5	1.10	.36
ABDE	2.2, 54.5	.86	.44
ACDE	2.2, 54.5	2.89	.058
BCDE	2.2, 54.5	.52	.62
ABCDE	2.2, 54.5	1.42	.25

Key : A = Facilitation of relevance
B = Facilitation of retrieval
C = order in which information was presented, either first or second
D = type of information, alternatives or additionals
E = type of inference, either modus ponens, tollens, affirmation of the consequent or denial of the antecedent.

Table 4

The effects of the four-way analysis of variance conducted on reponses to the first-presented blocks in Experiment 2

Variable	df	F	p
A	1, 24	5.08	.03
B	1, 24	.93	.34
C	1, 24	.50	.49
AB	1, 24	.50	.49
AC	1, 24	.00	.95
BC	1, 24	1.20	.28
ABC	1, 24	.70	.41
D	2.3, 54.7	12.69	.0000
AD	2.3, 54.7	1.4	.264
BD	2.3, 54.7	.21	.84
CD	2.3, 54.7	.64	.55
ABD	2.3, 54.7	.44	.67
ACD	2.3, 54.7	.42	.685
BCD	2.3, 54.7	5.04	.007
ABCD	2.3, 54.7	9.16	.0002

Key : A = Facilitation of relevance
B = Facilitation of retrieval
C = type of information, either alternatives or additionals
D = type of inference, either modus ponens, tollens, affirmation of the consequent or denial of the antecedent.

Table 5

The effects of the four-way analysis of variance conducted on the certainty judgements of the first-presented blocks in Experiment 2

Variable	df	F	p
A	1, 24	.09	.76
B	1, 24	3.40	.08
C	1, 24	.02	.88
AB	1, 24	6.30	.02
AC	1, 24	.38	.55
BC	1, 24	2.36	.14
ABC	1, 24	.07	.80
D	1.97, 47.4	2.30	.11
AD	1.97, 47.4	1.56	.22
BD	1.97, 47.4	.30	.76
CD	1.97, 47.4	2.09	.14
ABD	1.97, 47.4	.94	.40
ACD	1.97, 47.4	.26	.77
BCD	1.97, 47.4	3.47	.04
ABCD	1.97, 47.4	3.68	.03

Key : A = Facilitation of relevance
B = Facilitation of retrieval
C = type of information, either alternatives or additionals
D = type of inference, either modus ponens, tollens, affirmation of the consequent or denial of the antecedent.

Table 6

The results of the three-way analysis of variance conducted on the responses to Experiment 3

Variable	df	F	p
A	1, 42	.82	.37
B	2, 42	.96	.39
AB	2, 42	1.11	.34
C	3, 126	2.92	.04
AC	3, 126	6.59	.0004
BC	6, 126	1.82	.10
ABC	6, 126	2.22	.045

Key : A = situational information, either not given, or given
B = durational information, either not given, long or short
C = inference type, either modus ponens, tollens, denial of the antecedent or affirmation of the consequent

Table 7

The results of the three-way analysis of variance conducted on the certainty judgements in Experiment 3

Variable	df	F	p
A	1, 42	3.90	.055
B	2, 42	.32	.73
AB	2, 42	.89	.42
C	3, 126	.94	.42
AC	3, 126	.18	.91
BC	6, 126	1.87	.09
ABC	6, 126	.86	.53

Key : A = situational information, either not given, or given
B = durational information, either not given, long or short
C = inference type, either modus ponens, tollens, denial of the antecedent or affirmation of the consequent

Table 8

The results of the four-way analysis of variance conducted on the responses to Experiment 4

Variable	df	F	p
A	1, 54	.77	.39
B	2, 54	.56	.57
AB	2, 54	3.53	.04
C	1, 54	.25	.62
AC	1, 54	.09	.76
BC	2, 54	2.11	.13
ABC	2, 54	.85	.43
D	1, 54	1.48	.23
AD	1, 54	2.76	.10
BD	2, 54	.93	.40
ABD	2, 54	1.91	.16
CD	1, 54	.01	.9
ACD	1, 54	.62	.43
BCD	2, 54	2.50	.09
ABCD	2, 54	1.88	.16

Key : A = length of argument, either extended or standard
B = argument form, either all affirmative, a negative at the antecedent of the second conditional, or a negative at the consequent of the first conditional
C = inference direction, either forwards or backwards
D = inference form, either the same as, or different from the component referred to.

Table 9

The results of the four-way analysis of variance conducted on responses to Experiment 5

Variable	df	F	p
A	1, 26	1.53	.23
B	1, 26	1.03	.32
AB	1, 26	5.20	.03
C	1, 26	14.20	.0009
AC	1, 26	3.13	.09
BC	1, 26	.01	.94
ABC	1, 26	2.95	.10
D	1, 26	12.80	.001
AD	1, 26	.37	.55
BD	1, 26	.01	.93
ABD	1, 26	.25	.62
CD	1, 26	1.66	.21
ACD	1, 26	1.33	.26
BCD	1, 26	.40	.55
ABCD	1, 26	.15	.70

Key : A = certainty elimination
B = representation elimination
C = inference direction, either forwards or backwards
D = inference form, either the same as, or different to, the component referred to.

Appendix C

Table 10

The results of the four-way analysis of variance conducted on responses to the first items in Experiment 5

Variable	df	F	p
A	1, 26	.60	.44
B	1, 26	.60	.44
AB	1, 26	5.16	.03
C	1, 26	3.95	.0575
AC	1, 26	2.74	.11
BC	1, 26	.11	.74
ABC	1, 26	.02	.90
D	1, 26	33.69	.0000
AD	1, 26	.03	.86
BD	1; 26	.77	.39
ABD	1, 26	2.63	.12
CD	1, 26	13.81	.001
ACD	1, 26	1.20	.28
BCD	1, 26	.43	.52
ABCD	1, 26	.22	.64

Key : A = certainty elimination
B = representation elimination
C = inference direction, either forwards or backwards
D = inference form, either the same as, or different to the component referred to.

Table 11

The results of the four-way analysis of variance conducted on the responses to Experiment 6

Variable	df	F	p
A	1, 26	2.40	.14
B	1, 26	.06	.81
AB	1, 26	.13	.72
C	1, 26	4.31	.0479
AC	1, 26	3.53	.07
CB	1, 26	1.32	.26
ABC	1, 26	.32	.57
D	1, 26	25.53	.0000
AD	1, 26	.10	.76
BD	1, 26	.42	.52
ABD	1, 26	1.86	.18
CD	1, 26	4.51	.04
ACD	1, 26	.00	.95
BCD	1, 26	.00	.97
ABCD	1, 26	.09	.77

Key : A = argument length, either extended or standard
B = argument form, either a negative at the antecedent of the first conditional, or a negative at the consequent of the first conditional and at the antecedent of the second condition
C = inference direction, either forwards or backwards
D = inference form, either the same as, or different to, the component referred to.

REFERENCES

References

Adams, M.J., (1980). Inductive deductions and deductive inductions. In R.S. Nickerson (ed.), Attention and Performance VIII. Hillsdale, NJ, Erlbaum.

Anderson, A.R. and Belnap, N.D. (1975). Entailment: The Logic of Relevance and Necessity. Volumn 1. Princeton University Press.

Bracewell, R.J. and Hidi, S.E. (1974). The solution of an inferential problem as a function of stimulus materials. Quarterly Journal of Experimental Psychology, 26, 480-488.

Braine, M.D.S. (1978). On the relation between the natural logic of reasoning and standard logic. Psychological Review, 85, 1-21.

Braine, M.D.S. and Rumain, B. (1983). Logical Reasoning. In J.H. Flavell and E.M. Markman (eds.), Handbook of Child Psychology: Cognitive Development Volumn 3. New York, Wiley.

Braine, M.D.S., Reiser, B.J. and Rumain, B. (1984). Some empirical justification for a theory of natural propositional logic. In G.H. Bower (ed.), The psychology of Learning and Motivation, Volumn 18. New York, Academic Press.

Bucci, W. (1978). The interpretation of universal affirmative propositions. Cognition, 6, 55-77.

Cheng, P.N. and Holyoak, K.J. (1985). Pragmatic reasoning schemas. Cognitive Psychology, 17, 391-416.

Cheng, P.N., Holyoak, K.J., Nisbett, R.E. and Oliver, L.M. (1986) Pragmatic versus syntactic approaches to training deductive reasoning. Cognitive Psychology, 18, 293-328.

Chrostowski, J.J. and Griggs, R.A. (1985). The effects of problem content, instructions and verbalisation procedure on Wason's selection task. Current Psychological Research and Reviews, 4, 99-107.

Cohen, L.J. (1981). Can human irrationality be experimentally demonstrated ? Behavioural and Brain Sciences, 4, 317-370.

Copi, I.M. (1982). Introduction to Logic. Sixth Edition. MacMillian.

Collins, A. (1978). Fragments of a theory of human plausible reasoning. In D.L. Waltz (ed.), Theoretical Issues in Natural Language Processing - 2. New York, Association for computing machinery.

Cox, J.R. and Griggs, R.A. (1982). The effects of experience on performance in Wason's selection task. Memory and Cognition, 10, 496-502.

D'Andrade, R., cited in Rumelhart, D.E. and Norman, D.A. (1981) Analogical processes in learning. In J.R. Anderson (ed.), Cognitive Skills and their Acquisition. Hillsdale, NJ. Erlbaum.

Downing, C.J., Sternberg, R.J., Ross, B.H. (1985). Multicausal inferences: evaluation of evidence in causally complex situations.

Journal of Experimental Psychology: General, 114, 239-263.

Evans, J.St.B.T. (1972a). Reasoning with negatives. British Journal of Psychology, 63, 213-219

Evans, J.St.B.T. (1972b). Interpretation and "matching bias" in a reasoning task. Quarterly Journal of Experimental Psychology, 24, 193-199.

Evans, J.St.B.T. (1977). Linguistic factors in reasoning. Quarterly Journal of Experimental Psychology, 29, 297-306.

Evans, J.St.B.T. (1982). The Psychology of Deductive Reasoning. London, Routledge and Kegan Paul.

Evans, J.ST.B.T. (1983a). Linguistic determinants of bias in conditional reasoning. Quarterly Journal of Experimental Psychology, 35A, 635-644

Evans, J.St.B.T. (1983b). Selective processes in reasoning. In J.St.B.T. Evans (ed.), Thinking and Reasoning: Psychological approaches. London, Routledge and Kegan Paul.

Evans, J.St.B.T. (1984). Heuristic and analytical processes in reasoning. British Journal of Psychology, 75, 451-468.

Evans, J.St.B.T and Beck,M.A. (1981). Directionality and temporal factors in conditional reasoning. Current Psychological Research, 1, 111-120.

Evans, J.St.B.T. and Lynch, J.S. (1973). Matching bias in the selection task. British Journal of Psychology, 64, 391-397.

Evans, J.St.B.T. and Newstead, S.E. (1977). Language and reasoning: a study of temporal factors. Cognition, 5, 265-283

Falmagne, R.J. (1975). Overview: reasoning, representation, process and related issues. In R.J. Falmagne (ed.), Reasoning: Representation and Process. New York, Wiley.

Fillenbaum, S. (1975). IF: some uses. Psychological Research, 37, 245-260.

Fillenbaum, S. (1976). Inducements: On the phrasing and logic of conditional promises, threats and warnings. Psychological Research, 38, 231-250.

Fillenbaum, S. (1978). How to do some things with IF. In J.W. Cotton and Klatzky, R.L. (eds.), Semantic Factors in Cognition. Hillsdale, NJ, Erlbaum.

Geis, M.L. and Zwicky, A.M. (1971). On invited inferences. Linguistic Inquiry, 11 561-566.

Gellatly, A.R.H. (in press). Human inference. In K.Gilhooly (ed.), Human and Machine Problem Solving. London, Plenum.

Gilhooly, K.J. and Falconer, W.A. (1974). Concrete and abstract terms and relations in testing a rule. Quarterly Journal of Experimen-

Instituto Superior de Psicologia Aplicada
BIBLIOTECA

tal Psychology, 26, 355-359.

Grice, H.P. (1975). Logic and conversation. In P.Cole and J.L.Morgan (eds.), Syntax and Semantics. Volumn 3: Speech Acts. New York, Academic Press.

Griggs, R.A. (1983). The role of problem content in the selection task and THOG problem. In J.St.B.T. Evans (ed.), Thinking and Reasoning: Psychological approaches. London, Routledge and Kegan Paul.

Griggs, R.A. (1984). Memory cueing and instructional effects on Wason's selection task. Current Psychological Research and Reviews, 4, 3-10.

Griggs, R.A. and Cox, J.R. (1982). The elusive thematic materials effect in Wason's selection task. British Journal of Psychology, 73, 407-420.

Griggs, R.A. and Cox, J.R. (1983). The effects of problem content and negation on Wason's selection task. Quarterly Journal of Experimental Psychology, 35A, 519-533.

Guyote, M.J. and Sternberg, R.J. (1981). A transitive-chain theory of syllogistic reasoning. Cognitive Psychology, 13, 461-524.

Haack, S. (1978). Philosophy of Logics. Cambridge, Cambridge University Press.

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

Henle, M. (1978). Foreword. In R.Revlin and R.E. Mayer (eds.), Human Reasoning. Washington, D.C., Winston.

Hoch, S.J. and Tschirgi, J.E. (1983). Cue redundancy and extra logical inferences in a deductive reasoning task. Memory and Cognition, 11, 200-209.

Johnson-Laird, P.N. (1982) Thinking as a skill. Quarterly Journal of Experimental Psychology, 34A, 1-29.

Johnson-Laird, P.N. (1983). Mental Models. Cambridge, Cambridge University Press.

Johnson-Laird, P.N. (in press). Conditionals and mental models. In C. Ferguson, J.Reilly, A. ter Meulen and E.C. Traugott (eds.), On Conditionals. Cambridge, Cambridge University Press.

Johnson-Laird, P.N., Legrenzi, P. and Sonino Legrenzi, M.S. (1972) Reasoning and a sense of reality. British Journal of Psychology, 63 395-400

Johnson-Laird, P.N. and Steedman, M. (1978). The psychology of syllogisms. Cognitive Psychology, 10, 64-99.

Kahn, G.S. and Rips, L.J. (1983). Norms, competence and the explanation of reasoning. Behavioural and Brain Sciences, 3, 501-502.

Kahneman, D. (1981). Who shall be the arbiter of our intuitions ? Behavioural and Brain Sciences, 4, 339-340.

Keane, M.T. (1985). On drawing analogies when solving problems: A theory and test of solution generation in an analogical problem solving task. British Journal of Psychology, 76, 449-458.

Keane, M.T. (in press). On retrieving analogues when solving problems. Quarterly Journal of Experimental Psychology.

Levinson, S.C. (1983). Pragmatics. Cambridge, Cambridge University Press.

Maier, N.R.F. (1931). Reasoning in humans. II. The solution of a problem and its appearance in consciousness. Journal of Comparative Psychology, 12, 181-194.

Manktelow, K.I. (1981). Recent developments in research on Wason's selection task. Current Psychological Reviews, 1, 257-268.

Manktelow, K.I. and Evans, J.St.B.T. (1979). Facilitation of reasoning by realism: effect or non-effect ? British Journal of Psychology, 70, 477-488.

Mates, B. (1972). Elementary Logic. Second Edition. Oxford, Oxford University Press.

Marcus, S.L. and Rips, L.J. (1979). Conditional reasoning. Journal of Verbal Learning and Verbal Behaviour, 18, 199-223.

Marcus. S.L. (1982). Recall of logical argument lines. Journal of Verbal Learning and Verbal Behaviour, 21, 549-562.

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

Markovits, H. (1985). Incorrect conditional reasoning amoung adults: competence or performance ? British Journal of Psychology, 76, 241-247.

Mason, E.J., Bramble, W.J. and Mast, T.A. (1975). Familiarity with content and conditional reasoning. Journal of Educational Psychology, 67, 238-242.

McDermott, D. (1982). A temporal logic for reasoning about processes and plans. Cognitive Science, 6, 101-155

O'Brien, D.P. and Overton, W.F. (1980). Conditional reasoning following contradictory evidence: a developmental analysis. Journal of Experimental Child Psychology, 30 44-61.

O'Brien, D.P and Overton, W.F. (1982). Conditional reasoning and the competence-performance issue: a developmental analysis of a training task. Journal of Experimental Child Psychology, 34, 274-290.

O'Brien, D.P., Costa, G. and Overton, W.F. (1986). Evaluations of causal and conditional hypotheses. Quarterly Journal of

Experimental Psychology, 38A, 493-512.

Oakhill, J.V. and Johnson-Laird, P.N. (1985). Rationality, memory and the search for counterexamples. Cognition, 20, 79-94.

Osherson, D.N. (1975). Logical Abilities in Children, Volumn 3. Hillsdale, NJ, Erlbaum.

Osherson, D.N. (1976). Logical Abilities in Children, Volumn 4. Hillsdale, NJ, Erlbaum.

Pollard, P. (1981). The effect of thematic content in the Wason selection task. Current Psychologcial Research, 1, 21-29.

Pollard, P. (1982). Human reasoning: some possible effects of availability. Cognition, 12, 65-96.

Pollard, P. and Evans, J.St.B.T. (1980). The influence of logic on conditional reasoning performance. Quarterly Journal of Experimental Psychology, 32, 605-624.

Pollard, P. and Evans, J.St.B.T. (1981). The effect of prior beliefs on reasoning: an associational interpretation. British Journal of Psychology, 72, 73-82.

Piaget, J. and Inhelder, B. (1973). Memory and Intelligence. New York, Basic Books.

Reich, S. and Ruth, P. (1982). Wason's selection task: verification, falsification and matching. British Journal of Psychology, 73, 375-405.

Rich. E. (1983). Default reasoning as liklihood reasoning. In <u>Proceedings AAAI - 83</u>. Washington, D.C.

Rips, L.J. (1983a). Reasoning as a central intellective ability. In R.J. Sternberg (ed.), <u>Advances in the Study of Human Intelligence</u>. Hillsdale, NJ, Erlbaum.

Rips, L.J. (1983b). Cognitive processes in propositional reasoning. <u>Psychological Review</u>, <u>90</u>, 38-71.

Rips, L.J. and Conrad, F.J. (1983). Individual differences in deduction. <u>Cognition and Brain Theory</u>, <u>6</u>, 259-285.

Rips, L.J. and Marcus, S.L. (1977). Suppositions and the analysis of conditional sentences. In M.A. Just and P.A. Carpenter (eds.), <u>Cognitive Processes in Comprehension</u>. Hillsdale, NJ, Erlbaum.

Roberge, J.J. (1970). A study of children's abilities to reason with basic principles of deductive reasoning. <u>American Educational Research Journal</u>, <u>7</u>, 538-595.

Roberge, J.J. (1971). Some effects of negation on adults conditional reasoning abilities. <u>Psychological Reports</u>, <u>29</u>, 839-844.

Roberge, J.J. (1974). Effects of negation on adults comprehension of fallacious conditional and disjunctive arguments. <u>Journal of General Psychology</u>, <u>91</u>, 287-293.

Roberge, J.J. (1978). Linguistic and psychometric factors in propositional reasoning. Quarterly Journal of Experimental Psychology, 30, 705-716.

Rosch, E. (1978). Classification of real-world objects: origins and representations in cognition. In P.N. Johnson-Laird and P.C. Wason (eds.), Thinking: Readings in Cognitive Science. Cambridge, Cambridge University Press.

Rumain, B., Connell, J. and Braine, M.D.S. (1983). Conversational comprehension processes are responsible for reasoning fallacies in children aswell as adults: IF is not the Biconditional. Developmental Psychology, 19, 471-481.

Rumelhart, D.E. and Norman, D.A. (1981). Analogical processes in learning. In J.R. Anderson (ed.), Cognitive Skills and their Acquisition. Hillsdale, NJ. Erlbaum.

Staudenmayer, H. (1975). Understanding conditional reasoning with meaningful propositions. In R.J. Falmagne (ed.), Reasoning: Representation and Process in Children and Adults. Hillsdale, NJ, Erlbaum.

Staudenmayer, H. and Bourne, L.E. (1978). The nature of denied propositions in the conditional sentence reasoning task. In R. Revlin and R.E. Mayer (eds.), Human Reasoning. Washington, D.C., Winston.

Schank, R.C. (1982). Dynamic Memory Cambridge, Cambridge University Press.

Schustack, M.W. and Sternberg, R.J. (1981). Evaluation of evidence in causal reasoning. Journal of Experimental Psychology: General, 110, 101-120.

Sperber, D and Wilson, D (1986). Relevance: Communication and Cognition. Cambridge, Cambridge University Press.

Sternberg, R.J. (1985). Beyond I.Q.: A Triarchic Theory of Human Intelligence. Cambridge, Cambridge University Press.

Tversky, A. and Kahneman, D. (1973). Availability: a heuristic for judging frequency and probability. Cognitive Psychology, 5, 207-232.

Van Duyne, P.C. (1973). A short note on Evans' criticism of reasoning experiments and his matching bias hypothesis. Cognition, 2 239-242.

Van Duyne, P.C. (1974). Realism and linguistic complexity in reasoning. British Journal of Psychology, 65, 59-67.

Van Duyne, P.C. (1976). Necessity and contingency in reasoning. Acta Psychologica, 40, 85-101.

Wason, P.C. (1983). Realism and rationality in the selection task. In J.St.B.T.Evans (ed.), Thinking and Reasoning: Psychological approaches. London, Routledge and Kegan Paul.

Wason, P.C. and Johnson-Laird, P.N. (1972). Psychology of Reasoning: Structure and Content. London, Batsford.

Wason, P.C. and Shapiro, D. (1971). Natural and contrived experience in a reasoning task. Quarterly Journal of Experimental Psychology, 23, 63-71.

Winer, B.J. (1971). Statistical Principles in Experimental Design, Third Edition. New York, McGraw-Hill.

TRINITY COLLEGE
- 6 MAY 1987
LIBRARY DUBLIN

approaches. London, Routledge and Kegan Paul.

Wason, P.C. and Johnson-Laird, P.N. (1972). Psychology of Reasoning: Structure and Content. London, Batsford.

Wason, P.C. and Shapiro, D. (1971). Natural and contrived experience in a reasoning task. Quarterly Journal of Experimental Psychology, 23, 63-71.

Winer, B.J. (1971). Statistical Principles in Experimental Design, Third Edition. New York, McGraw-Hill.

TRINITY COLLEGE
- 6 MAY 1987
LIBRARY DUBLIN

Instituto Superior de Psicologia Aplicada
BIBLIOTECA