Looking for logic in all the wrong places:

an investigation of language, literacy and logic in reasoning

Marian Counihan

Looking for logic
in all the wrong places:

an investigation of language, literacy and logic in reasoning

ILLC Dissertation Series DS-2008-10

INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

For further information about ILLC-publications, please contact

Institute for Logic, Language and Computation
Universiteit van Amsterdam
Plantage Muidergracht 24
1018 TV Amsterdam
phone: +31-20-525 6051
fax: +31-20-525 5206
e-mail: illc@science.uva.nl
homepage: http://www.illc.uva.nl/

Looking for logic in all the wrong places:

an investigation of language, literacy and logic in reasoning

Academisch Proefschrift

ter verkrijging van de graad van doctor
aan de Universiteit van Amsterdam
op gezag van de Rector Magnificus
prof. dr. D.C. van den Boom
ten overstaan van een door het college voor promoties
ingestelde commissie,
in het openbaar te verdedigen in de Agnietenkapel
op vrijdag 17 oktober 2008, te 12.00 uur

door

Marian Elizabeth Counihan

geboren te Johannesburg, Zuid Afrika.

Promotor: prof. dr. M. van Lambalgen
Co-promotor: prof. dr. M.J.B. Stokhof

Overige leden promotiecommissie:
dr. J.J.H. Kurvers
prof. dr. G. Politzer
dr. V.A. de Rooij
prof. dr. K. Stenning
prof. dr. F.J.M.M. Veltman

Faculteit der Geesteswetenschappen

Cover design by Helen Counihan.
Printed and bound in the Netherlands by
PrintPartners Ipskamp, Enschede.

ISBN: 978-90-5776-185-0

For my parents

Contents

Acknowledgements

I first became drawn to the idea of doing research into human reasoning while a student in the Master of Logic programme at the ILLC, and following Michiel van Lambalgen's course on the psychology of reasoning. With a background in mathematics and analytic philosophy, most of the topics in psychology and semantics were new to me. Perhaps more 'new' for me though, was Michiel's research style. This involves original combinations of tools and insights from different fields to make a seemingly total study of the subject. It may not be to everyone's taste, but I found it exhilarating. For the first time in my brief academic life, the idea of research become appealing. By a fortuitous combination of timing, blind chutzpah and work, I was awarded a PhD position with Michiel and got the chance to pursue my growing interest. For this opportunity I was, and remain, exceptionally grateful. I am also grateful to Michiel for allowing me the freedom and making it possible for me to go my own way in research.

That was how it started. Of course, it didn't always go so smoothly. A few years later I was in the PhD-dip, demotivated and unproductive. We called in help, in the form of Martin Stokhof. Martin's supervising capacities are fairly legendary within ILLC/philosophy circles, and so it proved in my case too. His calm and encouraging presence in supervising our meetings helped me to understand better that doing research is not a high-jump competition, but a sculpting process, involving continuous formulating, adjusting, questioning, readjusting. This complemented the input from Michiel very well.

So with the inspiration and guidance provided by Michiel and Martin I gradually found my feet in research, and wrote this dissertation. I am very grateful for their patience and confidence, especially at times when I lacked both, and for their probing questions which raised the level of the arguments in the dissertation significantly.

There were other people who taught me much along the way. Firstly, I would like to thank Keith Stenning, whom I visited in Edinburgh in my first year of the dissertation. His long experience (and good memory!) of the field of psychology

of reasoning, and always thoughtful approach to experimental work, provided me with an invaluable introduction into the ways of psychology and enabled me to make a reasonable first stab at empirical work by myself. I feel privileged to have been able to learn from him.

Further experimental work for the dissertation was conducted in Hamburg, on the eastern coast of South Africa, with the women from the Keiskamma Art Project. The visit to Hamburg was a very restorative time for me, and although I didn't know it at the time, the work I did there became the basis for most of the rest of my dissertation, so it was a turning point in more ways than one. I would like to thank first and foremost, the wonderful Hofmeyr family: especially Graeme, for getting me down there in the first place, and Carol and Justus for making me feel so welcome. I'd like also to thank Jackie Downs, Nozeti Makubalo, Novuyani Peyi and Caroline Nyongo for helping me organise inteviews, and all the ladies from the art project and other residents in Hamburg and Ntilini for participating in my study. The interviews wouldn't have been possible without Zukiswa Pakama and Kwanele Ganto who translated for me, sometimes under rather difficult circumstances, and I am very grateful to them for their careful work. Finally thank you to Susan, Jacqueline, Florence and Maggie for looking after for me so well while I was there.

In my working environment in Amsterdam I experienced a lot of support from contact with my fellow PhDs, especially my roommates, Hartmut Fitz, with whom I had many lovely conversations and whose company helped me keep a sane perspective on things, and Katrin Schulz, whose quiet diligence initially intimidated and later galvanised me into working harder than I otherwise would have when in the office. Talking to fellow PhD students at the philosophy department helped me to know better what I expected and wanted from research, and for this I am grateful to everyone there. On the practical side of things, the ILLC has been a great institute to work in. It's a well-organised and friendly research environment and provides students with valuable cross-disciplinary exposure.

My more recent workplace is the University College Utrecht and over the last two years I have very much enjoyed and been buoyed by the supportive colleagues I have there. Being involved in teaching and tutoring activities was beneficial to my research, and being part of the university community there has made me feel much more comfortable with academic life. In particular thanks to Christel Lutz, not only for putting me in touch with UCU but also for helping me visualise post-PhD life.

Slowly developing in the background during the whole PhD project was my life in Amsterdam, and the quality of this life has contributed enormously to my energy for the dissertation. Because of this my friends deserve a special place in these academic acknowledgements. I can't mention everyone but I would like to say particular thanks to a few. Firstly, I would like to thank my two paranymphs, Lien de Jong and Martine Prange, who have supported and inspired me continuously over the last years, both academically and personally. They have,

in their different ways, taught me so much about Amsterdam, and made me feel welcome here. Two other friends should really also be my paranymphs: Hanna Schösler, who has been such a warm and loyal friend to me since my first year in Amsterdam, and Lotta Weckström, whose company on Wednesday afternoons probably did more than anything to get me through my PhD.

For help with the manuscript, I owe a debt of gratitude to Martine Prange for carefully reading many chapter drafts. Her comments made the last six months of writing a lot more manageable than they would otherwise have been. Also many thanks to Helen Counihan for help in designing a beautiful cover, Rik Nouwen and Frank Veltman for very helpful suggestions regarding the work covered in Chapter 2, and Hedde Zijlstra for help with translations in Chapter 2.

Finally, I would like to thank my family, who have supported me so much during difficult times. It is extremely sad for me that my parents are no longer alive to see me through to the end of the dissertation. I am happy and honoured that my siblings, Catherine, Helen and Jonathan, and my extended family, have rallied around so much and will be able to celebrate with me. I am very grateful to have such a family. Thanks finally to my new family, Jan-Willem, who has given me so much joy and love over the past seven years and who continues to do so every day.

Introduction

> "[O]nly irrationality is newsworthy."
>
> (Krueger & Funder, 2004, p. 318)

A browse through the psychology shelf at your local bookshop makes clear that bad reasoning is big business. Popular titles include

- "Inevitable Illusions: How Mistakes of Reason Rule Our Minds" (Piattelli-Palmarini, 1996, a best-seller in Italian),
- "Don't Believe Everything You Think: The 6 Basic Mistakes We Make in Thinking" (Kida, 2006, ranked number 38 on Amazon's list of over 5000 cognitive psychology titles[1]), and
- "How We Know What Isn't So: The Fallibility of Human Reason in Everyday Life" (Gilovich, 1991, number 74 in the same ranking).

Such titles are not very encouraging about the state of human reasoning ability. It seems that at every turn we are tricked by 'cognitive illusions' into drawing compelling but invalid conclusions. We might wonder: are things really that bad? Or is this just populist hype to sell books?

Actually, if some researchers in the psychology of reasoning are to be believed, things really *are* that bad. Byrne, Espino and Santamaria (1999) go so far as to blame the disaster at the Chernobyl nuclear faculty on the failure to draw the fairly simple *modus tollens* inference:[2]

[1] On www.amazon.com on 23 January 2008.

[2] The authors are here referring to the *modus tollens* inference based on the premises "if the test is to continue, the turbine must be rotating fast enough to generate emergency power" and "the turbine is not rotating fast enough". Using deontic logic, one could draw the conclusion that "the test must not continue".

> . . . [P]eople have difficulty making some inferences. The power plant workers in Chernobyl did not [make the *modus tollens* inference], and as a result partly of this inferential difficulty, the Chernobyl disaster occurred, with worldwide implications. . . . If we are to avoid repeating the mistakes of the past, we need to examine those mistakes carefully and establish how they came about (Byrne et al, 1999, p. 347).[3]

Indeed, if things are that bad, we need to examine very carefully if, when and how such mistakes come about.

So, what is the scientific research behind the popular titles? What kind of evidence is there for the claim that "mistakes of reason rule our minds"? There are several well-known research programmes feeding such ideas; perhaps the most well-known is the programme initiated by Tversky and Kahneman, investigating heuristics and biases, and associated with probabilistic reasoning. But another major source of empirical support is the research into logical reasoning, often termed the 'deduction paradigm' (Evans, 2002).

In the deduction paradigm subjects are asked to assess the validity of arguments. They are presented with a set of premises, and then either asked to decide whether or not a given conclusion follows from them, or to generate their own conclusions on the basis of the premises. This method is intended to evaluate the ability of subjects to reason logically. A central motivation for the research undertaken in this dissertation was to evaluate whether in fact studies using such a method have successfully done so, and if not, what kind of approach provides better access to the logic in reasoning.

Also in the deduction paradigm, the dominant characterisation of reasoning has been in terms of logical deficiencies and nonlogical influences. Indeed, from the first, studies of reasoning have produced negative findings. Wilkins (1928) found subjects to illicitly convert the quantifiers *all* and *some . . . not*. Woodworth and Sells (1935) reported so-called 'atmosphere effects' in syllogistic reasoning tasks: a negative premise increases the chance of a negative conclusion, and a particular (i.e. non-universal) premise increases the chance of a particular conclusion. Around the same time (the 1930s) Luria was conducting studies of reasoning with illiterate peasants in remote areas of Soviet Russia, and found "[t]he most typical responses of the subjects, therefore, were a complete denial of the possibility of drawing conclusions from propositions about things they had no personal experience of, and suspicion about any logical operation of a purely

[3]Byrne et al go on to explain that workers at Chernobyl might have not drawn the inference because they considered additional requirements from their background knowledge, and as such, "whether safety procedures had to be followed to the letter in such a case" (p. 347). If additional requirements are considered, then the *modus tollens* inference is not longer deployable. Human error is thus actually to be located in the judgement of whether or not the turbine's rotation speed provides sufficient reason to stop the test or not, and thus whether the *modus ponens* inference is appropriate or not; it has nothing to do with a fault in drawing the inference itself.

theoretical nature" (Luria, 1977, p. 108). In more recent syllogistic reasoning research, the negative findings have continued. Reasoners are subject to belief bias (Oakhill, Johnson-Laird and Garnham, 1989), typicality effects (Sloman, 1998), and other supposedly non-logical content-effects.

The field is usually seen as becoming fully-fledged in the 1960s and 1970s, initiated by the very well-known selection task, as developed by Peter Wason (1968a). Wason's very first presentation of the selection task (1966, p. 145) is tellingly headed "Errors in deductive reasoning", and his later conclusions about his subjects' performance in this conditional reasoning task (described in detail in Chapter 5) are just as pessimistic:

> The results, however, are ... disquieting. If Piaget is right ... then subjects in the present investigation should have reached the stage of formal operations. A person who is thinking in these terms will take account of the possible and the hypothetical ... But this is exactly what the subjects in the present experiment singularly fail to do. ... Could it be that the stage of formal operations is not completely achieved at adolescence, even among intelligent inviduals? (Wason, 1968a, p. 281)

Perhaps the most widely accepted conclusion arising from Wason's selection task is that reasoners suffer from confirmation bias; that is, they seek to verify, to confirm, rather than test and possibly falsify, their beliefs, expectations, or salient hypotheses (Evans, 1989, Nickerson, 1998). Wason himself saw this as the major finding resulting from his series of reasoning experiments (Wason, 1966, 1968b, 1972).[4] Also in conditional reasoning research, the findings have persistently been expressed negatively. Apart from confirmation bias, many other erroneous tendencies have been proposed. Evans long pursued an explanation of conditional reasoning in terms of very superficial matching bias (see Evans, 1998, for an overview); Byrne (1989) presents experimental evidence that subjects "suppress" valid inferences in certain contexts, while in a recent paper, Johnson-Laird and Savary (1999) report on "Illusory inferences: a novel class of erroneous deductions".

This litany of reasoning errors contrasts with a parallel, though less prominent, realisation in the field, that interpretation of the materials and construal of the task situation play a vital role in determining subject performance, thereby

[4] In another experiment in which subjects were asked to generate hypotheses regarding a series of numbers, Wason seems to despair at what he finds:

> There would appear to be compelling evidence to indicate that even intelligent individuals adhere to their own hypotheses with remarkable tenacity when they can produce conforming evidence for them. What makes people so narrow minded and so cognitively prejudiced? Why did they find these trivial games so difficult? (Wason, 1968b, p. 172)

undermining claims as to the significance of results from studies which do not take this into account. In fact, early papers by Wason (Wason, 1968b, Wason & Johnson-Laird, 1970) emphasise that the way the subjects 'structure' the task determines their card selection. Henle's (1962) paper is well-known as claiming that *all* seeming errors of reasoning can be explained by subjects' interpretation of the premises in combination with their construal of the task situation. Evans (1972) refers to this in arguing that reasoning research is too concerned with classifying behaviour only as correct or 'erroneous' (p. 382) and recognises that interpretation of the premises play a role (p. 373)[5]:

> In order to understand the psychological basis of subjects' behaviour it is suggested that at least two types of influence must be distinguished: Those relating to the subjects' interpretation of the sentences constituting the logical premises of the problems; and those arising from the nature of the mental operations required on a given task.

Interest in subjects' interpretative behaviour is thus present from early on in the field, but curiously such processes are often considered to be unrelated to reasoning, and hence to logical processes. For example, Evans associates the focus on deviation from the norm with an over-reliance on classical logical as a normative-theoretical model of behaviour. This leads him to formulate a neglect of interpretation as a neglect of 'non-logical factors': in arguing for more attention for interpretative processes, he says experimenters "have tended to overlook psychological explanations in terms of factors quite unconnected with logic" (1972, p. 374). Thompson (2000) explains the preference for 'abstract' materials, as an attempt to prevent interpretative processes from interfering with reasoning. Abstract materials, "which presumably are not subject to the same type of interpretative analysis as more 'realistic' materials", serve as a measure to "control or eliminate the role of interpretative processes in [theorists'] experiments" (Thompson, 2000, p. 212). It is unclear whether Thompson herself endorses this association of abstract material with absent or unproblematic interpretation; at any rate let me be clear that I do not. If anything, more abstract materials usually contribute to interpretational complexity. And conversely, much can be explained in subject behaviour by means of the normal, everyday use of constructions found in premises. Both Chapter 2 and Chapter 5 provide case studies of this phenomenon.

[5]The realisation of the importance of interpretation is sometimes even present in articles which report on reasoning errors. For instance, Byrne (1989) summarises the implications of the suppression-effect findings by saying "The results suggest that the interpretation of premises plays an even more central role in reasoning than has previously been admitted" (1989, p. 61). Later on, she says "The moral of these experiments is that in order to explain how people reason, we need to explain how premises of the same *apparent* logical form can be interpreted in quite different ways" (1989, p. 79, my emphasis). I will later return to the precise status of that 'apparent'.

Because many findings were characterised in terms of deviations from the classical norm, many theorists saw that the need for a positive account meant avoiding simply comparing behaviour with this norm. Since the norm was seen to represent the whole of logic, that meant jettisoning logic from an account of reasoning. Such a rationale, I think, lies behind the new wave of theories which aim to characterise reasoning positively, for instance, in terms of adaptive behaviours (Gigerenzer, 1999), and evolutionarily-based accounts (Cosmides & Tooby, 1989), but which explicitly set themselves off against logic-based theories of reasoning. Chapter 4 describes the assumptions which led to the jettisoning of logic from theories of reasoning, and shows how they stem from an overly-simplistic view of the relation between natural language and logic. When we take a more accurate view of the relation between natural language and logic, as the current work describes, it becomes clear that the negative findings in the psychology of reasoning provide no reason to dismiss a role for logic.

An existing seam in the literature of research *does* make the connection between interpretative processing and logical reasoning. This research aims to explain inferential behaviour in terms of more general semantic and/or pragmatic considerations. For instance, Hilton (1995, p. 248) understands that "Failure to recognise the role of conversational assumptions in governing inference processes can lead rational responses to be misclassified as errors and their source misattributed to cognitive shortcomings"; Thompson (2000, p. 212-3) makes the more general claim that "in order to have any explanatory power, a theory of reasoning must contain a theory of interpretation, which specifies how information is derived from the problem environment and applied to a given task domain. In a non-trivial sense, therefore, a theory of reasoning is a theory of interpretation." Commitment to a comparable view underlies the work of Politzer on the topic of human reasoning (Politzer, 1986, Politzer & Noveck 1991, Politzer & Macchi, 2002, Politzer, 2004), as well as the research programme of Stenning and van Lambalgen (2001, 2005, 2008). The work of both Politzer and Stenning and van Lambalgen are discussed in greater detail later on.

The current dissertation fits squarely into this tradition – that is, it propounds a semantically-grounded approach to reasoning. As such, the main aim is to contribute to our understanding of how semantic considerations shape performance in reasoning tasks. This is achieved by investigating and, when necessary, redescribing the relations between natural language, logic, and reasoning, as these notions are employed in experimental studies of reasoning. New experimental results from diverse subject groups contribute to the analysis and allow us to reinterpret earlier negative findings. By the end of the dissertation I hope to have revealed more of the intrinsic connections between interpretative and inferential processes, and, in doing so, to have undermined the evidence presented in service of negative conclusions about reasoning. In a sense the whole dissertation serves as a counterargument to the belief that "mistakes of reason rule our minds".

Some key terms

Before we start it is useful to sketch the role of some of the central concepts that are employed in the dissertation. I then briefly motivate and evaluate the empirical approach employed here, and finally give a brief outline of what the reader can expect.

Logic, semantics, interpretation

The terms 'logic', 'semantics' and 'interpretation' will occur frequently in the coming chapters. What do I mean by them and how do they relate to each other? Perhaps most distinct from other psychology of reasoning research is the conception of logic upheld here. In a very general sense, a logic can be defined as a collection of expressions of a language, a collection of structures (usually models), and a satisfaction relation between the two. A structure satisfies an expression if the expression is true of the structure. The final independent definitional parameter is validity, which can be expressed in terms of satisfaction. A conclusion is *classically* valid in case it is satisfied by *all* structures which satisfy the premises; there are however many alternative notions of validity, in which a conclusion need be satisfied by some preferred subset of these structures, or even a different set of structures (as in statistical reasoning). A structure consists of a domain and an interpretation function. Intuitively, the domain is what the language is about, and the interpretation function tells us what the expressions of the language mean, in the sense that it assigns suitable denotations to the non-logical parts of the language.

Each specific logical system studied by logicians reflects different aspects of the structure of the world in its structures. For example, first-order logic considers domains of individuals. It can talk about properties of and relations between these individuals in terms of sets which serve as interpretations of the predicate symbols in the language. It cannot talk about relations between sets. Modal logic is designed to capture notions of necessity and possibility; as such, it has a more complex language, including operators representing possibility and necessity, and more complex modelling structures, often including relations between possible worlds. In a way similar to modal logic, deontic logic seeks to represent the intuitions we have about obligations and permission. It is unimportant for the reader to know the details of these logical systems What I want to emphasise is that logic is essentially about meaning. It is about framing a situation in a certain way; abstracting away from or ignoring certain aspects of its structure and focusing on others.

This is a much broader and less monolithic conception of logic than that employed in many of the existing studies of reasoning. Chapter 4 presents arguments as to why the earlier, narrower conception of logic, more or less equating it with classical logic, is outdated, and does not have the sovereignty often attributed to

it in the realm of human reasoning.

Semantics is most broadly the study of meaning. In the area of linguistics it is usually bracketed off from syntax, which is concerned with grammatical form, on the one hand, and pragmatics, which is concerned with language use, on the other hand.[6] In studies of reasoning, as we will see in Chapter 4, a simplistic reading of the notion that logic is concerned with form as opposed to content, and the attribution of logical form solely on the basis of grammatical form, has often meant that logic is understood to be almost antithetical to semantics, where the latter is understood to deal with content.[7] This might partially explain why many researchers have been reluctant to take subjects' interpretations of task materials seriously, i.e. because it understood to be quite a separate issue from the logic of the materials – as the earlier comment from Byrne, "we need to explain how premises of the same apparent logical form can be interpreted in quite different ways" (1989, p. 79), demonstrates.

In this dissertation the opposite view is taken. Here I align myself with Stenning and van Lambalgen (2008) in claiming that the 'logic' of the task is *always* relative to an interpretation of the premises, since this establishes what it means to assume the premises are true, the inferences which can be made on the basis of them, and the kind of validity these have. Logic is thus essentially indexed to semantics. In fact, logic, as sketched above, can be understood as what is in linguistics termed 'formal semantics'.

Finally, it should now be clear that the interpretative processes are intimately related to logic because they are concerned with attribution of logical form to natural language expressions. Consider the example "All bears in the north are white". Interpreting such an utterance means at least determining extensions for terms such as 'bears' and 'white', for the present purposes (non-trivial, because, for instance, dirty yellow might count as white when 'white' is interpreted in combination with 'bear'), as well as figuring out what the domain is, over which the 'all' ranges, and the effects of changes to the domain on the truth of the generalisation. Does one brown bear falsify it? Or does its truth allow for some exceptions? There are many choices to be made in settling on an interpretation. The experimenter's interpretation of the task materials is only one among many, as we shall see, and is often highly contrived, in a sense which I make more exact in Chapter 2. Moreover, the experimenter's assumption that their interpretation is privileged, because it relies only on 'literal meaning', representing the truly 'logical' interpretation of premise materials, is flawed and thus ultimately untenable, as Chapter 4 shows.

[6]The semantics-pragmatics boundary is addressed in Chapter 4, where it is argued that traditional priority of semantic factors above pragmatic factors, in determining 'what is said', has been shown not to hold.

[7]This is a slight exaggeration, since of course the meanings of 'logical' elements of language, that is, 'and', 'or', 'if' and 'not', are understood to contribute to logical form. But such elements *are* considered as distinct from 'content', and it is in this sense that the claim is intended.

Reasoning

What is meant by reasoning? It is hard to find a general but substantial definition in the literature. This is by no means just due to laziness on the part of researchers to define their topic. Everyone will happily concur that reasoning is based on making inferences. But inference is ubiquitous: we infer emotions from facial expressions, infer shapes from outlines and shades, infer body position from middle ear fluid levels, even male toads, "roaming through the swamps at night, use the pitch of a rival's croak to infer its size when deciding whether to fight" (Gigerenzer & Goldstein, 1996, p. 650). All cognition involves "going beyond the information given", to use Bruner's (1957) phrasing. What then, is the proper research area for studies of human reasoning?

Until now, reasoning has chiefly been operationalised as assessment of simple arguments viewed as part of syllogistic and propositional logic. This involves drawing or judging a single conclusion from a limited set of premises, in one-off interaction (although the selection task is a notable exception to this format). Such an operationalisation is probably mostly historically-determined. Since one of the goals of this dissertation is to assess the previous research based on such an operationalisation, I likewise investigate reasoning in this format. But it should be clear that the use of this form of evaluating reasoning is the theorists' choice; actual behaviour does not respect our disciplinary boundaries. As will become clear, for a large part the reasoning in the context of reasoning tasks is found in the process of settling an interpretation. This has led Stenning and van Lambalgen (2008) to distinguish reasoning *to* an interpretation and reasoning *from* an interpretation, instead of distinguishing reasoning as such, from interpretation. The point is that reasoning is a concept much more broadly applicable than the above operationalisation suggests. This narrow definition serves a pragmatic purpose and need not be harmful, unless we assume that reasoning is restricted to reasoning *from* an interpretation, in Stenning and van Lambalgen's terms, in which case one runs the risk of making the artificial dissociation of interpretation from logic identified above.

Empirical access to reasoning

Given the above discussion of the intimate relation between interpretation and reasoning processes, it should come as no surprise that purely quantitative studies, in which subjects indicate their evaluation of an argument by ticking from a pregiven selection of conclusions, or being asked to say "what follows" from a set of premises, are here considered insufficiently penetrating. Such set-ups basically take interpretation for granted. As an attempt to counterbalance the neglect of interpretation, the empirical approach to reasoning employed in this dissertation is extended tutorial interviews with a relatively small number of subjects. Such an

approach is very unusual in contemporary studies of reasoning and thus deserves an extended introduction.

Quantitative studies are far and away the norm in psychology of reasoning. For example, the experimental set-up which generated Byrne's (1989) suppression-effect findings was as follows. There were three groups of subjects, and each group was assigned to a different type of premise set (e.g. one or more conditional premises; for more details, see Chapter 1). The subjects were given a booklet with a number of different versions of the premise sets of the same form. Versions differed only according to 'content', as in "If she meets her friend then she will go to a play" has the same form but differing content to "If it is raining then she will get wet".[8] Note here that it is simply taken for granted that sentences with different content, but superficially the same grammatical form, are indeed exemplars of one and the same logical form – a matter taken up more fully in Chapter 4. Each premise set was accompanied by three possible conclusions. Subjects were instructed (on the front page of the booklet only) to assume the premises were true and then to choose one of the conclusions, "whichever you think follows from the sentences" (p. 68). The conclusions chosen by subjects for each item were then counted and statistical tests (ANOVAs, one-tailed planned comparisons) reveal significant effects – such as, in this study, main effects of types of premise sets, and conclusion types, and an interaction between types of premise sets and conclusions chosen.

There are several reasons why such experimental methods are inadequate. Firstly, as mentioned above, the most problematic aspect to a set-up such as Byrne's is that it takes interpretation of the premises for granted. No attempt is made to find out what it means for the subject "to assume the premises are true". Under what circumstances is a conditional premise such as "If it rains then she will get wet" true? What counts as a falsifying instance to it? Questions like this remain open and very substantial, as Chapter 1 and Chapter 5 both make abundantly clear. Complexities at the level of interpretation are plausibly responsible for, among other things, the main effect mentioned, of differences in the rate of inference from the various premise sets. Yet Byrne's method (which is typical for reasoning studies) is unable to access interpretational matters. Related to this, there is no check on broader task construal issues – what does the subject see as their role, in the setting of the task? In less educated subject groups this matter comes forcefully to the fore, as the next chapter demonstrates. The 'laboratory' situation is, unfortunately for the quantitative researcher, just as meaningful to the subject as any everyday linguistic interaction. This makes control of the stimulus difficult to achieve, and impossible when interpretative variety and, more broadly, the subjects' own perception of the task, is not taken into account. Note however, that interpretative processes can be probed with quantitative methods, such as paraphrase judgements. These have been successfully

[8] In chapters 1 and 2 differences between such conditionals are elaborated.

used by researchers (e.g. Stenning and Cox, 2006, Stenning and van Lambalgen, 2004) to couple interpretative tendencies to reasoning behaviour.

Secondly, each response is treated as if it came from a separate subject, thereby losing potential information about individual tendencies and any resultant 'trait' differences between individuals. Presumably this is justified by the aim of Byrne's study, which is cast in terms of determining whether reasoning depends on "formal rules of inference" or "mental models" (as described by Johnson-Laird, 1983) and as such is concerned with finding a universal mechanism by which reasoning occurs. But, as Stenning and van Lambalgen (2008) have cogently argued, reasoning is not a 'mechanism' which operates in a contextual vacuum. Such conceptions are in any case severely challenged by findings from subject groups other than the undergraduate population in which quantitative studies are invariably conducted. In less schooled groups, the fact that simple reasoning tasks prove very difficult to subjects argues against the idea that such tasks provide neutral access to some universal reasoning mechanism. Additionally, there are many studies which have found significant differences between individual reasoners. Research within literate subject groups shows that performance on the selection task correlates with SAT scores (Stanovich, 1999). Other reasoning studies which look for individual differences found that reasoners can be classified according broad groupings in their interpretative and reasoning strategies (Stenning and Cox, 2006, Ford 1995).[9] Politzer (1981, reported in Politzer 2004, p. 99) even found a difference in interpretational strategies between arts and science students at a university! In sum, simply collating data across subjects, when subjects drawn from a narrow band of the population, and finding significant differences across experimental conditions, does not indicate empirical access to universal reasoning behaviour has been achieved.

Finally, reliance on quantitative studies has meant that most explanations concentrate only on the modal answer[10]; for example, in Wason's selection task, this means explaining the choice of the A and the 4 card (for more details, see Chapter 5). In fact, there is a much bigger range of responses present, albeit often in smaller numbers. These rarer occurrences tend to be sidelined. There can even be divergent rationales behind a single choice – as the conjunctive condition of the selection task, reported in Chapter 5, illustrates: some subjects choose a certain card because it potentially falsifies the rule; others because they think the rule is already false. In many forced-choice studies, a full range of plausible responses is not even offered, as is the case in Byrne's original study of the suppression effect

[9]Ford's 1995 study of syllogism-solving strategies indicates that reasoning subjects can be divided into two groups, namely 'verbal' and 'spatial', and inasmuch as this difference can be cast as that between subjects who focus on sentential form and those who don't, Byrne's treatment of all subjects as equal is potentially undermining her own argument for one or other type of 'mechanism'.

[10]The modal answer is that which is chosen by the highest percentage of subjects. It represents just one point on the distribution over the space of possible answers.

task described above. As Dieussaert, Schaeken, Schroyens and d'Ydewalle (2000) showed, subjects will spontaneously generate a much broader range of responses than the three simple propositions offered in Byrne's original study; this being so, a forced-choice study runs the risk of seriously underestimating the categories of response which need to be explained.

With the above arguments I hope to have given the reader a sense of why a purely quantitative approach is an inadequate instrument with which to access the complexities of reasoning behaviour. Before phenomena can be counted, classified and explained, the categories for classifation must be established. In short, there is still exploratory work to be done.

One might well wonder why there has been such a reliance on quantitative studies in the psychology of reasoning. A partial reason might be the bad name that 'think-aloud' protocols received around the time the field was burgeoning. Nisbett and Wilson's (1977) famous article suggested that we are often not conscious of, and thus not in control of, the influence that certain stimuli have on our behaviour. As such, verbal reports of reasoning might just be *ex post facto* rationalization, unrelated to the original automatic reasoning processes driving our behaviour.[11] But even before this, there was a preference for the laboratory stimulus-response method. The abstract which opens Wason's first publication about the selection task is telling (Wason, 1966, p. 135):

> One of the curious things about the earlier, introspective studies of thinking was that they demonstrated more than anything the inadequacies of their own methods. The course of thinking is affected by factors which are not available to introspection. Modern experimental work has avoided some of the issues by restricting itself to studying what people do when they solve problems.

Why "studying what people do when they solve problems" should be limited to asking subjects to tick boxes, as if thereby accessing "factors which are not available to introspection" is unclear to me.

Another reason that the stimulus-response method might have been viewed as sufficient, would be that experimenters have operated from the assumption that interpretation is a straightforward process. After all, one might argue, we are communicating successfully all the time, so interpretation can't be going that badly. What's more, spoken interaction often consists of incomplete or seemingly vague utterances, which are augmented by context and pragmatic conventions, yet in general this proceeds in a smooth and uniform manner. Why wouldn't it do so in the context of a reasoning task, one might ask, where the premises are written in full sentences, thus presumably exceptionally clear and unproblematic?

[11] Incidentally, such a distinction between conscious, verbal processing and unconscious, automatic processing has provided important impetus for dual-process theories of reasoning (Evans, 2003).

Well, as will become clear, as soon as we examine the specific linguistic material used in these tasks, the meaning of the premises is by no means clear or unproblematic. The language used is far less unequivocal as it might seem – especially, as registered above, in the case of 'abstract' materials. Moreover, in the typical laboratory set-up, the context needed for interpretation and assessment of the linguistic material that is used, is often found lacking. It will become clear that successful performance in most reasoning tasks requires a delicate combination of contrived and everyday interpretation of linguistic materials, which it takes quite some effort to maintain.

Perhaps Wason's comments should be understood as part of a broader *zeitgeist*, in which cognitive science was trying to establish itself a science. Bruner diagnoses this as resulting in an over-reliance on a computational metaphor for understanding the mind (Bruner, 1990, p. 4):

> Very early on [in the cognitive revolution], ... emphasis began shifting from "meaning" to "information", from the *construction* of meaning to the *processing* of information. The key factor in the shift was the introduction of computation as the ruling metaphor and of computability as a necessary criterion of a good theoretical model. Information is indifferent with respect to meaning.

It should be obvious that this criticism is especially pertinent in cases where meaning is central to the cognitive task being investigated. In studies of perception, by contrast, the "construction of meaning" may not play such a central role. But as we will see, constructing meaning is integral to reasoning.

It may be impossible to pinpoint why the emphasis has been so firmly on quantitative studies, but it is important to see the shortcomings of such a one-sided empirical programme and to be open to multi-method studies, which also allow space for in-depth qualitative studies.

Tutorial interviews: profits and perils

In this dissertation, empirical investigation took the form of extended interviews with subjects covering several reasoning tasks. With both schooled and unschooled subjects, structured interview protocols are invaluable for several reasons. Firstly, the considerations a subject takes into account in reaching an interpretation can be very revealing of the sometimes antagonistic considerations and resulting tensions that subjects deal with in understanding the task, both their role in it and the intended interpretation of the materials. These varying considerations are only identifiable in an extended interaction with the subject. Secondly, as mentioned above, studies which allow subjects to generate their own conclusions have revealed the great variety of responses that reasoning subjects make. If we are to take our subjects seriously, we need to understand what is behind this variety, and this can be aided greatly by a conversational interaction

with the subject. Once the varying rationale behind the full range of responses has been identified, we are better able to understand the range of choices generated in the quantitative studies of tasks, and not only the modal choice. Thirdly, tendencies in individual subjects can be better identified; this can then be used to design quantitative studies aimed at testing the validity of individual profiles. Finally, a spin-off value of the experiments with unschooled subjects, described in Chapter 1, is their ability to 'make the familiar strange', allowing us to explore assumptions about language functioning which in most cases are shared between the experimenter and the subjects of the experiment, but are nevertheless assumptions, only occasionally warranted in the experimental situation.

There are two obvious perils of an interview method: firstly, the inherent subjectivity of the interviewer/experimenter in the interviewing situation. The experimenter is bound to focus on responses they find interesting or relevant, and ignore others. The experimenter might inadvertently influence the interviewee in this, encouraging them to follow one or other train of thought, and dismiss some others. However, using a structured or semi-structured protocol somewhat counteracts this problem as it brings a measure of objectivity into the topics covered. Furthermore, the topic of the interview is a set of cognitive tasks, with clear goals, which further curbs the range of relevant responses, and as such, I think that subjectivity is a relatively minor problem here.

A second, related, problem is the richness and open-endedness of the data, which arises even in structured, goal-oriented interviews such as used here. In the stage of analysing transcripts, the experimenter constantly makes choices about which phenomena to discuss and which to ignore. As such, it depends on the discretionary ability and interest of the experimenter which semantic phenomena are identified and analysed, and how they are classified. This remains a hazard. Given the fact that the research is intended to be exploratory, however, and as long as no claim is made to exhaustivity, this aspect does not fatally detract from the value of the approach.

Perhaps the most basic problem associated with structured interview studies in general is the difficulty in creating the same meaning *in situ* across participants (Hill & Anderson, 1993), and this was something I was especially aware of in the study reported in Chapter 1, where the specific demands of the interview situation were foreign to the subjects. University undergraduates, by contrast, have probably quite rich ideas about what is expected from them in such an experiment (which is not to say this helps). But this problem becomes a strength for our current purposes: a unique aspect of the current research is that the process of interpreting the experimenter's words is part of the interview itself; it is not something which must be neutralised or 'controlled for'.

The empirical data of this dissertation are generated by very simple argument forms, but they reveal a wealth of complexity and depth in human semantic interactions. With this dissertation I aim to show that such interactions belong at the heart of a theory of reasoning, not outside it.

Outline

The dissertation is structured as follows. In the next chapter, the findings of an interview study with subjects with varying but low education levels, similar to those conducted by Luria (1976) and Scribner (1997), are presented. The aim of the chapter is to engage with and challenge the interpretation of illiterate reasoning behaviour as typified in Luria's negative conclusions. This is achieved by reanalysing the typical responses garnered in interaction with less educated subjects, but with heightened attention to semantic and pragmatic factors which are shaping their responses. The increased understanding of discourse contexts and their impact on the attribution of logical form as reflected in modern semantic and pragmatic study enables us to do this.

Chapter 2 takes as its starting point the finding that conditional premises of a certain format proved easier to the less schooled subjects than certain universally quantified premises. An explanation of this is sought in everyday usage of these constructions. To this end, a preliminary analysis is made of occurrences of *all* in a corpus of spoken English, and its equivalent in a small corpus of spoken Xhosa, the language spoken by the subjects of the experiment reported in chapters 1 and 2. Findings here support a semantic analysis of *all* generalisations as primarily contingent as opposed to law-governed, which in turn explains the difficulty especially unschooled subjects have in using such sentences as premises in the context of a reasoning task. A parallel analysis of the kinds of conditionals used in reasoning tasks explains why conditional premises, by contrast, are less problematic for subjects.

In Chapter 3 these findings are placed in the context of the broader debate on the cognitive consequences of literacy. A critical survey of empirical findings and general theories in this area indicates that much work has not been self-critical or precise enough. The work of David Olson is described and assessed in more detail, and an adapted version of his literal meaning hypothesis is proposed to explain the difference between schooled and unschooled subjects' performance in reasoning tasks.

Consideration of the notion of 'literal meaning' feeds directly into Chapter 4, the chapter which forms the analytic backbone for the whole dissertation, as it explains how much previous work in the psychology of reasoning has relied on an oversimplified picture of the relations between natural language, logic, and reasoning, giving rise to confusions about logical form, its relation to meaning in general and to literal meaning in particular. Here the main thesis is that a notion of 'literal meaning' of premises is uncritically and mistakenly used as a basis for normative judgements about logical reasoning. This notion of 'literal meaning' has served as a proxy for active investigation of interpretative processes, but it is imported from an outdated theory of linguistic meaning. I use recent arguments of philosopher of language François Recanati to show this.

Finally, Chapter 5 reports on an interview study on Wason's original selection

task and several variations of it. The analysis builds on the work of Stenning and van Lambalgen regarding the selection task paradigm, in identifying and describing the plausible interpretations of the presented rule which explain the large range of responses recorded (also in previous quantitative studies) and the task-level semantic parameters which interact with these to explain subject behaviour. This chapter is intended to show that standard reasoning experiments, although they have previously been aligned primarily with explanations of reasoning in terms of deviation from a norm, provide rich data for a positive theory of human reasoning.

Chapter 1

Logic, language and Khamrak

E: In the Far North, where there is snow, all bears are white. Novaya Zemlya is in the Far North and there is always snow there. What color are the bears there?

S: I don't know what color the bears are there, I never saw them.

...

E: But what do you think?

S: Once I saw a bear in a museum, that's all.

E: But on the basis of what I said, what color do you think the bears are there?

S: Either one-colored or two-colored ... [ponders for a long time]. To judge from the place, they should be white. You say that there is a lot of snow there, but we have never been there!

Subject: Khamrak., age forty, miller from remote village, illiterate.
(Luria, 1976, p. 111).

This is the transcript of an exchange which took place far away and long ago: in rural Soviet Russia in the 1930s, between the Russian psychologist Alexander Luria and his experimental subject, Khamrak. On reading it, one might well think that it comes from too far away and too long ago to any longer be of academic interest. One might write it off as an experimental anachronism, with entertainment value but not worth further serious study. That would be a mistake, as I hope to persuade you in this chapter. Such exchanges are highly relevant for contemporary theories of reasoning and, at the same time, yet to be fully understood.

The first notion we should disabuse ourselves of is that the above exchange was a one-off. In fact, the subject Khamrak's response is in key respects strikingly similar to other responses, not only those gathered by Luria, but also in more recent similar studies in Liberia (Scribner, 1997, Cole, Gay, Glick & Sharp, 1971), elsewhere in Soviet Russia (Tulviste, 1991), North America (Hamill, 1990), Zambia (Willemsen, 2001) the Netherlands (Kurvers, 2002) and, as reported here, South Africa. It is thus no experimental oddity. Rather it reveals something

general about an unschooled response to logical reasoning tasks. Now the sceptical reader might grant that it reveals something general about performance on these tasks but not something general about human reasoning, since the task is an induced reasoning situation and it is not clear how performance in such an 'artificial' situation relates to general reasoning behaviour. Again that would be wrong. I will argue that, although work in this area has, as yet, failed to connect adequately performance in logical tasks to general reasoning and linguistic behaviour, the connection *is* there. A more precise description of what it might be is precisely what this chapter seeks to contribute. I hope to show that for the most part illiterate responses on logical tasks are on a continuum with those given by schooled subjects; and that for both groups we can learn much about general reasoning behaviour on the basis of their performance, if we only we consider the meaning constructing processes which generate the transcripts, so to speak. Both the continuities and the discontinuities with schooled subjects are highly instructive for theorising about reasoning.

This connects up with another dimension of the study of reasoning which makes unschooled reasoning, as illustrated above, pertinent. This dimension is the representativeness of the existing experimental sample. Almost all empirical work in the psychology of reasoning has been conducted with a subject population of university undergraduate students. In terms of interaction with the presented material we might well wonder how representative this group is of the rest of the human population. Which parts of what they do are an artifact of their specific, and perhaps specifically literate, approach to the problem? Which aspects are driven by truly universal human cognitive traits? Mostly, we can't yet tell. As such, the validity of their response profile on any one task for a general theory of reasoning is unclear. Moreover, there is a more subtle confound in work with undergraduate subjects. Reasoning researchers themselves come from this same sub-population, at the high end of the literate scale, and thus there might be many shared assumptions about linguistic material which both the experimenter and his subject share, but which are not more widely shared in the general population. Towards this end, research with less literate groups helps to 'make the familiar strange' and enables us as researchers to see, and thereby evaluate, our implicit assumptions about language and how it relates to reasoning.

Apart from this more general motivation for studying unschooled subject populations, there are internal motivations. The first of these is the narrow range of materials used in reasoning studies with unschooled subject populations.[1] Data has mostly been collected in syllogistic-type tasks (Luria, 1976, Scribner, 1997, Hamill, 1990, although Cole et al, 1971, and Willemsen, 2001, also investigated conditional, conjunctive and disjunctive arguments), with a specific focus on the points of divergence from a normative/competence model. Roughly, when pre-

[1]The range of cognitive tasks investigated is broader, however, in other areas of study of cross-cultural cognition (Berry, 1971, Rogoff, 2003, Norenzayan, 2006, for instance).

sented with these particular materials, the responses of illiterate subjects differ from those given in the stipulated normative logical model, whereas literates' responses coincide with it. But we know very little about *why* we hear the responses we do and whether they are related only to syllogistic-type arguments, or to more general tendencies in reasoning. We may say, the representativeness of the reasoning material has yet to be tested.

This yields a second internal reason to investigate further unschooled reasoning behaviour: as yet, we have no positive account of reasoning for this group. What we do know is, as mentioned above, mostly expressed in terms of limits, of divergences. Early theorising tended to focus on discrepancies with schooled subjects – as we'll see here, while much later theorising has moved away from this experimental paradigm altogether (in this vein, see for instance Rogoff & Lave, 1984). There are certainly valid reasons to choose for this direction (which we touch on in the next section), but it does leave unexplained the reasoning phenomena picked up by the more traditional experimental task.

The experimental study here reported thus aimed to redress these imbalances in the following ways: 1) by broadening the subject sample, 2) by gathering data on a wider array of tasks across subjects in less literate groups, and 3) by submitting the data to semantic analysis which has been lacking in previous studies, thus contributing to a positive account of reasoning in unschooled subjects.

1.1 Previous work on logical reasoning

The seminal empirical study of the effect of literacy level on cognition was conducted by Alexandr Luria in 1930's Soviet Russia, at Vygotsky's suggestion. Luria saw his research as a means to verify the social-psychological thesis that "all fundamental human cognitive activities take shape in a matrix of social history and form the products of sociohistorical development" (1976, Preface, p. v). To conduct his study, Luria traveled to remote regions of Uzbekistan and Kirghizia where radical social and economic restructuring was going on at the time, part of the implementation of Stalin's first five-year plan. Included in the plan's aims were the elimination of illiteracy, and a transition to a collectivist economy, both of which required the large-scale introduction of schools with adult literacy programs and short-term courses for specific skills training, in rural areas. Luria wanted to see what the effects of these practices of literacy and of new economic activity, would be on individual thinking. He tested his subjects on a range of tasks, intended to chart diverse cognitive activities such as perception, generalisation, classification, deduction, reasoning, imagination, and self-awareness. Because of the politically sensitive nature of his findings, he did not publish them until much later, in the seventies, after he came into contact with Michael Cole.

Cole and his colleagues were by then participants in the burgeoning field of cross-cultural psychology (see for instance Cole & Scribner, 1974 or Berry

& Dasen, 1974) and conducted experimental studies aimed at replicating and furthering Luria's results. Cole et al (1971), for instance, reports on tests with logical problems with conjunctive, disjunctive and conditional premises.[2] The results were mixed but in all studies conditional premises generated the most correct responses – a finding which anticipates the results of the current study. One interesting deviation to this is in subjects' explanation of their answer once they had given it. Here, Cole at al found justifications to conjunctive-based conclusions to be more often correct (58% in non-literate adults) than those for either conditional (24%) or disjunctive based conclusions (29%). Unfortunately Cole et al do not specify the criteria by which responses were judged correctly so it is difficult to interpret these results further. Cole et al's conclusion was rather less equivocal: "The subjects were *not responding to the logical relations contained in the verbal problem.* Rather they were ... responding to conventional situations in which their past experience dictated the answer" (Cole et al, 1971, p. 188). To anticipate slightly: in much the same line of thinking as Luria, Cole at al suggest that reasoning based on logic and reasoning based on past experience are mutually exclusive. They go on: "In short, it appears that the particular verbal context and content dictate the response rather than the arbitrarily imposed relations among the elements in the problem." (*ibid.* p. 188).

Cole conducted much of his later experimental work with his colleague Sylvia Scribner, in Liberia, among the Kpelle and Vai peoples in the 1970s and 1980s. Scribner in particular focussed her investigations on reasoning skills and the tasks used to test them. In what follows I will analyse in greater detail the studies reported in Luria (1976) and Scribner (1997), and will refer to Cole et al (1971) only inasmuch as their conclusions diverge from Luria's. Other references, such as Hamill (1990) and Willemsen (2001) are not discussed further because they are judged to have sufficiently similar approaches and results to Luria (1976) to not warrant separate study.

Both Luria and Scribner tested their subjects on a range of cognitive tasks but within reasoning focussed on syllogistic-type tasks with a quantified or generalised 'major' premise and a particular statement as the 'minor premise', followed by a question.[3] An example is:

> All bears in the far north are white. Novaya zemlya is in the far north. What colour are the bears there?

Scribner reports on a slightly broader range of materials, including major (initial) premises of the form,

[2] A sample (disjunctive) item is:

> Flumo or Yakpalo is in the house. Flumo is not in the house. Where is Yakpalo?

[3] Note that what in this literature is labelled a syllogism is a far broader class of problem than strictly understood by the term. For him, a pair of premises, in which the first is "in the nature of a general judgement" – for example, *'Precious metals do not rust'* and the second is a particular proposition, such as *'Gold is a precious metal'*, qualify as a syllogism.

A dog and a horse are always together,
If Sumo or Saki drinks palm wine, the town chief gets vexed.

For discussion on the significance of these different premise forms please see the next chapter. In the experimental study reported here all syllogistic tasks used quantified statements in the premises.

	type of syllogism	
group/solution	associated with experience	not associated with experience
illiterate peasants *		
– immediate solution	9 (60%)	2 (15%)
– after conditional assumption †	6 (40%)	4 (30%)
young people ‡		
– immediate solution	15 (100%)	15 (100%)

Table 1.1: Luria's results (1976, p 116):
*from remote villages (15 subjects);
†"from your words I gather that ...", and note that this category cumulates with those solved immediately;
‡with short-term education, farm activists (15 subjects)

	% type of justification		
	'theoretic'	'empiric'	'arbitrary'
Kpelle villagers	22.3	68.1	9.6
Kpelle students	75.0	21.9	3.1
American students	82.3	3.1	14.6

Table 1.2: Scribner's results (1997, p112)

As for their findings, tables 1.1 and 1.2 give the precise percentages garnered in Luria's and Scribner's studies. The tables should be read with the following in mind. Scribner reaches her categorization by classifying the justifications given to initial yes/no answers as follows: included in the category **theoretic** are "statements explicitly relating the conclusion to the information contained in the premises"; **empiric**, "statements justifying the conclusion on the basis of what the subject knows or believes to be true"; and finally, **arbitrary** covers irrelevant, idiosyncratic and "don't know" responses (in further studies, Scribner absorbs the 'I don't know' answers into 'empirical' – for example in Scribner, 1997, p. 130). Luria is not explicit about his classification, although the role of justifications for yes/no answers is certainly a factor in determining whether the subject 'solves' the task or not. One would assume that 'solution' means giving what is in Scribner's terms a 'theoretic' justification, but note that *all* syllogisms

'associated with experience' are solved when the experimenter stresses the qualification "but what can you conclude, *on the basis of my words*?". This suggests that 'empiric' justifications might also be counted as solutions for Luria in this class of problems. A reason for the more equivocal results garnered by Scribner might be her wider range of premises – the significance of which we will come to later on in the chapter. What is interesting is that Scribner, and Luria as far his categories align with Scribner's, take for granted that the categories "theoretic" and "empiric" are mutually exclusive. Not only do some subjects give 'mixed' answers which make reference both to personal knowledge and to information given in the premises – see Khamrak's last turn in the opening dialogue, repeated below – but, as is detailed further on, most answers are necessarily mixed, since personal, and sometimes not-so-personal, knowledge influences how subjects interpret the premises.

Let us take another look at the exchange opening the first chapter and see how it illustrates the type of response Luria and Scribner reported.

Khamrak, aged 40 years, illiterate:

E: In the Far North, where there is snow, all bears are white. Novaya zemlya is in the Far North and there is always snow there. What color are the bears there?
S: I don't know what color the bears are there, I never saw them.
...
E: But what do you think?
S: Once I saw a bear in a museum, that's all.
E: But on the basis of what I said, what color do you think the bears are there?
S: Either one-colored or two-colored ... [ponders for a long time]. To judge from the place, they should be white. You say that there is a lot of snow there, but we have never been there!

The first thing which springs out is an obvious discrepancy in understanding of the discourse: what the experimenter intends for the subject to understand is patently different from what the subject understands is required from him. This is addressed in the next chapter. Notice however, that perhaps even inadvertently, the subject gives the 'right' answer the experimenter is looking for – 'they should be white' – albeit with a justification which is ambiguous as to how exactly the inference is drawn from the premises.[4]

As already mentioned, the responses show remarkable patterning when considered together. In both Luria's and Scribner's studies, the majority of responses to the syllogistic problem fall into one of two groups. Most common, according to Luria, was a refusal to give a positive answer because of a lack of personal knowledge of the premises ("I don't know what color the bears are there. I never

[4]Does the subject infer for himself that because there is snow, the bears should be white, or does the subject rely on the experimenter's relating of the snowy environment to the colour of the bears?

saw them."). Scribner also reports this type of response as most common. To the problem

> 'All Kpelle men are rice farmers; Mr Smith is not a rice farmer. Is he a Kpelle man?'

she heard the following answer:

S: I don't know the man in person. I have not laid eyes on the man himself.

E: Just think about the statement.

S: If I know him in person, I can answer that question, but since I do not know him in person I cannot answer that question.

Refusal to answer is the initial response for the majority of subjects. Luria summarizes thus:

> "The most typical responses of the subjects, therefore, were a complete denial of the possibility of drawing conclusions from propositions about things they had no personal experience of, and suspicion about any logical operation of a purely theoretical nature" (1976, p. 108).

I will argue that the characterisation of the 'I don't know; I can't answer that'-type response as refusal to engage with premises is inaccurate and that such answers are the product of a very 'normal' engagement with the premises. The label 'refusal' should be read not as refusal to engage with the premises, but as 'refusal to give a definite answer', for good reason, in a way that is laid out in more detail in the discussion section of this chapter. For now, it bears noting that 'refusal' does not indicate that there is no reasoning going on in the subject's head – as the last turn in the excerpt from Scribner amply illustrates!

A second type of response observed was an engagement with the premises, on the subject's own terms so to speak: that is, by expanding and/or fitting them to (presumably) known and accepted conventional situations. Here it should be noted that the subject often constructs sophisticated logical argumentation with a combination of the given premises and their own additions to these. This type of response doesn't seem to arise immediately, but only after repeated questioning by the experimenter. This is illustrated in the above transcript of Luria's conversation with Khamrak (p. 108), as well as in the following, in response to the problem

> 'Cotton can only grow where it is hot and dry. In England it is cold and damp. Can cotton grown there?'

a subject responds (after first answering 'I don't know'):

> **Abdurakhm, 37, illiterate.**
>
> S: I've only been in Kashgar country; I don't know beyond that.
>
> E: But on the basis of what I said to you, can cotton grown there?

> S: If the land is good, cotton will grow there, but if it is damp and poor, it won't grow. If it's like the Kashgar country, it will grow there too. If the soil is loose, it can grow there too, of course.

Here the subject appears to realise the experimenter is expecting him to draw conclusions beyond his personal knowledge, and does so, but by means of the introduction of adapted conditional premises, which presumably can be accepted on the basis of personal experience ('If it's like Kashgar country...'). Luria has a dim view of such performances: "Frequently they completely ignored the premise and replaced the inferential process by considerations of their own, for example ... they would introduce general, rumor-based opinions ... In short, in each case, they would avoid solving the task" (1976, p. 107)

At this point I can't help but to anticipate the current analysis: this is a serious under-evaluation of the subjects' performance, as will become evident in the rest of this chapter and the next. The subject engages with the premises – but, again, has to interpret them first. Luria's argument has a hiatus where the step from presented material to interpreted material, the 'logical form' that the subject reasons with, should be. Such a hiatus is more generally evident in the reasoning of researchers in this area, and revealing and repairing it is really one of the main aims of this dissertation.

Getting back to the matter at hand, the findings so far need to be summarized. According to Scribner and Luria's conclusions, there are two main characteristically unschooled responses to the reasoning problem. The most common initial reaction is one of refusal to answer on basis of the premises because of a lack of personal acquaintance with them. The second type of reaction is engagement with the premises on a 'personal' level: adapting the premises to align with one's own knowledge, or conventional wisdom, even including them in a narrative form of discourse. For Luria (1976), the results showed that illiterate subjects did not grasp the logic of the syllogism: they "are not perceived by these subjects as unified logical systems." `[p. 106]` These subjects are limited to concrete, situational thinking, incapable of abstract thought. On the other hand, subjects with "well-established forms of theoretical thinking", those with even a short (1-2 years) time in school education, "tend to grasp the over-all logical structure" (pp. 103/4). For Luria, these findings provided evidence of the deep impact of literacy on forms of thinking. He concludes his study with the following summary: "as literacy is mastered, and a new stage of social and historical practice is reached, major shifts occur in human mental activity. These are not limited simply to an expanding of man's horizons, but ... radically affect the structure of cognitive processes" (p. 161). Crucially, though, Luria's conclusion is unable to explain why in some conditions, and accompanied by certain assumptions, subjects can and do give responses which reflect a grasp of the logical structure of the syllogisms.

Scribner's (1997) conclusion is more qualified: "In the present analysis, formal

evidence in a problem controls performance of the schooled groups. The non-schooled groups show no such homogeneity: some respondents appear at either end of the spectrum, handling all problems empirically, or, in fewer numbers, handling them all theoretically. The great majority have a mixed strategy, relying now on the formal information in the problem, now on the evidence external to it" (1997, p. 134). This outcome casts doubt on the unequivocal effect of schooling/literacy on reasoning proposed by Luria, and suggests the impact lies more in styles or preferences when approaching a cognitive task. As such, the main achievement of schooled subjects would be merely to make more consistent use of an already (i.e. pre-schooling) available interpretative strategy for drawing an inference. But this does not provide an explanation of what is driving choices in the so-called 'mixed strategy' used by unschooled subjects. Rather, our locus of interest is the conditions under which unschooled subjects adopt one or other interpretation of the premises, what these various interpretations are, how they are related to task material, and to more general issues of task construal.

Furthermore, we are interested in whether the response profiles generated on syllogistic-type tasks give a true representation of the logical reasoning ability of unschooled thinkers. Is it really so that unschooled subjects typically refuse to draw conclusions about situations they have never experienced themselves? Or is the difference in the response profile identified by Luria and confirmed by Scribner somehow amplified by the peculiarities of the particular tasks they used? What we really need to know is *why* such responses are elicited. Why *this* type of reaction as opposed to some other?

Even so, at this juncture, given the rather non-constructive nature of Luria's conclusions, and to a lesser extent Scribner's, the cognitive researcher might be inclined to give up on this experimental paradigm altogether. See for example Dasen's recommendation (1977, p. 197) that researchers should rather focus on tasks "which test the same cognitive structures, but which are directly relevant to the daily activities and interests of the subjects." There has been a large research programme in this vein; that is, one that focuses on tasks moulded to the everyday environment in which cognitive development happens, and going so far as to argue that the individual cannot be studied as a separate entity from the culture (Rogoff, 2003, p. 42). Such an approach is certainly valuable and yields very different knowledge but leaves unexplained the differences, and, more importantly, the *similarities*, between groups on the 'standard' cognitive tests. As we will see, the concerns of unschooled subjects in these tests remain similar across very different environments, and this itself argues for their validity.

The current work aims to combine the best of both approaches. The original reasons for the set-up used in the classic studies of Luria, Scribner and colleagues, still stand, namely to use measures of thinking that bore little relation to everyday life, so as to tap inferential ability independently of background knowledge or convention. This may *seem* to have been unsuccessful, but these studies have not yet had the benefit of a modern treatment based on pragmatic and seman-

tic theory; insights from these areas allow the common basis of response across groups to be seen more clearly. Our conviction is that there is a common fundament to linguistic 'culture' which supersedes other cultural boundaries; hence that the standard reasoning task *is* relatable to the everyday linguistic activities and interests of the subjects – to paraphrase Dasen – and hence still worth studying.

1.2 Experimental set-up: subjects, protocol, materials

The participants in the current experiment were of varying age and education level, all resident in the small town of Hamburg on the coast of South Africa's Eastern Cape province. The majority of the subjects were employed as beaders or embroiderers in the local art project. Three of the women worked as domestic help; several were unemployed. Of the 29 subjects, six had had no education at all, thirteen had less than ten years of education, and had been out of the education system for more than ten years, and ten had completed high school within the last 20 years. I will maintain these three groupings when discussing my data:

Group 1: No education (6 subjects)

Group 2: Between four and ten years of education, left the educational system more than ten years ago (13 subjects)

Group 3: Graduated from high school within the last twenty years (10 subjects).

All subjects were interviewed individually, in Xhosa. Each interview lasted approximately 30 minutes and participants were paid R20 for taking part. The translator, Zukiswa Pakama, was a Hamburg resident, well-known to most of the subjects. Pakama is a native Xhosa speaker but attends university (in English) in nearby East London and speaks fluent English. The interviews were video-taped and later transcribed. At the beginning of each interview, the subject was asked about their age, language skills and educational history. Then the tasks were introduced by saying that the experimenter wanted to see how the subject used language and would describe situations and ask them questions on the basis of these. It was emphasised that there were no right or wrong answers and that the primary interest was how they used Xhosa.[5]

The materials used in the interviews were drawn from the range of tasks previously used in psychology of reasoning studies, including syllogistic-style material *à la* Luria and Scribner's version of the syllogism, such as:

[5]Not being a Xhosa speaker myself I thought this would help mitigate against potential authority issues in the situation.

Suppose there's a faraway country called Markia.
All the woman in Markia are married.
Fatma is a woman who lives in Markia.
Is she married?

and conditional reasoning material, such as used in the suppression effect task, viz:

Suppose there's a girl called Thembi living in Hamburg.
If Thembi wants to see her boyfriend, then she goes to East London.
(If Thembi has enough money, then she goes to East London.)
She does want to see her boyfriend.
Does she go to East London?

These two sets of materials were considered of primary interest, given their importance in previous and current theorising about reasoning behaviour (Luria, 1976, and Byrne, 1989). Two other tasks were conducted: the quantifier interpretation task, and a thematic version of the Wason selection task. These are not reported on here since the issues they raise differ considerably from those we are currently concerned with.

1.3 Results and discussion

In this section the aim is to get behind the categorization given by Scribner and Luria of unschooled reasoning performance. This is tackled by a two-fold strategy for both the syllogistic and conditional material. As to be expected, observed phenomena mostly do not restrict themselves to one or other type of premise material, so the consideration of the two separately is somewhat artificial. However, this is counteracted by mentioning explicitly when a phenomenon occurred predominantly in response to only one or other type of material.

Firstly, new data from syllogistic-type tasks is analysed, primarily by identifying to what extent it is comparable to data gathered by Luria and Scribner. Once this is established, I set out to explain the categories defined by Luria and Scribner in terms of reasonable semantic interaction with the premises. Secondly, I present and analyse data from the conditional premise set, by the same strategy. This enables us to get some insight into difficulties associated with specific task materials, an issue which is taken up more fully in the next chapter.

Although some quantitative results will be given, these should be seen as a means to better contextualise the qualitative results, and not an end in themselves. The aim is not to provide a comparison of accuracy between subjects, as defined by some predetermined competence model, but to uncover the range of factors influencing subjects' interpretations of the material and the kind of role these play in generating responses. A full understanding of subject behaviour is

considered to be only possible when such interpretative factors have been taken into account.

There are several remarks to be made at the most general level. Firstly, the distinction between groups was not as sharp as Luria found it to be, especially with conditional reasoning premises. As the table indicates, qualified or elaborated answers occurred consistently across all groups with conditional premise material. This contrasts with the response pattern for syllogistic-type material, which was much more differentiated by group. (See table 2.1). This finding is in line with that of Scribner's study (1997, p. 112) which also included some conditional material as described above.

Secondly, and related to this, there was heterogeneity within groups, and particularly within the schooled subjects (group 3), more so than identified by either Luria or Scribner (e.g. Nomalungisa, Nokulula). Recall Scribner's summary of her finding: that "The great majority [of unschooled subjects] have a mixed strategy, relying now on the formal information in the problem, now on the evidence external to it." (Scribner, 1997, p. 134) The current results are consistent with this finding, but do not align as well with Scribner's finding that "formal evidence of the problem controls performance of the schooled groups." But this is not the first registration of a more heterogeneous response within the schooled group. Tulviste's (1991) experimental study, also with a more varied range of premise materials, among subjects from a similar range of educational backgrounds, reports remarkably comparable findings. Although Tulviste found the correlation between number of correct responses and educational level of the subject to be statistically significant , he also found that "In many cases, the protocols of the subjects who had attended school for quite a long time and those who had not attended school at all were practically identical" (Tulviste, p. 134). As an example he cites the responses of a 26-year old subject with 10 school grades to the following problems:

> Every morning Asan plays on the kamuz [a Kirghiz musical instrument]. Did Asan play his kamuz yesterday morning or not?
>
> S: How should I know?
> *The problem is repeated.*
> S: Maybe he did play.

> Asan and Kenesh always drink tea together. Asan is drinking tea now. Is Kenesh drinking tea now or not?
>
> S: No, he's not.
> E: Why do you think so?
> S: Because he may not be there now.

These two responses fall squarely into the categories given by Luria to describe illiterate performance: to paraphrase him, the subject "ignores the premises or replaces them with considerations of his own." Yet this is a subject with ten years or more of schooling. Scribner already observed that villagers who had

attended school more than ten years previously and has since returned to with rural life responded more like unschooled subjects; so this is not surprising with respect to my group 2, nor with respect to group 3. Both Luria's and Scribner's schooled subjects were either still attending or had very recently left the educational system, whereas in the current study many of the Group 3 subjects left school several years ago. But this is by no means a complete explanation, as one of my subjects, Thembakazi, illustrated. Thembakazi was still in high school when she participated in the experiment, yet gave responses usually associated with unschooled subjects, namely changing the status of the antecedent from hypothetical to actual and elaborating on the given premises, viz:

Thembakazi, group 3:

(After simple condition: If Ntombi wants to see her boyfriend then she goes to East London. And she wants to see her boyfriend. Does she go to East London?)

E: And what if we know that if Ntombi has enough money for taxi fare, then she goes to East London? And she wants to see her boyfriend.

S: She will go, because she's got the taxi fare.

E: And what if she doesn't have the taxi fare?

S: If Ntombi doesn't have the taxi fare, and she wants to see her boyfriend, then she will borrow some money from somewhere else, and go to East London, if she really wants to see her boyfriend.[6]

In this case, Thembakazi was sitting in the room with her grandmother and principal guardian Susan, and I later found out that there have been heated domestic discussions about exactly such situations. It seems likely therefore, that Thembakazi was addressing her defiant answers not only to the experimenter but to her grandmother too! In the majority of her responses Thembakazi gave 'theoretic' responses, which sometimes contradicted factual knowledge – 'snow is black' in one condition – thus illustrating that the adoption of a particular mode of response is fairly shallow and dependent on particular aspects of a condition. Such cases remind us of the matter of task construal, i.e. what kind of discourse the subject assimilates the material to, and the contribution that contextual and individual factors may make to this.

[6]There were other school-going subjects who gave such responses – for example:

Abongile, group 3:

E: Ok next one. Ntombi's boyfriend lives in East London. If she wants to see her boyfriend then she goes to East London. So today she wants to see her boyfriend. Do you think she goes to East London?

S: I don't know.

E: Is it a strange question?

S: No, it's not but I don't know.

See also transcripts from Nozuko later on.

	% theoretic answers with . . .			
	quantified premises		**conditional premises**	
Group	ratio	percentage	ratio	percentage
1	7/23	**30%**	9/18	**50%**
2	25/38	**66%**	28/32	**81%**
3	24/33	**73%**	28/33	**85%**

Table 1.3: Comparing initial responses across premise forms

For both Luria and Scribner the choice for one or other strategy – "theoretic" or "empiric" – is considered to be a function of the subjects' epistemic relation to the material: whether she or he is 'familiar' with the described situation, and if so, whether it contradicts reality as known. Luria's presentation of results attests to this; and Scribner sums up her findings with: "Adoption of a particular mode is influenced [by] . . . especially the factual status of the information supplied in the premises" (1997, p. 134-5). As will become evident, the relation is somewhat more complicated than this allows, since the subjects' response is shaped by the specific semantic structure they discern in the material. Subjects respond to material first and foremost by interpreting it, not just by accepting or rejecting it.

On the other hand, as we will see, several subjects exhibited consistent patterns of response, indicating that they attributed a stable semantic structure to an argument form with varying content. In some cases the subject gave consistently "theoretic" answers (Novuyani, Sebenza, Notuthuzelo); other subjects consistently volunteered different kinds of extra information in their responses (Zukiswa, Nomhle) – behaviour which Luria might have classified as "replacing the inferential process by considerations of their own" but which will here be interpreted more favourably. Individual subjects who repeatedly gave answers of a similar form are discussed in more detail as we get to each of these forms.

For comparative purposes a table summarising one aspect of the data is presented: the proportion of initial responses which are characterisable in Scribner's terms as 'theoretic'. By initial responses, it is meant: only the responses to the initial *modus ponens* question or the basic syllogistic question (All A are B, x is A, is x B?). This is judged to be most comparable to other data, and least affected by the vagaries of a specific interaction. Everyone has heard the same story at the point when these answers are given. Once a conversation has started the task of the subject becomes more open-ended and to compare later responses to a posited norm would lead to distortions in the data, if the foregoing conversation is not taken into account.

The rate of immediate elicitation of correct answers for the syllogistic task in the unschooled group was 30% – comparable to Luria's and Scribner's findings

of 40%[7] and 22.3% respectively. As for the conditional elicitation rates, the only figures with which to compare them would be those from studies with undergraduate students, which have been typically somewhat higher (e.g. 96% in Byrne, 1989), but this might again be a 'literacy' effect. But most interestingly, the results tabled show a discrepancy between the groups according to the type of premise they had to reason with. Conditional premises proved to be easier than quantified premises across the groups, and, as we'll see when looking at justifications, show less variation across groups – something which might explain the more muted group differences found by both Scribner and Tulviste.

Further what does a table such as 1.3 really tell us about illiterate reasoning? Very little. In fact, the main goal of this chapter is to deconstruct this and similar presentations of results to reveal the similarities of category of concern across schooled and unschooled groups when it comes to determining the semantic structure of the premises. The eventual aim would be to construct another table based on just these categories, what might be called 'semantic parameters for reasoning', which underlie the results presented in the form of tables 1.1, 1.2 and 1.3. These will range from general concerns with the hypotheticality of the premises, the informational structure of the task, epistemic access to the premises, to more specific concerns such as the appropriate interpretation of a generalisation. After all, as Scribner already pointed out,

> Before drawing conclusions about the subject's reasoning processes, then, the investigator must determine what problem the subject is actually attempting to solve. (Scribner 1997, p. 108)

1.3.1 Syllogistic-style task results

In this section the data is analysed by addressing the following questions, generated by the new data or left unanswered by previous studies:

(1) to what extent are the categories of response above evident in syllogistic-type tasks in the current study?

Regarding the 'refusal or elaboration of premises' phenomena:

(2) why do subjects refuse to reason with given premises? do only unschooled subjects do so?

(3) do individual subjects consistently refuse to reason with the premises or is refusal related to particular materials and inferences?

(4) why do subject go beyond the given premises when reasoning, assimilating them to their own experience? again, is this phenomenon associated only with certain groups, materials or inferences?

[7]This is a composite figure, based on 9/15 for syllogisms "associated with experience" and 2/13 for those "not associated with experience" (Luria, 1976, p116).

As to the specific formulation of premises:

(5) what is the interpretation of the quantified statement assumed by the subject? can this explain the 'elaboration' phenomena mentioned above?

(6) does the interpretation of the quantifier vary across materials – for instance more or less plausibly law-like generalisations – and if so, does it do so consistently across subjects and or/groups?

And with respect to more general task construal issues:

(7) to what kind of discourse do subjects assimilate this kind of material? is there evidence that they reflect on the purpose of the discourse and the 'naturalness' of the premises?

The syllogistic materials were comparable to those used by Luria and Scribner, and the subjects were comparable in terms of literacy levels, but in other potentially operative respects (language, economic activity, cultural milieu – e.g. access to 'literate' mediums such as television) the subjects in the current study differed from both Luria's and Scribner's subjects. Note that these groups are, in turn, also different from each other. The circumstances of the three testing situations differ along several dimensions, such as historical and geographical setting (more or less isolated communities, familiarity with visiting researchers), but also the experimenter's relation to the subjects, and the available recording equipment. Given these discrepancies in the subject groups and testing situations one might expect the data to be barely comparable across groups. What is striking is the extent to which the data *are* comparable. Scribner herself considers an even wider range of studies and concludes:

> "the consistency of the basic findings is impressive. Not only are the quantitative results strikingly uniform ... but certain qualitative aspects of performance are so similar that it is often difficult to distinguish the translated interview protocol of a Uzbekistanian from that of a Vai – cultural and geographic distance notwithstanding."

In the current study it is even more remarkable that the interviews yielded comparable data because the subjects' daily activities and cultural milieu was so different from that of either Scribner's or Luria's subjects. For starters, the subjects in the current study live in a mixed and dynamic community, with much urban contact – the city of East London is only an hour's travel away – and are familiar with modern media such as newspapers, television, even mobile phones. Think again of Luria's hypothesis, that "all fundamental human cognitive activities take shape in a matrix of social history and form the products of sociohistorical development". Now even though each subject group considered here has had a different social history, their cognitive strategies as measured by these tasks prove to be comparable. This suggests that there are but a few factors which have

a non-zero value in the matrix that Luria talks of – the foremost of these being education, and paired with this, literacy. The mechanisms, through which these factors operate to influence cognition, are explored further in the next chapter.

The applicability of Luria's/Scribner's classification

As table 2.1 preliminarily indicates, the data obtained in the current study are highly comparable to Luria's and Scribner's data. Specifically, many subjects exhibited what Scribner called 'empirical bias', although mainly in the second sense of assimilating the premises to own experience. Few outright refusals to reason with the given materials were observed. The few came from the older subjects who had had no education at all. This second aspect of the empirical bias was much more prevalent and was also present in all groups. This is a point at which my data differs from Luria's but aligns with Scribner's results. Also, recall Tulviste's (1991, p. 134) reporting that "In many cases, the protocols of the subjects who had attended school for quite a long time and those who had not attended school at all were practically identical". As mentioned above, the prevalence of 'mixed' responses even in the most schooled group in the current study is probably partly a recency of schooling effect, as described in Scribner, 1997, and partly a matter of the expanded test material. An example illustrates:

Nozuko, group 3:

E: Suppose all the women in Nigeria are married. Now there's a woman called Connie and she's not married. Can we say she lives in Nigeria or not?
S: What kind of clothes do they wear in Nigeria?[8]
E: Just suppose the world is a strange one in which all the women in Nigeria are married.
S: We can say she's a Nigerian but she hasn't got married yet.

The classification of such a response as an unschooled-type response, with attendant connotations of defective reasoning, is fitting in Luria's terms – recall his description of subjects who "replaced the inferential process by considerations of their own". But this is also a fine example of what will here be argued to be a very reasonable, even pressing, tendency of subjects to reckon with or enquire after a basis on which the quantification is justified, and thus what kind of semantics it should get. Nozuko's last turn in this exchange can be seen as making a case for why a law-like reading universal quantification doesn't make sense: at any point there will be young women who are yet to be married, who are nevertheless falling into the domain of quantification by living in Nigeria. Compare this to the intended reading of the quantification for correct performance in the task. The

[8]Note that in Xhosa custom a married woman indicates her status by wearing a specific kind of dress: she always covers her head, and wears an apron. There is also special language only for married women: some objects (e.g. cattle, stones) have two names, one for use by married women, one for use by everybody else.

premise "All women in Nigeria are married", like the premise "All bears in the far north are white", needs to be interpreted on a universal domain, and strictly – not tolerating exceptions. As will become clear in the next chapter, this is a highly uncommon and contrived use of the term. In its everyday use, quantification is explicitly or implicitly subject to domain restriction. Sentences like "Everyone came to the party" illustrate: "everyone" is clearly intended to refer to a very restricted set of people; an unrestricted-domain reading is as good as unintelligible. So the rejection of a universally quantified statement might be better viewed as the attempt to clarify domain restriction. Further, it can either be read as a law-like or accidental generalisation, but both of these are problematic as they either result in the correct response being uninformative or unfounded, as will be elucidated in the next chapter. What is key is to understand that the correct answer is *not* informative in the standard sense, but only informative about the knowledge state of the respondent. We should understand the exchange not as a failure of the subject to "accept the logical task" (Scribner, 1997), but as a negotiation between subject and experimenter as to what the logic of the task *is*.

Indeed, stripping the evaluative layer from the categorization given by Luria and Scribner reveals the concern for semantic factors which has been identified elsewhere and in different subject groups. Luria lists three factors which he judges to limit capabilities for theoretical, verbal-logical thinking, and which lead to the responses as categorised. Firstly, subjects "mistrusted" the initial premise. Under this he subsumes also refusal to answer and ignoring of the premise altogether. Secondly, subjects do not accept the premises as universal. "Rather they were treated as particular messages reproducing some particular phenomenon." Thirdly, subjects did not treat the premises as forming a unified set, but as "three independent and isolated particular propositions with no unified logic." Observe that the first two of these factors can equally well be uniformly explained by the concerns described in the previous paragraph, in that instance articulated by a schooled subject. That is, "mistrust" of the initial premise and failure to treat it as universal is actually a result of uncertainty about the appropriate domain and type of quantification intended, including the means to verify it. This last could explain the many responses along the line of "I don't know; I've never been there/met the man", that is, under a contingent reading of the quantifier, in which all instances need to be verified *before* the generalisation can be. Allied to this, taking the contingent reading of the generalisation decouples the intended relation between the two premises, and thus could explain what Luria diagnoses as failure to treat the premises as a uniform set. Please see the following chapter for a full description of this phenomenon.

The 'refusal' phenomenon

There were very few cases, as mentioned, of flat refusal to reason with the premises.[9]When this did occur, it was after the first question, right at the beginning of the interview, in both cases with conditional premises. It only happened with unschooled subjects. In one other case, the subject at first answered, and then refused to answer at a later point in the exchange:

Susan, group 1:

E: Another one. Remember that it doesn't matter if it's true or false. Just listen to the words. Suppose all the birds in Cape Town are penguins. (translator has to explain what they are; subject nods). If someone sees a bird in Cape Town, what kind of bird will they see?
S: When she sees it in Cape Town?
E: Yes. Repeats question.
S: It could be a dove, or a raven, or a swallow.
E: But ignore what the real world is like, and just pretend that all the birds are penguins in Cape Town, then if you see a bird, what kind of bird is it?
S: *It will be a bird, maybe a dove or any other kind of bird.*
...
E: Imagine we are talking about the North Pole, and I tell her that all the birds there are penguins. If you see a bird there, what kind of bird will you see?
S: *I don't know.*
E: And if I tell you all the birds there are penguins?
S: I don't accept that.
E: Why not?
S: *I don't know these penguins, I've never seen them.*

[9]Here are two examples. On being asked whether Ntombi would go to East London, Susan replied:

Susan, group 1:

S: How will I know? I don't know.
E: If you just listen to the words, and to the situation they are describing, an (repeats question), will she go then? We are not referring to a specific person here.
S: I don't think she will go because her mother is watching her and not letting her go.

And another example, this time in a story about a girl called Ayanda:

Vulelwa (group 1):

S: I don't even know this Ayanda, and I don't know if she will go to East London. Does she live in Hamburg, this Ayanda?

This last turn resembles the answers garnered by Luria and Scribner as initial responses and here forms an interesting contrast with the response offered in the more familiar setting of Cape Town.

What also might be contributing to Susan's answer is a lack of the relevant taxonomic knowledge: that penguins are a type of bird, distinct from say doves, ravens and swallows. Although the translator explains this to her, describing the different appearance of the birds, she might still not 'get' the categorization. This might also have been a factor in Luria's study: with premises such as "Precious metals do not rust" we might well wonder whether subjects have the concept "precious metals" at all, and thus what meaning such a premise might have for them. See further discussion on this topic in the section on familiarity issues, and in Chapter 4 on literal meaning.

More generally, characterisation as refusal to reason is only accurate when the subject is assumed to have taken the intended interpretation of the premises, but, as the next chapter will show, this assumption is not generally warranted. The intended reading, in which the universal generalisation gets a law-like reading, is highly contrived and infrequent in spontaneous speech. Moreover, taking this reading results in a question-answer pair with an abnormal epistemic structure. This kind of question-answer is possibly dominantly used in schooling environments, so that unschooled subjects may dismiss the intended reading, if they consider it, because they fail to recognise the task as best fitting into the *test* discourse genre, with its peculiar epistemic structure.

What unschooled subjects are likely to be doing is interpreting the universal generalisation, in the way it would be in everyday discourse, and the question in line with normal question-answer structure. If this is the case, then 'refusal' answers are very much answers *on the basis of* the given premises, and as such evidence of reasoning as much as a schooled "yes" would be. Luria's equivocation of refusal to give a definite answer with refusal to engage with the premises stems from his failure to take interpretative processes into account, so that he cannot see a gap between the presented premises and his own (or someone else's) interpretation of them. In fact, the premises need a highly specific interpretation to get to the 'correct' answer.

There were few further subjects who gave an outright refusal to answer the question as posed. This might be a matter of politeness, of cultural norms, familiarity with other media, or of the relatively higher tolerance of my subjects to strange questions – the majority worked on an art project which had a continuous stream of foreign visitors and a documentary had recently been made about the town, for which some of the older subjects had been interviewed and thus could have gotten used to answering strange questions! More seriously, outright refusal was, as described, given as an initial response but subjects mostly consequently participated further. As such, refusal should perhaps primarily be regarded as a feature of task construal in the broadest sense, which, although influenced by the premises, had more to do with the whole setting of the experiment. (Reasons for

refusal related to the exact premises are discussed in greater detail below, in the paragraph 'The role of experience'.)

"Failure to accept" the given premises

The second feature of unschooled subjects' reasoning was broadly characterised by Luria as "failure to accept the premises" as given. This includes ignoring, rejecting, distorting, and elaborating on given premises. Note that, once again, this is only an accurate description when the subject is assumed to have the intended interpretation of the premises.

Ignoring the premises occurred in extreme varieties in the responses of only a couple of subjects, but within these quite consistently. For example, Susan was interviewed twice. The first reply in her first interview was reported above. She continued to give such answers throughout the first and the second interview, often seeming to ignore the given premises altogether. Here is another example of Susan ignoring premises, this time from the second interview:

Susan, group 1:

E: Next one. It's about the moon and on the moon all the stones are green. And a man goes to the moon and he finds a stone there. What colour is that stone?
S: *It's white and shiny.*
E: How do you know?
S: I just think that.
E: But remember, I said that all the stones are green there.
S: Yes.
E: So do you still think it's white?
S: It's only this one that's white.

In this case the subject's response cannot so easily be explained by a case of lack of taxonomic knowledge. She seems to truly not engage with the premises as a logical unity, and at first appears to ignore the first one altogether. My honest impression with this subject was that she simply wasn't paying much attention at this point. She often seemed bored and indicated that she found the questions rather bizarre. A lack of understanding of the experimenter's intentions and the aim of the interaction could certainly play a role in many illiterate subjects and cause them to be less attentive than they might otherwise have been, had the goal of the interaction (and their role in it) been clear to them.

Another subject who consistently gave Luria-style refusals to reason with the given premises or reasoning on the basis of personal knowledge was Vulelwa – see the following two excerpts:

Vulelwa, group 1:

E: Now suppose there are no cows in the whole of the Eastern Cape. And Hamburg is in the Eastern Cape. Will there be cows in Hamburg?
S: Because of what?

E: Just suppose that something happens, and there are no cows, in the whole of Eastern Cape. And Hamburg is in the Eastern Cape. So will there be cows in the Eastern Cape?
S: There may be cows.
E: Why?
S: There may be cows. If you say there are no cows in the Eastern Cape, there may be cows in Hamburg, even though there are no cows around Hamburg.
E: Even though Hamburg is in the Eastern Cape?
S: Yes.

...

E: Now suppose there are no cows in England. And there is a place in England called Fawley. Will there be cows there?
S: I don't know.
E: But what if I tell you there are no cows in the whole of England, and Fawley is in England. Will there be cows there?
S: This question is so difficult.
E: OK. I'll repeat it. In England, and the whole of England is an island, there's no cows. Now there's a small town on the island, called Fawley. Will there be cows there?
S: No, there are no cows there.
E: Why?
S: *Because you say it's a small town and there's no grazing fields there.*

We now address these types of responses in detail, again with the aim of uncovering semantic reasons for such responses, thus allowing us to understand them more constructively and on a continuum with schooled responses.

The role of experience

Content effects are well-documented in reasoning research. Belief bias phenomena (Oakhill, Johnson-Laird & Garnham, 1989) and typicality effects (Sloman, 1998) are two-well known examples of highly literate subjects reacting to quantified premises on the basis of their own understanding of them. With illiterate subjects content effects also operate, seemingly to an even greater degree, under some conditions. Luria's assessment is that there was typically "complete denial of the possibility of drawing conclusions from propositions about things [the subject] had no personal experience of" (Luria, 1976, p. 108). His results indeed show a marked difference in response on familiar versus unfamiliar materials: of the illiterate subject some 60% solved familiar problems immediately; only 15% solved those "not associated with experience" immediately (p. 116). Familiar problems involved "experience transferred to new conditions" – e.g.

> Cotton grows well where it is hot and dry. England is cold and damp. Can cotton grow there or not?

Unfamiliar problems were those "not associated with experience"; an example is the 'white bears' syllogism reported in the transcript above:

> In the Far North, where there is snow, all bears are white. Novaya Zemlya is in the Far North and there is always snow there. What color are the bears there?

Young subjects with some education solved both types of problems 100% of the time.

How is this familiarity/unfamiliarity effect comparable to 'content effects' identified in other psychology of reasoning studies? Scribner labels it an extreme form of content effect, and labels it "empirical bias" – that is, as it is occurs in reasoning studies with literate subjects, the effect of problem content which 'distracts' the reasoner from the formal task. In unfamiliar cases the bias is acting to such an extent as to function as what Scribner calls an 'organiser', resulting in a judgement that the problem is "*in principle* unanswerable", so that the subject does not engage with the premises at all, except to explain why s/he is not able to accept them. Scribner reports that on some problems 'empirical bias' entered into 75% of the responses; in others it fell to as low as 30% (1997, p 135). She remarks that this was related to the "factual status of the information supplied in the premises" but does not specify further.

It is important to keep considerations of factual knowledge and familiarity with material distinct, although of course they are related. With familiar materials subjects are simply likely to know more and so factual status of the premise or conclusion becomes a potentially salient consideration for the subject. In contrast, with unfamiliar material there is less danger of coming to a valid but known to be false conclusion, which could moderate the interference of own knowledge, but there is clearly also sometimes an independent reluctance to draw any conclusions about exotic material. Increased familiarity leads factuality to be an issue; less familiarity seems itself to be an obstacle to reasoning with the premises.

In my study I found this pattern in unschooled subjects. Consider, for example, the responses of Malinge (group 1) on two items: one about a truly unfamiliar situation, the stones on the moon, the second about a new object in a familiar domain, washing clothes.

Malinge, group 1:

E: This one is about the moon. And on the moon all the stone are green. A man, he goes to the moon and he finds a stone.
H*: Is he a Shangaan?[10] (general laughter)
E: OK so this man goes to the moon and he finds a stone there. what colour is that stone?
S: I won't know.
E: But suppose all the stones on the moon are green. And the man, he finds the stone on the moon. What colour do you think it is?
S: *If there are green stones, for sure then the man got a green stone.*

[10]The Shangaan are a southern African tribe – it's not clear what reason there is for saying this, unless they are purported to have magical powers enabling them to do things like fly to the moon.

...

E: Suppose there's a new kind of washing powder. It's called Cillit Bang. and suppose that if you use this new kind of washing powder you don't need to use water. And suppose you are going to wash something with this new kind of washing powder. Will you use water?
S: To wash clothes with this powder?
E: Yes.
S: *I don't think there is something you can wash with,without water.*
E: So you always have to use water?
S: In my knowledge, it's like that, you can't wash clothes without water.

* H = Subject's husband

So we see unfamiliarity does not always prevent reasoning according to the premises. Indeed, when the premises refer to situations about which the subject could reasonably have no own knowledge, it might seem less odd to rely on what the experimenter says – to take their word for it, so to speak. The subject has only that (linguistic) information as a source for inference. The experimenter is thus unproblematically the provider of the relevant information. In some of Luria's examples of "unfamiliar situations", on the other hand, the subjects might have some idea of what happens up north, and of what kinds of bears they are, or infer that the experimenter thinks they do – after all, *they* are being asked about it. This is exactly what we see happening in the excerpt from Malinge's transcript. Whose information takes precedence? Moreover, these "unexperienced" conditions are still about their world, and could plausibly be imbued with some political significance or used for political ends (although that's hard to imagine with the white bears example!). All this would influence how the subject responds to the task: they don't have the benefit of the tester's perspective, in which similarity of the test items is apparent. See further discussion of this in the questions and informativity section.

Moreover, there are several ways in which task material can be "not associated with experience": in some cases it's an unfamiliar situation being described (white bears in the north of Russia); in other conditions, the unfamiliarity centres on the objects or concepts under discussion. What Luria calls unfamiliar is of the former kind. However, Tulviste (1991) found that subjects with some schooling are inclined to respond to task materials of the latter kind 'logically'. That is, when asked to draw conclusions about unfamiliar kinds, they draw their conclusions from the given premises, more so than when familiar kinds are referenced. With premises like 'All precious metals do not rust. Molybdenum is a precious metal. Does molybdenum rust or not?' school-going children offered "theoretical"[11] bases for their answers which refer only to information given in the premises

[11]Tulviste maintains Scribner's categorization of responses; theoretic means "the subject draws only on the data contained in the problem and on the logical correctness of the conclusion from the given premises" (Tulviste, p. 120).

– much more commonly than in conditions where the premises described familiar, everyday situations. Another problematic aspect of the "associated with experience" and "not associated with experience" distinction used by Luria, is that it's not clear how he sorted his test set into these two classes. He gives an example of each – as mentioned above – but where would "Precious metals do not rust. Gold is a precious metal" fit in? Did his subjects know what gold is or not? Did they know the category 'precious metals'? If the answers are "yes" and "no" then we have a Tulviste-style unfamiliar condition in which we might expect more logical answers.

Questions and informativity

As we've seen above, the epistemic relation of the subject to the test material is an important variable in determining subject response. Comparing the question-answer pairs for the intended interpretation of these tasks with normal question-answer structure is illuminating in this respect. First, let us examine the epistemic structure of an everyday question and answer. Usually, the questioner indicates, by the act of asking, that they themselves do not know the answer to their question. Also, they signal to whoever they address the question, the addressee, that they expect *them* to know the answer. Consider the example (with polite forms stripped off for convenience):

> Q: Which way is the train station?
> A: Carry on over the intersection and then take the next right.

This is an unremarkable, if curt, exchange. But now imagine being asked the question by a train conductor (for the railway company operating out of that station). It would be very strange indeed to hear her ask such a question, because we expect her to already know the answer, even if we do too. It would be even stranger if the train conductor asked a more specific question, say:

> Q: Which way is the station staff room?[12]

because here we expect her to know the answer sooner than we, the passing pedestrian, could be expected to. In normal question-answer pairs, thus, there is epistemic asymmetry between questioner and addressee, which motivates the asking of the question in the first place.

This contrasts with the situation in the reasoning tasks as described above. Let us consider a sample question-answer pair – on the intended competence model – in this situation:

> Q: In the Far North, where there is snow, all bears are white. Novaya Zemlya is in the Far North and there is always snow there. What colour

[12] Or, even more distressingly:

> Q: Where's the brake?

are the bears there?
A: White.

Here the questioner has provided information before asking the question, information *with which* one can in principle answer the question. By giving this foregoing information, the questioner signals that he has the necessary knowledge to answer his own question, and more than this, creates common knowledge of it. Yet he asks the question of the subject. Now the subject is expected to answer *on the basis of information given by the questioner*. To our school-trained eyes, this might be quite normal. But imagine that you haven't got the benefit of school-trained eyes which immediately recognise that this is one of those circumscribed contexts in which the epistemic asymmetry doesn't hold. Then you would find it odd. You might find it as strange as the train conductor asking you the way to the station, after telling you how to get there. Note here that also subjects' familiarity to the material will play a role here in determining the sense of this.

Where do question-answer pairs which violate epistemic asymmetry occur? Certainly they are a large part of formal schooling, since any kind of test forms such a case. Are they restricted to school contexts? Not necessarily; plausibly any learning situation involves similar questions directed to the learner on the part of the instructor. But in learning a skill, especially a physical one, such questions would be restricted to displays of skill, of *know-how*, not of knowledge, *know-that* (even if this is in turn to show *know-how* in thinking).

About those bears: who knows what colour they are? Who cares? No-one in the situation need *really* know, or care. We are being asked to display our knowledge, as evidence of our skill in manipulating information. The question is not a question after information as such, but only after information as much as it shows that we have understood the intended coherence of the foregoing premises with the premise. However, if you are interpreting the question as one after information, then you would try and offer your own knowledge of the situation as an answer. Could this be explaining the 'personal' interpretation of the premises Luria saw?

Often a subject seems to employ a mixed strategy, using the premises to draw a conclusion but when asked for justification being informative by suggesting justification for the premises themselves and not *using* them as justification. This was common especially with conditional premises, and specifically subjects often seemed to want to offer justification for the conditonal relationship; for instance

Nomhle, group 2:

E: If Ntombi wants to see her friend then she goes to East London. And she does want to see her friend. Will she go to East London?
S: Yes.
E: Why will she go to East London?
S: Is it because she wants to see her friend?
E: Yes, she wants to go.

S: *Maybe she can talk to her parent and tell her that she wants to go and see her friend. Maybe Ntombi is interested to go and to see her friend but now she doesn't have enough money. So maybe her parent will give her money to go to East London to see her friend.*

See further discussion of this in the section on conditional premises.

Problems with generality

Luria found that subjects had specific difficulty with "the universal nature of the premises", which was "in all cases ... not respected". In the current study we also found that universal premises proved tricky for subjects, but for a variety of reasons. Why this is so is explored further in the following chapter but here there are several varieties of generalisation which deserve attention.

Problematic generalisations 1: Law-making statements

Certain premises in the syllogistic task materials lend themselves to a deontic reading: for example, "all people who own a house pay house tax", as used in both the current study, and also in Scribner's study with the Kpelle (Scribner, 1997, p. 131). So the house tax generalisation was often interpreted as "all people who own a house *should* pay house tax" or "all people who own a house *must* pay house tax". Indeed, the descriptive generalisation would presumably only hold because of an underlying decree – no-one except the rich and philanthropic would elect to pay house tax. As will be argued in the next chapter, a descriptive and law-like reading of the generalisation would be highly contrived for the subjects of the current study; the same would have held for Luria's and Scribner's subjects. In the following transcript, a high-school student explicitly ties the strict/law-like reading to the deontic reading of the generalisation:

Thembakazi, group 1:

E: More school problems. This one says suppose that all lawyers smoke cigarettes. And suppose that all people who smoke cigarettes also drive fast cars. Then, according to this problem, can we say that all lawyers drive fast cars?
S: (repeats problem) No, they don't all drive fast cars.
E: Why not?
S: Because *it's not law, people are not forced to drive fast cars.*
E: But just listen to the words, it doesn't matter if it's true or not. Suppose it's happening somewhere else, and suppose that in this place, that all lawyers ... [premises repeated]
S: Do I have to believe this is happening?
E: Yes, and if you pretend this is true, then can you conclude all lawyers drive fast cars?
S: Yes.

Assuming subjects do go for a deontic reading of the generalisation – as paraphrasing supports – this changes the semantic structure of the task, because, for example, someone not paying house tax *violates* but does not *falsify* the law. In some cases, the deontic statement supports the descriptive statement and subjects are happy to use the latter to draw conclusions about individuals; but when the descriptive generalisation is required as a premise to conclude something about a whole group – often leading to an unlikely scenario – then the subject points out the gap between deontic and descriptive statements. This can be understood as a default or generic variant of the descriptive generalisation, i.e. tolerating counterexamples, with the strict law-like reading applying only to the deontic reading – "all people *should* pay house task and *generally people do* – but not always". This explains the following excerpts:

Nonkululeko, group 2:

(preamble) All people who own houses pay house tax. Sabelo does not pay house tax. Does he own a house?
S: He doesn't have a house if he's not paying.
E: And now suppose that none of the people in Cape Town pay house tax. Do they own houses or not?
S: They have houses.
E: Why?
S: They can have houses because there are places where you don't pay tax, like the squatter camps.
E: So they can have houses and not pay?
S: They may, they can live at the squatter camps.

Rosie, group 2:

E: Suppose that all the people who own houses pay house tax. And suppose Luazi owns a house in Hamburg. Does he pay house tax?
S: He has to pay if he's got a house.
E: And suppose Sabelo doesn't pay tax. Does he own a house?
S: He doesn't have a house, if he's not paying tax.
E: Now suppose none of the people in Hamburg pay house tax. Does that mean that they own their houses or not?
S: Those who have houses, they will pay tax and those who don't have houses, they won't pay tax.
E: So if no-one in Hamburg pays tax, does that mean that no-one owns a house in Hamburg or not?
S: They may have the houses, but they decide not to pay.

Sometimes the deontic is used to justify why a descriptive generalisation *would* result:

Nokulula, group 2:

E: OK. Now suppose that no-one in Hamburg pays house tax. And remember that everyone who has a house does pay house tax. Does that mean that people in Hamburg own houses or not?

S: If they are supposed to pay and don't pay, then the houses will be taken anytime, because they are breaking the law. They are not paying, even though they are supposed to pay. So.
E: So they break the law basically?
S: Ja.

As is obvious from the above transcripts, subjects are often taking a deontic reading of the descriptively intended premises of the form "All people pay house tax". A deontic reading has a different semantic structure to a descriptive interpretation of the generalisation, and, in particular, does not support the required inferences without an additional premise which connects the law to actuality i.e. one which states that everyone actually does obey the law. Precisely the plausibility or justification of this necessary but implicit additional premise is what subjects are concerned with in the above transcripts.

Problematic generalisations 2: Generic statements

Generic statements are characterised by their tolerance of putative counterexamples, which turn out to be mere exceptions to the rule. The statement "chairs have four legs" can be accepted along with the existence three-legged barstools which, strictly, falsify the statement. In fact, it is sometimes surprising how tolerant generics are taken to be, as the following example illustrates.

Headline: **'Women long for plastic surgery'**
Subheader: 'One in two young women are so dissatisfied with their appearance they would consider plastic surgery, a new poll has revealed'.

(ITV.com news website, Jan 29, 2007)

According to the British ITV's news website, just half of a set is enough to justify a generic claim!

This issue is relevant because Luria complained that his subjects never maintained the 'universality' of the quantified premise but in fact he ignores the possibility of a generic reading of the premise. In fact, many of his subjects' responses contain a *more* generic variation of the premises than those originally presented by Luria: "If the land is good, cotton will grow there", "Each locality has its own animals" (pp. 108–109).

Positing a generic reading for the universal premises weakens the link between that premise and a particular second premise, because it comes to depend on the specificities of the protagonist. The possibilities for the relation between protagonist and the generalisation widen from just exemplification to include also exceptionality. Being asked to justify your conclusion then becomes a request for a justification of the choice between an exemplar and an exception interpretation of the protagonist. Merely mentioning a protagonist is singling them out, in a sense, and we might thus even expect a tendency towards the exception reading.

In the following two excerpt we see a subject twice switching to an exemplar reading when the justification for an exception reading is ruled out:

Nokulula, group 2:

E: Suppose all people who own houses pay house tax. And suppose Luazi owns a house in Hamburg. Will he pay house tax?
S: No, he doesn't pay.
E: Why?
S: *For example, I've got a house in Hamburg, and I built the house myself, so I don't pay house tax.*
E: OK. But now suppose there's a new law, that all people who own houses must pay house tax. Then would you have to pay house tax?
S: *Yes, then you pay.*
...
E: Ok. And now suppose I tell you all Xhosa people own cattle. And Peter is some man, we don't know whether he's Xhosa or not. But suppose we know that he doesn't own cattle. Then can he be Xhosa?
S: *He can be Xhosa because there are Xhosa who don't have cattle.*
E: But suppose that all Xhosa own cattle.
S: *He won't be Xhosa if all of them have cattle.*

Reassessing Luria and Scribner's findings

In the foregoing we have explored aspects of the semantic structure of reasoning tasks with quantified premises with the aim of evaluating whether Luria was right to conclude that unschooled reasoners "are limited to concrete, situational thinking". Similar to Scribner it was found that less schooled reasoners tended to employ a "mixed strategy", but it was argued that the seeming 'mixed' character of their responses stems from a unified set of semantic concerns, about, amongst other things: the epistemic structure of the task; related to this, the epistemic standpoint of the subject and experimenter relative to the task materials; the relation of the protagonist to the generalisation, and the intended interpretation of the generalisation. In the upcoming chapter, when we compare quantified and conditional formulations of premises, it will become clear that especially this last matter illustrates why unschooled subjects are more justly seen as 'normal' language users than as non-logical reasoners.

To sum up, it has become clear that Luria's neat distinction between schooled and unschooled subjects reflects only a very superficial understanding of the reasoning processes of subjects. Not only were typically 'unschooled' responses present also in schooled subjects' performance in the current study, but, prefiguring terminology from Chapter 4, once we take into account subjects' processing in reasoning *to* an interpretation, the responses of unschooled subjects can be seen to be very reasonable and understandable even from our highly literate point of view. In the following we extend the analysis to conditional reasoning tasks.

1.3.2 Conditional reasoning results

Some background: the suppression effect task

The suppression effect task was first reported in Byrne (1989). Subjects are presented with a set of sentences comprising either one or two conditional sentences ('If she meets her friend she will go to a play') and a simple sentence ('She meets her friend'). The second conditional sentence is judged to bear either an "additional" or an "alternative" relation to the first one. For instance, 'If she has enough money, she will go to a play' would be labelled "additional" to the first conditional because it suggests an extra requirement to make the consequent clause 'She will go to a play' true. Another kind of conditional is labelled "alternative" because the antecedent contains another requirement which is by itself sufficient to make the consequent true. See sample sentences in table 1.4.

Premise label	**Example premise**
simple	*If she meets her friend, she will go to a play.*
additional	*If she has enough money, she will go to a play.*
alternative	*If she meets her family, she will go to a play.*

Table 1.4: Labels for the different conditional premises in the suppression task

This set-up provides the three different sets of materials used by Byrne. The first set is of "simple arguments": one conditional and one simple sentence per test item. The others sets, "alternative" or "additional arguments" have an additional conditional premise, respectively an alternative or additional one. Note that the categorization into "alternative" and "additional" is based not on structural features of the conditional, but on prior (to the task) interpretations. The conditional relationships are just as open to interpretation during the task by the reasoning subjects. The experiment is premised on idea that logic is monotonic – i.e. for instance, an inference drawn in the simple argument condition should also be drawn when new premises are added – even if those specify extra requirements to make the consequent true. The link to monotonicity can only be made, however, by in turn supposing that a logical interpretation of the material treats the conditional premises individually – so that no 'compound' conditional premise would be generated. We return to this point in Chapter 4.

The 'suppression effect' is the label given to a pattern of responses; namely that the presence of certain types of conditional premises leads to lower elicitation rates for some conclusions, than in the condition where no 'extra' premises have been added i.e. the simple argument condition (see above).[13] So for instance, in

[13]In the context of this set-up, Byrne's description of her finding as a "suppression effect"

Inference type	single conditional premise	+ alternative premise	+ additional premise
Modus Ponens (MP)	96%	96%	38%
Modus Tollens (MT)	92%	96%	33%
Denial of the antecedent (DA)	46%	4%	63%
Affirmation of the consequent (AC)	71%	13%	54%

Table 1.5: Rates of inference in the suppression effect task (from Byrne,1989)

the group which gets the 'simple' arguments, rates of MP are higher (96%) than in the group which gets the 'additional' arguments (38%), but the same as the group working with 'alternative' arguments (also 96%). A similar pattern can be observed in the MT inferences (92% $\searrow$ 33%). The lower rates of DA and AC inferences are further reduced only in the presence of 'alternative' premises: 46% $\searrow$ 4% and 71% $\searrow$ 13% respectively. This pattern of responses is summarised in table 1.5. Clearly it makes no sense to compare percentages with conversations, to measure how 'typical' the current subjects are compared to Byrne's. This is not possible, nor is it the goal of this study.[14] Rather, the current study is in-

is somewhat self-serving. This stems from her experimental set-up (see 'Experiment 1', 1989 p. 66): no subjects are given both the simple and the additional/alternative conditions. Thus, the 'suppression' is posited to explain the differences not between two different answers from one subject as material is added, but the differences between subjects presented with either one or two conditionals in the premises of the given arguments (the second premise being either 'alternative' or 'additional' to the first as described above). This can justly be called "suppression" only if we assume that subjects in the two-conditional case somehow generate, say *modus ponens*, on the basis of the first conditional and are then led to suppress it upon reading the second conditional. This would entail *modus ponens* being somehow automatically generated upon reading the first conditional. This is a possibility. But it might also be the case that subjects read and assimilate all the presented material into a single semantic structure – which does not mimic the grammatical/textual structure – and then draw their conclusions. For instance, they might assimilate both antecedents into one complex conditional with either a conjunctive or a disjunctive antecedent, before making any inferences. Byrne's results are elicited from a set-up which, without further arguments about discourse processing, do not warrant her suggestively labelling the patterns she observes, as "suppression". Here the terminology is maintained for ease of comparison with other studies, but the above point should be kept in mind.

[14]Byrne's results are garnered per data item, not per subject, so individual differences are lost. This is not the case in interview situations, where individual subjects' responses are generally collated. Another difference is that the two conditions (i.e. first presentation of a single conditional premise, then an additional or alternative one) can be investigated in the

tended to contribute to the body of work identifying the range of interpretations given to conditional premises and the relations between them in the suppression effect task. To this end, we can make use of existing work which has been conducted with undergraduate populations, e.g. that by Dieussaert, Schaeken, Schroyen and d'Ydewalle (2000), Lechler (2004) and Stenning and van Lambalgen (2008). These studies used sequential presentation, and/or production or interview (elicitation) techniques and as such are suitable to compare with my interview data.

The suppression effect task is a reasoning task in a similar vein to the much-used syllogistic-type task, but with different premise sentences. The original motivation for studying material presented in this form is the so-called 'suppression effect' it elicits in subjects (Byrne 1989). The phenomenon is so named because with the addition of certain types of conditional premises subjects are judged to 'suppress' an inference they would have drawn had the extra conditional premise not been added. Byrne analysed this phenomenon as evidence that subjects do not use logical rules in drawing inferences. Regardless of whether or not we agree with Byrne's analysis, the task, and resulting response pattern she identified, are interesting because they give us insight into how subjects collate information in a intuitively fairly natural discourse and how their inferences adapt as they do so. As far as I am aware, the existing data have been collected only in schooled populations. This analysis is intended to contribute to bridging the gap between schooled subject data and unschooled subject data, by examining 1) how unschooled subjects deal with conditional premises, and 2) to what extent their response in the suppression effect task resembles that of schooled subjects.

But note that we are now in a very different comparative situation than we were in for the syllogistic-type task, because, as mentioned, the task elicits neither 'correct' nor uniform responses from schooled subjects. Byrne's original experiment already shows us that undergraduate subjects' conclusions depend on the perceived relation between two conditional premises[15], and further studies, by Dieussaert et al (2000), Lechler (2004) and Stenning and van Lambalgen (2008) have discovered a wide range of responses that subjects give to combined premises. For example, within just the 'additional' case for the MP inference, when given the premises 'if p then r; if q then r; p' Dieussaert et al's subjects came to the following conclusions:

- r
- r if (also) q
- not-r if not-q
- both r if q & not-r if not-q
- nothing follows
- r and q
- other

course of an interaction with a single subject.

[15]Other studies (Cummins 1995, Byrne 1999) have shown subjects' sensitivity to the availability to counterexamples

It is thus much more difficult to compare the conditional premise data across groups, because the 'norm' for the task[16] is now split from the schooled response, and this itself has splintered into a much broader range of responses. The implication of this for the upcoming analysis is that we cannot do what was done in the syllogistic case, namely take the posited norm as the yardstick for the schooled response, and thereby derive a comparison of schooled and unschooled on the basis of it. As such, this study is exploratory, rather than comparative:

We first look at the 'simple' condition, to discern whether the (unitary) logical form proposed in studies with undergraduates (specifically Stenning & van Lambalgen) can explain the responses garnered here, and whether there are differences between groups in this condition. We then proceed to briefly examine combinations of conditional premises, here primarily with aim of ascertaining what range of responses is apparent, and if so, to what extent it varies by group. Sequential presentation of premises to ensure subjects did both the simple and additional/alternative conditions generated much data on single conditional premises, another reason to pursue this split analysis.

Data from the conditional premise set is analysed in the following, with the following questions in mind:

(8) to what extent are the categories of response identified by Luria and Scribner for the syllogistic-type task evident in conditional reasoning tasks?

(9) do unschooled subjects ratify the same inference patterns as those ratified by schooled subjects in this and other studies?

(10) is the interpretation of the conditional statement assumed by the subject recognisably the same as that assumed by subjects in other studies (e.g. Stenning and van Lambalgen)?

(11) does the interpretation of the conditional vary across materials – e.g.

[16] I don't mean to suggest the 'norm', as Byrne intended it i.e. the response suggested by the rather contrived classical logical reading of the conditional premises, is an appropriate normative response to the task. In the syllogistic task, the statistical norm in the schooled group coincided with the posited normative answer. This lends plausibility to the posited normative answer, i.e. that generated with a 'classical' reading of the quantifier and the relation between the premises entailed by it, under the assumption that people are in general logical. This is not so in the suppression effect task, where the range of responses given diverges widely from the intended normative answer. More than this, the normative answer only makes sense under implausible assumptions about the structure of the task – such as how the two conditionals combine. In the case of the Wason selection task, the situation is again slightly different, since subjects often seem to gain *insight* into the task when moving from the statistically common response to the posited normative answer.

This is not to say that the posited normative model is the *only* norm available for this task; or the most appropriate one. In the final chapter justification for various readings of the conditional in the context of that task are presented; in Stenning and van Lambalgen (2004) there is extended discussion of the alternative non-monotonic notion of validity represented by closed-world reasoning, and how this relates to choices in the selection task.

when interpreted as a habitual vs a singly-occuring relation – and if so, does it do so consistently across subjects and/or across groups?

(12) to what kind of discourse do subjects assimilate this kind of material? is there evidence that they reflect on the purpose of the discourse and the 'naturalness' of the premises?

A notable preliminary finding in the conditional premise data was the relatively smaller number of 'Luria-type' responses (that is, those that have earlier been labelled 'refusal' or 'personal' interaction with premises) within the unschooled subject group *and* the relative increase in such responses in the more schooled subjects; this finding immediately adds credence to the earlier suggestion that differences in reasoning behaviour across literacy levels might have been overestimated in earlier work because of the focus on syllogistic-type arguments. What we see here is that the scale of difference between subject groups varies according to task material, and that conditional premises yield less contrasting responses between groups. A possible explanation for this finding is outlined in the next chapter.

The rest of this section is devoted to examining in more detail the findings on the conditional premise set. There were several recurrent phenomena across all subject groups with this material. Many of these have been identified in other studies of interpretations of the suppression effect material with undergraduate subject populations (Stenning and van Lambalgen, Lechler). The extent to which the categories presented here are continuous with those identified in other studies will be discussed as we progress through them. We start with phenomena which at first sight are continuous with those identified by Luria as characteristically illiterate responses, and move onto the more general 'suppression' phenomena.

The applicability of Luria's/Scribner's classification

As we have seen in the previous section, the characteristics of unschooled reasoning identified by Luria, such as elaborating on the given premises or rejecting them as a basis for the conclusion, can be understood as adequate responses, involving extensive reasoning, once allowance is made for subjects' interpretative engagement with the premises (again, to anticipate, we can make use here of Stenning and van Lambalgen's terminology 'reasoning *to* an interpretation'). Here this understanding of subject behaviour can be extended in analysing reasoning with conditional premises. So although a start is made with the terminology used by Luria, this is only to facilitate a suitable replacement with semantically-derived descriptions. These are not only more accurate but also show how the findings here connect up with those from studies with undergraduate populations.

"Refusal to answer"

There were only two occasions in the conditional premise interviews where subjects seem to interpret the premises only relative to their own knowledge. In both cases it was the first item in the interview. On being asked whether Ntombi would go to East London, Susan replied:

Susan (group 1):

S: How will I know? I don't know.

This first assertion is quickly overridden and the subject gives an answer (albeit not the expected one – but this is something we'll discuss elsewhere). Another subject gives a similar initial response:

Vulelwa, group 1:

S: I don't even know this Ayanda, and I don't know if she will go to East London. Does she live in Hamburg, this Ayanda?

I avoid the question by suggesting Ayanda lives in Bodium and the subject proceeds to give positive answers. The refusal to answer for lack of knowledge of the characters/situation described, is thus, as it appears in these cases, a relatively easily discarded interpretative set, and the basis for it might be construal of the task as a genuine query for information, something which appears to be more generally the case. When this is ruled out ('we are not referring to a specific person here') the subject proceeds to answer. In the first example, it looks as if Susan takes the question to be a more general query about sensible behaviour.

Vulelwa elsewhere:

E: Now suppose there are no cows in England. And there is a place in England called Fawley. Will there be cows there?
S: I don't know.
E: But what if I tell you there are no cows in the whole of England, and Fawley is in England. Will there be cows there?
S: This question is so difficult.
E: OK. I'll repeat it. In England, and the whole of England is an island, there's no cows. Now there's a small town on the island, called Fawley. Will there be cows there?
S: No, there's no cows there.
E: Why?
S: Because you say it's a small town and there's no grazing fields there.

'Personal' interpretation or natural interpretation?

In Luria's study, unschooled subjects were judged to reason badly because they interpreted the premises 'personally', expanding or adjusting them to fit their own knowledge, or more general knowledge ("If the land is good, cotton will grow there."). Is this really an idiosyncratically unschooled response, and, more importantly, is it a sign of inability to reason with the premises? It seems more likely, from what we've seen

in the syllogistic task data, that subjects, especially unschooled ones, are inclined to search not so much for a 'personal' but for a 'natural' or common-sense interpretation of the premises, assimilating them to everyday discourse form, while schooled subjects access the intended interpretation by suspending the natural one. (At this point, I used 'natural' in an intuitive way, but the next chapter is exactly aimed at pinpointing a precise sense in which the discourses are more or less natural.) If this is indeed the case, then one might expect that in premise sets which resemble naturally occurring discourses this effect is reduced, because both schooled and unschooled subjects would use the readily available naturalistic interpretation.

How can this idea be made more exact in the context of the suppression effect data? Use can be made here of the research done by Stenning and van Lambalgen (2008) who have proposed that natural language conditionals hide an 'abnormality' clause, that is, a sentence of the form 'if A then B' is of the logical form

(13) If A, *and nothing abnormal is the case*, then B,

"where what is abnormal is provided by the context"(p. 163). Stenning & van Lambalgen (2008) demonstrates how attributing this form can explain many of the suppression effect phenomena. The conditional premises used in the current study lend themselves to such an interpretation, especially if the 'abnormality' clause is understood contrariwise as a marker of 'normality'. So, for example, "If Ntombi wants to see her boyfriend then she goes to East London" is certainly best understood in everyday conversation as expressing a more generic habitual relationship, which tolerates exceptions, and thus is adequately expanded by "If Ntombi wants to see her boyfriend *and nothing else is going on*, then she goes to East London."[17]

There is much evidence that subjects employed this abnormality-sensitive reading of the conditional. Firstly, in several cases the subject gave back a modified form of the conditional, in which it was explicitly marked by the introduction of a marker of (weak) habituality ('sometimes' – Thaboliwo and Florence, group 2, Zoleka, group 3);

Thaboliwo, group 2:

E: If she has to fetch water she goes down to the river. and she has to fetch water. Where will you look for her? Where do you think she is?

S: *Sometimes*, she has to go to the river to fetch water. ...

Florence, group 2:

E: One more story, about my friend Simon. Imagine I'm looking for Simon. I know that if Simon has homework to do then he will study late in the library. And I know that he has homework to do. Do you think he will be at the library?

S: I don't know, because you don't know for sure he's got homework, you haven't seen his homework.

E: Suppose I tell you that I know he's got homework.

[17]Note that a 'generic' reading of many of the syllogistic task premises is possible too: "All bears in the north are white" read as "In general, bears in the north are white".

> S: He might be at the library because *that's what he usually does, but we can't be sure.*

But secondly, and more vividly, direct evidence is provided by subjects' enquiries about relevant factors which would prevent the deployment of the conditional relationship, as well suggestions to how these may be overcome. This can be understood in terms of the logical form given above as concern for relevant abnormalities. Often the subject spontaneously mentioned factors which could serve as abnormalities, and which would thus prevent the consequent being fulfilled.

Sensitivity to abnormalities is not always the most accurate way to describe other ways in which subjects qualify the conditional relationship. Many of my subjects were concerned with qualifications to the conditional which were more like (necessary) preconditions than abnormalities. In fact, there is an alternative form which does treat qualifications as preconditions in Politzer (2004). Politzer argues that a range of results from conditional reasoning research can be explained with a single form for the conditional, as proposed by Mackie (1974) for causal conditionals:

$$[(\mathrm{A} \,\&\, \mathrm{A}_1 \,\&\, \mathrm{A}_2 \ldots) \text{ v } (\mathrm{B} \,\&\, \mathrm{B}_1 \,\&\, \mathrm{B}_2 \ldots) \text{ v } \ldots] \rightarrow \mathrm{C}$$

where A is the antecedent currently under consideration; B would be an alternative antecedent which in appropriate contexts justifies the assertion of *if B then C* In some cases relevant preconditions were spontaneously offered, sometimes before the *modus ponens* inference was granted. (This will come into play in explaining the suppression effect data.) Since in many cases the B's would be null, Politzer concentrates on the abridged form:

$$(\mathrm{A} \,\&\, \mathrm{A}_1 \,\&\, \mathrm{A}_2 \ldots) \rightarrow \mathrm{C}.$$

The key aspect of this form is the role of the A_ns, which are "separately necessary with respect to C" and combine to form the sufficient condition (A & A_1 & A_2 ...). The A_ns are what Politzer calls *complementary necessary conditions* or CNCs. The CNCs enter into a reasoning process because

> in asserting the conditional *if A then C*, the speaker assumes that the necessity status of the conditions A_1, A_2 ... is part of the cognitive environment, and most importantly that the speaker has no reason to believe that these conditions are not satisfied.

But crucially, in some cases the satisfaction of the CNC-clause is brought into doubt – typically when one has "high availability", presumably in terms of salience or, as shown in work by Cummins (Cummins et al 1991, Cummins, 1995), when there are many of these CNCs available ('disabling conditions' in Cummins' terms). This form of the conditional can also certainly be fruitfully applied to

suppression effect data, as Stenning and van Lambalgen did with their version. Here the two versions are treated as of a piece, since a positive precondition A can be captured by the implication 'not-A $\rightarrow$ ab', where ab formalises "something abnormal is the case".

Sometimes the subject first denies the conclusion and gives an additional requirement (or CNC if you like) as justification. Overwhelmingly, when this requirement is met the subject is happy to draw the inference. In the following excerpt, the subject infers the consequent but when asked for justification, she is more hesitant, asking first if the antecedent condition is the reason, and then anticipating obstacles to the conditional relationship obtaining, such as strict parents or not enough money, and suggesting ways in which these can be overcome so that the conditionality of the situation is adequately captured by the premise as stated. She is not so much concerned with abnormalities as preconditions.

Nomhle, group 2:

E: If Ntombi wants to see her friend then she goes to East London. And she does want to see her friend. Will she go to East London?
S: Yes.
E: Why will she go to East London?
S: Is it because she wants to see her friend?
E: Yes, she wants to go.
S: Maybe she can talk to her parent and tell her that she wants to go and see her friend. Maybe Ntombi is interested to go and to see her friend but now she doesn't have enough money. So maybe her parent will give her money to go to East London to see her friend.

This subject's responses were consistently of this form. Other examples, from all groups:

Nothabile, group 1:

E: OK next one. ... it's about a young man called Simon. He lives in a town near East London. Now suppose you want to know what Simon is doing today. You know sometimes he goes to East London.
S: He goes to East London if he wants something.
E: Yes, like if he wants to visit his sister then he goes to East London. So today he wants to visit his sister, do you think he will go to East London?
S: *He will make a phone call first to make sure if she's available.*
E: And if he calls and she's available?
S: And *if she's available, and she's his sister, and he's missing her*, then he'll go to East London.

Vulelwa, group 1:

(preamble) If Maria finds a job, then she will hire a maid. And she does find a job.
E: Will she hire a maid?

S: She may hire someone or she may not.
E: Why? Why would she hire one or why won't she hire one?
S: Sometimes it will be difficult for her to hire someone after she has just found a job, immediately hiring someone, because she won't have enough money, because she's just got a job.

Another example, in which the subject suggests the precondition is the friend's desire to see Ntombi:

Sebenza, group 3:

(preamble) If Ntombi wants to see her friend then she goes to East London. And she does want to see her friend.
E: Do you think she'll go to East London?
S: *I think she can phone the friend if she wants to see her.*
E: So she won't go to East London?
S: She'll phone, and hear from her, the friend, if she can go.
E: And if she's arranged it with the friend and it's ok?
S: If they arrange all this then she can go.

Other subjects – group 1 mostly – first gave an answer which belied extra conditions and afterwards checked whether these were met.

Nomvumisa, group 1:

(preamble) Ntombi wants to see her friend. If she wants to see her friend, then she goes to East London.
E: Will she go to East London, do you think?
S: When she wants to see her friend?
T: Yes.
S: No, she can't go to East London.
E: Why?
S: If she's here and the other person is in East London, does she have the right to go to East London?
T: Yes, that's no problem.
S: Well, if there's no problem then if she really wants to see her friend then she will go to East London.

This kind of response can be seen as fitting Luria and Scribner's diagnoses that the subject 'goes beyond the information given' in the premises, but again, it is not a peculiarly illiterate phenomenon, as my data illustrate, and more generally it is absolutely in line with the findings of for instance, Byrne 1999 and Cummins et al 1991, where schooled subjects' willingness to draw a conclusion has been shown to depend on the availability of disabling conditions, or in Stenning and van Lambalgen's terms, relevant abnormalities/preconditions. Concern for both abnormalities and preconditions is certainly influenced by task construal: offering possible counterexamples or necessary preconditions is very informative, under the assumption that the task is about establishing conditions for drawing the inference. See also the section on informativity for more discussion on this issue.

Ignoring premises or taking time into account?

Nonkululeko, group 2:

(preamble) Patricia is looking for her friend Susan. She knows that if Susan has an essay to write she works at school.
E: Where do you think Patricia will look for Susan?
S: She will look for her at home.
E: Why?
S: She will look for her at home because that's where she stays.
E: But suppose we know that she has homework to do, and if she has homework then she goes to school.
S: The first place she will go to is her home, and then afterwards she can go to the school if she's not at home.

This is a typical example of what Luria would have called rejecting or ignoring (before elaborating on) the given premises. But notice that the conditional relationship described in the premise "If Susan has an essay to write, then she goes to school" has a temporally bound character: it might be more fully expressed as "If Susan has an essay to write, then she goes to school at some point before the deadline to work on it". We don't expect her to be at school *continuously* until the essay is finished; in fact, she is probably at school for a relatively short amount of time during this essay-writing period. She could still be expected to spend the majority of her time at home, for instance. With this background, i.e. taking the 'base-rate' of time spent at home into account, looking first for Susan at home first is a better strategy than immediately going to the school. Another subject goes further in elucidating this:

Nomhle, group 2:

E: Where do you think Patricia will look for Susan?
S: She will look for her at her home. If she's not there she will look in the library.
E: Why won't she look in the library straight away?
S: To make sure she's already gone, to the library. That's why she'll go first to the home.
... The reason why she has to go first to the home, is that maybe Susan, this girl, maybe her parent, her mother, has asked her to do something first, before she goes to the library.

Why would a subject take this 'base-rate' into account? One reason would be if you aim is to give an optimal strategy for finding Susan. If you understood the purpose of the task to be to describe the best way to find Susan, and *not* to demonstrate your grasp of the logical structure of the premises, then a good answer *should* take the base-rate into account. This 'search strategy' reading of the question is also apparent when subjects were faced with an additional condition which might not be fulfilled, leading to a dead-end for the search viz:

Vulelwa, group 1:

E: And now suppose that if the library is open then Susan is working in the library. What do you think now? Do you still think she will be there?
S: She will go there.
E: And what if the library isn't open?
S: I got a problem now, if the library's closed.

The temporally-bound nature of conditional premises played a role in the responses of many subjects, especially when the question was phrased in this 'looking for x' fashion. For instance, with the premises "If Thembi has to fetch water then she goes down to the river. She has to fetch water. Where will you look for her?":

Sweetness, group 2:

S: If at home they said she's not there, I'll go to the river.

In the next excerpt the intermittent fulfilment of the conditional relationship is mentioned explicitly:

Thaboliwo, group 2:

E: Where will you look for her? Where do you think she is?
S: Sometimes, she has to go to the river to fetch water. Thembi sometimes goes to the river, maybe in the afternoon or the morning. When I see her going to the river, maybe in the morning, I'll go to her then and see her.

Is this typically unschooled behaviour? Absolutely not. It depends on task construal, and this varies also within schooled subjects. Lechler (p. 60) gives excerpts from undergraduate subjects' responses which show the same considerations about temporal ranges – and also note in the second to last turn the mentioning of 'an infinite number of possibilities' which could prevent the consequent from holding – again evidence that the subject is concerned with abnormalities/preconditions not mentioned in the premises. In this case note the two conditional premises were attributed to different sources in this condition, and the subject is also told the protagonist 'was quite often in the gym'):

(preamble) If she has an essay to write, she will be in the library. If the library stays open, she will be in the library. She has to hand in an essay next week.
Subject 6:

E: So, where would you look for her?
S: Um it is a good chance of finding her in the library. Maybe she could be by the gym as well.
E: Okay, so *what could prevent her from being in the library?*
S: *Well, she could be in the pub, you know.* (laughing). *There's a whole, an infinite number of possibilities.*
E: So is there any information you would need to decide where she is?

S: Well, it'd be interesting to know how conscientious a student she was. You know, if she doesn't give her ... *Say she is a first year student, and she's got a week for an essay deadline. She is probably not gonna be in the library. But if she is a final year student and she wants to stay on for a PhD, she will probably stay in the library, working really hard.*

Lechler (p. 102) sums up her finding as "Some subjects treat the presented statements in a similar way to everyday discourse, others regard them as some kind of logical task." In my data, Abongile, a high-school student, illustrates the former reading, where the conditional is understood to be temporally limited:

Abongile, group 3:

(preamble) If Thembi has to fetch water, then she goes down to the river. And you know she has to fetch water. So this is what we know: [Repeat premises.]
E: Do you think she's at the river?
S: (quiet)
E: repeats premises.
S: No.
E: She's not at the river?
S: No.
E: Why not?
S: *She'll fetch the water and go home.*
E: So she will go to the river and then come back?
S: Ja for sure she'll go to the river and then go home.

This contrasts with Mzikazi, also a schooled subject, who *does* treat the conditional atemporally, although note that 'being at home' is still apparently the default location, and can be understood as a switch to a temporal reading:

Mzikazi, group 3:

E: Where do you think Patricia will look for Susan?
S: In the library.
E: And what if I also tell you that if the library is open, then Susan is in the library, do you still think that Susan is in the library?
S: Yes.
E: *And what if the library is not open?*
S: *She's at home.*

The subject's assumption that the additional condition, such as 'library open', holds, was a recurring feature of her treatment of additional premises. It might be the case that an 'atemporal' reading of this kind of conditional *is* a literate default in this context – not necessarily a general tendency, associated perhaps with treatment of the task as 'some kind of logical task', but clearly it is not always triggered by material in which a temporally-bound reading is more natural.

One last excerpt provides a nice illustration of how a subject juggles with these possibilities – and note that the subject has never been to school. This excerpt is from Susan, whose responses have cropped up elsewhere, because of their close resemblance to those of Luria's subjects. When the translator intervenes, to suggest Vuyo is really a hypothetical character, it seems that Susan switches from giving advice to the 'logical task'.

Susan, group 1:

(preamble) If Vuyo has to look after the baby, then he stays at home.

E: And what if we know that Vuyo has to look after the baby today?
S: You're asking me what I would say? I would search for him, tell him to look after the baby.
T: You haven't met the person, you just have him in your mind. You haven't seen him. So would you know he's at home or not?
S: If I knew he had to look after the baby then I would know that he's at home.

Justification of the conditional itself

As a final example of interpretational mismatch between experimenter and subject, we look at the phenomenon whereby, on being asked to draw the *modus ponens* inference, several subjects gave an answer which included a justification of the antecedent of the conditional on which the inference is based. This is a very interesting response. As far as I know, it has not been reported in the undergraduate subject groups. In my data I had one incidence of such a response among the most schooled group:

Thembakazi, group 3:

(preamble) Thembi's mother is concerned about her. But she believes that: If a student works hard then they will pass. And if a student is clever then they will pass. And the teacher says that Thembi is clever.
E: Do you think the mother will think Thembi will pass?
S: Yes, the mother will think that Thembi will pass, because the mother has gone to the teacher and asked the teacher, and the teacher has told her she's smart, *maybe she has seen her books.*

A similar response comes from an unschooled subject, when presented with the same premises, who defends the (amended) conditional 'If the teacher says a student is clever, then the student is clever, (and thus they will pass)', but does not phrase it in terms of the mother's beliefs:

Maggie, group 1:

E: Then what do you think the mother will think – will Thembi pass her exams or not?

S: If the teacher says that Thembi is clever, then she will pass, *because the teacher has seen her performance at school.*

In some instances the subject's response seems to indicate a justification of why they themselves are inclined to accept the conditional:

Rosie, group 2:

E: If she finds a job, then she will hire a maid. And she does find a job. Do you think she hires a maid?
S: Yes, she's supposed to find a maid because she won't have someone to look after the kids.

An example of how this has been treated in earlier work is to be found in Cole et al (1971): the response from a group of village elders to the problem "Everybody who has a house must pay house tax. I have a house. Therefore, I must pay house tax" was unanimous agreement that the last statement was true "because *it had been decreed by the government that we have house tax.*" Cole et al call this "extraneous information"; I would rather call it "justification for accepting the major premise". If the purpose of the task is not clear to the participants they might well see this as a sensible response.

Offering justification for the premises are a counterpart to offering additional necessary conditions for its fulfilment: both responses answer the question "What further information would make this a reasonable inference to draw?" Luria and Scribner might have called these responses 'empirical bias'; but they are as above more accurately described as specifying grounds for the premise itself, a strategy which could be understood as resulting from concerns to be informative beyond demonstrating one's own cognition and therefore have to do with general task construal.

1.4 Summary, conclusions and outlook

In the current chapter, earlier experimental work with illiterate subjects employing syllogistic-type materials (Luria, 1976, Scribner, 1997) was replicated. The inclusion of suppression-effect task materials provided an extension of this earlier work. The results from the experiment showed remarkably similar responses to those identified by Luria as typically illiterate, albeit on a lesser scale. The conditional premises derived from suppression-effect materials provoked a more similar response between the schooled and unschooled subjects.

At this point it is worthwhile to reflect on the value of the foregoing analysis. The results of Luria's experiments often evoke one of two reactions. In Tulviste's phrasing (1991, p. 118):

> Some see in them evidence for the underdevelopment of thinking in people from traditional cultures, their low mental abilities. Others

> reject these results, maintaining that the methodology of the experiment is evidently not suitable for use in traditional cultures, that the subjects do not understand what is expected of them, etc.

One can certainly maintain that syllogistic problems are not suitable for use in the study of reasoning in illiterate subjects. The problems *are* strange to the subjects. Yet, as Tulviste rightly points out (p. 118),

> All investigators evidently sensed this inadequacy even before doing the experiments, for they used the most simple syllogistic problems, and not problems of the type "Some academics are parents. All parents are drivers. What conclusion can you reach?" (used by Johnson-Laird, 1983). But as we have seen, even these simple problems are not "simple" for the unschooled subject. It seems to us that it is exactly this inadequacy that is of primary interest.

In this chapter our primary interest has indeed been in uncovering some of the interpretational parameters which cause these seemingly simple problems to generate the range of 'wrong' responses Luria and Scribner both observed in their data. I argued that once we have more insight into the range of task interpretations available to a subject we see that their responses are well-argued and not lacking logic as Luria would have had.

In particular, in analysing the syllogistic task data it became clear that a logical answer does not preclude reliance on previous experience, and vice versa; we also saw that the reasoning task has a pathological epistemic structure; and that generalised premises can be read as generic or law-making statements, thus changing their relation to a particular premise.

Moreover, when subjects interacted with conditional premises such as those used in the suppression effect task, their interpretational tendencies aligned with those found in studies with literate subjects by both Stenning and van Lambalgen (2008) and Politzer (2004). Subjects can be seen to be employing a common interpretation of the conditional, which has been identified in other studies as, as including an abnormality clause. Many of their 'personal' responses make sense as reference to relevant abnormalities, again, something which has been shown to be more general behaviour in studies in other subject groups (Cummins, 1995). Also the temporally-bound nature of natural language conditionals features in the responses of both schooled and unschooled subjects. These factors, along with some allowance for the occasional caprice of a mystified interviewee, strongly undermine Luria's conclusions that subjects rejected or ignored the given premises or their logical structure.

The analysis thus provides corroborating evidence to the central claim, namely that illiterate and literate reasoning performance, for all its differences, should be seen as stemming from a common base of semantic concern. That is to say, all subjects must first interpret the given premises before they reason with them.

Differences in logical reasoning ability can only be claimed once the differences arising at the interpretational stage have been accounted for. As we have seen, this is not yet the case, and, as such, we have no grounds to claim that literacy brings increased logicality. In the following chapters this claim will be both strengthened and qualified.

Seeing the commonalities in reasoning behaviour across the subject groups brings the illiterate reasoning data back into the arena of interest for reasoning researchers. At the very least it should inspire us to look beyond the university for subjects, to seek out the full range of interpretational parameters which inform reasoning behaviour.

Chapter 2

'*if p then q*' . . . and *all that*

2.1 Introduction

From any given proposition, a multitude of inferences can be drawn. Yet most of them are not. Take a simple proposition, such as "Today is Tuesday". Would you expect anyone to conclude, on hearing that, "Oh, so it's Tuesday or Saturday"? No. More generally, for any p we may infer the weaker $p \vee q$. But this is not something we can expect to see commonly occurring. Why not? One explanation is that $p \vee q$ is less informative than p, and this violates Grice's maxim of quantity: be as informative as possible. As such, we would expect the use of p to conclude $p \vee q$ to be infrequent in reasoning and communication contexts. Undoubtedly there are contexts to be thought up, in which it would be a natural conclusion to draw – for example, when the disjunction is needed as input for further reasoning – but these are circumscribed.[1]

Now observe that reasoning research doesn't focus on the percentages of people who generate $p \vee q$ on presentation of p, nor on developing theories of why people don't draw such inferences. Though there is some work on reasoning *from* disjunctive premises – see for instance Van der Henst, Yang and Johnson-Laird, 2002 – reasoning research overwhelmingly focuses on premises with conditional and quantified phrasing. Why should this be so? There are at least two places to look for an answer.

One is in the history of the field. There are historical reasons why some inferences are studied above others, and the syllogism is just such a case. Aristotle

[1] Imagine a tax form, with Category B defined as the group of people who either earn less than a certain amount or don't work at all. Then if you are working as a PhD student, say, you would conclude that you fit into Category B. Under some description you've concluded the disjunction holds on the basis of one of the disjuncts, and used this to conclude that you belong in Category B. Compare this also with normal use of disjunction, which carries the implication that the speaker does not know which of the disjuncts is true: "When is Peter leaving?" "Monday or Tuesday." Note also that inferring a disjunction is a different discourse 'move' from that when a disjunction is offered as a correction of a simple proposition.

thought that all valid inferences could be represented in the form of a syllogism. Although medieval logicians were well aware of the limitations of the syllogistic system, Aristotle's view was definitively overruled only in the nineteenth century, when Boole and Frege laid down new formalisms, for propositional and predicate logic respectively, and which formed the beginnings of modern symbolic logic. Early studies of reasoning (such as Wilkins, 1928, Woodworth & Sells, 1936, but also Luria, who conducted his research in the 1930's) focused exclusively on syllogistic arguments.

A second reason for the dominance of certain types of premises might be that inferences based on them are more natural, recognisable, common, than others; that they 'make sense' to experimenters and subjects alike. A conditional premise would seem to be just such a candidate. And, we might assume, so would the basic syllogistic form. Yet, when one looks at the transcripts from Luria's study, one gets the distinct impression that for unschooled subjects the intended inference from syllogistic premises doesn't 'make sense' at all, as evident in the following transcript:

E: In the Far North, where there is snow, all bears are white. Novaya Zemlya is in the Far North and there is always snow there. What color are the bears there?

S: I don't know what color the bears are there, I never saw them.

...

E: But what do you think?

S: Once I saw a bear in a museum, that's all.

E: But on the basis of what I said, what color do you think the bears are there?

S: Either one-colored or two-colored ... [ponders for a long time]. To judge from the place, they should be white. You say that there is a lot of snow there, but we have never been there!

Subject: Khamrak., age forty, miller from remote village, illiterate (Luria, 1976, p. 111.)

We have already seen this transcript in the previous chapter, as an example of the characteristic illiterate responses found by Luria in syllogistic reasoning tasks. Such transcripts crop up all over psychological literature, and are commonly used in psychological textbooks (such as Gray, 1991, p. 389) in the section on cross-cultural psychological differences. It seems clear that their illustrative appeal stems from obvious, even comical, misunderstanding – from our point of view – on the part of the subject, of the purpose of the exercise. What to us looks like a rather boring schoolish enquiry is responded to with any manner of off-chart replies. The exchange is in a sense a *failure* of exchange, from the experimenter's perspective, and shared by us, because the subject does not answer the question as it is put to him. "Refusal" was how Luria put it: "refusal to resort to logical inference from the given premises" (Luria, 1976, p. 108); "refusal to draw conclusion because of lack of personal experience" (*ibid*, p. 110).

But what if the subject isn't refusing to answer so much as trying to figure out what the question is? This would occur if the subject could not easily discern the purpose or structure of the exchange. That illiterate subjects often give a non-answer, rather than a wrong answer, gives this idea initial plausibility. At the level of the discourse, this mismatch would be driven by the relative availability of discourse 'templates' or genres, trickling down to the level of the sentence where a mismatch could be caused by atypical use of the linguistic forms found in the premises. The idea that there are more or less natural ways of describing situations is by no means new – already in, say, Donaldson (1978) we find discussion of why it is much more 'natural' to say "The flowers are on top of the television set" than "The television set is under the flowers"; additionally, plenty of recent research within psychology of reasoning has exactly the aim of relating reasoning task performance to everyday language use (for example, Stenning & Cox, in press, Politzer & Noveck, 1991). Understanding the reasoning task as a linguistic structure, a discourse, with more or less similarity to typical language use, opens up the possibility of understanding that the so-called failure to reason on behalf of the subject can also, possibly more justly, be characterised as a failure of communication between two interlocutors.

In this chapter I explore the hypothesis that Luria was led to an overly negative conclusion regarding his subjects' reasoning ability because of his focus on syllogistic premises. This is suggested by the results reported in the previous chapter, where we saw that subjects in all groups fared better with conditional premises, and group differences were more muted with such premises. Support for the hypothesis is found in the use of quantified constructions in spontaneous speech; and a comparison of this with the use of conditional sentences in spontaneous speech.

Yet, on the other hand, the simplest formal analysis of quantified statements gives them a implicational structure, viz: $\forall x(Px \rightarrow Qx)$. Hence I aim also to explain why the apparent similarity between quantified forms and conditionals does not result in similar inferential properties for the two types of sentences. As we will see, there are subtle differences in the semantics and use of such sentences, in both spontaneous language, and in the context of the task, which go some way to explaining this phenomenon. It should be stressed that the present proposal is surely not the only one which explains the data; nor does that matter here – our aim is to showcase what a semantically-grounded analysis[2] of reasoning data looks like, and to hopefully thereby illustrate its value.

If it can be shown that subjects are in fact exhibiting normal language use in the tasks as explained, the charge of illogicality loses force. The claim can then be made that they are the normal conversants, and that in the schooled case the

[2]This should be understood to include some traditionally pragmatic concerns. The intensive interplay between pragmatic and semantic factors in reaching an interpretation – and indeed the very sense of drawing a sharp boundary between pragmatics and semantics – are discussed extensively in Chapter 4.

subject 'colludes' with the experimenter on a special kind of artful discourse. The chapter should not be understood as a comparison of two experimental conditions; if this were the case the premise sets would differ in only one regard, namely the use of conditional or quantified phrasing. Rather this is a comparison between two experiments, which were primarily conducted with the aim of replicating previous paradigms (namely Luria, 1976 and Byrne, 1989). The two test paradigms differ not only in the phrasing used in the premises but also in the content of the premises. A number of suggestions are made as to further empirical means to further assess the impact these factors have on reasoning performance.

The plan of the chapter is as follows. First, a pilot corpus study, which distinguishes categories of usage of *all*, is reported. The significance of the categories, especially with respect to the relation between the quantifier and its domain, is discussed in the light of formal work on quantifiers and domains. It is argued that different types of generalisation are associated with differential relations of *all* to its domain, but that this association is distorted in syllogistic reasoning materials.

Next, I aim to explain the better response to conditional premises despite the aforementioned apparent similarity with universally quantified forms. Analysis of the semantics of the (types of) conditionals used in reasoning tasks indicates that the way they are used in reasoning tasks is very similar to how they are typically used in spontaneous spoken language. On the basis of this their inferential properties can be contrasted with the more problematic inferences from universal generalisation.

2.2 Are *all* premises equally difficult?

As we have seen in the previous chapter, Luria tested his subjects' reasoning by means of what he calls 'syllogistic' problems. This means that the major premise takes the form of a universal statement, either expressed with the use of a universal quantifier, translated as the English *all*, as in "In the far north, all bears are white", or as a simple generalised statement such as "Precious metals do not rust". Moreover, we saw that Luria's subjects had great difficulty with reasoning from such premises, seemingly either reluctant to draw a conclusion on the basis of them or adapting them to their own version of the premise. This difficulty was to be observed also in the current study, especially as reflected in the differential rates of immediate assent to quantified premises, as compared to conditional premises, given in table 2.1. In other words, subjects in the current study had greater difficulty in reasoning with quantified premises such as "All birds in Cape Town are penguins", than with conditional premises such as "If Thembi wants to visit her friend she goes to East London". The question is what causes this discrepancy, and the aim of this chapter is to seek an answer in possible discrepancies between everyday language use and the use of premises

	% correct with ...			
	quantified premises		conditional premises	
Group*	ratio	percentage	ratio	percentage
1	7/23	**30%**	9/18	**50%**
2	25/38	**66%**	28/32	**81%**
3	24/33	**73%**	28/33	**85%**

* Recall the group division from the previous chapter is as follows:
Group 1: No education (6 subjects)
Group 2: Between four and ten years of education,
left the educational system more than ten years ago (13 subjects)
Group 3: Graduated from high school within the last twenty years
(10 subjects)

Table 2.1: Comparing initial responses across premise forms

in reasoning tasks. We start by comparing everyday use of universal quantifiers with their use in syllogistic tasks. Since only explicitly quantified generalised statements were used in the current study, I concentrate on them. Comments regarding difficulties with generalised statements are to be found in the previous chapter and also in the section on 'recall data' in the next chapter.

2.2.1 *all* in spoken discourse

To be able to judge whether the quantifier is being used in a natural or recognisable way, we need to know how it is used in contexts of spontaneous speech. Unfortunately, I could find no previous corpus-based research on this topic, so a small study was made for the purposes of this chapter. The reported results are preliminary, and were garnered with the express intention of exploring their power to explain the data reported in the previous chapter.

The use of English *all* in discourse

A small sample of *all* in the spoken British National Corpus formed the basis for the study.[3] First, a random sample of 50 occurrences was analysed (where *all* is being used as a determiner according to the corpus coding, excluding, for example, adjectival modifier use, as in 'all grumpy', 'all fired up'[4]) and at least four different usages were discerned. After the first fifty the classification was applied to a second fifty occurrences as a check on validity. The only change made after analysis of the second fifty was inclusion within the category 'forward quantification' of sub-categories for science, law and religion. Apart from this, all uses could be subsumed under the first categorization. The following numbers are based on a second pass through the samples with the amended classification,

[3]Located at `http://www.natcorp.ox.ac.uk/`.

[4]The vast majority are deemed determiners: 236 518 of the 277 147 in the corpus.

	Category	**%**
1.	stock phrases	24
2.	emphatic usage	22
3.	anaphoric or deictic use	23
4.	forward quantification	27
5.	miscellaneous	4
	total	100

Table 2.2: Categories of '*all*' usage

which is given in table 2.2. As will become evident, the categories often overlap: for instance, emphatic contexts are also often formulaic and could be grouped under stock formulations. It is thus worth bearing in mind that the treatment of the categories as disjunct is somewhat artificial, and the categories should be seen as points along a continuum rather than partitions of a space, for reasons which will be highlighted later. There were only a few overlaps, and these do not affect the contour of the findings, as will become clear.

A notable initial finding is that there were but a few occurrences of sentence initial *all*: three in the sample I looked at, all in subcategory of law-giving contexts. Clause-initial *all* was however much more common, especially in the anaphoric and forward categories. I first report results for the English term *all*. A similar study was made of the Xhosa equivalent, the suffix *-onke*, and the categorization applied equally well there, as we will see. In fact, there was an even stronger tendency towards anaphoric/deictic use, this being the biggest category at 30% of usage, with forward quantification accounting for only 11% of occurrences.

These categories are now discussed in more detail:

1. **Stock phrases** included conventionalised constructions such as 'all night', 'all the time', 'all around him', 'all his heart', 'that's all', 'all but impossible', 'all the same', 'and all', 'all in all', 'first of all', 'after all'. These are phrases which are to be interpreted figuratively – in many cases a 'literal' reading is not even apparent, as with 'all the same', or 'all in all'. By literal meaning I mean that the domain, over which the *all* quantifies, can be properly specified. In English this appears to be a very common usage, accounting for 14 of the 50 occurrences analysed in the first sample and 10 in the second one – averaging 24% overall. It remains to be seen whether this category is significant cross-linguistically. A rule of thumb to judge this category is to try replacing *all* with other quantifiers – try 'most' or 'some' – and see if the new phrase is useable. If it is, then it doesn't fit into this category. This heuristic also suggests that stock phrases are a kind of fossilised universal quantification. For this reason we'll exclude them from further analysis. Examples of stock phrases in context:

- 'Maybe I'm not too late **after all**' (FS1 1299)[5]
- 'Mr(sic) sent for an ambulance **and all**' (KDU 570)

2. **Emphatic** use. This is related to the above category in that there doesn't seem to be true quantification going on. Rather, *all* seems to serve to emphasise what's being said, often also in conventionalised formulations. The heuristic I used to assign this category was the following: if *all* is omitted, or replaced with a determiner (as in 'all three countries' becoming 'the three countries'), would the sentence get a different reading? If not, then it belongs here. Also fairly large, this category accounted for 22% of occurrences. Examples are:

   - '**Let's all** get the hell out of here' (B1N 459)
   - 'Joanne and her parents agree that having the baby has **brought them all** closer together as a family (FU1 1037)
   - 'The kouroi and early female figures **all** carry the aura of the block's four faces' (FPW 343)
   - '**Tell me all** about this woman' (GV8 2583)
   - '**First of all**' (F9D 729)
   - '20% of **all dialogue**' (GOW 2835)
   - 'in **all three countries**' (AP7 474)

   In the example about 'Joanne and her parents . . . ', the pronominal 'them' apparently refers only to Joanne and her parents according to the rest of the excerpt. If there were siblings or other family members mentioned in the previous discourse then the *all* would function as a means to let 'them' refer to the whole family and not just the closest 'Joanne and her parents' candidate for reference. In this case the *all* could be seen to be functioning anaphorically. Another example where the same situation might hold is in (A6N 1988): '**They all** went on to be priests'. This case also is counted as emphatic in the current tally, and addressed in more detail in the next category.

   Observe also that in the majority of cases of emphatic usage given above, the quantifier is 'floating' . It occupies the position normally occupied by an adverb. Adapting one of the corpus examples, compare

   (1) *All* the kouroi carry the aura of the block's four faces.

   (2) The kouroi *all* carry the aura of the block's four faces.

[5]This number gives the tag used in the British National Corpus to identify this excerpt.

A discussion of the relations between such sentences is beyond the scope of the current study; for our purposes it is enough to know that they are considered to be logically equivalent to each other in the linguistic literature. As quoted in Bobaljik (2002), their "quantificational properties" are considered "identical" (Sportiche, 1988, p. 426).

3. **Anaphoric** or **deictic** use. This is what I have termed the coupling of the quantifier with an anaphoric or deictic terms such as 'which', 'this', 'those', 'the others', so that the kind of entities to be quantified over are given either in the previous discourse, or by the context of utterance. Bare *all* counts here too – see examples below. Sometimes the phrasing is quite conventionalised, but replacement by for instance 'most (of)' is generally unproblematic – suggesting there's some live quantification going on. Anaphoric and emphatic usage serve a similar function: in many cases the emphatic could be seen as a reinforcing or contrasting with the default reference to a group; the anaphoric serves to do this where the anaphoric or deictic reference might not clearly distinguish between a group and subgroups of its members. Anaphoric use was also a big category – accounting for 23% of the sample (12 occurrences in the first sample and 10 in the second one). Examples from the corpus are:

   - 'Of course I know **all that's** no reason I can't have a bike' (G3P 1937),
   - '**all this** will take time to negotiate' (ABE 2666),
   - '... some of which have a less visible protestant ethos, but **all of which** have a loyalist ethos as well.' (A07 1375)
   - '**All** can be reached by public transport and offer quiet and relaxing woodland walks.' (FTU 808)

   In this category I have also included anaphoric/deictic reference which also have some descriptive content, such as the following:

   - 'In spite of acknowledging **all these factors**' (BLW 480),
   - 'where **all these views** would concur' (CD9 371),
   - 'But after the vote on Mr Craxi, parliamentary leaders agreed to take **all such decisions** by a show of hands (CR9 1860),

   There was again here a case on the border between emphatic and anaphoric/deictic, but which have been tallied as anaphoric, such as: '**All this industry** must be sign of things looking good' (HA6 1031).

4. **'Forward' quantification**. This is the category which most closely resembles how *all* is used in syllogistic arguments. The entities to be quantified over are explicitly mentioned after the *all*, as in 'all ethnic minorities' or 'all or part of the primary school cycle', 'all denominations' etc. I have splintered off a number of subcategories here, based on the clear context

for use of such sentences. These are: statements of policy or law; as part of religious doctrine; in writing about scientific research. I think that these subcategories deserve special attention because they are probably the only place where quantification truly 'globally' quantifies, and, perhaps related to this, they are more or less directly derived from textual discourses. We examine first some examples which fall outside these subcategories:

- 'Thanks were extended to **all the Kent teachers who had hosted the event**' (KAE 91)
- 'The bank gave consideration to **all matters relating to the company in question's affairs** (AHB 588)
- 'It has rendered untenable the simplistic belief that members of **all ethnic minorities** are part of one undifferentiated black mass . . . ' (A1T 40)
- 'And **all the water courses** been blocked up and then it was swampy as well.' (HER 499)

Then we come to the subcategories:

(a) Statement of **laws/rules/policies**:
   - '**All penalties** are cumulative, but penalties for disobedience depend on ...' (BPB 301),
   - 'A complementary excursion to the Dolomites for **all guests staying 14 nights**' (ECF 3790),
   - 'Issued to **all Gulf warmen** . . . ' (K1M 4027).

(b) **Ideological (religious, political) doctrine**:
   - 'It is the mystery of the Creation, the God of **all Jews** . . . who transcends **all beings**' (A3F 55)
   - 'Christ had died for **all**, **all men and women**' (CLM 268)

(c) **Scientific research**:
   - 'Patterning in **all systems** occurs in small groups of cells, . . . ' (ASL 992)

What can be said about this category? Although there is explicit description of entities to be quantified over in this category, in almost all instances there must be further domain restriction in order to pick out the appropriate group quantified over – the latter excerpt is a good example, where the 'systems' quantified over are clearly of a pre-specified sort, or range, given elsewhere, but which are currently under discussion. In fact the only cases for which quantification can truly be said to operate over a universal domain are the religious categories – what that means I'm not sure!

5. Miscellaneous:
   - **Proper names**: 'All Saints', 'All Angels'

These proper names could been seen as part of the 'stock phrase' usage since they represent fossilized quantification and not live quantification, so to speak.

The use of Xhosa *-onke* in discourse

The above classification is based on an English corpus study; ideally, the classification would be cross-linguistic. Obviously any argument to explain syllogistic reasoning data is greatly strengthened if the discourse function of universal quantification in Russian, Vai, Kpelle, (Turkish, Berber) and Xhosa is similar to that of English *all*, as reflected in the above categories. I make a start with a corpus study of spoken Xhosa.[6] I made use of the only electronically available corpus of spoken Xhosa[7], drawn mostly from telephone and face-to-face conversations and interviews, and which is still fairly small (around 60 000 words). Nevertheless, the above classification applied surprisingly well to Xhosa.

The Xhosa translation of *all* is the suffix *-onke*, which attaches to one of *z*, *y*, *s*, *w*, *l* or *b*, or stands alone, depending on the type of entities being quantified over – people, animals, or things (with for instance *l* prefixing quantification over locations, *b* and *s* prefixing to people). For example, '*sisebenza sonke*' translates as 'we all work together' (we+work all); '*lonke elo*' translates as 'the whole area' or 'all over (the place)'. The *-onke* suffix has a slightly wider usage than English *all*, as we see in the categorization.

The categories found for English sufficed here except for the fact that *-onke* also translates as *every* and *whole* as well as *all*. These cases are discussed below. More significantly, difference with English was found in the distribution of occurrences across the categories. Use in stock phrases and for emphasis was much less frequent, as was 'forward' quantification. By far the biggest category was the anaphoric/deictic usage of the *-onke* suffix. This is a very suggestive finding in the light of the current claims about the usual function of *all* in everyday discourse; possible implications are drawn out in detail below.

First, the use of *-onke* when translated with *every* and *whole*. These could easily be paraphrased with *all* in English by, say, 'all people' in place of 'everyone' and 'all the world' instead of 'the whole world', 'all day' instead of 'the whole day'. But I will analyse them as they have been translated; nothing rests on the choice for one or other translation since what we are interested in is the use of the constructions in discourses, and how they are related to categories already

[6]The excerpts from the corpus have been translated by Johannesburg-based Xhosa translator Amanda Blossom Bulelwa Nokwele.

[7]The data is drawn from the *Spoken Language Corpora for the Official Languages of Southern Africa* Project, a still-ongoing collaborative research project between the Linguistics Departments at the University of South Africa (UNISA) and the University of Göteborg. The aim is to create the first online textual corpora of "spoken and phatic language use in a variety of social activities in a natural environment" for the nine official African languages of South Africa, one of which is Xhosa.

Category	**no.**	%	% of '*all*' subset
1. stock phrases	5	8%	12%
2. emphatic usage	8	12%	19%
3. anaphoric or deictic use	20	30%	47%
4. forward quantification	8	12%	19%
5. miscellaneous	2	3%	5%
total *all*	43	65%	100%
'everything/one/where'	13	19 %	–
'whole'	11	16 %	–
total	67	100	–

Table 2.3: Categories of '-onke› usage

identified. In many cases of 'every' it is used in a stock or conventional phrasing which can only be figuratively interpreted. In other cases it adds emphasis, and would easily fit into the 'emphatic' use of *all*. Central to use of *every-* phrases (everything, everyone, everywhere) is the necessity to interpret them on restricted domains. This fits with the observation from English that interpretation of quantification is often accompanied by contextually-given domain restriction (about which more later). The following examples illustrate these points:

- '**zonke ke ziza** kwenziwa kakuhle enzele ukuba iimali zingene endaweni eyione amnike **yonke le nto** ayifunayo"
  all you see, SUBJ.will do.PASS well SUBJ.will.do so that the monies SUBJ.will enter at the place that is one
  ***everything** will be done so that the money is channelled to one place.'* (69)

- 'ukuze **izinto zonke** sizibone zihamba kakuhle kungoba ...'
  so that things all PL.SUBJ.see PL.go well it is because ...
  *'so that we were able to see that **everything** goes well, because ...'* (16)

- 'Kulo **lonke** eli lizwe kumdaka.'
  At this all this country there.be.filthy.
  *'**Everywhere** in this country it is filthy.'* (25)

In very few cases was the 'every' what could be called 'forward' quantification, and even then it is clear that there is an anaphoric aspect to it:

- 'and ikhona enyeinto evela ku msoma ethi makusubmithwe **zonke izinto** ezenzeka ezicenteni, ipersonnel, iimali, **zonke ezo zinto**'
  and it is there something else SUBJ.come from Msoma SUBJ.say there SUBJ.must.submit.pass all the things SUBJ.happen.PL at the centres, the personnel, the monies, all those things
  *'... and there is something else from Msoma, **everything** that is happening at the centres must be submitted, the personnel, money/funds, **everything.**'* (68)

The need for pragmatic restriction of the domain in the case of 'every' is made more explicit by the lack of precision in what entities are quantified over: the 'body', 'one', 'thing', 'where' are generic terms for people, objects or locations – more we cannot tell. In this sense the phrasing is anaphoric/deictic: e.g. the use of 'everybody' is functioning like 'all of them' or simply *all*.

The case of 'whole' is even more clear-cut: it's emphatic or stock.

- '**umhlaba wonke wonke** akho mntu ungamaziyo laa tata lowa.'
  the whole whole world there.NEG person SING.SUBJ.do.NEG.know that father there.
  *'in **the whole whole world** there is nobody who does not know that man (colloq.) there.'* (17)

- 'nay(e) udisappointed because kaloku **yonke laa process** kwathiwa mayibuyeleumva **yonke laa process**'[8]
  s/he is also disappointed because you see all that process it.SUBJ.said.PASS it.must.go back all that process
  *'s/he is also disappointed because, you see, it is said **the whole process** must be reversed, **the whole process**.'* (65)

- '**(i)mini yonke** nje kutyiwa idina'
  the whole day there is SUBJ.eat.PASS dinner
  *'dinner is served **the whole day**'* (63).

We now discuss those examples translated with *all*.

1. **Stock phrases** were much less frequent than in English but they were still present:

   - 'naku itishala zigcwele **yonke le ndawo** kunzima ... '
     there teachers PL.SUBJ.full all this place it.be.difficult ...
     *'there are teachers **all over the place**, struggling ... '* (11)

   - 'so ke lilonke sifuna ukuya kabini nje ngesiqhelo'
     so all in all PL.SUBJ.want to go twice as usual
     *'so **all in all** we want to go there twice as often'* (59)

2. **Emphatic usage**:

   - 'sisebenza sonke emsebenzini'
     PL.SUBJ.work all at work
     *'**we all** work together at work'* (52)

   - 'nihleli **ninonke** apha esikolweni'
     PL.SUBJ.sit PL.SUBJ.all together here at school
     *'you are **all** seated here at school'* (14)

[8]Sic. The English phrases which crop up in the Xhosa transcript are examples of code-switching, not typos!

3. The largest group: **anaphoric/deictic** use – and even larger as a proportion of occurrences which are translatable as *all*: anaphoric usage accounts for 47% of this subset. The group is classified by the same criteria as the English category, i.e. followed by pronominal or deictic elements. Examples are:

   - 'xa **bebonke** yithathe le nto uyise emapoliseni'
     when they all SING.SUBJ(2P).IMP.take this thing to police
     *'when **all of them** take this to the police'* (34)
   - 'ayiyeke **yonke laa nto** yoba bendize apha ndizosebenza'
     SING.SUBJ(3P.).MOD.ignore all that thing SING.SUBJ(1P.) SING.came here SING.SUBJ(1P).work
     *'they must ignore **all that**, I came here to work'* (35)
   - '**bonke** sebesifundisa enye into'
     all PL.SUBJ.OBJ.teach other thing
     *'**all of them** are now teaching something else'* (47)

4. And finally, **'forward' quantification** was much less frequent in Xhosa than in English. The law-giving and science writing contexts were not represented at all in this corpus. It is impossible to say conclusively whether this is because Xhosa is a predominantly spoken language or whether this arises from bias in the corpus. Xhosa certainly *is* much less a text-based language than English. The first Xhosa dictionary was compiled in the twentieth century, and in South Africa all tertiary education institutions are still English- or Afrikaans-medium; there are no textbooks beyond the school level written in Xhosa.[9]

   Examples are:

   - 'kuthiwa ngenxa yalo myalelo wenkundla onke amabhinqa akhulelweyo anentsholongwane ihiv atsho anethemba'
     there SUBJ.say.PASS because of this order of the court all females that that are pregnant that have the virus of hiv they have become having hope
     *'it is said that because of this court order **all expectant females living with HIV** are hopeful.'* (73)
   - 'kwindwe kwanyanzela **onke amakristu** ukuba abe ngamajoni' (6)
     personal SING.PAST.force all Christians that they be soldiers
     *'it was imperative that **all Christians** become soldiers'*

[9] We also see many English phrases in the excerpts which are concerned with science or political affairs:

'. . . hayi igovernment yonke . . . abanye individuals . . . '
'*no, the whole goverment, . . . others, individuals*' . . .

- 'umzimba unokuhlaselwa lula ziintsholongwane umzimba unokuhlaselwa zizo zonke iintlobo zezifo
  the body OBJ.MOD.attack.PASS easily by viruses the body OBJ.MOD.attack.PASS by all types of deseases
  'the body can easily be attacked by viruses, it can be attacked by ***all kinds of diseases***' (75)

The Xhosa corpus analysed in the foregoing was small but nevertheless the use of the universally quantifier *-onke* was found to be similarly classifiable to the English quantifier *all*. Differences lay only in the distribution over the categories, *-onke* being used predominantly anaphorically or deictically and much less frequently in other contexts. This finding bolsters the claim *all* or equivalents are typically used 'anaphorically' in combination with a previously or pragmatically specified domain, especially in spoken language, and supports our explanation of the difficulty in reasoning with such premises, as outlined over the next two sections.

Summarising the corpora data

Categorising the occurrences of *all* and the equivalent Xhosa *-onke* in spoken discourses has yielded a perhaps unexpected picture of their typical usage, one that suggests that the "quantifier" is often not actively quantifying. As already mentioned, in most of the stock phrases it is impossible to interpret it literally as quantifying. In most examples it is used in a metaphorical way: 'all his heart'. But these stock phrases are language-specific and hence this category might not have as much significance cross-linguistically. Similarly emphatic usage seems to represent an ossified version of quantification, in which *all* cannot be traded in for any other sort of quantifier – in fact, it can usually be done without. And also here, it is unclear what the cross-linguistic significance of emphatic usage of universal quantifiers would be.

Discounting the appearance of the quantifier in stock phrases and emphatic usage, we are left with two large categories where *all* or its translation is used: anaphoric/deictic and what we've called 'forward' quantification. In the anaphoric/-deictic category, *all* performs a summarizing function, acting as a kind of fishing net for aforementioned or contextually given referents, about which something *further* can then be said. This is more plausibly a universal usage – witness the frequency of this usage in Xhosa. Here *all* may be truly quantificational, but it quantifies over an extremely restricted domain: that given by the previous discourse or the context of utterance.

By contrast, the cases that have been labelled 'forward' quantification don't seem to sum up anything previously mentioned, but are 'forward-looking' in the sense that what they quantify over is introduced after *all* is: 'Thanks were extended to all the Kent teachers who had hosted the event', 'Christ had died for all, all men and women'. It is in these cases that *all* is most likely to be available

as a premise (the argument 'Christ died for all men and women, therefore Christ died for you' sounds vaguely like Catechism class).

Observe, however, that in both the large categories of *all* usage there appears to be a division of semantic labour between linguistic and contextual input (either from the previous discourse or the non-linguistic environment) by which the domain of quantification is determined. In some cases – the legal, religious, and science contexts – the burden lies more on the linguistic side, and can potentially be contained totally in the quantifying sentence. Although 'anaphoric' and 'forward' usages have been presented as separate categories, they could be better viewed as points on a slide of determining quantification – the one side being quantification items being determined by 'new' linguistic information, the other side by previously or contextually given information. The large category of quantifier usage for which the domain is determined 'purely' anaphorically should thus be understood as an articulated continuation of what extends below the threshold of linguistic explicitness – determination of the domain by context.

Contexts in which 'forward' quantification was used were often derived from written discourses – think about the religious and law contexts.[10] What is interesting is that the quantification in these subcategories is the closest thing to strict quantification. Any exceptions to the universality would at least have to be specified. For example, consider (ECF 3790) the following, 'A complementary excursion to the Dolomites for all guests staying 14 nights'. We might well expect exceptions: guests who benefit from some other special offer, or stay on reduced rates, might not be entitled to their free trip to the Dolomites – but this would have to be explicitly mentioned (the small print!), as a caveat to the rule 'all guests get a free trip'. And of course, 'all guests' is understood to apply only to a certain group of guests – those who stay at whichever hotel made the offer.

As this illustrates, the common condition in uses of *all* as quantifier is the pairing of its usage with determination of its domain. Quantification always functions over a domain. When we say 'The burglar took everything' we take it that 'everything' ranges only over the valuable objects in a certain house.[11] When we say 'All students sat the exam' we mean all students who were registered for a particular course at a particular university in a given term. In these cases the exact domain will be either explicitly or implicitly given by the previous discourse or by the context of utterance: in the first example, the speaker will have introduced the topic by saying who the burglary 'happened to', say, Mr and Mrs Bloggs, and the hearer infers that the 'everything' in the sentence refers to the valuable objects in the Bloggs' household. In the second sentence, the identity of the conversants might be enough to deduce what *all* quantifies over: if the speaker is the teaching assistant for Maths 101 at the University of Cape

[10]Sometimes phrasing even betrays a written text underneath the spoken one: 'all of the above' (CET 1734) (and occurs again in FTB 1394), a visual metaphor only appropriate for written texts (possibly the transcript is from a text being read aloud).

[11]Example from Recanati (1996).

Town in the second semester of 2006, and is addressing the professor teaching the course, then they could reasonably conclude that the registration list for this course comprises the domain of quantification.

These informal remarks are made more precise in the following, in which we explore the matter of fixing an appropriate domain of quantification and how exactly the type of quantification interacts with its domain of quantification.

2.2.2 The semantics of *all that*

The most intriguing category of *all* usage is that of anaphoric/deictic usage, where what is quantified over is not made explicit in the quantifying statement but is indicated by demonstrative elements presumably referring to the previous discourse context or extra-linguistic context, as for instance in

(3) 'Of course I know **all that's** no reason I can't have a bike' (G3P 1937)

(4) '**all this** will take time to negotiate' (ABE 2666)

Now given the fact that *all* so often combines with anaphoric or deictic elements, we might wonder what kind of formal machinery is needed to enable this to function well. What is needed to fix in each case 'this' or 'that' which is being quantified?

Domains are given by context sets

The pervasiveness of the anaphoric/deictic category highlights a distinction which has already been proposed in formal work on quantifiers. For instance, Westerståhl (1985) argues that in providing a semantic analysis for *all* one needs to distinguish three types of universe – as opposed to just two as is usually proposed in the 'flexible universe' strategy, in which pragmatic processes are assumed to continuously adjust the discourse universe appropriately.[12]

In more detail, his account is as follows. Model-theoretic semantics routinely makes reference to a universe of models, or *discourse universe*, M, in a model $\mathcal{M} = < M, [[.]] >$[13] as well as the denotation of the noun in the model (where NP = determiner + noun), which can be viewed as the *NP universe*. But Westerståhl (1985) argues that semantics should distinguish also a *context set*, a contextually selected sub-universe of M. The role of selecting this sub-universe is usually assigned to pragmatics; as Westerståhl says, "in practice this means identifying context sets with (temporarily chosen) model universes" (p. 46). He offers two types of argument why context sets cannot be identified with discourse universes, the first methodological, and the second via concrete examples. Methodologically

[12] He goes on to show how this three-way distinction can be implemented in an enriched version Barwise and Cooper's Logic for Generalised Quantifiers.

[13] [[.]] is an interpretation function assigning interpretations to natural language expressions.

the two do not match because discourse universes are large and constant over pieces of discourse while context sets are not, and determiners are 'universe-independent' (pp. 48 – 51) in the sense that their interpretation is not affected by the discourse universe in which they occur, as long as it's large enough (a property referred to as EXT in many logic textbooks[14]).

The vivid second argument offered by Westerståhl is in the form of examples for which there is simply "no way to make sense of [the] sentences if the discourse universe is identified with the context set" (p. 49), viz:

(5) The English love to write letters. *Most* children have *several* pen pals in many countries. (emphasis mine)

Now this sentence only makes sense if we understand *most* to be quantifying over children in England but *several* to be quantifying over children *anywhere* in the world. Were we to identify the context set with the discourse universe, either the first NP-universe would be 'most children in the world' or the second one would be 'several children in England' – clearly neither of which is intended in the context of utterance. In fact, the first NP does not operate on the discourse universe but on a restricted context set, given by the previous sentence, while the second NP operates over a bigger set. Obviously, the discourse universe must include this bigger set of children anywhere. The possibility remains open that in certain sentences the context set coincides with the universe of discourse M – for instance in the amended example below the context set remains the same for both NPs and can be taken to coincide with the discourse universe:

(6) Children love to write letters. Most children have several pen pals in many countries.

Specifying a role for the context set as apart from the NP-universe in determining the domain of quantification can explain how anaphoric and deictic usage works. In these cases, the context set gives the sub-universe in which the anaphoric or deictic elements are interpreted – and the NP-universe restriction is minimal. Moreover, the size of the category anaphoric/deictic illustrates the importance of the context set in determining the domain of quantification in everyday use of the quantifier *all*. For demonstratives such as 'that', 'this', 'those' can only get the appropriate denotation when interpreted in a restricted sub-universe of the discourse universe, i.e. in a context set. Westerståhl proposes just such an explanation for bare or pronominal use of *all*, arguing that "the lack of an argument is a visible *context set indicator*, which signals the implicit occurrence of a context set" (*ibid*, p. 49, his emphasis).

[14]Formally expressed by

(EXT) If $A, B \subseteq M \subseteq M'$ then $D_M AB$ iff $D_{M'} AB$

To place Westerståhl's proposal in context, more recent work in linguistics makes this same kind of distinction albeit with different terminology. Apart from explicitly quantified formulations such as

(7) All the girls jumped in the lake.

there are 'bare' plural noun phrases such as

(8) The girls jumped in the lake.

Yet these are often taken to imply quantified formulations such as

(9) Every girl jumped in the lake.

This phenomenon has lead many theorists to posit a so-called "D operator", an implicit distributivity operator which introduces quantification in the denotation of the plural 'the girls' (Brisson, 2003). Note though that the first sentence above is slightly 'weaker' in that, as will be discussed below, its truth can in many situations withstand a few non-swimming girls. In linguists' lingo, the quantification is nonmaximal. The influential paper by Brisson (2003) uses Schwarzschild's (1996) idea "that a D-operator has a domain variable in its restriction *whose value is contextually specified*" (p.130, my emphasis), to propose that *all* is in fact not a determiner-quantifier, "but rather interacts with the quantification introduced by the D operator to rule out the nonmaximality that a D operator normally allows" (p. 141). 'Nonmaximality' here means allowing for exceptions, so in this respect resembles genericity. Brisson's proposal is that *all*, as witnessed in (7), functions as a means to adjust the domain to make it maximal, *where the domain is contextually-selected by a variable within the D operator*. In this function *all* is different from quantifiers such as *every* or *each*, which also cannot combine with plural noun phrases, viz:

(10) *Every the girls jumped in the lake.

Although Brisson's proposal is very suggestive in the light of the current analysis, the further details go beyond our purpose and scope. Westerståhl's relatively simple proposal suffices make the general idea of local domain determination precise, but it should be clear from the above that alternative formulations are available.

To return to Westerståhl's terminology: as we will see, in the context of the reasoning task the context set is considered identical with the discourse universe, from the experimenters' point of view, but in certain conditions the two could – and even should – sensibly be distinguished, and this leads to a divergence in denotation for the determiner NP for the experimenter and the unschooled subject. But under which conditions is the domain of quantification given by context set as different from the universe of discourse, and how does the quantification relate

to the context set in such cases? We now turn to address this question.

Law-like and contingent generalisations

Two main interpretations for *all* sentences are considered and discussed in relation to domain determination and the role of the context set. The first type of interpretation considers statements of the form '*All x are y*' as expressing some kind of theory-supported or causal generalisation. This does not necessarily mean that the quantification is interpreted strictly. Certainly in English *all* can be used to make a generic statement.[15] Consider the sentences

a. All doctors wear white coats.[16]

b. All bears have four legs.[17]

One can read the sentences as expressing something about the prototypical doctor, or bear, and thus equivalent to the bare generic formulation

a.′ Doctors wear white coats.

b.′ Bears have four legs.

This only goes through in the case of a conventional or causal, that is, a law-like, connection. In such cases, the generalisation can withstand counterexamples, or rather, exceptions, so that the odd three-legged bear, having perhaps been caught in a hunter's trap, does not shake our belief that 'All bears have four legs'. This does not work for contingent generalisations. Borrowing from an example in Pelletier & Asher (1997), if by some cruel twist of fate all bears in the world lost a leg, we would not assent to 'Bears have three legs', but would maintain the generic 'Bears have four legs' – even while admitting that 'All the bears have three legs'. In other words, the generic reading of a universally quantified statement is not available when it expresses a mere contingent generalisation. Here genericity is expressed with a bare noun phrase – a point to which we shall return. In general, the degree of robustness to exceptions seems to vary with the degree of theoretical basis for the generalisation.

Notice however that in certain contexts the generic reading of *all* is ruled out and it gets a strict reading while paired with a law-like connection. This is the case for statements expressing scientific or theory-based laws, such as 'All bears are mammals', and Goodman's example 'All butter melts at 150° centigrade', as will be discussed in more detail further on. In these cases one can read the quantification as applying to *kinds* (of bears, or butter) rather than individuals (single bears or pats of butter). Combination with deontic modals in imperative

[15]Later on in this chapter it is suggested that such a reading of *all* statements might be an artifact of linguistic theorising and not one common in naturally occurring language use.

[16]Example from Partee (1985).

[17]Example from Pelletier & Asher (1997).

statements also yields an exceptionless reading: 'all doctors must wear white coats', although an exception here does not falsify the rule but violate it.[18] As we saw in the previous chapter, many subjects seemed to take a deontic interpretation of some reasoning premises, such as 'All people who own houses pay house tax'.

A second available interpretation for an *all* sentence is that of a contingent generalisation, where the connection described by the predication is accidental. Imagine, if you will (after Goodman, 1947), that upon checking my pockets this morning at home, I found only silver coins in my right pocket. Then the universal generalisation 'All the coins in my right pocket are silver' is true. But now suppose that on the way to work I bought a coffee and slipped the change into my right pocket. Does the generalisation still hold? Maybe, maybe not. It depends what I put in my pocket. Certainly we wouldn't say that the additional coins *became* silver on being put into my right pocket. If there was a copper coin among my change then it's no longer a true generalisation. The generalisation may be 'universal', but only in the small and rather gloomy universe comprising the inside of my pocket on a particular winter morning. It's a very circumscribed and contingently constituted domain.

These can be compared to universal generalisations which express physical laws, such as

(11) All butter melts at 150° F

This is a statement of a law-like relationship, and can be distinguished from true contingent generalisations like

(12) All the coins in my pocket are silver

by the fact that the first statement can be accepted as true before all cases of it have been determined – these undetermined cases being predicted to conform with the law. In contrast, a statement like (12) "is accepted as a description of contingent facts *after* the determination of all cases, *no prediction of any of its instances being based upon it*" (Goodman 1947, p. 124, second emphasis mine). In other words, there are different criteria of acceptance for the two kinds of statement. For the case of a law, a few positive instances may lead us to accept the statement as true, but in the case of an accidental generalisation we need to have tested all instances before we can accept it as true.

Why did I assent to the sentence 'All the coins in my pocket are silver' this morning? I could do so precisely because the sentence is taken to refer to all coins that were in my pocket *at the time of utterance* – and not the coins which have been or will be in my pocket. This is what makes it a contingent generalisation. It's a description of a part of the world at a particular point in time. We might say, the generalisation operates on a context set which can be, and probably is, much smaller than the discourse universe. The only instances which counted

[18] See Stenning and van Lambalgen (2001) for an extended discussion of this in the context of the Wason selection task.

were those known of in the specific situation. Sometimes this can also involve past instances – 'All my boyfriends are short' – but essentially nothing can be claimed about future instances which will fall under the generalisation, i.e. future coins in my pocket, or future boyfriends, or even instances which fall beyond the domain in some other sense – like coins in the trousers hanging over a chair in my bedroom. So we see that, in cases where we do accept contingent, un-lawlike generalisations it is because there is a limited range of instances to which they apply, and these have been checked.

There is a further difference in the examples offered by Goodman contrasting law-like with contingent generalisation: the latter is expressed with a definite article. In English, the use of the definite article 'the' after *all* supports a contingent reading because it suggests an identifiable and finite domain. For example, compare

(13) All women have two children

to

(14) All the women have two children.

In the latter it is clear we are talking about a specific group of women to whom the generalisation applies; in the former this reading is not available without considerable contextual support. But although the bare version does not get a contingent reading, the reverse is not always the case. For example, 'All the bears in the North shed their winter coat' can be read as generic, but, again, in that case it does suggest quantification over *types* of bears rather than individual animals. Note that an accompanying definite article rules out tolerance of counterexamples, so that a single woman in the relevant domain with only one child falsifies the generalisation 'All the women have two children'.[19]

For English materials the use of the definite article would create a lurking confound when testing syllogistic premises, if it is indeed associated with contingent generalisations, because, as we'll see, a contingent generalisation makes for an awkward premise. In the current study, we can ignore this issue because in Xhosa, there is no distinct part of speech corresponding to the definite article in English (Tsolwana, 1996), so that, for example, 'all women' and 'all the women' both get translated as 'bonke abafazi'. This means that the difference between generic/law-like and contingent generalisation is not expressed by the use of the definite article. This will turn out to be an important point in our later analysis.

The two kinds of generalisations can be seen as opposite poles on a scale of domain-sensitivity. On the law-like end, we have generalisations which are often based on causal relations, and which can therefore be judged on the basis of few instances. They are in this way true universal generalisations, beyond any given

[19]This is often called the 'maximising' character of *all*, contrasting with the nonmaximality of 'The women have two children (each)' (Brisson, 2003). See previous section for more discussion of this feature.

domain (the lab, say, where the experiment is conducted). There are two things to note about this: 1) even, or rather, especially, here, ceteris paribus clauses hold, qualifying extrapolation to new cases; 2) these generalisations are usually not explicitly quantified: for example, 'butter melts at 150° F' expresses the same law as that expressed in (1). Simply by virtue of being butter it falls under the generalisation. In fact, law-like regularities are probably not expressed in natural language with universal quantifiers, but with generic statements. Goodman's distinction between the two types of generalisation might only be relevant in the domain of scientific language. Should this be so, it only strengthens the current claim that universal statements are usually used contingently and with a prespecified domain.

What about *all* statements which are read generically such as described above, but which are not interpreted strictly and are thus not falsified by single 'counterexamples'? These are the generalisations which describe stereotypes, patterns, habits, typicalities, which are more predictable than simple accidents, but which are not as reliable as laws of nature. These can cross domains but the ceteris paribus clause will become more difficult to enforce and the presence of instances which serve as exceptions will increase. A key aspect of these exceptions is that they are not real counterexamples because they need not negate the quantification. As illustration here consider the example

(15) All chairs have four legs

or the previous example

(16) All doctors wear white coats.

One might agree that these statements hold in some general sense while also granting that in certain contexts – respectively, say, an *avant-garde* design exhibition and a surgical operating theatre – exceptions will abound.

On the other end of the scale we have these descriptions of states of affairs which express entirely contingent or accidental generalisations, which do not extend beyond a known domain. This is what we might call local universal generalisation. Here quantifying expressions do real work because there is no inherent property of the entities involved which makes them fall under the generalisation: there's nothing causal connecting my pocket and silver coins. As such, should anything change about the domain – such as new coins get added – we can say nothing about the status of the generalisation. It's hyper domain-dependent in a way that laws aren't. Another way to see this is to compare the quantified formulation of generalisation with conditional formulation of it: 'if something is butter, it melts at 150°F' holds, under ceteris paribus clauses, for undetermined entities. But 'if something is a coin in my pocket, it is silver', apart from being a very awkward paraphrase of the quantified statement, we can only be sure to be true for entities in the original domain. (See the next section for more discussion of conditional formulations.) Here especially, we might expect context sets

to play a key role in domain determination – i.e. contextually-given, surveyable sub-universes in which instances of the generalisation are known or visible.

If we grant that quantifiers are usually interpreted on a contextually-given sub-universe of the discourse universe, then we also have a means to explain the difference between law-like and contingent generalisations. The current proposal is that contingent generalisations are interpreted on a context set which is varying from context to context, while law-like generalisations do not rely on a context set for their interpretation: they can be interpreted on any universe of discourse. The phenomenon that Goodman observed, that one can extrapolate to new untested instances which were not necessarily part of the original domain, stems from the theoretical basis for the generalisation.

The next section shows why, in experimental conditions, the failure to distinguish a role for the context set might lead the experimenter to project an inaccurate – and perhaps unwarranted – interpretation for the quantified premises onto unschooled subjects' reasoning.

2.2.3 When *all* sentences make lousy premises

In the syllogistic task, the subject is first presented with a universal generalisation, say "All bears in Novaya Zemlya are white." One can interpret this either as a strict law-like, or a generic, or a contingent generalisation. Next, the subject is presented with a possible instance of this generalisation, in the second premise "My friend saw a bear in Novaya Zemlya." Then, depending on which reading of the generalisation is taken, the subject should answer the question: "What colour was the bear?". Herein lies the anomaly:

Option 1: Strict law-like. All instances of bears are predicted to conform, so answering that the bear is white is minimally informative, given that the subject presumes that their interlocutor knows what she's asserting (i.e. the premises). The answer would be more informative if there was some reason to believe that the bear was a counterexample to the law, or if the question is understood to be about the basis for a law-like connection, including specifying what the ceteris paribus clauses involves – i.e. what determines the certainty of prediction. A response which fits this reading is found in the following excerpts (see especially emphasised parts):

Abdurakhm., age 37, illiterate.

S: There are different sorts of bears.
[The syllogism is repeated.]
S: I don't know; I've seen a black bear, I've never seen any others ...
Each locality has its own animals: if it's white, they will be white; if

it's yellow, they will be yellow.[20]

Ishankul, age 63, illiterate.

S: If you say that they are white from the cold, they should be white there too. Probably they are even whiter than in Russia.

Recall also our discussion of the strictly unnecessary use of *all* for expression of law-like regularities, something which might contribute to the improbability of the law-like reading of the premise.[21]

Preliminary evidence that subjects have problems taking a 'law-like' reading of the generalisation are given in the following two excerpts. Firstly, the sheer impossibility of a truly universal reading of the quantification is voiced:

Florence, group 2:

E: OK. So one more question. So suppose that all lawyers are alcoholics. And all alcoholics smoke cigarettes. Do all lawyers smoke cigarettes?
S: No, they can't all smoke.
E: But suppose we make it true.
S: *But even though we make it true, there will be others that don't smoke.*
E: It's impossible to make it true?
S: Yes.

The exchange indicates a tension in her to accept the generalisation on a universal domain because of its inherent implausibility. She seems to be saying: it can only be true in a limited domain.

The second excerpt contains an incisive query about the sense of a law-like reading – recall the coins which surely do not turn silver on being slipped into my pocket. This query comes from a schooled subject.

Nontembeko, group 3:

E: So suppose there's this imaginary country called Markia. And all the women who live in Markia are married. And Fatma is a woman who lives in Markia. Is Fatma married?
S: If Fatma is a lady then definitely she's married because all the women who live in Markia are married.
...
S: *Is it the law that all the women there are married? Is that possible if it's me, I'm going there, and I'm not married?*
E: You can go there as a visitor. If you visit, you don't have to be married.

The outcome, as seen in the last turn, is that the experimenter is forced to qualify the generalisation to exclude visiting women!

[20]Where I read 'it' as referring to the locality. An intriguing suggestion of natural selection? (Compare with the case of the English Peppered Moth.)

[21]Obviously, this only applies to generalisations expressed using *all*.

Option 2: Generic. As discussed above, these tolerate counterexamples. So for example, bears in Novaya Zemlya might usually be white, and the odd roaming brown bear wouldn't threaten this generalisation. Any individual bear could thus be an exemplar or an exception. It might even be that singling one out suggests exception rather than exemplification. Support for this is given by Clark & Bangerter's (2004) review of research which shows that subjects identify referents according to salience against the common ground, where common ground includes for instance 'given' information. Witness, for example, Nofezile from my study, upon presentation of the problem 'In Markia all women are married, and Fatma is a woman who lives in Markia. Do you think she's married?':

Nofezile, group 2:

S: Does she stay alone?

E: We don't know. All we know is that all the women in Markia are married.

S: I don't think she's married if she stays there.

E: Why?

S: *I will say so because you said all the women in Markia are married and then you say her, living in Markia alone.*

There are two ways to interpret the subject's initial response, 'Does she stay alone?' Firstly, one can read this as a question about the law-like basis for the generalisation: on what basis can we go beyond the known instances? This suggests a law-like reading is available to the subject, but note that we can get away with positing a deontic law-like reading, something like: 'All women in Markia must marry their live-in boyfriends'. Alternatively we can understand the question as one after further specification of the domain, trying to establish a context set so to speak. The last turn indeed suggests that the subject has introduced her own context set, and decided that the protagonist must fall outside the putative – restricted – domain of application, that is, cohabiting women in Markia. And indeed, why would we mention Fatma unless she was somehow exceptional?

Option 3: Contingent. Now either this is an untested instance about which we strictly can say nothing, or it's a known instance in which we're in situation of Option 1. The countless refusals along the lines of 'I don't know, I've never been there' can be understood as one of two possibilities on this analysis: either the subject is saying that he has interpreted the premise as a contingent generalisation, in which case it's a moot point whether any new instance falls under it or not – he would have to know what the context set was as it were – or that he has no grounds for a theoretical basis for the generalisation, which would enable him to predict the colour of the next bear simply on the basis of the given generalisation. An informative answer here would be to give conditions under which

you could determine which of the cases you're in – i.e. to specify whether or not and why the bear in question falls under the domain of the generalisation or not. This is what we find: witness Nozuko seemingly resorting to a limited-domain interpretation when no theoretical grounds for a law-like reading are forthcoming:

Nozuko, group 3:

E: Suppose all the women in Nigeria are married. Now there's a woman called Connie and she's not married. Can we say she lives in Nigeria or not?

S: What kind of clothes do they wear in Nigeria?[22]

E: Just suppose the world is a strange one in which all the women in Nigeria are married.

S: We can say she's a Nigerian but she hasn't got married yet.

This last turn – especially the 'yet' – can also be understood as a plea for a temporally-delimited interpretation of the concept "woman" as females of marrying age.

After surveying these three options which are open to a subject on hearing the question 'What colour is the bear?', we see that answering 'white', is not as straightforward as it might seem. In fact, the given range of options suggests that the subtext of the question is a question after the applicability of the generalisation. Especially the generic reading of the generalisation suggests the question be best interpreted thus. On the other hand, to simply use the generalisation as a premise, thus taking a 'straight' reading of the question, seems to result in either being uninformative, or rash. If this is indeed the case, we would expect reactions to include further suggestions or queries about specification of the domain; or about the relation of the named particular to the domain. This type of reaction is indeed present in a good many responses.

An alternative means to capture the oddity of the syllogistic question is in terms of the difference between an instance being known to conform and predicted to conform. Using the given premises as intended implicates that the conclusion is foregone because the instance is known to conform, but this is uninformative. On the other hand, assessing whether the instance is predicted to conform results in an informative response, but involves assessing the domain and type of quantification, and thus going beyond the question as asked.

This mismatch between domain and interpretation can explain the frequent asking for or volunteering specification of the limits of the domain of quantification. This occurred in all groups but was less common in the schooled group. An example:

[22]In Xhosa custom a married woman indicates her status by wearing a specific kind of dress: she always covers her head, and wears an apron. There is also special language only for married women: some objects (e.g. cattle, stones) have two names, one for use by married women, one for use by everybody else.

Nonkululeko, group 2:

Preamble: 'All people who own houses pay house tax. Sabelo does not pay house tax. Does he own a house?'

S: He doesn't have a house if he's not paying.

E: And now suppose that none of the people in Cape Town pay house tax. Do they own houses or not?

S: They have houses.

E: Why?

S: *They can have houses because there are places where you don't pay tax, like the squatter camps.*

E: So they can have houses and not pay?

S: They may, they can live at the squatter camps.

The subjects' response to the question, 'Do [the people in Cape Town] own houses or not?' might be taken as a Luria-style 'ignoring the premises' response, or even as inconsistent with her previous turn, if it were not for her next response, in which she justifies her conclusion, that people in Cape Town do have houses, by explaining that the given generalisation does not hold in Cape Town – that is, we've gone beyond a restricted domain – the context set – in which the quantification was strict. She is in essence saying, in her second turn, 'I am free to assert they have houses in Cape Town because the given statements do not apply there.'

So if the generalisation is so problematic for some subjects, how is it supposed to function from the experimenter's point of view? The domain restriction is intended to be 'self-contained', i.e. given only by the NP, with no further domain restriction given by a previously determined context set. Now we can connect this with the insights of both Westerståhl and Goodman with regards the nature of domain-relations. In anaphoric use of quantification this means that all elements of the domain are 'known' – or at least they have been previously indicated by the context. 'Forward' quantification, so long as it introduces the elements to be quantified over only in the generalisation itself, leaves open the extent of the domain, thus allowing for previously unmentioned or unknown elements to still fall under it. I would like to speculate that interpreting the generalisation on an unspecified context set (or, as a putative default, the universal domain), which is required for correct response in the syllogistic task, might be more common in scientific or literate discourses, and the assumption that the unschooled subject understands it as a law is thus problematic. For instance, witness Luria's observation that: "In all instances, when a subject repeated the premises he did not give them the character of universal assertions" (1976, p. 106). As we have seen, taking the strict (unrestricted) universal reading of *all*, where a ceteris paribus clause suffices for domain restriction, is atypical for everyday spoken language use of universal quantifiers. In fact, I would speculate that the whole idea of

a universal domain is a literate idiosyncrasy. Yet this is the default domain in syllogistic tasks.

If this is indeed the case, then the difficulties that unschooled subjects have with quantified premises should not be attributed to their faulty reasoning ability, but rather to the experimenter's chauvinistic view of what interpretations are available and plausible for such premises. The implicit theory of language supporting such chauvinism, one which is maintained in much reasoning research, is articulated and critically evaluated in Chapter 4.

The most obvious prediction that follows from the account given here of difficulties with *all* premises is that such premises would become unproblematic when the nature of the generalisation is clarified. So, for instance, if the context set was specified then subjects should be able to answer questions about particular instances of the generalisation. Recent work (Haan, 2007) addresses this question and verifies that this is indeed the case. Haan conducted two reasoning tasks with an explicit context set, designed to test whether an explicit domain aids unschooled subjects. The results from the tasks confirm that this is the case.

In the first task, the 'box task' (versions one and two), subjects were presented with a tray containing three red boxes and shown that each box contained a ping-pong ball. The content of the boxes was then hidden and the subjects were asked the following questions:

i. 'Is it true that all red boxes contain a ball?'

ii. (On one of the boxes being produced again): 'What is in this box?'

and in a third version which rules out *yes*-bias the subject was first shown three blue boxes, only two of which have a ball in them, and asked

iii. 'Is it true that all blue boxes contain a ball?'

The rates of correct answers for the tasks were respectively 100%, 69%, and 100%, thus confirming our prediction that using universal generalisations in a typical fashion (that is, contingently on a specified domain) removes difficulty for unschooled subjects in reasoning with them. Such a task set-up looks to be very fruitful for further investigation because the multi-modal sources of information avoid many of the problems of informativeness often present in purely linguistic tasks.

We should however give some attention to the lower elicitation rates for the second question, which Haan labels the second "version" of the Box task. In fact, given that the order of questions was not counterbalanced, we should see the second question as a follow-up on the first, and thus as part of the same 'version' of the task. Asking the question 'what is in this box?' when the previous question 'is it true that all red boxes contain a ball?' certainly changes the informativeness of various answers from the situation in which it would be asked first-off. For instance, the subject might read the second question as some kind of trick;

a 'straight' reading of the question puts us back to a question-answer situation comparable to the purely linguistic tasks, with their attendant pathological information structure. But even if the question had been asked first, uncertainty remains as to what kind of answer is required. What level of detail would be informative? Was the box empty to start with? Any of all of these factors could explain the lower rates of correct answer to the second question.

Haan's second task designed to test the role of domain-specification did not yield such high performance. In this task subjects were presented with the following set of premises[23]:

> I have three brothers and one sister. All of my three brothers live in Rotterdam. Jan is one of my brothers. Where does he live?

On this task, perhaps surprisingly, subjects gave the answer 'Rotterdam' only 25% of the time. However there are several differences to the box task which make the 'brothers' task more problematic. Firstly, the question as phrased is underspecified. 'Where?' could mean 'which city?' 'which street?' or 'where in relation to the rest of the family?'. Given the fact that the premises already state the brother lives in Rotterdam the subject could judge it uninformative to give this as an answer and would then be stuck for a specific answer. A further issue with this question-answer set-up is that of authority: the experimenter tells of his own situation, so he clearly is the authority on the matter and once again we have anomalous epistemic asymmetry (the questioner knowing more than the person he's asking) as described in the previous chapter. Moreover the naming of brother Jan separately and after the quantified statement leads to the same unclarity of the relation of the named individual to the domain: does he belong to the original domain or not? Although on the standard reading 'brothers' is a closed set, a metaphorical interpretation of the concept, something which is certainly very common in South Africa, would allow new members to be added indefinitely. If the same reading is available in Moroccan or Turkish culture then this would introduce another interpretation of the premises. For these reasons I take the outcome of the 'brothers' task to be non-consequential for the predictions outlined here.

2.3 What *if* conditionals are easier?

In the foregoing, I identified aspects of the use and interpretation of the universal quantifier *all*, which can explain the difficulty unschooled subjects have with syllogistic-style premises. In this section I focus on conditional premises with aim of finding candidate explanations for the relative ease with which subjects reasoned in the simple condition of the suppression-effect task materials.

[23] A variation on this, the 'daughters' task, is judged to suffer from the same ambiguities as discussed above so I will not discuss it further.

As described above, unschooled subjects were much less inclined to give 'non-answers' (what Luria termed 'refusals') when presented with this type of conditional premises. Perhaps even more interestingly, within the schooled group subjects were more inclined to give elaborations to such conditional premises than they were with quantified premises. In other words, the gap between the two groups narrows, from both sides, when subjects are presented with conditional premises.

There are several factors which may play a role in explaining why this is so. Firstly, the conditional premises were presented within a 'story' context, so that subjects were, for instance, first told,

> 'This is about a girl Thembi. She lives in Hamburg.'

Then the premises are presented:

> '*If Thembi wants to see her boyfriend she goes to East London*, and *Thembi does want to see her boyfriend.* So will she go to East London?'

Now this embedding in a narrative context might partially explain subjects' comfort with the material; it has been argued that narrative discourse is the most basic type of discourse (Bruner, 1991). It would indeed be interesting to compare this presentation of premises to unschooled subjects with the original premises used by Byrne (1989) which opened cold with: 'If *she* has an essay to write ...' (Who's she? The cat's mother?) In the conclusion I return to this point when suggesting new experimental conditions.

However, simple narrative embedding does not tell the whole story, because in the syllogistic task subjects also got some sort of introduction to the material before the premises are presented:

> 'Suppose there's an imaginary country called Markia. And in this country *all (the) women are married. Fatma is a woman who lives in Markia.* Is she married?'

Moreover, as we have just seen in Haan's (2007) 'brothers' task (and even more so for the 'daughters' task), simply embedding syllogistic premises further does not improve performance on the task. The difference lies rather in the quality of the embedding. In both problems the subject is asked to draw a conclusion about a specific named individual, a 'protagonist'. One factor in the conditional task is that the protagonist is introduced *before* the generalised premise – here in the form of a conditional – is given. By contrast, in the syllogistic task, the particular individual about whom a conclusion must be drawn is introduced *after* the generalisation. As discussed in the previous section, this might make a subject more inclined to view the protagonist as somehow contrastive, and therefore deviating from the generalisation. But the plausibility of such a reading of the protagonist's position is more dependent on the form of the generalisation than the order in which it is presented – as will become clear as we further

our analysis. Given these considerations, the difference between conditional and quantified formulations of a generalisation remains a potentially significant factor in the different tasks, especially given the foregoing analysis of *all* and its 'misuse', so to speak, in syllogistic problems. It is to this which we now turn.

This variation across premises is by no means obvious because, as mentioned, both quantified statements and conditional statements are typically given a logical form containing an implication. Both can express generalisations and thus say something about more than one entity or event, meaning that in both cases 'domains', understood in a loose sense (and which we'll make more precise later on), are important. It has been argued in the previous section that *all* premises are difficult because they are used in reasoning tasks with hanging semantic parameters; more specifically, an unspecified context set. No distinction is made between the universe of discourse and context set even though, especially in the case of contingent generalisations, this is usual. In Chapter 1 it was argued that conditional premises also have hidden elements, namely abnormality clauses, which have not been taken into account in the analysis of experiments. We might just as well expect subjects to stumble with such premises, because of the unaccounted-for elements, as they did with *all* premises. In this section the aim is to distinguish reasons why the conditional sentences used in the study are easier for our reasoning subjects. As with *all*, a start is made by investigating how *if* phrases are used in everyday spoken language. This then guides a semantic analysis of experimental materials.

2.3.1 *if* in conversation and reasoning

Unlike the situation for universal quantifiers, there is much existing research on the use of *if* conditionals in practice. Firstly, conditionals have consistently found to be more common in spoken than in written discourse: Hwang's (1979) analysis of a corpus of spoken (of 63 746 words) and written English (357 249 words) found 4.2 *if* conditionals per 1000 words in speech against 2.7 per 1000 in text, and Ford & Thompson (1986) found 7.2 per 1000 words in speech versus 4.6 per 1000 in text. Secondly, initial conditional clauses – that is, a preposed *if* clause – outnumber final conditional clauses by about three to one, or even more in spoken language; Ford & Thompson (1986) found initial *if* clauses in 82% of spoken occurrences of the conditional. Comrie (1986), in a typological study, found no counterexamples to a preference for initial conditionals. So far so good: *if*-clauses, and especially pre-posed ones, are common in spoken language.

Conditional constructions are ubiquitous, complex, and varied. Apart from *if*, conditionality can be expressed in English with many other expressions, including *when*, *whenever*, *whether*, even *and*, *or* ('Do it and/or I'll punish you!') or simply co-subordination ('The more I work, the more I earn')[24]. As for typologies of

[24]Both examples are from Declerck & Reed, 2001. They distinguish 14 different syntactically

conditional constructions, there are several available (Comrie, 1986; Declerck & Reed, 2001, but also Haegeman, 2003, and from a reasoning perspective, Bonnefon & Hilton, 2005). So for instance, Comrie (1986) offers a cross-linguistic typology of conditional phrases, and judges clause order, marking of conditionality, degrees of hypotheticality and time reference to be relevant parameters in classifying conditional constructions.

Here we are interested in conditionality as expressed with *if* clauses. Since Comrie is dealing with a much wider class of constructions than those, not all of his parameters are relevant for us. Only varying degrees of hypotheticality and the matter of time reference will be a feature of the typology offered here. Similarly, Declerck and Reed's analysis is broader than that needed here. For this reason, the taxonomy of *if* phrases given in Athanasiadou & Dirven (1995, 1997) will be employed.

As for what *kind* of *if* conditionals are frequently used, Hwang found the structure 'if + pres., pres' to be most commonly used (19.2% of spoken conditionals, 16.5% of written), closely followed by 'if + pres., will/going to' (10.9% spoken, 12.5% written). Elsewhere the latter construction, often known as future or predictive conditionals, has been found to be most common (Comrie, 1986). We now go into these kinds in more detail.

Categorising conditionals

Athanasiadou and Dirven's (1995) analysis of 300 instances of *if* clauses from the COBUILD corpus yields a category classification into, principally, event conditionals and marginal conditionals. This split becomes one between event conditionals and pragmatic conditionals in their 1997 analysis. For reasons described later on, however, we should not just rely on the 1997 study categorization. We focus initially on the largest category as reported in the 1995 study, event conditionals, which itself clefts into course-of-event and hypothetical conditionals. These are both considered event-based conditionals and differentiated from marginal conditionals, which divide further into logical and conversational *if* clauses. By far the largest group is the course-of-event sub-type, covering 44.7%, and followed by hypothetical conditionals at 36.9%. Examples of these two types are, respectively:

> If there is a drought at this time, as happens so often in central Australia, the fertilised egg in the uterus still remains dormant.
>
> If there is no water in your radiator, your engine will overheat immediately.

In a later analysis (Athanasiadou & Dirven, 1997) which was extended over five corpora, the authors found hypothetical conditionals to be the largest group, at 42%, and course-of-event conditionals slipped to 30%. The latter had an

marked conditional structures!

especially low presence in the Leuven Drama corpus of modern British plays (9%), which perhaps skewed the outcome; plays are a very specific literature genre. By contrast, the COBUILD corpus represents a large variety of English registers. At any rate, the significance of the precise percentages should not be overestimated; suffice to say that both course-of-event and hypothetical conditionals represent sizeable categories of usage. As will become evident, these categories represent the vast majority of conditionals used in reasoning tasks.

The other change from the 1995 to the 1997 study is that the category 'marginal conditionals' from the earlier analysis is replaced by that of 'pragmatic conditionals'. Pragmatic conditionals are characterised by their discourse-interactional function, as in,

> If you are thirsty, there's beer in the fridge.

This subtype entails its consequent, and is given the generic form "If there is a need x, let me give you information y, so that you can arrive at solution z" by Athanasiadou and Dirven. Haegeman (2003) characterises pragmatic conditionals as 'discourse-structuring'. The broad distinction between event conditionals and pragmatic conditionals is one which has been consistently upheld in the literature on conditionals even since before Athanasiadou and Dirven's study – see for example Haegeman (2003) for an overview of authors who have proposed such a distinction, albeit with different labels.[25]. Pragmatic conditionals are not central to our current interests because they are not the type typically used in reasoning tasks. However, for just this reason, they provide a key means to test experimental predictions generated by our explanation of reasoning behaviour with event conditionals. They are discussed further when suggestions for further experimental work are made.

How are course-of-event conditionals (henceforth CECs) and hypothetical conditionals (HCs) characterised? There are three main differences between them.

1. **Recurrence.** CECs refer to "generally or occasionally recurring events", while HCs mostly refer to a singly-occurring event. This is highlighted by the fact that "in CECs we can always substitute *if* by means of the temporal conjunction *whenever*" (1995, p. 617).

2. **Seriality.** CECs can refer to simultaneous or consecutive events, whereas the events of an HC must be consecutive.

3. **Immediacy.** CECs refer to real time, while HCs refer to "hypothetical time": events situated in the future, combined present and future, or in the imagined past or present (pp. 612–613), as is illustrated in the following two formulations (adapted from Athanasiadou and Dirven, 1995, p. 628) :

[25]For instance, Sweetser (1984) differentiates between content and epistemic conditionals; see also Bonnefon & Hilton (2005) who make use of this terminology

> If there is no water in the radiator your engine overheats immediately.
>
> If there is no water in the radiator your engine will overheat immediately.

These differences come together in a difference in the level of commitment to the realisation of the events mentioned in the antecedent and consequent, only CECs being committed to both of these actually occurring. Apart from tense differences, there are other means to indicate the speaker's "epistemic distance" to the events described in the conditional. Indeed, the authors follow Langacker (1991) in assuming that tenses "do not in the first place denote time but the more abstract distinction 'immediate' vs. 'non-immediate' " (pp. 616–617). So for instance, modal auxillaries such as *will*, *can* and *may* can also be used to create epistemic distance from the described events. CECs are distinct from HCs in that past tense cannot be used to create distancing effects there but must be interpreted as past time (such as in telling a story), while the modal auxillaries can occur in both types of conditionals. In CECs especially those evoking a near reality such as *going to* or *can* are in evidence, for example in the following:

> If there are distance problems, when engaged in conversation, then there are clearly *going to* be even bigger difficulties when people must work privately in a shared space. (p. 620, emphasis mine)

Also adverbs such as *normally*, *always* and even *sometimes* are used in CECs, as a means to simultaneously establish "an effect of generality and reality" (p. 619).

Within CECs, Athanasiadou and Dirven distinguish three further sub-types. These are descriptive, inferencing and instructional CECs, illustrated respectively in the following examples (p. 616):

> But if there has been rain and there is good pasture, then the egg now restarts its development.
>
> He looked at his watch: if the soldier was coming, it was nearly time.
>
> It is wise to call the doctor in all cases of sore throat, especially if there is a fever of 101°.

In descriptive conditionals both the antecedent and consequent have been observed in reality, while in inferencing conditionals although both events are 'real', the second event of the conditional has been inferred from the occurrence of the first event, rather than observed directly. In instructional CECs the consequence has imperative force, forming an instruction of what is to be done in case of the situation described in the antecedent. Presumably such CECs are to be distinguished from pragmatic counterparts by the recurring nature of the described situation, indicated by the phrasing "in all cases" . This last example does however serve to highlight that Athanasiadou and Dirven's classification depends also on contextual factors, not simply on the form or content of a given conditional, and as such sometimes delivers seemingly overlapping categories.

Conditional constructions in Xhosa

In an ideal world the classification of *if* clauses used here would be based on typological data. This is unfortunately not possible, but most relevant for the current argument is the comparison with Xhosa, the language in which the experiment was conducted. This fortunately is feasible.

All conditional sentences in the task materials were translated with the *xa* construction. According to Mncube's 'Xhosa manual' (n.d.[26]) *xa* is used for the simple conditional expressing a present or future oriented relation. The consequent can take several forms in this construction. Here the antecedent is introduced by *xa* and followed by an indicative consequent, as in the following:

(17) *Xa umntu etshayela kuqhuma uthuli.*
when a person SING.SUBJ. sweep there rise dust
When one sweeps, dust rises.

The *xa* construction can also be used for a conditional expressing doubt or 'inexpectancy of fulfilment in present or future time', viz:

(18) *Ndingambetha, xa ungandinika ikhulu leeponti..*
SING.SUBJ.MOD.OBJ.beat if SING.SUBJ(2nd pers.).MOD.OBJ+give one hundred pounds
I would beat him if you would give me a hundred pounds

Mncube describes a third construction for expressing conditionality involving doubt or inexpectancy of fulfilment: the consequent is introduced by the optative *nge*, in participial mood, followed by the antecedent in the indicative mood introduced by *ukuba*:

(19) *Nge ehlala ukuba ndiya funa.*
SING.SUBJ.MOD.stay if SING.SUBJj(1st pers.).want
He would stay if I wanted him to.

The *nge* conjunctive can also be used with *ukuba*, as above, to express counterfactuality. According to Mncube the *ukuba* construction is usually used to express a counterfactual, what he calls an 'unfulfilled past condition'. The antecedent is introduced by *okokuba* or *ukuba* and generally followed by the participial mood.

A more recent source (Pinnock, 1994) also gives *xa* as the equivalent to English *if*, *when*, *whenever*, and offers *noxa* as a translation of 'even if' and *nangona* for 'although'.

[26]The book must have been published between 1931 and 1961. This can be deduced because the 'settlement of the orthography question' in 1931, and the work of a late Xhosa scholar on grammars and dictionaries in the new orthography are mentioned in the foreword, so the book is definitely published some years after 1931. Also, there are sample sentences referring to the pound. The rand was introduced as the currency of South Africa in 1961, so the book was published before then.

So the two main conditional operators in Xhosa are *xa* and *ukuba*. As we've seen in the analysis of the transcripts from the current study, the majority of conditionals are interpreted as habitual or CEC conditionals, which can in English also be expressed with *when* or *whenever*, and in Xhosa with the *xa* construction. The *xa* construction was indeed the predominant translation used in the current study.

2.3.2 Why some conditionals make natural premises

The majority of the conditionals used in reasoning studies (and the few used in Scribner's 1977 and Cole et al's 1971 materials) can straightforwardly be classified as course-of-event conditionals, recognised by the fact that *if* can unproblematically be substituted with *when* or *whenever*. It is also most plausible to interpret them as describing recurring events, consecutively occurring and probably both observed, hence falling into the sub-type descriptive CECs. For example, I used:

> If Ntombi wants to see her boyfriend then she goes to East London.

And from Scribner, 1977:

> If Sumo or Saki drinks palm wine, the Town Chief gets vexed.

A second type used in reasoning studies uses the future tense in the consequent clause, and can be characterised as inferential CECs:

> If Thembi works hard then she will pass her exams.

This last example might also be adequately characterised as descriptive, or even hypothetical. Luckily this does not matter greatly for the current analysis. Again, it does however illustrate that the categorization into descriptive CEC or hypothetical conditional can depend on contextual factors. This applies equally to Byrne's original (1989) suppression effect materials, such as:

> If it is raining then she will get wet.
>
> If she has an essay to finish then she will study late in the library.

The first example given can be doubly classified as hypothetical and inferential CEC, since the latter express inferences about time-bound relationships and often contain modal auxillaries. The second example from Byrne can, despite the future tense, be read as a descriptive CEC, as above, as it can be seen to instantiate a stable relationship rather than express a singular connection. In fact, other studies based on this material (Lechler, 2001, and the current one) used a present tense formulation of the conditional premise in some of the conditions, viz:

> If Rosa meets her friend Liz then she goes to the cinema.

and

> If Maria studies hard then she gets high grades in the exams.

In summary, the materials used in conditional reasoning tasks are classifiable as expressing either descriptive or inferential course-of-event conditional relationships, and in some cases they may also be viewed as hypothetical. Recall that CECs are recognised by the fact that they describe the "normal course of events", simultaneous or consecutive events which are both assumed to occur on a recurring basis. We now turn to semantic analyses of such conditionals.

The common characteristic of both course-of-event and hypothetical conditionals lies in the stability of relation between the antecedent and consequent they seem to convey, at least in unmarked usage. The presumption of a relevant relation is illustrated by the unacceptability of conditional statements in which the antecedent is irrelevant to the consequent, such as 'If 2+2=5, then Utrecht is the capital of the Netherlands'.[27] Otherwise expressed: it is difficult, if not impossible, to read event conditionals (including CECs and HCs) 'contingently'. Rather, event conditionals seem to tell us something about stable relationships playing out in the world. Marked usage plays with exactly this setting by connecting events which have nothing to with each other, such as in 'If <unlikely event>, then pigs can fly' or ' <unlikely event> when hell freezes over'.

The ummarked event conditional is to be contrasted with the other major category of conditionals, premise conditionals, in which a stable connection between antecedent and consequent is not necessarily evident, such as in

> If you're hungry, there's pizza in the fridge.

or, to echo an earlier example:

> If you're looking for a silver coin, check my pockets.

In these examples there relation between the antecedent and the consequent is extremely context-dependent: there's no intrinsic connection between one's hunger and the contents of the fridge (or a type of coin and the inside of my pocket). Rather, in a specific context this connection could be made for the specific purpose of alleviating one's hunger (or need for a silver coin!). I would argue that this is a marked use of conditionals and that the unmarked usage is in course-of-event conditionals where some kind of intrinsic connection, often but not necessarily causality-based[28], between antecedent and consequent. This is emphatically not the unmarked usage of quantified phrasing. To illustrate, consider the contrast in

[27]Example from Veltman (1986). Such statements are nevertheless true if the conditional is attributed the semantics of the material implication – a phenomenon often labelled as a 'paradox of material implication'. More discussion on this can be found in Chapter 4.

[28]This allows for the possibility of a third event as common cause; also convention may explain the connection, such as in

> If it's Tuesday, I go swimming.

.

All the coins in my pocket are silver
vs
If a coin is in my pocket, it is silver.

In the second phrasing, the conditional, we feel the need for a story about *why* it would be so – it cannot be a mere accidental fact. By contrast, the quantified version is most naturally read as a happy accident. Here also Fillenbaum's (1978) data – as reported in Chapter 3 – indicates the seemingly inexorable tendency to attribute a sensible intrinsic relation to the events conveyed in event conditionals.

Haan (2007) makes the suggestion that the difference between quantified and conditional generalisations may originate in what they generalise over – namely, typically entities and events, respectively. Perhaps events are inherently more causally related to each other than predicates? The following example, in which only events are quantified over, argues against this.

Consider first the sentence

Every time this button gets pushed, the alarm goes off.

We can imagine the first sentence being uttered *to* the alarm system technician, perhaps from a disgruntled user. The correlation between the button-push and the alarm going off is puzzling, unintended, accidental, even indicative of a mistake. By contrast, the conditional variation,

If this button gets pushed, the alarm goes off.

lends itself much more to utterance *from* the alarm system technician, as explanation of how the system works. The alarm in this case is *supposed* to go off. Very clearly here, it is the conditional phrasing which brings out the causal, or law-like, nature of the relation, and not the mere fact that the relation described is between two events.

A note on formalisms: dynamic semantics for the conditional

An issue we've come across in evaluating subjects' reasoning is the mismatch between the argument structure viewed from a classical logical perspective, and the argument as a discourse. One of the limits of classical semantics is that it works on a sentence-level translation of terms from natural language into a formal language. This means, for example, that classical logical translations will not distinguish between

(20) Bill fell and John hit him.

(21) John hit him and Bill fell.

But this limitation has been overcome by recent developments such as dynamic semantics which enable us to deal with sequences of sentences (Veltman, 1986). Within a dynamic semantics framework it is possible to elucidate formally the

reasoning steps that happen against an epistemic background, a 'hearers cognitive state', as opposed to a background of no information.

In this brief note I give some indications how this could be applied to modelling the data from reasoning tasks from a reasoner's perspective. In the above a speaker's reasons for the choice of one or other conditional construction, such as course-of-event, or hypothetical, was discussed, but once a construction has been chosen and needs to be interpreted the focus shifts to that of a hearer's perspective. We now turn our attention away from the speaker to the hearer, as it were.

Much work has been done on understanding the semantics of conditional constructions from a hearer's perspective; attributing the truth conditions of the material implication to indicative conditionals in natural language (as has commonly been done in psychology of reasoning) leads to the apparent truth of highly unintuitive if not downright unacceptable constructions such as 'If 2+2=5, Utrecht is the capital of the Netherlands'. In such cases, the antecedent seems irrelevant to the consequent. One approach which is aimed at overcoming this paradox is that of dynamic semantics, where the meaning of a sentence is analysed to be an operation on the hearer's cognitive states (i.e. beliefs) – in keeping with the slogan: 'meaning is change in information'[29]. On this approach, a conditional is interpreted as a step-wise 'test' on the current cognitive state, rather than the incorporation of new information into it. The 'test' works as follows: it returns the current state, if the adding the antecedent p to the current state generates a new state which supports the consequent q, and the absurd state otherwise. A state supports a sentence if adding that sentence to the state does not change it – that is to say, if the sentence could already be deduced from that stock of beliefs. This interpretation of the conditional captures the intuition behind the Ramsey test for conditionals, which is the hypothetical adding of the antecedent p to one's stock of beliefs and subsequent check for the truth of the consequent q as a means to evaluate the conditional as a whole. This account predicts a difference between course of event and hypothetical conditionals as the former involve actual addition of the antecedent and consequent to the state, according to Athanasiadou and Dirven's analysis. So the reasoning process is augmented by belief revision in this case.

Pragmatic constraints on quantity and quality of utterances ensure that the normal context in which to assert an indicative conditional 'if p then q' is one in which both p and q are uncertain – i.e. none of p, q, or their negations $\neg p$, $\neg q$ has previously been asserted and taken up into the listeners' stock of beliefs. This is because if $\neg q$ had already been asserted (and accepted), it would be impossible to have a successful 'test' of the conditional – the second step would generate

[29]Or more fully, "You know the meaning of a sentence if you know the change it brings about in the cognitive state of anyone who wants to incorporate the information conveyed by it" (Veltman, 1994).

the absurd state. Likewise for a prior assertion of $\neg p$ – here the first step would already generate the absurd state. These two cases nicely explain the workings of constructions such as 'If <unlikely event>, then pigs can fly' or ' <unlikely event> when hell freezes over' – they rest on forcing accommodation (of the presupposition) of the negation of the unlikely event, if we take 'pigs will fly' or 'hell freezing over' to represent the absurd state. This also explains why the above example is an abnormal case, because it is uttered in a context where the antecedent is known to be false. Further, if either p or q would already be known, then asserting the conditional would be non-informative because, in both cases, q would already be deducible from the initial belief state.

Note that this analysis is intended for a series of assertions from a single source; a dialogue would introduce other constraints again, involving the intentions of each interlocutor regarding moves in the foregoing discourse as well as assumptions about the hearer's belief states.

The dynamic account of conditionals can be related to the abnormality-sensitive reading of the conditional given in the previous chapter[30], by considering the relevant abnormalities, or necessary preconditions, as salient elements of a hearer's belief state which would impact on the support for q in the updated-with-p state – thereby determining whether or not it is felicitous to utter the conditional in the context, according to the hearer.

As we did with quantified premises, let us consider the position of the reasoning subject, this time in a conditional reasoning task. First the subject is presented with a conditional premise, say "If Thembi wants to see her boyfriend then she goes to East London." This describes two events, which can be understood as simultaneous or consecutive, but also generally recurrent. The speaker indicates they are committed to the realisation of these events by the use of the simple present tense. Then the second premise is presented: "Thembi does want to see her boyfriend", followed by the question, "Does she go to East London?"

Recall that course-of-event conditionals suggest that the speaker is committed to both of the events described actually occurring. To give the answer 'Yes, Thembi goes to East London' is for the hearer to indicate acceptance of the conditional. There are two possible reasons that this is a sensible response, perhaps mutually strengthening. First, it is general conversational convention to indicate comprehension of the speaker's assertions (often accomplished by non-verbal indications such as nodding or murmuring 'mm-hmm'). Secondly, the dynamic analysis of the conditional suggests that a question after the consequent is a question after the success of the test expressed by the conditional. To answer 'yes' is to answer 'pass' to the conditional test. The hearer agrees, as it were, that if one adds p to one's set of beliefs, one can infer q.[31]

[30]This was the reading for a sentence of the form 'If A then B' as 'If A, *and nothing abnormal is the case*, then B'.

[31]The given account is very compatible with the broadly-accepted discourse-functional thesis that conditionals provide topics, as first proposed in Haiman (1978), concisely expressed, the

However, once the hearer has indicated the conditional test has 'passed', we might expect them to volunteer reasons why this is possible, that is, what other conditions have to be met for the conditionality to be captured in the premise as given. This is where the abnormality clauses become explicit, and as mentioned in the previous chapter, this was a commonly occurring phenomenon in my data – seen in excerpts such as:

Nomhle, group 2:

E: If Ntombi wants to see her friend then she goes to East London. And she does want to see her friend. Will she go to East London?
S: Yes.
E: Why will she go to East London?
S: Is it because she wants to see her friend?
E: Yes, she wants to go.
S: Maybe she can talk to her parent and tell her that she wants to go and see her friend. Maybe Ntombi is interested to go and to see her friend but now she doesn't have enough money. So maybe her parent will give her money to go to East London to see her friend.

On this account, the *modus ponens* condition in these reasoning tasks is more a test of communication, or acceptance than anything else. This would not necessarily be the case when other conditional-based inference schemas are tested, nor when types of conditionals are used – for instance premise conditionals. Some suggestions for investigating other types of conditionals empirically are given below.

As for other schemas, if acceptance of a CEC conditional means commitment to both the antecedent and consequent occurring, then *modus tollens*, i.e. reasoning from the negation of the consequent, becomes problematic.

Nothabile, group 1:

After premises 'If Thembi has to fetch water then she goes down to the river. If she washes her clothes then she goes down to the river.'
E: And she doesn't go down to the river today, does that mean she will fetch water or wash clothes or not?
S: She can't wash her clothes if there's no water, and *she can't stay at home if there's no water*, so she must go to the river if she has to fetch water or wash clothes.

Now this excerpt is difficult to interpret, but it is at least clear that the subject resists drawing the conclusion, "Thembi doesn't fetch water or wash clothes" on the basis of the fact that she doesn't go to the river. Although the subject short of asserting the consequent and antecedent, as a CEC classification predicts, her words do seem to imply she feels unable to reject either of them. She emphasises

thesis is that: "A conditional clause is (perhaps only hypothetically) a part of the knowledge shared by the speaker and his listener. As such, it constitutes the *framework* which has been selected for the following discourse" (1978, p. 581).

again the connection between the situations, as if to say, "You can't have the one without the other" and resists taking on board what is intended as a premise, i.e. the negation of the consequent. It is as if she feels caught in a contradiction. This is certainly consistent with a CEC interpretation for the conditional in this case. With a similar argument we can predict affirmation of the consequent to proceed automatically as *modus ponens* does; and denial of the antecedent to be difficult for subjects with CEC premises.

This is of course a far too brief introduction to the possible application of a dynamic semantic approach to modelling the conditional-based inference; the reason for including it is to indicate that there are available formalisms which are equipped to deal with some of the phenomena associated with a discourse-based view of reasoning. Unfortunately a fuller analysis is beyond the current scope.

2.4 Summary, conclusions and outlook

In the first half of the chapter I explored the use of the universal quantifier *all* and its Xhosa equivalent *-onke* in spontaneous speech and was able to categorise them in four main categories. Stock and emphatic usage was not further analysed; the focus of the analysis was on the differences between anaphoric or deictic usage and so-called 'forward quantification'. It was argued that this latter usage comes closest to the usage of *all*-sentences as premises, but that it is atypical for spoken usage, and possibly even derived from written contexts.

Reasons for these findings were sought in a semantic analysis of *all*. I argued for the finding that *all* sentences are chiefly used with a pre-given context set, to express contingent generalisations. I proposed that reasoning with *all* sentences is difficult for the unschooled subject because the context set is not specified; when this is absent the subject cannot make inference about individuals, because their relationship to the domain is not known. An alternative law-like reading is less context-set dependent, but this is specialized use of *all*, and moreover results in an uninformative answer. We can summarise by saying that *all* is used primarily to pick out a pre-given domain and that this makes *all* sentence unnatural first premises.

In the second half of the chapter I investigated features of conditional sentences, focussing on the type used in reasoning tasks so far, which can be categorised in Athanasiadou and Dirven's (1995, 1997) classification including course-of-event conditionals, hypothetical conditionals, and premise conditionals. It was found that conditional premises used in for example the suppression effect task closely resemble common course-of-event conditionals. Relevant features of this type of conditional are the recurrent nature of the events described and the commitment to their occurrence implicit in use of the conditional. I argued that this makes this type of conditional sentences natural bases for drawing inferences and can at least partially explain the greater inclination of subjects to assent to

the *modus ponens* conclusion in the conditional reasoning tasks studied. Brief suggestions were made regarding available appropriate formalisms for these and other response patterns observed in the data.

This chapter is a first attempt to related reasoning behaviour to everyday language usage and remains very much a first attempt. Corpora data was matched with a range of more theoretical semantic work to explain the differences between kinds of premises, specifically those between quantified and conditional formulations. Although the findings in the corpus study were preliminary and there remains much to address regarding the match between data and formal work, it seems reasonable to conclude that the approach taken here is a fruitful one and one which deserves further study.

2.4.1 Experimental predictions and evaluations

If the account given here about difficulties with *all* premises is correct, then we should expect subjects to reason more easily with the following variations of quantified premises:

- *all* sentences for which the context set is specified. As we've seen, Haan's work (2007) bears this prediction out.
- *every* sentences, since these do not lend themselves to a generic reading, but rather to a purely contingent one. Here specification of the domain would be necessary; otherwise the task would become purely a matter of prediction, based on similarity judgements between the particular instance named and those known to belong to the domain.
- conditional phrasings of generalisations, such as 'If a bear lives in the north, it is white', to encourage law-like readings.

Regarding conditional premises, event conditionals are to be contrasted with pragmatic conditionals in terms of their role in a discourse. Specifically, pragmatic conditionals play a discourse-interactional role – a very different discourse function from event conditionals. Hence we might expect subjects to exhibit different patterns again when reasoning with such conditionals. In particular, as Athanasiadou and Dirven (1997) point out, in premise conditionals the truth of the consequent is always presupposed. Given this, we can predict that the *affirmation of the consequent* and *denial of the antecedent* inferences will be much less common. See if you find yourself committing a fallacy with either of the following:

> If you're hungry, there's pizza in the fridge.
> There's pizza in the fridge.
> What follows?

and

> If you're hungry, there's pizza in the fridge.
> You're not hungry.
> What follows?

Also worth further exploration is the depth of the posited distinction between course-of-event and hypothetical conditionals. In several examples offered by Athanasiadou and Dirven it was clear that the distinction between these categories is only to be determined by contextual indicators. However, if experimental differences could be found in more clear-cut examples of one or other type, this would lend credence to the suggested categorical distinction between them. For instance, it could be investigated whether subjects do commit to both the antecedent and consequent occurring in CEC formulations, but not in hypothetical conditionals. Suitable materials would differ only in the tense phrasing, for instance,

> If it is raining then she will get wet.
> If it rains then she gets wet.

In my data, as mentioned above, the fact that some subjects had resisted the *modus tollens* inference with conditionals expressed in the simple present ("If Thembi has to fetch water then she goes to the river") suggests the distinction is a semantic one. If however, this turns out not to be the case, Athanasiadou and Dirven's categorization should be simplified accordingly.

Other avenues

The idea that illiterate reasoning behaviour can only be judged within the communicational context of the task forms the backbone of the current chapter. Indeed, the 'meta' goal was to emphasise of the value of usage-based studies of language in helping us to understand inferential behaviour.

But the 'flesh' given to this idea here is a kind of experiment, there are other possible explanations which might fit the data just as well. In this chapter I have pursued some aspects of the materials used in reasoning tasks related to their role in discourses. There are aspects which have been left unexplored so far but which could well yield just as rich an explanation of behaviour as that given here. For example, I have no more than touched on the positioning of the protagonist in the premises, and the ways that this could contribute to the ease or difficulty of drawing a conclusion from conditional or quantified premises. A preliminary consideration of the varieties of relation the protagonist can have towards a generalisation suggests this would be worth further analysis: they may be examplar, exception or counterexample, all of which would impact on the conclusions one can draw about them and the relation of these to the generalisation.

Neither have I explored the widely-accepted idea that narratives are a 'basic' type of discourse (Bruner, 1991). The relative strength of fit of the various premise sets to a narrative structure might also be a fruitful means to explain variation in responses.

Finally, a feature of these tasks which has persistently cropped up both in the category analysis of the previous chapter and in the information-structure sections of the current chapter is their peculiar epistemic structure. Usually, when one asks a question, it is because one doesn't know the answer. Rhetorical questions are no doubt a universal phenomenon, but whether or not the subject understands that the questions asked in the context of a reasoning task are of a sort with these is unclear. One reason to suppose that the subject did not understand the question as rhetorical is that such questions don't require answers, whereas in the experimental setting an answer is clearly expected from them. In many cases, the subject replied with a question, asking for confirmation from the experimenter for a 'correct' answer. In other cases it was clear that unschooled subjects simply couldn't understand the question they were supposed to be answering – they are looking for the 'question under discussion' as it were. The idea that this 'QUD' is determined by the 'genre' of the task is touched on in the next chapter, but further exploration of the contrast between this type of question-answer situation and that occurring in spontaneous dialogue would certainly be worthwhile.

Chapter 3

The farmer *vs.* the undergrad

3.1 Introduction

The past two chapters reported the responses in reasoning tasks, of subjects with varying educational levels, and endeavoured to show some of the ways in which the response profile across subjects is a function of semantic interaction with the material of the task. It was argued that schooled and unschooled subjects are concerned with the same semantic factors; yet in the case of syllogistic reasoning tasks, schooled subjects might have learned to suspend everyday interaction strategies with the linguistic materials and to employ a contrived interpretation of the premises necessary to solve the task as intended, whereas unschooled subjects do not as readily reach the required, but contrived, interpretation. We concluded that to subjects' varying willingness or ability to ignore such semantic factors explains much of what has been reported as differences in reasoning abilities.

The goal of the current chapter is to relate these findings to the broader theoretical context of 'cognitive consequences of literacy'. In particular, what are the operative aspects of 'literacy' which can explain the above results? This incorporates sub-questions such as: What are the theoretical proposals already advanced in the anthropological and historical literature on the subject, and are they of use here? How do the findings relate to other data from psychological studies with illiterate subjects, and can similar explanations be given for these results? On the explanatory side, what is meant with the terms 'literacy' and 'schooling', and are they employed consistently? It goes without saying that coverage is by necessity highly selective; none of the topics mentioned here can be done justice, while they certainly bear mentioning.

Qualifications aside, the plan of the chapter is as follows. Firstly, I briefly mention existing theoretical approaches, and delimit how the current findings can be related to them. Secondly, I report on some other cognitive tasks, to see whether results were similar to those found in reasoning, and whether similar, seemingly deflationary, semantics-based explanations of illiterate performance can

be applied there. This helps us to assess the sources of difference in literate and illiterate performances more generally. Then I reflect on the complex of relation and dissociation, both empirical and conceptual, between literacy and schooling. Finally, the work of David Olson is singled out and evaluated for its adequacy in providing a theoretical embedding of the current findings, since Olson is the most prominent theorist in this area to address experimental findings such as those presented in the previous chapters. On the back of this, some new experimental approaches are proposed, to help further understanding of the impact of literacy on cognitive performance.

Terminological trickery

The reader might have noticed the equivocation of schooling and literacy in the last chapters. Of course, schooling and literacy are *not* synonymous. Nor are they unrelated. But the proper locus of investigation (i.e. schooling or literacy) is a discussion in itself. Often, the way the term 'literacy' is used makes it akin to 'Western-style education', while school activities are often text-based but also carry specific social and cultural values. I will, for the time being, be rather casual about using 'literacy' and 'schooling' interchangeably. In section 3.4, the relation between schooling and literacy is addressed explicitly with the goal of differentiating retrospectively, and as far as it is possible, the appropriate use of the terms.

3.2 Background theory

Cultural anthropological debates on rationality

Since the topics of literacy, culture and cognitive ability have been the topic of study in so many disciplines (including history, anthropology, educational science, cognitive psychology, sociolinguistics) it is helpful to first delineate which debates are relevant to the current investigation. For starters, there is, predating Luria's pioneering empirical study on the impact of literacy, a long-running debate amongst anthropologists about differences in cognitive ability between different groups. The debate concerns the possibility and consequences of a cultural relativist account of rationality and has, roughly, been about whether different cultures have different rationalities, or whether cultural differences should be understood as stemming from different conceptual frameworks, often called "world views" (Lukes, 2000) to which the same universal standards of rationality are applied (see for instance Lévy-Bruhl, 1926/1910, Winch 1964, Wilson 1970, Lukes, 2000).

The debate is relevant for our current interests as far as the traditional *vs* modern distinction coincides with the literate *vs* illiterate boundary – which is

rather a lot – but also inasmuch as it addresses the same kind of empirical phenomena – which it largely does not. The debate in cultural anthropology is about differences in whole-sale belief systems, at the level of societies rather than individuals, and which certainly direct the behaviour of individuals but only within contexts in which cultural beliefs play a part. Although I will later touch on some issues addressing the possible social bases of rationality/logicality, I further judge the literature in this area to be of limited relevance for the current study. Interestingly, though, fuelling much of the discussion is a type of belief held by members of traditional societies which seemingly confound 'modern' semantic analysis. Statements such as "twins are birds", "a cucumber is an ox" (first reported in Evans-Pritchard, 1956, of the Nuer tribe of the Sudan), "corn is deer" (held true by Huichol of Mexico, reported in Myerhoff, 1974), and "we are red parrots" (asserted by the members of the Bororo tribe, quoted by Vygotsky & Luria, `[date]`, – p. 70 – without a reference) are emphatically *not* understood metaphorically, according to the anthropologists' enquiries, but are taken simply to be true by the tribe under study. But they are patently false! – or is that just our blinkered Western world-view? These examples have served as fuel in the debate on the possibility of a cultural relativist account of rationality in cognitive anthropology. A relativist account has been expressed in various forms, the most extreme being the view that there is a 'primitive' mentality, different to ours in being driven by mystical, as opposed to logical principles.[1] We are inclined to accept Evans-Pritchard's evaluation that Western scholars have tended to misinterpret statements like the ones mentioned above because "it was not recognized that they are made in relation to a third term not mentioned in them but understood ... A cucumber is equivalent to an ox in respect to a God who accepts in place of an ox" (Evans-Pritchard, 1956, pp. 147-8), that is, as a sacrifice. Strangely enough the work which is often quoted as sparking the debate is

[1]Lévy-Bruhl was one of the first theorists to treat as central the influence of historical and cultural factors on thinking. Early in his career he proposed that the psychology of primitive societies is fundamentally different from that of modern, 'civilized' societies (Lévy-Bruhl, 1926). The key contrast for Lévy-Bruhl was in the *type* of thought, as being magical versus logical. According to this account, primitive societies are 'prelogical'; their thinking is magical (or mystical). Magical thought is primarily characterised as being insensitive to contradictions; the basic feature is the "law of participation", in which the same object may participate in several different forms of being. In contrast, the logical nature of the thinking of 'civilized' societies means it has as its basic feature the law of non-contradiction. The transition from magical to logical thinking is a process of development: the modern mode of logical thinking is a more advanced one than that of the magical mode of thinking found in primitive societies. Lévy-Bruhl intended his comparison to refer to belief systems at the level of societies rather than individuals, and later formulated a much more nuanced position. Unfortunately for him, his name is still associated with the rather simplistic mystical-logical dichotomy just sketched. Theorists such as Levi-Strauss stressed the commonalities between primitive and modern thought. Levi-Strauss (1966) concluded that primitive/magical and modern/scientific thought require the same sort of mental operations but function on a different level, primitive thought being rooted more in the concrete.

Evans-Pritchard's 1937 study of the witchcraft and magic of the Azande tribe of central Africa (and Winch's 1964 paper responding to it). When considered in the normal context of utterance, i.e., a ritualistic one (Gellner, 1988), such statements are directly comparable to singularly Western statements such as "this bread is my body", the simple, non-metaphorical acceptance of which is a central tenet of the Roman Catholic faith! This analysis only strengthens the case for paying careful attention to contextually driven aspects of semantics when assessing linguistic phenomena across cultures.

Historical theories on literacy

More directly relevant is the theoretical discussion on the differences between societies with or without literacy. The nature and impact of written language has been a subject of contemplation almost as long as writing has been around, and, just as for much western scholarship, we find ideas germinal to later theories on these themes in Plato and Aristotle. One such idea is the supposed supremacy of the written word when it comes to matters of reason. Scribner and Cole point out that in Plato's *Republic* dramatic oral poetry is considered not to appeal to reason but to emotion, a precursor to the theme of "pitting certain oral modes of discourse against reason," ... "a theme that has never disappeared from critical studies on the psychology of reasoning" (1981, p. 6).

Contemporary exemplification of this thinking is found in Goody and Watt's influential paper *The consequences of literacy* (1963). Goody and Watt defended the thesis that the invention of alphabetic script was a necessary precursor to the emergence of such key western social institutions and practices as democracy and logic. This comes about because of the permanence of script, which, by fixing content, turns myth into history. The possibility of historical enquiry engenders scepticism. This sceptical attitude leads to the kind of analysis at the heart of the intellectual tradition in modern literate societies today. Regarding logic, their view is succinctly formulated thus: "the kinds of analysis involved in the syllogism, and in the other forms of logical procedure, are clearly dependent on writing, indeed to a form of writing sufficiently simple and cursive to make possible widespread and habitual recourse both to the recording of verbal statements and then to the dissecting of them" (Goody & Watt, pp 344 – 5). The second part of this quote hints at the importance of the alphabet for Goody and Watt, in contrast to for instance a syllabic script, as a catalyst for development – a view also seen often in early theories of writing. For instance, a dramatic expression of this is found in Rousseau's *Essay on the origin of language* (1754–91/1966, cited by Olson, 1994):[2]

> These three ways of writing correspond almost exactly to three dif-

[2]The French and Dutch terms for 'illiterate': respectively 'analphabète' and 'analfabeet' also illustrate this script chauvinism.

> ferent stages according to which one can consider men gathered into a nation. The depicting of objects is appropriate to a savage people; signs of words and of propositions, to a barbaric people, and the alphabet to civilized peoples.

Goody and Watt's theory is an example of what Finnegan termed 'Great Divide' theories: those which posit a principled and fundamental difference between literate and oral societies (see also Ong 1982, McLuhan, 1962). Such theories ally the development from traditional to modern, from simple to complex, from primitive to civilised, to the shift from oral to literate. For Goody and Watt, the essential difference lies in the permanency of written language. For Ong, the process of writing is "completely and irremediably artificial" depending on "consciously contrived rules" whereas speech is a "natural process" making use of "speech organs" (p. 199, 1982). For McLuhan, it's the visual nature of writing which explains its role in a new kind of thinking, linear thinking, and creates 'typographic man'; he goes so far as to claim that "by the meaningless sign linked to the meaningless sound we have build the shape and meaning of Western man" (McLuhan, p. 50, 1962).

While specific notation systems surely do afford different means of expression, Goody and Watt's (1963) alphabetic hypothesis is undermined by for instance much anthropological work criticising grand general theories such as theirs (Jahandarie, 1999). For instance, anthropologists (Finnegan, 1973, Street, 1984, Akinnaso, 1992) have stressed that the diversity of illiterate or 'oral' societies is underestimated in Western-centric theories; conversely, 'literacy' is not one thing, but a variety of skills loosely grouped "under a modernist rubric" (Stock, 1990, p. 141, quoted in Jahandarie, p. 279). As such the significance of bare literacy, when divorced from the cultural milieu in which it is practised, can be overestimated. On the other hand, there are many literate traits to be found in oral societies. For these reasons anthropologists plead for specific ethnographic studies to further our understanding of the differences between literate and oral cultures.

Historical theories about the consequences of literacy are in a sense irrelevant to the current investigation, because we are interested in individuals who inhabit the same cultural-historical world, whereas historical accounts describe changes played out in historical, and cultural time. But these theories are informative about what we take the effects of literacy to be – in a general sense. The ramifications of literacy at the level of society will inevitably impact on individual literacy. This is because there are multiply different activities in which and for which the individual might use reading and writing. Moreover, these activities, and understanding of their significance, changes over time. For instance, Carruthers (1990) describes how for medieval monks reading functioned primarily as an aid to memory in the rote learning of holy texts, which were read aloud. Note that these texts are absolutely fixed. Also, there was no such thing as reading silently. Reading and memorizing were taught as a single activity and

writing was used to check one's memory, and for transcription and translation of holy texts, as opposed to being used for communication. It seems obvious that a medieval monk must have had a very different understanding of reading than, say, a modern-day school-going child, who is interacting with many kinds of texts: narratives, other- and self-generated, fictional and non-fictional; learning aids, textbooks, and, nowadays, a whole new welter of textual genres through the internet and digital communications. Clearly any differences in conceptualising would be not only a function of involvement in different reading activities, but also the social, cultural and technological milieu. As such, when discussing the effects of literacy, we should consider not only an individual's literacy-based activities but also the environment in which they are being practiced.

Thus we see that historical theories *are* relevant to theorising about the impact of literacy on the individual, since an individual's literacy is shaped by social and cultural factors. Conversely, all these theories make use of the individual as 'middle man' by which changes to society or culture come about, so understanding of the impact of literacy on the individual will enable us to better assess their credibility, and to identify where other variables are needed to explain cultural-historical differences.

Having indicated some of the kinds of theories that are out there, and delineated where they are and are not useful to understanding changes at the level of the individual, it is time to approach from the other direction, from the data. What kind of evidence do we have for the changes brought on by literacy in an individual?

3.3 Some data on the impact of literacy

There are at least three kinds of data which can help to answer the question: what are the cognitive consequences of literacy? Firstly, we can look at the differences in adult populations with and without literacy. There have been few large-scale empirical studies of this kind. This is probably due to the difficulty of finding suitable conditions, i.e., those where populations differ only in their level of literacy. Indeed, literacy is typically associated with changes in socio-economic status and participation in an industrialised and technologised world. This makes a single-factor comparison of literate and illiterate hard to achieve. We give a (very) selective evaluation of the work in this field below.

Secondly, we can look at changes in children's cognitive development associated with literacy. There is some work on this – such as Greenfield's studies reported below – and of course much research in education can be understood as describing the effects of literacy. But comparing literate with pre-literate children is problematic, because literacy is acquired in tandem with processes of general cognitive development, as well as (spoken) language learning.[3] As such, in prin-

[3]For instance, some semantic structures are acquired only when children are already learning

ciple the 'purest' way to test the effects of literacy would be with adult groups which newly have access to literacy programmes. However, this does not necessarily generalise, because literacy achieved in adulthood might well have different effects to literacy acquired before cognitive maturation. Indeed, much of what we regard as 'normal' development is probably heavily influenced by the process of a literacy-based education. This will turn out to be an important point, one to which we shall return.

Thirdly, we can look at the effects of specific writing systems on cognition, in much the same vein that comparative cognitive linguists looks for differences in conceptualising arising from cross-linguistic differences in grammatical structuring. There are some elegant studies addressing this for different writing systems, one of which we discuss below, but, again, suitable situations in which a single-factor analysis can be achieved are rare.

Quite apart from any practical difficulty in measuring the effects of literacy, there is a more fundamental conceptual difficulty of specifying what is meant by 'becoming literate' or 'acquiring literacy'. For instance, a distinction is commonly made between *primary literacy*[4], i.e. the skills of reading and writing, from *secondary literacy*, meaning longer term engagement with varieties of texts, involvement in writing activities, understanding of the conventions around texts, and more generally, extended participation in the literate world. This latter stage of literacy is plausibly responsible for much more cognitive change than a circumscribed introduction into the skill of reading and writing, not least of all because it encompasses such a diffuse range of activities. Perhaps we should rather say there's a distinction between '1.n-literacies', each of which refers to a specific interaction with texts, and all of which piggy-back on primary literacy (Street, 1984). The next step would be to associate these literacies with specific sets of conditions which define them.

For the moment we put aside these concerns to examine some of the data that have been gathered on the response of illiterate adults in tasks designed to tap cognitive ability, since this is our chief interest in the current study. Later on in the chapter we'll return to the other two sources of data and discuss how they can help us to interpret the data from illiterate adults.

3.3.1 Illiterate performance on cognitive tasks

As mentioned in previous chapters, the seminal experimental study into the effects of literacy was conducted by Alexander Luria in the 1930s Soviet Union. Luria examined performance on a battery of cognitive tasks: classification, generalisation,

to read and write (see Moore, Pure & Furrow, 1990, for evidence that children master modal expressions of speaker certainty after 5 years of age).

[4]This is not the same as functional literacy. Functional literacy is the use of primary literacy skills in highly circumscribed situations, for non-literate aims, such as filling in forms, checking schedules etc.

definition, arithmetic, imagination, self-awareness and perception tasks, among others. Many of the tasks investigated by later theorists are found in Luria's experiments and as such it is worth following the framework he laid down. On all the tasks he tested, Luria found differences between literate and illiterate subjects' responses. Here we concentrate on Luria's findings in the tasks which are most easily classified as purely 'cognitive', that is, without perceptual components, although, of course linguistic perception is also perception!

Luria's classification task results

In his classification task Luria found that the vast majority of illiterate subjects did not classify on the basis of categorization, but on the basis of what he calls 'situational thinking' (or, elsewhere, the 'graphic method of grouping'). The associated percentages were: 80% of the illiterate subjects classified solely according to 'situational thinking'; 4% solely on the basis of categorical thinking, and 16% used a mixed strategy. All 12 literate subjects classified using categorical thinking. The fact that the categorical approach is used by 20% of the illiterate subjects implies that literacy isn't a necessary precursor to the ability. It would seem literacy is at most creating a bias towards a categorical strategy for classifying. Yet Luria's conclusion is that "we failed to get these subjects to shift to a logical plane of thought" (p.64).

On being shown drawings of a hammer, a saw, a hatchet and a log, and asked to pick the odd one out, the 'categorical' response would be to choose the log, since it does not fall into the category *tools*. An example of 'situational thinking' would be to choose one of the tools instead of the log. When reminded that "a hammer, a saw and a hatchet are all tools", Luria's subjects replied with answers such as "Yes, but even if we have tools, we still need wood – otherwise, we can't build anything" (Luria, 1976, p. 56). This was a commonly-occurring response and is labelled as 'situational thinking' because it involves the subject introducing "a concrete situation in which the objects could function together" (*ibid*, p. 64).

A second excerpt provides an illustration of a mixed strategy (*ibid*, p. 66):

> E: ...Is an ax like a sickle in some way – is it the same type of thing?
> S: Yes, they're both tools.
> E: What if I were to put some barley here?
> S: No, that wouldn't be right. Barley is food, it's not an asbob [tool].
> E: Would the group be alike if I put the barley here?
> S: It would because you can chop with the ax, reap with the sickle, and eat the barley.
> E: Suppose I would put the saw here?
> S: Yes, that would fit. A saw is also a tool.

For Luria this excerpt showed that (p. 67):

> even when a subject appeared to have learned the principle of abstract classification, his grasp remained far from firm. As he proceeded to think through a problem, he would revert to his habit of constructing imaginary situations in which objects functioned together.

Let us examine this extract more closely. The subject first classifies according to category *tool*, and reiterates this by rejecting the barley. But the experimenter persists, repeating the question about the barley. Now the subject offers another basis for classification – roughly, 'useful things for man' – which does include the barley. It should be observed here that as long as we consider *useful things* a category, this is just as much a 'categorical' classification as the previous one.

The above extract brings to the fore the fact that there are multiple possibilities for classifying according to category, something which Luria does not acknowledge. For instance, on being presented with the series 'bayonet-rifle-sword-knife' a subject responds "there's nothing you can leave out here!" – which is true if one uses the category *weapons*. Luria does not classify this as categorical classification.

Kurvers' multiple bases for categorical classification

What kind of results do you get if you *do* take different categorical bases for classification into account? Kurvers (2002, p. 111-116) did exactly this, and her results differ markedly from Luria's on this task. A majority (55.3%) of her illiterate subjects gave 'categorical' answers in the classification tasks compared with 77% for the literate group. Kurvers criterion for 'categorical' was that the subject named a common characteristic of three items which the odd one out lacked. These are illustrated in responses to the series 'rabbit-cow-fish-dog'. In response to the question: 'what doesn't belong here?' all of the following responses were counted as categorical:

> dog – 'because you can eat the others' (category *animals we eat*)
>
> rabbit – 'because it's wild, not useful for people' (*bred for people*)
>
> fish – 'because it lives in the sea; the rest live on the land' (*land animals*)
>
> fish – 'because the others are all mammals' (*mammals*)

The last option would have been the only one to be labelled 'categorical' under Luria's criterion; it's the only one relying on the concept *mammal*.

Using this broadened criterion for categorical classification does not eliminate 'non-categorical' responses. A sizeable group of Kurvers' subjects (25.5%) also responded with a 'situational' classification strategy. This is similar to Luria's label, so that for instance in the series 'chair-stool-television-sofa' the odd one out is chosen to be chair, "because that belongs in the kitchen." A third group of subjects used what Kurvers calls an 'idiosyncratic' strategy, which involves reacting to a feature of *one* of the objects (choosing a saw because "the others

can't saw") or choosing an item for personal reasons ("because I like it"). This accounts for 19.1% of the subjects' responses. These examples illustrate that the difference between categorical and non-categorical is a matter of degree: 'situational' categorization plausibly reflects non-subjective stereotypes – chairs belong in kitchens for instance. This could be experimentally investigated, and if such choices are stable across many subjects then this brings 'situational' classification closer to 'categorical' choices than to idiosyncratic subjective classification described above.

We can sum up the results so far as follows: according to these studies, classification tasks yield a mixed response in illiterate populations along two dimensions. Firstly, there is variety in the choice of strategy for grouping, primarily either categorical or 'situational' functional grouping. Then, within categorical choices, there is variety in the basis for categorization – much more so than in literate subject groups (Kurvers, 2002, p. 112). The heterogeneity in the illliterate response more generally contrasts with that in literate subjects, who overwhelmingly respond on the basis of categorical classification according to concepts such as *mammal*, *tool*, *furniture*, etc.

How do Kurvers' results compare to Luria's findings? They suggest a much more qualified and subtle distinction between literate and illiterate classification behaviour than allowed for by Luria when he says that his subjects "do not employ verbal and logical methods to group objects but reconstruct graphic situations in which the latter can function" (p. 91). Instead it seems that illiterate subject do use categorical bases for classification, although many more so than literate subjects, but that these categorical bases were just as 'abstract' as those intended by the experimenter. The most problematic aspect of Luria's reading of his results is the association of 'logical' or 'abstract' thinking with categorical classification; he has no independent story about why proposing a hypothetical situation, in which objects bear a functional relation to each other, is not evidence of logical or abstract thinking. In fact, his definition of situational thinking belies the flaw (p. 49):

> objects are grouped together not according to some general principle of logic but for various idiosyncratic reasons. Any such group can be extended to include the most diverse objects (all of which, however, apply to a given situation).

This is not an accurate representation of the transcripts where subjects classify 'situationally'. As much as the subject grasps what's being asked of them, they often gave an eloquent explanation in terms of stereotypical situations, surely just as abstract as categories. Moreover, they are *capable* of doing categorical classification, just as well as their literate counterparts. The following excerpt illustrates both aspects (p. 71) – see especially the first and last turns:

> *Subject is given the series* tree–ear of grain *to match with one of* bird–rosebush–house.

S: *There should be a house next to the tree and the flower* (ear of grain).
E: But is a house really like a tree in any way?
S: If you put a rosebush here, it won't be of any use to a person, but if you put the house here, *a person could live in it and have beautiful things around him. ...*
E: But are trees and a house alike in any way?
S: They don't look alike but they go very well together. *If you want to pick the one that's alike, you've got to pick the rosebush.*

To label such a response any less logical than a simple 'rosebush' answer seems absurd.

Greenfield's data on classification behaviour

Another widely cited source of evidence for differences between schooled and unschooled subjects on classification tasks is Greenfield's work with Wolof children in Senegal (1966, 1972). Rather surprisingly given the high co-occurrence of citations of Greenfield and Luria, the results presented by Greenfield at least partially controvert what Luria reported for classification tasks. To wit, her unschooled subjects did categorise according to abstract categories such as *redness*. Greenfield analyses her data in terms of Vygotsky's definition of advanced conceptual structuring, *superordination*, in which objects are grouped by sharing a single common feature. She summarizes her findings thus: "superordination became more frequent with age in all three cultural milieus [unschooled rural, schooled rural, schooled urban]" (1972, p. 174). This contradicts what Luria reported among illiterate Siberian adults, where he found categorization to be done on the basis of functional relations, as we've seen above. Possibly this is a function of different testing materials: in Luria's examples the 'odd object out' was often still functionally linked to the others, while in Greenfield the objects did not have such an obvious functional relation. In her first study, Greenfield (1966) presented unschooled children with ten objects bought at the local market, which could be sorted into four round objects, four red objects, or four articles of clothing. Her results were as follows:[5]

10% of 6–7 yr. olds formed one of these groupings
30% of 8 & 10 yr. olds did so
100% of 14–16 yr. olds did so

Despite the apparent dissimilarity in results, Greenfield interprets hers similarly to Luria – that is, lack of literacy is characterised by lack of abstraction.For Greenfield this is manifested in the different ways the subjects expressed superordination: either by a holophrase (a single word) or by linguistic predication (a full sentence). So, for instance, when grouping red objects together and asked

[5]One notable phenomenon not discussed here was the predominance of the 'colour' grouping among unschooled children. An exploration of the implications of this finding is beyond the scope of the current discussion.

to supply a reason for doing so, you can either reply by simply saying 'red' (holophrastic), or by indicating each object and saying something like 'this is red' for each one, or saying 'they are red' (linguistic predication). The latter case "involved an explicit statement of the connection between attribute and group members". Greenfield contrasts this with the holophrastic approach, whose communication value "is more dependent on the situational context". This now becomes the key point of divergence of the unschooled group from the schooled: "While the unschooled children became increasingly systematic with age in their object groupings, they continued to express the attribute basis for their groupings in a single word". Big deal, you might think. But in fact this is taken to be very significant by Greenfield. The significance of this finding rests on her remark that the holophrastic expression "demands greater knowledge of the concrete situation – in this case the experimental stimuli – to have communication value for a listener" (1972: 174). If you grant this, then you'll agree that "embedding of a label in a total sentence structure (complete linguistic predication) indicates that it is less tied to its situational context and more related to its linguistic context" (pp. 174-5). From here it is but a short (theoretical) step to Greenfield's central thesis, that "context dependent speech is tied up with context-dependent thought, which in turn is the opposite of abstract thought" (p.169).

Greenfield's extrapolation of this thesis from her findings is precipitous to say the least. The basic findings were that unschooled Wolof children were more likely to justify categorization with single word answers, while schooled children were more likely to give complete sentences as answers. This is taken to be evidence that unschooled children have more "context-dependent" thought. There are several remarks to be made here. Firstly, what does "context-dependent" actually mean? The "context" of the task is probably very different from the schooled subjects' point of view; once we take this into account it might be that their thinking is just as "context-dependent" as unschooled subjects' thinking. Secondly, when we compare the findings to Luria's earlier findings, we see that they are different, but are nevertheless being taken to have the same significance as Luria's; namely that illiterate subjects are less capable of 'abstract' thought. But the means to measure 'abstractness' has shifted – it's no longer how 'abstract' the basis for categorization is, but how 'abstract' the justification for giving the categorization is, as measured by some ill-specified sense of context-dependence.

The lesson here is that until there is a precise understanding of the range of construals of a specific cognitive task – for both the experimenter and the subject – we cannot make evaluative judgements about performance in the task. We need to find out what the task *is* before we decide what's significant in performance and what's not. In classification tasks with illiterate subjects this has not yet been achieved. Instead the results show variation in the difference between subject groups according to the study, and hence it is difficult to identify a clear pattern characterising illiterate behaviour as different from literate.

Reasoning tasks

The most widely-known findings in this area are those on syllogistic reasoning, as described in the previous chapters. For Luria these findings provided decisive evidence for his conclusion that illiterate people cannot reason in abstract or logical terms. But as we've also seen in the previous chapters, results from syllogistic reasoning tasks reveal greater differences between groups than results from, for instance, conditional reasoning tasks. This tells us that focussing solely on syllogistic reasoning tasks gives a skewed picture of the reasoning abilities of illiterates. Even within the syllogistic format, differences vary according to material used. For instance, Tulviste (1991) found that with unfamiliar material group differences were minimized (see previous chapter for details). When we consider a broader range of reasoning materials it becomes clear that there are in fact many similarities in reasoning behaviour across differently educated subject groups; in some settings they are practically indistinguishable.

Yet the differences remain. Like Luria, we found that syllogistic premises often do not yield 'logical answers' from illiterate subjects. In the previous chapter it was argued that, because some premises more closely resemble natural discourses in certain ways, they are relatively unproblematic for all subjects. More importantly here, they are treated in the same way by schooled and unschooled subjects. In contrast, unschooled subjects are in some conditions reluctant or unable to draw conclusions from syllogistic premises, because these premises *don't* resemble naturally occurring discourse. But for schooled subjects these premises are as easy and unproblematic as the naturally occurring ones. And this is exactly the observation we want to explain: schooled subjects manage just as well with the *un*natural premise sets.[6]

Further work on the interaction of reasoning and literacy is minimal – probably, as Stanovich (1993) suggests, "stifled" by the widespread acceptance of the outcome of Scribner and Cole's (1981) study aimed at separating schooling from literacy effects. As we'll see later on, there is plenty of scope for disagreement and qualification of their findings.

We can sum up the reasoning research covered here as indicating that logical aspects of illiterate performance have been overlooked because of a lack of attention to the semantic structure of reasoning tasks. When a semantic analysis of, for instance, the syllogistic task is given, we see that the difference between subject groups can be explained in terms of their ease in ignoring certain parameters of 'normal' interpretation, such domain specification preceding *all*-usage. In this case, then, the impact of literacy would seem to be a broadening of ways of interpreting linguistic materials. How general this is, is unclear, since it might

[6]We should remember that they don't manage that well, as plenty of results from 'mainstream' psychology of reasoning show. Sloman (1998) provides a nice illustration here. See Chapter 4 for more discussion on this and the seeming jump in the gap between the norm and response from here to there.

just be an artifact of the similarity of syllogistic tasks to story sums used in school settings. How significant this it is in terms of cognitive development is also as yet unspecified – we would need an account of how semantic flexibility features in cognitive processing to be able to judge this. We return to this issue up further on.

Other tasks covered by Luria

As well as investigating syllogistic reasoning and classification, Luria (1976) compared literates and illiterates on perception, self-awareness, calculation and definitional tasks. The former two are not purely 'cognitive' so we won't discuss them here.

The definition task was apparently an unmitigated failure: illiterate subjects couldn't be enticed to give definitions. For example, on being asked, "Try to tell me what a tree is", the response was "Why should I? Everyone knows what a tree is, they don't need me telling them" (Luria, p. 86). Luria does not give quantitative results for this task.

In the 'problem-solving' tasks, Luria presented his subjects with basic mathematical problems such as

> It takes thirty minutes to go on foot to a certain village, or five times faster by bicycle. How long will it take on a bicycle?

Luria reported similar types of responses as for the reasoning tasks: subjects rejected or ignored premises, answering on the basis of their own knowledge, calculations or guesswork ('One minute!').[7] One subject in this condition asked for the problem to be converted to 'versts' (a Russian unit of length), and then, when Luria refused, said "We don't reckon in hours; I had better reckon in days", but then did solve an equivalent problem posed only in terms of buttons to be divided among five people. This was typical. Again, though, Luria's synopsis is damning: "The subjects can be made to solve the problem when they operate with concrete entities (versts). But when the problem changes to an abstract level (time), the subjects are incapable of reasoning about conditions divorced from practical experience, and they slip back into arguments based on experience" (p. 130), and further on: "All this shows that the formal operation of problem-solving presents major, sometimes insurmountable difficulties for these subjects." (p. 132) I disagree. Subjects *are* solving the same formal task, i.e. making the same calculation, merely with different labels for the basic units, and they are not familiar with calculations in time. Which European can calculate 11 stone less 6 pounds 2 ounces? Again, a lacking specification of the informational task

[7] In the particular condition mentioned, Luria offered thirty buttons to the subject to help him solve it – something that would only work if the subject understood that each minute was represented by a button and that 'five times faster' should be represented as division by five in the problem. Needless to say the subject didn't use the buttons!

the subject is supposed to be solving allows for different interpretations of the results.

Interim conclusions

How to sum up the findings from this selection of empirical studies? Results in this area have been often incomparable or mixed. As such, it is hard to generalise beyond saying that subjects' performance suffers because of their unfamiliarity with the task, in some cases in terms of what's expected of them, in some cases in terms of the materials used; and these factors also overlap. Sometimes diminished performance seems to be an artifact of the task analysis – think of Luria's too-narrow definition of categorical classification; but in other cases it plausibly *does* indicate a lack of cognitive ability, because the skill is inherently bound up with experience with a symbolic system – for instance, as is the usually the case for complex arithmetic. In general, it is unclear whether results are *merely* a result of different uses of language usage – from the level of words up to the level of discourse genre – or whether, perhaps even because of different language usage, they do tap a difference in cognitive processing. But to get further in determining this we need a theoretical account of the relation between symbolic systems and cognition.

More serious though, is the inconsistency with which results have been interpreted – evident in Greenfield's interpretation of her subjects' categorical classification as nevertheless showing more "context-dependent" thought, despite the fact they met Luria's criterion for 'literate' thought. As we have discussed in the previous chapters, Luria's interpretation of his own data is seriously undermined once subjects' interpretations of the task materials are taken into account. It would seem that researchers have such strong expectations, that they overlook aspects of the findings which do not support their apparently foregone conclusions. One wonders what would have happened if the studies could have been conducted double-blind.

At this point it is worthwhile to reflect on what such tasks are supposed to be testing. What are we hoping to learn from them? It has just been argued that Luria's unschooled subjects were as capable of making calculations in versts as their schooled compatriots. But of course Luria's subjects would have been stretched by more complex problems – no-one would expect them to do long-division. Everyone knows you need to learn maths at school.[8] Similarly, one can certainly maintain that syllogistic problems are not suitable for use in the study of reasoning in illiterate subjects. The problems *are* strange to the subjects. Yet we needn't shy away from a seeming mismatch of material and response; we should use it to learn what we value in terms of cognitive processing and why. The

[8] Although see Saxe's (1988) fascinating study of how Brazilian street children developed a complex mathematics for candy-selling, which they incidentally also go on to use in school!

mismatch itself can be very telling.[9] We can get information about the range of task interpretation available to a subject and then we are in a position to better judge how our cognitive development piggybacks on interactions with externally developed symbol-systems, like the arithmetic used in school mathematics, but also language-genres, and how specific the link is between a skill and its context of use. In the case of reasoning tasks, it seems we have neglected exactly this aspect of reasoning behaviour. Results have not been sufficiently 'vetted' for semantic factors; nevertheless they have often been taken to be convergent about the idea of a deep difference between literate and illiterate thought. We now explore the context in which literacy is overwhelmingly acquired: school.

3.4 Are we barking up the wrong tree?

Most people learn to read and write at school, or at least, most people who read and write have also been to school. Acquisition of literacy skills is thus usually paired with a process of enculturation in an educational system; any comparison between literate and illiterate subject groups really should take account of the confound effect lurking here. This is very difficult to do, though, precisely because the two factors (literate and schooled) almost always occur together. There are rare opportunities to study the two separately. Scribner and Cole had one such rare opportunity, in 1970s' Liberia, where literacy and schooling could be tested separately – to some extent. The results of their large-scale empirical study within the Vai population of Liberia are reported in Scribner and Cole (1981).

3.4.1 Literacy without schooling

The opportunity arose through a fortuitous diversity in scripts and learning contexts within the Vai tribe. The situation was as follows: schools (usually mission-based) were wide-spread in Liberia and had English as medium of instruction. Additionally, there were many religious schools for study of the Qur'an, where the Arabic script was used. Attendance at these religious schools was often an after-school activity – so the English and Arabic literacy were commonly paired. But the Vai also have a native syllabic script, which is taught and used in private

[9]For instance, Kurvers tested illiterate subjects on the Raven Standard Progressive Matrices task (Kurvers, 2002, p. 116). All illiterate subjects scored very low on the task – the average was 16.32, from a maximum of 41 points. More interestingly, however, was the narrow range of scores in this group – outliers aside, the highest score was just above 20, the lowest just above 10 points. In the literate (but relatively unschooled) group, the average score was 25.88 but the scores ranged from 10 to 41 points. In the latter group only, then, it would seem that the task is achieving its purpose: differentiating intelligence levels. This cannot be said of the test for the illiterate group. Note that these results were garnered on a non-verbal task; the fact that illiterate subjects showed 'illiterate' behaviour here highlights that literacy does not apply only to texts but to pictorial representations too.

contexts.[10] Although originally (150 years ago) otherwise, when Scribner and Cole visited Liberia there was no formal instruction for learning the native script. Rather, it was taught largely in a one-on-one manner; in many cases the student lived with the teacher. Students of the script were mostly young men, who chose, usually for personal reasons, to learn it. Scribner and Cole (1981) report one man's experience:

> [H]e began to learn when he went to work at a sawmill in the high forest of the Vai Koneh district, where a fellow sawyer was able to read and write Vai script: "All of us were living together ... Every time he received letters he read and answered them, so I too got encouraged and decided to ask him to teach me." Over several months, as they worked together, they would have a session, "sometimes five minutes when we met, because we never used to spend to long talking about it".

This seems to have been a fairly typical way to learn; very few reported learning on their own or in a more organized context. Many of these Vai 'students' had never been to school. This then is the key group: they are literate, in the Vai script, but have no experience of a formal schooling environment.

Scribner and Cole conducted a large-scale study of the effects of the different literacies and learning contexts on performance in a range of experimental tasks, including abstraction, classification, memory and reasoning tasks. The question was, who would unschooled literate subjects look more like: schooled (literate) subjects, or (unschooled) illiterate subjects? If, in any specific task, the former was the case, then that would suggest that indeed, the skill tapped in that particular task was advanced by skills in reading and writing per se. If, on the other hand, unschooled literate subjects looked more like illiterates on a task, then it would suggest that familiarity with school practices (which might include some specific literacy practices) made the greater difference to performance in the task.

The results from the testing were striking: no specific effect of (Vai) literacy was found on a number of cognitive tasks, including geometric sorting tasks,

[10] The Vai's native syllabic script is used, as stated above, principally in a private capacity: for letter-writing, record-keeping and accounting. Scribner and Cole report one case in which a Muslim association had a constitution and bylaws written in Vai script. The script is composed of approximately 210 signs, representing all possible combination of consonant and vowel (most syllables have CV form), plus seven vowel symbols and a 'syllabic nasal sign. Precise origins of the script are unclear. What is known is that it was first developed early in the nineteenth century, and may originally have had a more pictographic character. A well-known origin story tells of a man 'Duala Bukele', who was visited in a dream by a tall white man in a long coat who brought a book to show Bukele. On waking, Bukele could not remember the signs, so gathered together with friends and made new signs. At that time the Vai had already had contact with Roman and Arabic scripts (both alphabetic): Portuguese traders had established connections with them in the mid-fifteenth century, and the Vai themselves were migrants from the Mande region along the Niger river and had much contact with Islam.

classification tasks, recall tasks and syllogistic reasoning tasks. Scribner and Cole summarize their findings thus: "The most impressive finding is that formal schooling with instruction in English increased ability to provide a verbal explanation of the principles involved in performing the various tasks ... neither syllabic Vai script literacy nor Arabic alphabetic literacy was associated with what are considered the higher-order intellectual skills" (1981, pp. 130-132). The authors themselves argued that their findings "lay to rest some misconceptions about the psychology of literacy that went unchallenged in the past for lack of empirical data. ... The small and selective nature of Vai script and Arabic influences on cognitive performance precludes any sweeping generalisations about literacy and cognitive change" (p. 132). And expressed even more strongly elsewhere "Our results are in direct conflict with persistent claims that 'deep psychological differences' divide literate and non-literate populations." (p. 250). Specific effects of Vai literacy were found only in tasks more tightly tied to literacy: grammaticality judgements, rebus reading, and integrating syllables.

Not surprisingly, this study has popularly seen as the 'death blow' to Great Divide theories (Stanovich's phrasing, 1993, p. 138). Jahandarie (1999) says of the impact of Scribner and Cole's study: "[that] literacy did not have any of the general cognitive consequences attributed to it by the literacy theorists [has] become received wisdom in many discussions of the topic" (p.267). Greenfield (1983) expresses the 'general view' (according to Olson, 1994, p. 20) in saying that Scribner and Cole (1981) "should rid us once and for all of the ethnocentric and arrogant view that a single technology suffices to create in its users a distinct, let alone superior, set of cognitive processes." (Greenfield, 1983, quoted in Olson 1994). As such, it deserves a closer look.

Interpreting the results from Scribner & Cole (1981)

Jahandarie (1999) takes Scribner and Cole to task for over-interpreting their own results. He questions their measures in some tasks, but more worryingly, identifies several points at which Scribner and Cole offer a 'summary' or evaluation of the results which according to him incompletely represents, or is not justified by, the actual results. An example is the word definition task. Scribner and Cole report the following outcome:

> After considerable experimentation, we coded each definition on a binary basis. ... On this measure we obtained *no noticeable population differences*, but there was striking evidence that definitional adequacy was controlled by the semantic properties of the words being defined. ... Words that were more familiar or concrete in meaning elicited the most adequate definition, while level of description dropped for words at a greater distance from concrete experience. (my emphasis)

Note that this is a finding which applies across all groups. Yet Scribner and Cole continue:

> This outcome is consistent ... with Luria's distinction between mundane and academic concepts – a distinction that literacy in Vai script or Arabic does nothing to diminish. (p. 150)

... And neither does schooling in English, a crucial omission at this point.

Another set of findings for which Jahandarie questions Scribner and Cole's conclusions is based on the syllogistic reasoning task, where he reports that although all the literate groups (Vai, Arabic and English) performed equally well, Scribner and Cole summarize as follows: "Taken together, these studies of logical-verbal problem solving cast doubt on hypotheses that implicate literacy directly in the acquisition of metalinguistic knowledge about the properties of propositions" (Scribner and Cole, p. 156, reported in Jahandarie, p. 270). Jahandarie concludes: "Once again, a finding not justified by the actual findings." Now in fact it's Jahandarie who's over-interpreting his findings. Scribner and Cole conducted two sets of tests on syllogistic material. In the first round of testing, conducted with familiar content materials, they found that:

> Of all the survey tasks, logic problems proved the most predictable and demonstrated the strongest effects of schooling. Not only did amount of school increase the number of correct answers, but it contributed to the choice of theoretical explanations, over and above correct answers. Schooling was the only background characteristic to improve performance; neither Vai script nor Arabic literacy had an effect on either measure."

The *second* round of testing was conducted with new materials, this time with unfamiliar content, i.e. content about which subjects could not have had own experience of – for example "All stones on the moon are blue". With these materials, group differences disappeared[11] although exactly how the tests went is not clear from Scribner and Cole's reporting of them. As well as this change in materials, they distinguish between a metalinguistic survey and a replication survey; the difference between these two is not explained. The one variable that functioned as a predictor was the order in which the tasks were given: "when logic problems followed all other tasks [including conversations about grammar and words], the rate of theoretical responses was significantly higher" (p. 156). Scribner and Cole are pursuing the hypothesis that familiarity with the relevant discourse genre accounts for the shift in performance on the tasks – not literacy per se – and their reported results are certainly consistent with their conclusion as presented above, which they continue as follows: "In moving from one study

[11] See also Tulviste (1991).

to another we found greater variability arising from differences in materials, procedures, and experimenters than in literacy or other background factors."

This discussion shows the importance of reporting others' results in a complete way. The failure to do so – but this time in the other direction – has resulted, as Stanovich (1993) points out, in Scribner and Cole being over-interpreted in the literature in literacy, resulting in the 'death blow' interpretation of their results for Great Divide theories. "A major issue that is often glossed over – but that, interestingly, was raised by Scribner and Cole themselves – is whether the nature of Vai literacy was such that it provided a valid test of the claims of Great Divide theorists." There are several ways in which Vai literacy do not meet the criteria for the literacy intended by theorists: for Goody & Watt (1963), it would be the lack of access to accumulated knowledge, for Olson (1994), no essayist tradition, even Havelock (1982) would quibble that it's not alphabetic.

But more generally, the Vai script might fail to provide a means to test literacy because it was used for highly circumscribed purposes: primarily for letter writing and, to a lesser extent, record-keeping. In several ways these texts are 'atypical' according to the criteria given by many literary theorists. For instance, letter-writing is highly personalised, assuming much shared knowledge between writer and addressee. The letters usually follow a fairly prescribed format, and are used for only two communicative purposes: making requests, and giving information, such as announcements about births, deaths, or reasons why money has not yet been repaid, for example (Scribner and Cole, pp. 71-75). Moreover, the script, as described above, is learnt much later in life and in a one-on-one environment. Another relevant factor is the lack of 'literature' in the Vai script – there's almost nothing to read in Vai.

Interim conclusions II

To give a short answer to the question we started with: no, we are not barking up the wrong tree. These considerations show that the empirical results garnered by Scribner and Cole by no means strike a 'death blow' to literacy hypotheses. There are at least three main reasons why this is so: firstly, their results are not as clear-cut as they could be, as we've seen above. Secondly, and more importantly, Vai literacy is a very restricted form of literacy; most theorists would be happy to only make claims about a more diversified literacy, for instance, involving at least reading texts from unknown authors. Related to this, a third key problem is that a simple separation of schooling from literacy as Scribner and Cole claimed to have is not tenable: schooling is centred around literacy-based activities. In fact, Scribner and Cole's study could be understood as investigating the effects of say *primary* versus *secondary* literacy. Another factor which should be mentioned is bilingualism; this is known to affect metalinguistic knowledge, but in this case the contribution of bilingualism is in all likelihood only applicable to the English schooled group, who spoke both English and Vai, and not to the

Qu'ranic schooling group, as in the latter case the students just had to rote learn the text without understanding what it meant.

Moreover, that literacy and schooling are not easily separable is a conceptual as much as an empirical matter. School in many senses forms our concept of what it is to be literate and what our understanding of language is – as the next section indicates. On the other hand, we've seen that the term 'literacy' is used where 'schooling' or 'schooled literacy' could often more accurately be substituted, since the term is used to refer both to the basic skills in reading and writing and to deeper interaction with literate culture. In some cases, the sense in which literacy is attributed to someone makes it tantamount to meaning: this person has passed through a Western-style formal education system. The identification is even made explicit by some authors. We see this, for instance, in Ogbu (1990), addressing difficulty with literacy in minority groups in the United States: "I define literacy as the ability to read, write, and compute in the form taught and expected in formal education. Put differently, *I consider literacy to be synonymous with academic performance*" (1990, p. 520). Indeed, schooling *is* a heavily literacy-oriented institution. The 'three R's': reading, writing, and arithmetic, are all about becoming literate in different systems of notation. Much of school activity involves interaction with texts; in this sense, schooling is a vehicle of literacy, and we could say someone is 'schooled' when they are literate in the right ways. Describing the literacy-based aspects of education would mean specifying what kinds of interactions are valued, not only with texts, but also with other notational or symbolic systems .

Another terminological issue we should bear in mind is the varying application of the term 'literacy' at the level of society. For instance, the term 'literate society' is sometimes used when referring to a society in which some small elite is literate (as in classical Greece); sometimes to a society in which it is widespread and needed to negotiate every day life (as in modern society where for instance official forms, road signs, prices, etc are written). Moreover, literacy effects at a societal level can interact with those at an individual level. If a literacy-based activity/technology catalyses some or other insight on language then it might well be widely absorbed by a culture without specific individuals gaining reading and writing skills. So, for instance, 'word magic' could die out in a society before literacy becoming widespread. On the other hand, if reading and writing skills are used in a very narrow range of activities, in an otherwise non-literate environment they might not be accompanied by the consequences otherwise associated with the acquisition of literacy. This phenomenon is reported in, for instance, Narasimhan (1991), which describes the continuing dominance of oral practices in India despite long literate traditions, and, as we've seen, precisely the same phenomenon is uncovered by Scribner and Cole in their 1981 book. This point has led some theorists to suggest that it is technologies and practices associated with scripts which have more impact than just the presence of the script in a society – for example, for both Finnegan (1973) and McLuhan (1962) the invention of the

printing press was more important in shaping modern society than the invention of scripts.

Thus we see that the task is to specify which activities, conventions, forms of knowledge, etc. engendered by the schooling environment are the relevant ones for the tasks at hand (see Ceci & Roazzi, 1994 for evidence of why this is important). It might turn out that the aspects of schooling which advance cognitive performance are *not* directly related to literacy – for instance the social environment. Results of this kind would be far more damaging to literacy hypotheses. We need to delve into the schooling group, to ascertain literacy-based *vs* non-literacy-based effects on cognition, in order to evaluate the literacy hypothesis. A way to investigate the social environment aspect, for instance, would be to compare home-schooled subjects with normal schooled subjects. Comparing subjects schooled under different educational philosophies would also help to separate these effects somewhat.

3.4.2 Social factors in cognitive performance

The results presented and reported in earlier sections of this chapter and in previous chapters make it clear that cognitive development cannot adequately be viewed as an individual-bound process of maturation, whereby cognitive skills or structures simply unfold with age. As well as this, Scribner and Cole's results show that the school environment has a role over and above literacy skills in determining cognitive performance. The message is clear: we need to seek to explain the effects of literacy relative to their contexts of acquisition and use. The cultural context of acquisition and use is *especially* pertinent when a task involves cognition on symbolic material which needs to be interpreted. The ways in which cultural and social factors may shape cognition are manifold; in this section we explore ways in which cultural and social factors can shape cognition via the imposition of norms in a school environment.

Schooling as a source of cognitive and linguistic norms

In Cole, Gay, Glick and Sharp's (1971) study with the Kpelle they had little success in getting unschooled subjects to group objects according to categories in categorization tasks – until, that is, they asked subjects "how stupid people would do it" (Glick, 1975, p. 636). Then suddenly, unschooled subjects used categories to group objects – as schooled subjects were doing all along.

This example illustrates the importance of the value we place on *ways of approaching* problems. Cognitive and social development is paired with the acquisition of values. Values can attach to both activities, goals, knowledge, skills, and to *ways of* performing activities, achieving goals, acquiring knowledge, practicing skills, approaching problems. For instance, calculating with numbers is a skill worth acquiring. On top of that, there is a correct way to calculate. There's

no value in calculating if you do it incorrectly. We might say the skill is both valued and 'normed', that is, practice of the skill is always constrained by normative standards for how it *should* be practiced. As Goodnow (1990, p. 259) puts it, "We do not simply learn to solve problems. We learn also what problems are worth solving, and what counts as an elegant rather than simply an acceptable solution." The role of school in this process is more apparent in some domains than in others. Enforcing appropriate behaviour in the social realm is explicitly part of a school's prerogative. But in the cognitive domain, just as in the social domain, there is a very clear indication of what kinds of activities, knowledge, goals are valued – the dual authority of the teacher and the syllabus determine what is taught in class. And the 'norms' attaching to correctly acquiring and practicing are given by judgement of performance, for instance, in grading exercises and tests. More generally, the norm in the cognitive domain is expressed in terms of 'intelligence': some behaviours are smarter or more intelligent than others.

The example from Cole et al suggests that each cultural group has its own definition of what are the interesting and worthwhile problems.[12] Along similar lines, Luria (1976) remarks that his experimental subjects often found it very strange that he was interested in such trivial tasks.

There is another sense in which school imposes norms, but this time in the sense of 'normalisation': fostering a sense of homogeneity. This occurs simply by the standard classroom set-up: with the possible exception of the naughty ones having to sit at the front, everyone is treated the same. Everyone has to do the same work and come up with the same answer. This engenders the notion that everyone has the same basic set of cognitive capacities – especially in the early years. In fact this 'normalisation' is accomplished in some sense by the common norms applied to performance, which generates the idea of a class of 'like-minded' children.

In the norming of cognitive activities and the normalisation of the members of the class, the child learns that there is a correct way to do things, and that it's the same for everyone. It seems to me that this very basic way of approaching school-like problems is present only in literate performance on reasoning tasks. Illiterate subjects in my study often seemed very concerned to 'get it right', but more out of a sense of wanting to understand, to have a successful communication, rather that from a sense of oneself as a representative of a kind of thinking, or thinker.

[12] A result of cultural valuing of knowledge, as Goodnow observes, is the phenomenon that particular groups take ownership of some skills and areas of knowledge: these skills and knowledge can seem to 'belong' to some people more than others. She gives the example of mathematics in our culture being considered more relevant for males than females – although this is changing rapidly. Another illustration of this is Fordham & Ogbu's (1986) fascinating account of resistance in black students in the States to academic performance, out of fear of being accused of "acting white".

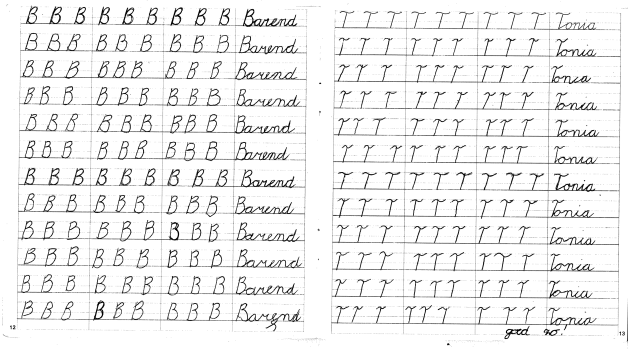

Figure 3.1: Letter-shape norms in learning to write

Schooling as a source of norms about language

A central argument in this dissertation is that linguistic material used in reasoning tasks is not a neutral catalyst for reasoning processes, but is integral to the reasoning itself. We've also observed that educational background of a subject influences how they interact with the linguistic material. Where can we find a theoretical home for such an observation? Mainstream linguistic theory cannot provide it. As the sociologist Pierre Bourdieu so adroitly puts it, linguistic theory suffers "the illusion of linguistic communism" (1991, p. 43). All speakers have equal and untrammelled access to a shared language. For Bourdieu, this stems from the conception of language as a kind of 'universal treasure', and linguistic competence as the deposit of this treasure in the individual.[13] Bourdieu sees this conception maintained in Chomsky even though Chomsky makes explicit that he is concerned with an ideal speaker-listener. Distinguishing between competence and performance, à la Chomsky, does not rid us of lingering communist ideals, precisely because this vocabulary hides a *fictio juris*, "converting the immanent laws of legitimate discourse into universal norms of correct linguistic practice". Bourdieu's concern is that this move "masks [the] social genesis [of the language as object]". In other words, the use in linguistic theory of an idealised speaker-hearer bolsters the sense of a universally given system of norms for linguistic practice and masks the social nature of these norms. It identifies 'official language' with 'ideal language'. This would not be so damaging if, say, the Chomskyan picture of language had not been so influential. As it is, his ideas have wide currency beyond linguistics.

There is strong empirical evidence that in fact the 'treasure deposit' is differ-

[13] A good example of the use of this assumption is the reliance on native speakers' intuition in theories of linguistic structure. A single speaker, any speaker, will do, since they are assumed to have access to the same linguistic treasure trove as any other speaker.

ent for different speakers, and, more importantly here, that it differs according to educational background. For example, Karanth and Suchitra (1993)'s research review indicates that the ability to judge grammaticality develops over a long period, increasing dramatically in the age-group 6-7 yrs, is qualitatively different at different stages of development, and develops at different rates for different aspects of grammar (for example, sensitivity to plural markers is early, sensitivity to participial markers is late, reaching ceiling after 11 yrs). The timing of this development is consistent with the proposal that grammaticality judgements develop out of school activities. Corroborating this, Karanth, Kudva and Vijayan (1995) report that illiterate subjects "generally refused to perform the task" or "gave indiscriminate responses" (p. 304) on grammaticality judgement tasks. This result is limited in the sense that it suggests the grammaticality judgement task does not effectively tap any intuitions about grammaticality of the subjects, but it does raise the question: does the illiterate subject have the same intuitions about language norms? Dabrowska's (1997) comprehension study overcomes the limitations of a grammaticality judgement task, thus going further to answering this question. And it would seem the answer is no. Dabrowska conducted her study in a group of university employees of varying educational levels. The least-educated group consisted of cleaners and janitors, most of whom had had no more than ten years of schooling. The next group consisted of first and second year undergraduates from different faculties. The third and fourth groups were graduate students and university lecturers, respectively. The latter group "had a professional interest in language": they came from literature or foreign-language departments. All these respondents had spent a substantial amount of time in education. Still, Dabrowska found a clear progression in their scores on comprehension tasks. Adults' ability to process complex syntactic structures thus depends on their level of education. There is no universal competence to which we all have equal access.

So much for the experimenter's reliance on a non-existent norm. But what about the reasoning subjects? They are no Saussure or Chomsky readers. Where does their sense of a linguistic norm originate? And for the linguistic theorist, what is his source for the 'immanent laws of legitimate discourse', giving the 'ideal language'? The key to this lies in education: schooled subjects have been taught to talk properly, as it were. As Bourdieu says, "[i]n the process which leads to the construction, legitimation and imposition of an official language, the educational system plays a decisive role" (p.48). This could be achieved by any or all of several aspects of the educational system: firstly, the teacher, the authority in the classroom, teaches in the standardized language, inclining the students to see it as this. Related to this is the dynamic of the classroom situation itself, in which all pupils are treated similarly by the teacher, again fostering a sense of sameness between pupils. Secondly, of course, there are grammar lessons, in which the standard language is explicitly taught. Even if children do not use this language in the playground, the perception that there *is* such a language is strengthened.

In this way, the notion of 'ideal language' and the 'official language' coincide.

But if the idea of such a universal language competence is a theoretical and cultural construct, what should we replace it with when conceptualising language between groups? According to mainstream linguistic theory the only way to interpret variation in interaction with linguistic materials is as divergence from a universal norm. As *mis*use of language. As mentioned above, this conception of divergence extends beyond the borders of linguistic theory, into theories of cognitive behaviour grounded in the use of language. But this is clearly inadequate to explain the range of behaviour and, as we've seen, premised on the mistaken assumption of "linguistic communism". The ramifications of different language competences resurface when we sketch a theoretical proposal later on in the chapter.

3.4.3 Literacy improves intelligence: a tautology?

The professed aim of the current chapter is to explore the consequences of literacy on cognition. But the data presented are all about *illiterates*' performance, and the ways they do not match up to a posited norm. The implication is a functional near-equivalence of 'literate' response with 'correct' response. This is interesting, because it indicates an implicit bias towards literacy in the selection of test materials. Researchers in this area have looked overwhelmingly at tasks in which 'we' do well, leaving only the possibility to discover deviation from the norm as a source of difference between literates and illiterates, even before the tests have been administered. There may well be good reasons for this – one being that one needs to understand the logic of a test to be able to administer it – but this is very different from the situation elsewhere in reasoning research, where highly literate subjects are found to be systematically diverging from the norm. It seems there is a kind of relative distance between subject and experimenter, which widens or narrows depending on what kind of results are interesting in the context.[14] The equivocation of 'literate' with 'correct' behaviour can also lead to circularity in explanation – for instance, as we will see when we come to discuss the work of David Olson, where logical reasoning comes to be dependent on literal interpretations of language, a literate achievement, because that's what's required in the experimental tasks used thus far to test it.

Within studies of illiterate reasoning, this reduces to taking for granted the representativeness of the experimental tasks. In other words, it is assumed that the experimental setting is suited to elicit behaviour characteristic of the way people process information in the world. This reveals an issue which we have not yet touched on: the matter of generality. How do results generalise, if at all? Do the observed phenomena provide a representative portrait of unschooled

[14]The so-called 'negativity bias' inherent in focussing on norm divergence is addressed in more detail in the concluding chapter.

versus schooled subjects' ways of thinking, or are we getting a random snapshot which has little significance outside the experimental setting? This is a complex issue; in unravelling it you can learn much about the logic of any experimental endeavour to understand cognition.

To get insight into the case at hand, Cole, Sharp and Lave (1976) suggest considering the consequences of learning in a totally different domain. They use the example of learning to be a carpenter:

> Sawing and hammering are instances of sensorimotor co-ordination. Learning to measure, to mitre corners, and to build vertical walls requires mastery of a host of intellectual skills which must be co-ordinated with each other and with sensorimotor skills to produce a useful product. ... To be sure, we would be willing to certify a master carpenter as someone who had mastered carpentering skills, but how strong would be our claim for the generality of this outcome? Would we want to predict that the measurement and motor skills learned by the carpenter make him a skilled electrician or a ballet dancer, let alone a person with 'more highly developed' sensorimotor and measurement skills?

One might readily grant that the master carpenter has 'more highly developed' sensorimotor and measurement skills – within the domain of carpentry skills. Whether this could also be claimed about the carpenter's ability in other areas depends on how skills and tools demanded for performance in both areas are related. We might say: it depends on the extent to which the assembly of sub-skills involved in the skill of say, measurement or co-ordination, overlap. The master carpenter can measure with a measuring rule; he can't therefore measure temperature with a mercury thermometer, because they require different assemblies of sub-skills. He can co-ordinate his hand in the act of shaping a piece of wood with a lathe; he can't thereby co-ordinate his body parts in a dance move, because this requires a totally different set of sub-skills.

But surely the performance in the experimental tasks have a much tighter relation to general cognitive skills? Well, this is because general cognitive skills are defined in terms of just such tasks. As Cole et al (1976) point out, some version of practically every experimental task on the effects of education can be found in Alfred Binet's early work on the development of tests to predict children's success in school.[15] Notice the circularity here. The effects of education are measured with tests originally designed to predict success in education. As such, Cole et al say, "the correlation between successful performance on Binet's tasks and success in school was a tautology" (p. 227).

The way out of this circularity would be, of course, the purported inherent transferability of the skills taught in school: they are designed to be general,

[15]Binet is considered to have developed the first intelligence test; his tasks form the basis of modern IQ tests.

high-level, transferable skills. That would be why they are taught at all. There is some experimental evidence that this is indeed the case, i.e. that schooling fosters the development of flexible problem-solving ability. For instance, Cole and Scribner (1974) found that schooled subjects treated groups of problems as cases of a single type, and applied common operations to solve them. The problem is, the relations between these problems are probably very evident to the subjects from their school experience. Problems unlike those encountered in school might not be solved so easily. More broadly, just how 'flexible' school-learned skills are depends on what we take to be the domain of problems to be solved. Usually the domain is defined by academic performance. Less attention has been paid to the cognitive demands of everyday life – although contemporary cognitive scientists are beginning to tackle this (for instance, Gigerenzer and his colleagues in Gigerenzer and Selten, 2001).

The above point is important because it will influence what the effects of literacy are hypothesised to be. If the literate response is too easily taken to equate with the correct response then the effects of literacy could be overstated or misconceived. Tomasello has made this point in the domain of language acquisition: what the child acquires should be modelled in terms of adult linguistic behaviour – not in terms of a theoretical linguistic structure. Admittedly, in the domain of cognition there are much more compelling reasons to posit a theoretical norm which might not be achieved in practice but is nevertheless correct for independent reasons – think here of complex arithmetic calculations. This rests on the status of the theoretical models of mathematics, as opposed to say models of language. The norm/practice distinction is not parallel to the competence/performance distinction in language use because a competence model describes, and doesn't prescribe, linguistic behaviour.

On the other hand, this is also a point which suggests a widening of the experimental approaches would be fruitful. If a broader range of cognitive tasks would be considered, the results might be more varied. Cole et al (1971) purposefully tried to include tasks where literates *don't* necessarily do well, and found many cases where illiterates did. They report that unschooled Kpelle farmers were much more accurate in estimating quantities of rice, for instance, than their schooled urban counterparts. Note that there is no inconsistency in concluding that indeed schooled subjects 'do more' with the information given in a task, and yet that their skills are not widely transferable.

Interim conclusions III

The underlying issue here is what we take the relation between symbol use and cognition to be. This also has implications for when behaviour is appropriately 'normed' and when a competence model is more appropriate. In western society, the cognitive achievements derived from literacy are potentially over-represented in our palette of intelligent behaviour. As an example, see Stanovich (1993) list

"vocabulary, syntactic knowledge, metalinguistic awareness, verbal short-term memory" as aspects of cognition which contribute to intelligence. Now the question is: is this a contingent, chauvinistic view to take on cognition? Or is human cognition fundamentally shaped, or even advanced, by engagement with symbol systems, of which the written word just happens to be a particularly potent example? Insofar the latter is the case, we should be prepared to grant linguistic capacities agency in the cognitive domain, and reassess whether semantics-based explanations of cognitive performance are indeed deflationary. But this, again, is tied to general theory. An answer rests on a specification of what constitutes a cognitive ability and the kind of agency which can bring it about.

One should, however, be wary of singling out literacy as a cultural product. In the literacy debate, the temptation is to suppose that written language is 'learned'; spoken language is 'given'. As we've seen above, such a principled distinction lies behind much theorising about the effects of literacy. In particular, the physicality of written language often stressed: it is a material thing (which is spatially-extended), in contrast to speech (which extends in time[16]). Thus text can be analysed, broken up and re-interpreted in countless ways – "in short, it comes under the language user's control" (Scribner, 1997, p. 166). In this way the language becomes an object of further analysis and not just an instrument which we use for communicative (or other) purposes.

Is the distinction so sharp? Ingold (2000) argues that it is not; and does so by comparing the contrast between speaking and writing, to the contrast between walking and cycling. For, as Ingold (2000) points out, "it is generally accepted that bipedal locomotion is a universal human characteristic, whose evolution entailed a distinctive suite of anatomical adjustments. ... Cycling, by contrast, is an acquired skill which has appeared relatively lately in some, but not all human populations." The supposition is that we are born to walk, but we learn to cycle. Walking is innate; cycling is acquired. And although it is accepted that certain environmental factors (such as an attentive and supportive caregiver) are necessary for the ability to emerge, the sense remains that we are bound to walk (provided all environmental conditions are met), whereas we are not bound to cycle, and that the body is hard-wired, or 'ready-made' for walking, but not for cycling. But these distinctions are one of degree, not category, says Ingold. Certainly cycling requires more specific environmental preconditions (a bike, for starters, and probably the ability to walk) than walking, and is practised much less widely as a result of this. And certainly there is not such a short critical window period for learning to cycle as there is for learning to walk. These are differences, but they are a matter of "extent, rather than principle". The similarities are much greater: "if walking is innate in the sense – and only in the sense –

[16] And thus comprehension depends on the peculiarities of the auditory modality, as opposed to sight. An interesting discussion on characteristics of different modalities can be found in Ingold (2000).

that given certain conditions, it is bound to emerge in the course of development, then the same applies to cycling. And if cycling is acquired in the sense that its emergence depends on a process of learning that is embedded in contexts of social interaction, then the same applies to walking."[17]

Ingold's expressive analogy suggests that we should be wary also of splitting speech and writing into 'given' and learned'. Speech conventions are as much subjected to cultural conventions as textual interactions are; this point makes it clear that there may be no single defining feature of the impact of literacy which separates literates from illiterates. Nevertheless, it is worthwhile attempting to connect specific aspects of literacy with specific cognitive insights or behaviours. This is the aim of the next section, namely, investigating what aspect of literacy might explain the differences in reasoning performance as reported in the previous chapters..

3.5 Olson on the effects of literacy

As we have seen, most of the theoretical work on the consequences of literacy predates experimental results and/or engages only with the findings of anthropological fieldwork studies. An exception to this is found in the work of David Olson. In developing his account of the consequences of literacy, Olson endeavours to account for experimental results from psychology of reasoning, language acquisition and reading studies, amongst others. Already in his 1977 article "From utterance to text: the bias of language in speech and writing", one of the most widely-cited in the literature, he references reasoning studies such as Cole, Gay, Glick & Sharp (1971), Henle (1962) and Wason & Johnson-Laird (1972) (more recently, in this vein see Olson, 1993, and Olson, 1994). With colleagues he has conducted his own experimental work, especially on children's development of language understanding and adult metalinguistic knowledge (see for instance Lee, Torrance & Olson, 2001, Astington & Olson, 1990, as well as Olson & Astington, 1990). The aim in this section is to assess how useful Olson's ideas are, in explaining the impact of literacy as evidenced in the reported experimental work. Recently Olson has published a comprehensive formulation of his views in the book, *The World on Paper* (1994), so I will primarily respond to his ideas as outlined there.

3.5.1 Central themes

No illocutionary force in texts?

A central theme throughout Olson's writings is the increasing "explicitness" in the transition from spoken to written language. In his earlier work (1977), he states that "language is increasingly able to stand as an ambiguous or autonomous

[17] Living in the Netherlands lends this citation special resonance!

representation of meaning" (p. 258). One consequence of this is that societies move towards greater explicitness as they become more literate; another is that children's language and thought will become more explicit as they become more and more encultured in literate habits, particularly through schooling.

The increasing explicitness hypothesis goes against the commonly-held conception that writing is simply the transcription of speech. This idea dates back to Aristotle: "written words are the signs of words spoken" (from *De interpretatione*, quoted by Olson, 1994, p. 65) and endured to be explicitly upheld by seminal linguistics scholars such as Saussure and Bloomfield (Olson, 1994; Harris, 2000) – although note their goal was to emphasise the legitimacy and centrality of spoken language as a subject for linguistic study. The validity of the transcription view has been challenged by contemporary literacy scholars, such as Roy Harris, as well as Olson. In his book *Rethinking Writing* (2000, and see also Harris, 1986), Harris sets out to deconstruct the Western view that writing is "depicted speech". This view is both what he calls "phonoptic", that is, sees speech as basic, and biased towards considering (alphabetic) writing as progress on speech. Harris sets out to disabuse us of these mistaken notions. Olson (1994), like Harris, also aims to provide a more convincing alternative view of the relation between writing and speech; it is this view which we now explore in more detail.

The crucial flaw in the conception of writing as the transcription of speech, is the teleology implicit in it. A transcription account requires projecting modern conceptions of languages backwards onto the inventors of scripts, and understanding the history of scripts as one of development: each change is an advance towards an optimal transcription. In other words (Olson, 1994, p. 67): "[transcription views] assume the inventors of writing systems *already knew* about language and its structure – words, phonemes and the like, and progress came from finding ways to represent those structures unambiguously."

In keeping with the view of writing as increasing explicitness of meaning, Olson (1994) rejects the transcription account of writing, and argues instead that the direction of influence goes the other way: writing provides a model for speech. This occurs because writing brings aspects of language into consciousness which are not available to language users in an oral society. These newly-explicit aspects of language then *provide* "the concepts and categories for thinking about the structure of spoken language" – the opposite of transcription views. This applies at several levels of linguistic structure, from phonemes, to words, to sentences, to discourses. Expressed generally: "whatever is represented in the script becomes an object of knowledge or awareness to the person literate in that script" (Olson, 1994, p. 91).

The reason for the greater explicitness found in written language is that it has, more than spoken language, to stand by itself. Maintaining his earlier (1977) distinction between "utterance" and "text", Olson (1994) argues that texts fail to represent certain aspects of meaning, "such as the indications provided by a speaker and by the shared context as to how what is said is to be taken by

the listener or audience" (1994, p. 91). In other words, texts only represent the *locutionary act* or 'what is said', but do not represent *illocutionary force*, which tells us 'how they are to be taken'. This is in contrast to utterances, which portray both aspects of meaning. In Olson's words (*ibid*, p. 93):

> All oral utterances are composed of both what is said and some indication of how they are to be taken – as statement, question, command, promise or whatever; writing, capturing only what is said, represents only the former. How it is to be taken is underspecified and hence becomes the central problem in interpreting written texts and a critical problem in composing them.

This is division of labour is surely overly simplified, for written language does come with a context and indications from the speaker/author. First, consider the case of context. Texts do not appear out of the blue, but in a newspaper, a novel, on a sign in a railway station. Whoever writes the text can make use of this to get his message across. Beyond contextual indications, we can imagine all manner of cues within the text itself which indicate further aspects of illocutionary force. These cues include punctuation, formatting, even font. Punctuation often functions as an explicit marker for an illuctionary act – a full stop indicates an assertion, a question mark a question, quotation marks indicate reported speech, and so on. Furthermore, we can discern how the text is to be taken by asking: What is the surrounding text? How is the page on which it appears laid out? Are there accompanying diagrams or illustrations? What register is the text written in? In the case of a letter, is there an official letterhead? Is it hand-written or typed? Is it signed? All these cues help us to deduce aspects of illocutionary force. In sum, the means and manner of conveyance of illocutionary force may change in writing; that is not to say it is absent. Moreover, we risk understating the force of text if we suppose what it does afford in terms of transmitting intentions is only a manque attempt to convey what speech naturally does. Text rather also offers novel ways of conveying meaning.

Examples offered by Olson as illustration of the distinguishing line between speech and text fail to convince. For instance, he suggests that the illocutionary force of the statement "You're a real friend", which may be uttered sincerely or ironically, can only be distinguished in spoken language. Which of these readings is intended would be conveyed, in speech, by extralinguistic factors such as tone of voice, facial expression, and context. But a writing system "which simply transcribed what was said would capture neither tone nor context. Yet the tone and context convey part of the meaning of an utterance" (p. 91). This is a problematic illustration; the statement "You're a real friend", when written, has a context just as the spoken version would have, as described above. This context *would* provide cues, such as the response of the interlocutor, and the events leading up to the utterance, to help determine illocutionary force. Of course one might miss the cues – but the cues accompanying the spoken version could just as easily

be missed.[18] A second illustration given by Olson is Herbert Simon's assertion that the mind is a computer. "That expression, however, gives no indication of how it is to be taken" (p. 92). As Olson himself writes, however, this applies equally well to the spoken and written versions of it. Consequently, Olson's examples not only fail to illustrate his claim but actually weaken it, since he relies on the assumption that speech gives unambiguous, and stable representations of a speaker's intentions. Rather the given examples can be used to argue that illocutionary force is in fact neither unambiguous nor stably given in spoken language; it may be absent or be dynamically given.[19] Finally, transcribing dialogue is neither sole nor primary function of writing, as Olson has emphatically told us. Yet his example relies on a text "which simply transcribe[s]" an utterance. This is a recurring weakness in this part of Olson's thesis: the writing as transcription view, and inherent problems, are implicitly preserved and surface at crux points in his account.

What Olson will go on to argue is that the history of writing has been one of developing means by which to convey illocutionary force, primarily by lexicalising it. This is probably the best way to take his claim that writing does not represent illocutionary force: i.e. as an historical one, that writing did not originally capture as much illocutionary force as speech does. Early writing systems were limited in their scope to capture the complexity of a spoken linguistic interaction – indeed, the original functions of written language were record-keeping, mnemonics, and such-like, rather than transcription of spoken language – but gradually devices have been developed to improve on this function, such as a massive expansion of communication and mental state verbs, the development of punctuation devices such as question marks and exclamation marks, the differentiation into genres, and the conventions such as reported speech. Lack of illocutionary force in early written language and the development of means to represent it in historical time give his theory a depth it does not have in a synchronic reading of it. In general, Olson's work suffers for the continuous mingling of historical and developmental changes associated with the growth of literacy.

However, it is still unclear why Olson places so much emphasis on this aspect in the history of writing, as evident in his assertion that "[t]he history of literacy, in other words, is the struggle to recover what was lost in simple transcription" (1994, p. 111). Why should the lack of illocutionary force, at least when it is understood to be a feature primarily of spoken language, be so dominant an impetus in the development of writing systems? Other devices, such as graphic accompaniment – think of the illuminated manuscripts of the middle ages – con-

[18] British humour is rife with dry remarks whose wit relies on the subtlety of the ironic reading, and is all too easily missed.

[19] For instance, rising intonation in spoken English is generally taken to signal a question, but in certain sub-populations – such as some groups of young British women – it is used at the end of declarative sentences, indicating that phonological cuing of the speech act is dynamic. This phenomenon has been labelled "up-speak" (Bradford, 1997).

vey the significance of the text, but we should be careful about equating their function with that of providing 'illocutionary force', especially if we follow Olson in taking this to a be primarily a property of spoken language. In fact, as already mentioned, we risk missing novel functions of writing which are not traceable to speech functions if we focus only on the ways in which writing captures aspects of a verbal interaction. Text is a representational system on its own, not just a derivative of speech. The confusion comes because (non-pictographic) writing only *means* through its relation to spoken language, unlike, for instance, pictures, which represent their meaning directly. But that writing represents meaning *via* spoken language does not mean it represents *in the same way* as spoken language.

Writing as a model for speech

Apart from these qualms, the idea that "writing provides a model" – only not just for speech, but for language in general – seems very plausible. In fact, the challenge is not to show this is the case but to make this claim more than a platitude. There is bountiful evidence that literates have a very different conception of language to illiterates, and it is extremely likely that acquisition of literacy changes the individual's perspective on language and makes them more aware of the properties and structure of it. More specifically, research focussing on the classroom situation suggests that it does draw attention to language functioning, and thus generate metalinguistic awareness or knowledge. For instance, Dolz & Erard (1999) report on the use of reflection on language as part of teaching practice in language class, and Allal's (1999) study of text-production in the classroom also shows evidence of awareness of interpretational divergence.

Claims of difference in language perception between literate and illiterate subjects, for instance at the level of phonemic awareness, are experimentally borne out in many studies. An example is the elegant study by Read, Zhang, Nie and Ding (1986). Read and his colleagues compared the ability of two groups of Chinese speakers to segment words into consonants. The groups were differentiated by their familiarity with alternative written scripts for the same spoken language, Mandarin. The customary script for Mandarin is logographic, or character-based, but an alphabetic counterpart to it was introduced in schools in China in 1958[20]. This means that only those Chinese schooled after 1958 are literate in the alphabetic script for Mandarin. Read et al took advantage of this educational juncture in their study: one group of subjects was literate in both the alphabetic and character based script, while the other (older) group knew only the character-based script. They tested the groups on their ability to delete or add single consonants to spoken words. Spoken Chinese is ideal for this kind of testing because syllables consist of a syllabic nucleus with an optional single initial consonant and optional final 'nasal' consonant so, for example, /a/, /da/, /an/ and /dan/, are

[20]This alphabetic script is known as *Hanyu pinyin*.

all possible syllables. Subjects were asked to delete or add a consonant from the initial or final position of a syllable. A response was deemed correct if the syllable nucleus remained the same as in the presented syllable, as judged by three phonetically trained transcribers. The results show a significant difference between alphabetic and non-alphabetic groups, and a significant difference between word and non-word target items, with no significant interaction between the two effects. That is, alphabetic literates were significantly better than their non-literate counterparts at the task, (83% vs 21% correct on nonword targets, 93% versus 37% on word targets), and both groups found the conditions with word target items significantly easier than those with nonword targets. Effects such as reported here are not exceptional: similar results have been found in a study which compared literate and illiterate Portuguese speakers (Morais et al, 1979), and several studies have found correlations between reading level and segmentation ability in children (see Ehri, 2000, for an overview).

At the level of words and sentences, a review of the research on this topic is to be found in, for instance, Kurvers (2002)[21] (but see also Olson 1994, and Scribner & Cole, 1981). Kurvers conducted her own research to investigate illiterates' awareness of language at the phonological level, the level of single words, and the level of sentences and texts, as well as their perception of scripts as such. She carried out three kinds of tasks to investigate language awareness at the level of sentence and above: sentence imitation; syllogisms[22]; and story-telling. On all these tasks she found the illiterate subjects to perform significantly worse than literate subjects. Here we report in more detail on the sentence imitation. In this task simple sentences were often repeated exactly; only 6 of the 24 illiterate subjects had difficulty with simple sentences. Repetition failed more frequently when sentences contained embedded clauses or two simple sentences joined by a connective like 'because', or 'but'. A paraphrase was often given; this usually maintained the meaning but often omitted words, such as the connective, and was often a more conventional formulation than the original one. This is reflected in the average word length of the answers across groups: 9.79 for illiterates versus 12.78 for literates. The absolute number of non-repetitions are 13 in the group of 15 literates; 60 in the group of 24 illiterates. Kurvers codes non-exact repetitions more precisely along these lines:

1. **words not repeated:** 'She wanted to go home' instead of 'Now she really wanted to go home'; omission of function words like 'but', 'because' (58.3% of mistakes)
2. **paraphrases:** 'It rains very often in the Netherlands' instead of 'In the Netherlands it rains very often' (29.9%);

[21] This is a Dutch book; all translations are my own.

[22] The fact that Kurvers includes syllogism tasks in her chapter on 'knowledge of language at sentence and text level' indicates just how much she sees the task to be indexing linguistic rather than cognitive skills!

3. **reactions to content:** 'Yes, it's eleven o'clock' when asked to repeat 'Do you know what the time is?'; 'That is your pen' instead of 'That is my pen' (8.3%)
4. **other reactions:** 'Come aeroplane?' when asked to repeat 'Did you come to the Netherlands by aeroplane?' (3.3%)

These results are very much in line with the puzzling recall data recorded in Luria (1976) and Scribner (1997). Both Luria and Scribner interleaved a recall task with the various reasoning tasks when interviewing their subjects. This was motivated by the realization that subjects' interpretation of the premises plays a role in what conclusions they will draw. One means of accessing the subjects' understanding would be to simply ask the subject to repeat the premises just heard. For example, Luria reports attempts at repetition of the following syllogism: "Precious metals do not rust. Gold is a precious metal. Does it rust or not?" (1976, p. 104–6). The following responses were recorded:

> Kurb., 18 yrs, peasant from remote region, illiterate.
> S: Do precious metals rust or not? Does gold rust or not?
>
> Gal., peasant from remote region, almost illiterate.
> S: Precious money rusts ... there was something else, I forget. (1)
> S: Do precious metals rust or not?(2)
>
> Iganberdy, 34 yrs, Kirghiz, illiterate.
> S: Precious metal rusts. Precious gold rusts. (1)
> S: Does precious gold rust or not? (2)
> S: Do precious metals rust or not? Does precious gold rust or not? (3)

We return to the significance of this data in more detail later on in the chapter.

In sum, studies like those mentioned above make evident that literates and illiterates have a different perception and conception of language. The task is to specify how this impacts on cognitive behaviour such as evidenced in reasoning tasks. Considering Kurvers' data above, the most obvious explanation of reasoning performance would be that unschooled subjects do not, or cannot, pay attention to exact wording, and this somehow influences their reasoning behaviour. Olson's explanation is of this form: he proposes that unschooled subjects do not make the distinction between 'what is said' and 'what is meant', since this distinction is a product of literacy. In the following section we assess the evidence for this claim and its relevance in reasoning tasks.

3.5.2 Saying *versus* meaning

Donaldson (1978) writes of a child learning to speak as becoming more sensitive to "sheer linguistic form" in determining the meaning of utterances. Olson

(1994) echoes this in proposing that a primary effect of literacy is awareness of the difference between the text and its interpretation, thus enabling the distinction between 'what is said' and 'what is meant'. The ability to distinguish what was said from what was meant is a result of the model of language engendered by writing, according to Olson, because it "makes interpretative divergence possible", by making the process of interpretation more explicit and thus the concept of interpretation available. It is easy to imagine how distinguishing 'what is said' from 'what was meant' could result from school activities. Firstly, and perhaps foremost, is the very basic observation that learning to read is a long and arduous process. The child's attention is focused for many months on squiggles on the page, which gradually become letters, words, and sentences. Exactness is of the essence. In this sense learning to read is partly a process of learning to pay attention to exact wording: meaning can only be gleaned by doing so. This sharpens the notion of 'what is said'. Secondly (or perhaps 1a), written text is invariant, candidates for intended meaning (both in what has been read and in what has been written by the child) are checked against it. Especially being corrected when the process is unsuccessful would be relevant here. In other words, the comparative and evaluative process should also serve to focus attention on exact wording. Thirdly, and here we employ a circumscribed version of Olson's earlier mentioned claim, the realization of the very possibility of differentiating between intended meaning and expression used to communicate that meaning might be prompted by learning to read. The suggestion here is that in written text the speaker's intention is apparent only through their words (and perhaps accompanying pictures), which might result in a less determinate or apparent speaker intention than in the case of spoken language. In the latter case extra-linguistic clues as to speaker intention: gesture, tone of voice, facial expression, context, and so on, work in concert with the uttered words to communicate speaker intention. The multi-channel nature of this process would presumably leave less doubt than the single- (at most dual-) channel process that reading is[23].

There are several studies which address the development of the discriminatory ability and report similar results (Robinson, Goelman and Olson, 1983; Torrance and Olson, 1987; Beal and Flavell 1984, Bonitatibus 1988, Lee, Torrance and Olson 2001): that is, that children under six years of age are strongly inclined to accept 'what was meant' as 'what was said', and become less so as they pass the six-year mark. This is just when the children are becoming readers at school. The results suggest that learning to read and write play a role in acquiring this explicit knowledge. But frustratingly, detailed information about what school activities their subjects have experience of, or of their literacy levels, is not given in any of these studies. Torrance and Olson (1987) do address this issue; if only to say that "it is a mistake to assume that [these interpretive distinctions] are strictly associated with the practices of learning to read and write." They suggest

[23]Being read to is again different.

that ambient affects of a literate culture might be just as effective in catalyzing the development of the distinction, such as how the parents talk to the child and the experience with books the child may have prior to reading themselves. Moreover, the direction of influence goes both ways: being able to distinguish a text from its interpretations is a key part of being a competent reader later on. All we can conclude at this point is that distinguishing intended from some as yet unspecified literal meaning is an ability which develops from the time when children are attending school and learning to read and write.

However, there is a caveat to be made, regarding song and poem contexts. Lee, Torrance and Olson (2001) explored the affect of discourse genre on awareness of the verbatim/paraphrase distinction, and found that children as young as 3 yrs recognize when a response is not an exact reproduction of the stimuli – but only in the context of a nursery rhyme. Also, given the fact that songs and poems rely on exact wording and are surely not restricted to literate cultures, we may conclude that the ability to recall exact wording is not especially a literate feat, but an attentional bias which is triggered in contexts of rhymes, songs and so on. Again, it's not just the skill itself but also the context of practice which matters.

A conservative hypothesis would be that attention to interpretational divergence is engendered by specific literacy practices in certain contexts. But here Olson is again prone to overstatement: the result of writing "is a kind of interpretational anarchy, each interpretation being taken by some individual or group as what was said." And again, I think that Olson's thesis has merit when it is applied to a restricted domain, such as perhaps the history of texts, but it does not necessarily apply to individuals. At the level of the individual, attention to interpretational divergence is heavily dependent on the context of acquisition of literacy skills. This is demonstrated by for instance the study by Scribner and Cole discussed above. But even if the context of acquisition does foster awareness of interpretation this doesn't mean that the man on the street is concerned with or even aware of the difference between understanding and interpretation when reading the sports page. In certain circumstances when it becomes pertinent – a typo in the text? a reasoning task? – it might be available to him as a means to resolve misunderstanding or discrepancy. In the following we explore how this might operate in such a situation.

The literal meaning hypothesis

As we've seen, a primary tenet of the reasoning task paradigm is that people reason, or rather *should* reason, on the basis, and solely on the basis, of the presented material. 'Solely on the basis of the given premises' can be taken to mean interpreting them strictly, that is, literally, verbatim, word for word, and contrasted with a 'paraphrase reading'.[24] The question is: how does the ability to distin-

[24]It is interesting to observe the terms 'literally', 'reading', here being used in a metaphorical sense, are terminology from literate practices.

guish verbatim phrasings feature in reasoning tasks? What consequences could a non-literal reading of the premises have? As we saw in Kurvers' and Luria's data, illiterate subjects were had difficulty producing verbatim repetitions of sentences or premise sets. This offers a new vantage point from which to understand many of the responses in the reasoning tasks proper. From this view, the instruction to reason solely on the basis of the given premises changes meaning. The given premises are converted into a reconstruction of what was meant, and it is *this* which serves as a basis for reasoning.

How exactly might literacy help with verbatim recollection? There are two types of hypotheses available here. Firstly, an undifferentiated concept of literal meaning can serve as an explanation for this phenomena; so that what was said is conflated with what was meant (or posited to be meant). Without concentrating merely on exact wording used, as distinct from what might be meant by it, the subject would have to rely on similarities (however scanty) between presented material and familiar forms of utterances. Interestingly, none of Luria's subjects reproduced the premises in the form: two assertions followed by a question; instead the majority returned either a pair of questions or a pair of assertions. Here the lack of relevant discourse 'genre', presumably gained in a school environment, could be crucial, because it prevents the subject from figuring out 'what is meant' but at the macrolevel – what the experimenter intends with his task. Familiarity with story sums, for instance, opens the possibility of reconstructing the experimenters intention as some such challenge. Secondly, being able to recall exact wording might be facilitated by being able to convert spoken language into a written representation, which in turn might be helpful in holding novel sentences in working memory and reproducing them: so that not only do the subjects not pay attention to wording, but that, even if they did, it wouldn't help them because they wouldn't have the computational aid of a visual representation of the sentences to hold in their mind, a suitable format to put the exact words into, if you like. The hypothesis here would be that literacy enables recall of exact words by providing a means to keep them in mind, namely in a textual form, writing as a mental mnemonic if you like. Recall of exact wording in song and poem contexts speaks against a strong version of this hypothesis; nevertheless literacy might play a facilitating role.

In the context of the reasoning task, non-reliance on 'literal meaning' might make it easier to adjust paraphrase readings according to what makes more sense in the situation, as it were. This is illustrated in the following excerpt:

> Nonkululeko, 56 yrs, illiterate. Presented material is
> 'All people who own houses pay house tax. Sabelo does not pay house tax.
> Does he own a house?'
>
> S: He doesn't have a house if he's not paying.
> E: And now suppose that none of the people in Cape Town pay house tax. Do they own houses or not?
> S: They have houses.

E: Why?
S: They can have houses because there are places where you don't pay tax, like the squatter camps.
E: So they can have houses and not pay?
S: They may, they can live at the squatter camps.

In this case the subject was initially using the first premise as a basis for reasoning, as can be seen in her first turn. Moreover, she would have to be reading it as a strict universal to be able to generate a conclusion about a random individual, Sabelo. But when reasoning with it as a strict universal would entail an extraordinary conclusion, namely that none of the people in Cape Town have houses, she can be seen to revise her interpretation to a generic reading, and thus avoids generating that strange conclusion. We can see this as an example of the tendency to attribute 'awry-ness' to language before the world, to anticipate Fillenbaum's formulation in the next section.

One might find the literal meaning hypothesis appealing without wanting to go so far as to identify logical reasoning ability with the ability to distinguish literal meaning in a text. Olson, however, does go so far. In his 1977 article, he claims that (1977, p. 274–5): "logical development in a literate culture involves learning to apply logical operations to the sentence meaning rather than to the assimilated or interpreted or assumed speaker's meaning. Development consists of learning to confine interpretation to the meaning explicitly represented in the text and to draw inferences exclusively from that formal but restricted interpretation." Although it is unclear what 'logical development' means here, this sounds quite uncannily like the argument in Goody & Watt (1963) which Olson so vehemently sets himself against. Olson maintains a variation of the claim throughout his subsequent work, where he says (1993, p. 177, based on arguments in his 1994 book): "Logic and literal meaning are, therefore, completely interdependent and both dependent on properties of language rendered explicit by writing." The logical (if you'll pardon my non-literal use of the term) outcome of this claim is that unschooled subjects are not in the possession of logical faculties. In fact, Olson is intending logic in a very specific sense; logic in his 1993 paper is reduced to a rather vague notion of 'logical proof', which relies on a distinction between the literal and metaphorical (p. 172). This notion of logic is not only specific but also circular, since later on he defines literal meaning in terms of logic (p. 177): "this constitutes a working definition of literal meaning; literal meaning of a statement is that meaning for which strict logical rules apply. ... ordinary language rarely achieves such logical purity". As the next chapter shows, this may well be the case under certain descriptions but the point is whether this is the correct notion of 'logic' being applied. The problem arises when logical reasoning ability is judged solely on the basis of the tasks such as reported in the previous chapter, which apparently do rely on some literate tendencies to be performed correctly. These conceptual confusion are not the only problems with the literal meaning hypothesis described thus far, as the next section shows.

Problems with the literal meaning hypothesis

There are two kinds of challenges to such a strong formulation of the literal meaning. The first stems from empirical work which suggests that the distinction between 'what is said' and 'what is meant' is neither exclusive to literate cultures nor widespread within them. Indeed, it proves under certain circumstances to be a very recalcitrant notion, as will become evident when we report on Fillenbaum's (1978) study of strange conditional constructions. Moreover, it is not clear that the recall data, adduced to show that illiterates cannot or do not access literal meaning because they do not produce verbatim reproductions of sentences, actually does the job in showing this.

The second type of challenge lies within the notion of literal meaning itself. This is addressed in the next chapter, where we address in greater detail the problems with assuming that logical subjects reason solely on the basis of the given premises, i.e. on the basis of the literal meaning of the premises.

This section is devoted to an evaluation of some empirical findings which are relevant for the literal meaning thesis as stated thus far. The bulk of the empirical research reported by Olson (including many he has done with colleagues) are developmental studies, so the subjects are children of school-going age. Apart from the possibility that the relevant distinction might be pertinent in exactly the phase of learning to read and discounted after that, these studies suffer from a potentially disabling confound between general cognitive development, including (spoken) language acquisition, and acquisition of literacy skills. Eradicating this confound may well be nigh impossible to avoid in such studies, but Olson should at least attempt to offset it by comparing results to those for populations in which these two factors come apart. He relies rather heavily on anecdotal anthropological evidence to achieve this but would strengthen his argument by exploiting cross-cultural research further.

Besides this, there is anthropological evidence which goes against his claim that awareness of interpretation and attendant ramifications are a uniquely literate achievement. For instance, Feldman (1991) investigates to what extent oral genres separate text and interpretation. Indeed, as she points out, there will be certain events characterised by their form of talk (Feldman calls them 'linguistic jobs') in every culture – story-telling, songs, conflict resolution, contract negotiation, ceremonial and ritual talk. We tend to focus on conversation when we think of oral genres, perhaps because written genres are more explicitly differentiated, or because we're in a literate society where written genres may have even taken over some oral forms. But the case can also be made that oral forms be "more varied and better defined in [oral] cultures because there is no competing written language". Feldman argues that artful genres which stress the difference between text and interpretation are very much present in oral societies. She cites the example of 'kiyori', a type of political poetry used by the Wana tribe of Indonesia. A kiyori has a set form and fixed pitch contour, and makes use of special am-

biguous and metaphoric language. When delivered it is usually repeated several times until the exact wording is stable. It might also be repeated by the person to whom it is delivered. Then it is replied to, either in the form of another kiyori or in a general conversation, in which interpretations of the kiyori are discussed. So kiyori are an oral genre in which text is first fixed, and subsequently interpreted. With this example Feldman wants to illustrate that oral forms can also be seen to make the distinction between a text and its interpretation. This in turn undermines what she calls the 'general claim' of literacy theorists, that writing is necessary for the development of consciousness of this distinction.

Upon closer analysis of the recall data mentioned in the previous section, it becomes apparent that also here the conclusion that illiterate subjects do not have the concept of literal meaning is too hasty. One major reason to believe so is the nature of the errors. Luria provides transcripts but does not analyse the errors, but Scribner (1997) does provide an analysis of the errors she found in her recall data. A "principal form of error" was the omission of a premise, according to Scribner. Both examples of omissions that she offers concern premises which could plausibly be taken as common knowledge. This might make them not worth repeating, to the subject, because they go without saying, as it were. For instance, with the premise "Mr Ukatu's store is in Kpelleland", presumably Ukatu is a Kpelle name, so we could take for granted that Mr Ukatu's store is in Kpelleland. It's even more plausible to consider that premises such as "all the people we know are in Liberia" are considered common knowledge, and thus doesn't bear repeating in the subject's eyes, who has most likely never been outside Liberia and had minimal contact with foreigners (the experimenters were also local Kpelle people). Perhaps most striking is the consistent omission of these kinds of premises in Luria's data too. Recollect that Luria reports attempts at reproduction of the following syllogism: "Precious metals do not rust. Gold is a precious metal. Does it rust or not?" (1976, p. 104). Only one of the five subjects quoted by Luria included the premise "Gold is a precious metal" in their reproduction of the premise. This is again a premise which could conceivably be considered so well-known that it does not bear repeating. Indeed, we can't tell from this data whether these 'common knowledge' premises are forgotten or rather taken for granted. At most the claim can be made that unschooled subjects are not inclined to pay attention to exact wording in the context of reasoning tasks. The claim that unschooled subjects can't pay attention to exact wording in general seems dubious anyway; recall for songs and poems speaks against it.

Literal meaning is contrasted by Olson with intended reading, in which one goes after what was *meant*, what was intended to be communicated by the sentences. That this latter method of interpreting is dominant, even among literate populations, is supported by several sources. One such source is Fillenbaum's 1978 study of conditional constructions.

In a paraphrasing task Fillenbaum identified what he called 'pragmatic nor-

malization'. Subjects were presented with a set of sentences and asked to paraphrase them. Instructions emphasized that the task was to preserve given meaning, that is, the subject was to (Fillenbaum, 1978, p. 187):

> paraphrase or rephrase "the sentences as accurately as you can conserving meaning as completely as possible" and that the subject was not "to improve the sentences or make them more sensible, but to *paraphrase* them, rewording each in a way that captures its meaning as accurate as possible."

Despite these exhortations[25], subjects were inclined to reorder any sentence which violated a somehow natural or expected relation (percentages ranged from 70% in conjunctive sentences violating an entailment, to 54% in 'perverse' conditional threats)[26]. This is what Fillenbaum terms 'pragmatic normalization'. An example illustrates the phenomenon (p.187):

> "Clean up the mess or I won't report you"

becomes

> "If you don't clean up the mess I'll report you".

The latter situation is easier to envision than the former, but it also does not take a massive stretch of the imagination to envision a situation for which the literal reading is fitting. All we need is a situation in which being reported is a desirable outcome – as in, being reported as a potential employee of the month. So there is a sensible literal reading available for the sentence[27]; nevertheless, subjects tended to paraphrase non-literally in these cases. The sentence structure itself is in the above example not unconventional, as can be seen in the following variation:

> "Eat your vegetables or you won't get any pudding".

Perhaps most surprising in this study is the responses of subjects after the fact. Once the subject had finished the task proper he was instructed to go through his paraphrases, and asked (*ibid*, p.187)

> "about the remaining differences IF ANY (sic)" between each sentence and its paraphrase, and that if he did "see some shred of difference" to say what seemed to be involved.

[25]Or because of them? That the sentences are *not* already clear, is presupposed by the suggestion that the sentences can be 'improved' or 'made more sensible'. This only makes the task pragmatically stranger.

[26]Note that literal meaning and plausibly intended meaning need to be distinct for the tendency towards one or the other to become apparent. Thats why only the 'disordered' sentences, that is, those for which this distinction holds, are relevant.

[27]Fillenbaum unfortunately does not specify for which sentences normalization was most common. There are other sample sentences – such as that mentioned above 'Don't print that or I won't sue you' – for which a literal meaning is more difficult to find. However, you can do it: in the above-mentioned, you can simply take out the unnatural negation, rephrase, and stick the negation back in.

The experimenter then worked through each item individually with the subject, asking "is there any difference?, and "what sort of difference?" if the subject answered affirmatively. Also, if the subject made changes to his paraphrases at this stage he was asked to explain the original answer. The results from this stage are mixed: according to the type of sentence the number of subjects who detected a difference of meaning, as against those who overlooked it, varied. But both groups remain substantial: for example, the figures for those who detected difference *vs* those who didn't are respectively 27% and 43%, for conjunctions, and 20% and 34% for conditional threats (types as described above) In the cases where subjects did detect a difference in meaning, and so were asked to explain their initial choice of paraphrase, the results are less equivocal. Fillenbaum reports that the comments of the majority of the subjects could be classified along five lines: that is, claims of the form

(a) that the paraphrases made things clear and more sensible
(b) that the paraphrases put things into natural order
(c) that the original sentences violated expectancies
(d) that the original sentences were illogical
(e) that they knew what the original sentences were trying to say so they said it.

Further, seven of the remaining subjects said they had "misread or read incorrectly" the original sentences (p. 190), which suggests they did not see the literal meaning in a first reading.

The picture we get is of subjects who often can see the difference between a literal and non-literal paraphrase, but who defend the latter on the basis that it was a more successful formulation of what was intended to be said. In Fillenbaum's words: "what is taken to be awry or extraordinary is not the world but the linguistic account of it. ... a difference detected is not so much one between descriptions of two different sorts of events but as one between two different descriptions of the same event, with the paraphrase expressing properly what is *intended and badly expressed* by the original sentence" (p. 190, his emphasis). That this occurs when university students are in a laboratory task which is very clearly about paraphrasing only strengthens the import of the findings, because if there is any situation which piques attention to literal meaning, this is one.

Fillenbaum's study illustrates that preserving a literally faithful but pragmatically strange complex sentence meaning is a difficult task for highly literate individuals; nonetheless in many cases they are able, if only when heavily prompted, to distinguish between a literal and non-literal paraphrase of a sentence. There are other studies reporting such 'corrective' reading of pragmatically strange sentences (Garnham & Oakhill, 1987 for instance), but others showing that in the case of simple passives sentences such 'corrective' reading doesn't happen (Macwhinney, Bates and Kliegl, 1984). The task settings under which these results were elicited differed: the former was a paraphrase task; the latter asked only for thematic role information ("who's the actor?" in sentences like

The eraser bit the turtle). So the phenomenon is related to the specific task setting, including the emphasis placed on a coherent holistic interpretation of the presented material, as well as the complexity of the presented materials.

The empirical findings presented above suggest that the claim that literate subjects are uniquely and unproblematically able to distinguish a text from its interpretation, or literal meaning from intended meaning, is too strongly formulated. The literal meaning hypothesis as stated in Olson (1993, 1994) needs to be scrutinised.

Qualifying the literal meaning hypothesis

Fillenbaum's sentences contained unexpected negations (for instance of a typically negative outcome in a threat, as above) and unusual use of other connectives in combination with these. The sentences are semantically/pragmatically incongruent. In this case paraphrasing can be seen as *correcting* this incongruency, in order to make the sentence more sensible for comprehension. In other situations, such as the say/mean studies conducted with young children (Robinson, Goelman and Olson, 1983, Bonitatibus 1988) the original sentence is ambiguous between two potential referents, resulting in a communicative failure. So a paraphrase here serves to *disambiguate*, thereby enabling the communicative task (i.e. indicate the right referent) to succeed. In the recall studies reported by Luria, the 'paraphrases' offered by the subjects (although the experimenter asked for repetitions) often omit information which can be considered general. So the paraphrases here can at least partially be described as '*repackaging*' the information contained in the original sentences, by backgrounding common knowledge and highlighting new/contentious information. When sentences are idiomatic, metaphoric or ironic, appropriate paraphrases would supply a new sentence which *elucidates* the non-literal meaning, for example 'his job is a jail' rephrased as 'his job is very demanding and inflexible' or something to that effect. Here paraphrases tell us about the use of a particular stock construction within a language community (such as that 'what's up?' is used as a greeting among American English speakers). Similarly with rhetoric devices, such as rhetorical questions, which are not to be understood as questions at all. An appropriate paraphrase of the question 'How I am supposed to explain that?', for instance, might (depending on the context) be the assertion 'I don't know how to explain that'.

The variety of paraphrasing possibilities indicates the different relations between literal and non-literal meanings, and that taking a sentence literally might involve different counteractions depending on the context, the goal of the interaction, and the common usage of the sentence. In fact, to prefigure what will be covered in the coming chapter, literal meaning, as it functions in Olson's theory, comes as part of a package, a specific theory of language which assumes a semantic core of meaning is determined prior to the input of pragmatic factors. This semantic core, however, turns out to be theory and context-dependent notion,

at best an *abstraction* over a range of contextually determined meanings. Once literal meaning loses its status as semantic core, it also loses its place in theories of logical reasoning. In the following chapter I address the sense in which the notion of literal meaning can be applied in reasoning tasks, and show that there is no simple connection between literal meaning and logic.

3.6 Summary, conclusions and outlook

As the first half of this chapter illustrated, it is difficult to say anything interesting *and* general about the consequences of literacy on the individual, further than that they impact on awareness and knowledge of language. There are various reasons for this difficulty; amongst others, these are that existing historical theories concentrate on whole-scale differences in societies, and that empirical research with individual subjects has been inconsistently interpreted because of a lacking standard approach to categorization of responses. For instance, we saw that Greenfield's data on categorization from illiterate subjects, although fitting Luria's criterion for 'literate' thinking, was judged according to Greenfield's own criterion and therefore interpreted as having similar significance to Luria's, namely that illiterate people are limited to situational or context-bound thought.

Fresh empirical research into the effects of literacy might well have been inhibited by the reception of Scribner and Cole's (1981) study which aimed to separately test the influence of schooling from that of literacy. The general perception was that Scribner and Cole's study found schooling to be responsible for more cognitive change than literacy. However, examination of their findings indicated that their results were by no means as unequivocal as this conclusion suggests, since they only tested a very restricted form of literacy. More generally, literacy and schooling are not to be readily teased apart, either practically or conceptually. Literacy is almost always acquired with schooling, bringing a whole complex of values, norms and practices with it, and the terms literacy and schooling have been differentially applied to aspects of this complex.

The task is to index consequences of literacy to contexts of acquisition and use, and especially to the social norms and conventions about language and cognition active in these contexts. This entails specifying which activities, conventions, forms of knowledge, etc. engendered by the schooling environment are operative for performance in the tasks at hand.

With this in mind, we pursued the hypothesis, found in the work of David Olson, that performance on logical reasoning tasks is a product of the literate ability to distill a 'sentence meaning' as distinct from 'speaker meaning', or 'what was said' as distinct from 'what was meant'. We reviewed the evidence that this is a peculiarly literate achievement and argued that it much be understood in a qualified form and that it the ability should not be too quickly associated with logical behaviour.

3.6.1 Experimental suggestions

Conceptualising literacy not as a unitary skill but as a host of skills leading into differential engagement with the world of text opens up new avenues for experimental research into the effects of literacy. It makes clear that literacy can be segmented into different levels and that at each stage we may look for cognitive change. As the inventive work of Dabrowska shows, there are huge differences between subjects within the literate spectrum. This avenue has also been begun to be explored by Keith Stanovich and his colleagues. For instance, Stanovich (1993) presents evidence for print recognition (briefly, how many book and publication titles a subject recognises) as a measure of 'secondary' literacy and argues that using such a factor provides a very efficient means to get a handle on the cognitive correlates of literacy while staying within literate populations. Additionally, different educational environments also engender different approaches to cognitive tasks, as Miller's work on mathematics learning in China and the United States reveals (Miller, Kelly & Zhou, 2005). Even within undergraduate subject in different disciplines (humanities *vs* sciences) we might expect to see differences in response on the kinds of task reported on here. Paired with further experimentation, ethnological studies of literate environments are needed. Case studies of particular interactions with texts, such as reported in Camps and Milian (1999), enable us to get behind the platitudes of saying that literacy impacts on language perception, and identify mechanisms fundamentally shaping cognitive behaviour.

Further research with unschooled subjects would benefit by being coupled with interrogation of the language perception of such subjects, as Kurvers (2002) did with her subjects. However, suitable tasks for this are not readily to hand and here borrowing from the anthropological literature might prove fruitful – for instance Feldman's (1993) extended exploration of the role of interpretative processes in different 'linguistic jobs' across cultures could yield as yet unexplored routes of access to such notions in further empirical work. This would enable experimenters to get a better grip on the possible interpretations their subjects are taking and thus have a more accurate semantic analysis of the task from the perspective of such subjects.

Someone who is sceptical of the whole linguistic basis of the tasks done here might want to see non-linguistic tasks developed, to tease out the consequences of literacy on cognition apart from those which we access via language. This seems to me to be a very difficult task, because, firstly, representational conventions may be literate although not textual, as we saw with the Raven Standard Progressive Matrices task (Kurvers, 2002) which is intended as non-linguistic and yet which stymies unschooled subjects and fails in its goal of measuring non-verbal intelligence in such a population. Indeed, the whole paradigm of an experimental task is closely linked to an educational testing paradigm, as such it might be inescapably bound with literacy acquisition (in school contexts). Moreover, some tasks are inescapably linguistic; I wouldn't know how to design a non-linguistic

equivalent of complex arguments, such as full syllogisms, with two quantified premises. Granted, these are also difficult to test with purely verbal interactions, but in combination with physical objects – such as in Haan's ball-in-the-box task, discussed in Chapter 2 – one can plausibly get much further than by aiming for a purely non-linguistic paradigm, insofar that can be achieved at all.

Chapter 4

What's logic got to do with it?

4.1 Introduction

> What counts as human rationality: reasoning processes that embody content-independent formal theories, such as propositional logic, or reasoning processes that are well-designed for solving adaptive problems?
>
> (Gigerenzer & Hug, 1992, p. 127)

This citation opens an article by Gigerenzer and Hug in which they go on to argue for the second alternative. Their position is representative of a widespread tendency to dismiss the role of logic in the psychology of reasoning: Fiddick, Cosmides and Tooby (2000) forthrightly claim that "[s]ocial contract theory . . . has no commitment to logical formulae", while Evans' (2002) assessment is that "[f]ew reasoning researchers still believe that logic is an appropriate normative system for most human reasoning" and "researchers have progressively de-emphasized the importance of logic in human reasoning". In large part this is due to the negative conclusions drawn from empirical findings in the field, such as those reported in the previous chapters, and which have been taken to show systematic non-logical behaviour on reasoning tasks. The feeling is this: 'Well, subjects clearly aren't using logic, therefore it has no place in our theories of reasoning'.

Here I set out not so much to root for the first option in the choice offered by Gigerenzer and Hug as to reassess the sense of such an 'either-or' understanding of logic's place in psychology. The professed aim of this dissertation is to contribute to a semantically grounded research programme in the psychology of reasoning, recently campaigned for by Keith Stenning and Michiel van Lambalgen (2001, 2004, 2005, 2008), but also advocated by Guy Politzer (1986, 2004), in Thompson (2000) and incipient as early as Henle (1962). In the first three chapters I dissected the performance of variably schooled subjects and argued that, in the case of less schooled subjects, their problems in interpreting various premise sets in reasoning

tasks plays a large part in explaining their difficulty in the task. I identified semantic factors for which they show concern, and proposed a reinterpretation of their responses as a function of everyday language usage. Implicit and sometimes explicit in the approach was the sense that their responses have been too quickly characterised as non-logical. In this chapter the aim is to make this sense explicit. Note that this is not to say that we expect to re-interpret all subjects as being logical all the time. The claim is rather that judgements of non-logicality have been wrongly directed, because of an understanding of logic which is by turns too narrow, too broad, and almost always too vague. I will show that views about logic's relation to reasoning employed by reasoning researchers in large part rely on implicit yet outdated theoretical assumptions about the construction of meaning.

So firstly, I claim that empirical findings have been interpreted in the context of a number of theses which form background assumptions about meaning and its working in reasoning tasks, and that these theses should be challenged. Secondly, I would like to make the stronger claim that we are not in a position to dismiss logic's role in reasoning and that we should in fact expect its role to become only clearer as we understand more about the semantic strategies that subjects employ. These claims will be developed in the context of a critical discussion of Stenning and van Lambalgen's recent exposition of the relation between logic and psychology, to be found in their forthcoming book, *Human Reasoning and Cognitive Science* (2008). Most of the analysis presented in the previous chapters fits well into Stenning and van Lambalgen's programme, but there are also important points of divergence. As we go along, I describe the fit and the divergence of the current account relative to Stenning and van Lambalgen's work.

The plan of the chapter is as follows. I start by identifying the theses alluded to above. These are then examined one by one, starting with the assumption that logical form is betrayed by grammatical form, then addressing the dichotomy between logical and non-logical elements of language. This leads us to the central argument of the chapter, namely that it is untenable to equate logical reasoning to reasoning 'solely on the basis of the given premises'. Finally, I consider the implications of this for the normative status of logic in theories of reasoning and rationality.

4.2 Logic *vs.* logic

The rejection of a role for logic in the study of reasoning is premised on a number of interlinked assumptions about the relation between logic, natural language and reasoning. At least the main assumptions are listed below. For a sample of papers which express or embody all or most of them see Wason (1968a), Braine (1978, 1990), Griggs & Cox (1982), Cosmides (1989), Gigerenzer & Hug (1992), Newstead et al. (1992), Johnson-Laird (2001); for each thesis I try to offer an

illustrative citation from one of these publications.

1. The grammatical form of a natural language sentence directly corresponds to the logical form of the proposition expressed by that sentence (Wason, 1968a, Cosmides 1989), viz (Wason, 1968a, p. 273):

   > This investigation is concerned with the difficulty of making a particular type of inference from conditional sentences, statements of material implication of the form, "if P then Q" (P⊃Q).

   Sometimes, there is even a dichotomy made between syntax and semantics, as if logic was *only* to be found in syntax. This leads to confused formulations such as (Johnson-Laird, 2001, p. 434):

   > Many theorists suppose that the mind constructs syntactic representations of the logical form of assertions and applies the rules of a formal logic to them. There is another possibility: reasoners could rely instead on their grasp of meaning, their general knowledge, and principles akin to those for the semantics of a logic.

   This is a false dichotomy, because validity can be determined either via syntactic operations or by model checking i.e. semantically, and – crucially – these two approaches are equivalent in a sound and complete logic (which both propositional and predicate logic are).

2. Logic is content-blind and domain-independent, often termed *topic-neutral*; this in turn rests on the division of natural language sentences into logical and non-logical elements; for instance, Cosmides claims with respect to variations on the Wason selection task that

   > The correct formal logic response is P *&* *not-*Q, regardless of content. (Cosmides, 1989, p. 199)

   where the only specification of logical formalisms given in this article is to label P and Q 'logical categories'.

3. Reasoning logically means reasoning without taking world or general knowledge into account; logical inferences are those drawn *solely on the basis of the given premises* (Braine, 1978; Johnson-Laird, 2001; Newstead et al, 1992). The following excerpt from Braine (1978, p.2) illustrates:

   > Practical (i.e., everyday) reasoning uses all the information at a person's disposal, whereas formal reasoning is concerned with whether conclusions follow conclusions follow from premises. ... formal reasoning makes two demands not made in everyday reasoning:

> (a) Reasoners must compartmentalize information (i.e. *restrict the information used to that contained in the premises*) and (b) they must take a special attitude in comprehending the premises, by attempting to discover *the minimum commitments of the premises as they are worded*.

4. The base case (in a sense we'll make exact later on) for logic is classical propositional logic, or even logic *is* just classical propositional logic. This is not always made explicit; in many papers reference is simply made to 'logic' (e.g. Norenzayan et al, 2002) or to 'formal logic' (e.g. Cosmides, 1989[1]). In papers where a specific logical system is mentioned, such as propositional logic, the full machinery is not described (Johnson-Laird, 2001; O'Brien et al 1994). The fact that propositional logic doesn't to provide a formalism for quantified arguments doesn't seem to bother most researchers.

5. The normative standards of rationality are given by logic, or, as Evans et al (1993) pithily put it, "rationality=logicality".

This complex of theses, taken together with the empirical results in psychology of reasoning, do indeed form a formidable barrier to sensibly maintaining that experimental subjects are employing logic in their reasoning. We will see, however, that the above theses are, at best, in need of qualification, and at worst, untenable, mostly because they rely on language theoretical frameworks which have been successfully challenged. On closer analysis of the workings of language, it becomes apparent that casual reference to 'logic', 'form', and 'the information contained in the premises' as evident in the above research, is precisely part of the reason we are led to false dichotomies such as that suggested by Gigerenzer & Hug (1992)[2]. In the following sections the tenability of the each of the above theses is critically evaluated, and the resulting view is in turn related to Stenning and van Lambalgen's counter-suggestions to the theses.

But before we start, it is helpful to get an idea of a logician's contrasting perspective on logic. For this, see the opening comments of the widely-used introductory textbook, *Logic, Language, and Meaning, Volume 1* (Gamut, 1991, pp. 7 – 8):

> Whether logic is seen as the science of reasoning or the science of relationships between meanings, either way there is no such thing as a

[1] In numerous places reference is made to the predictions of 'formal logic' and 'logical structure': for instance on page 192 "*The logical structures* of these two Wason selection tasks are identical"; on page 197: "*From the point of view of formal logic*, only the combination on the same card of a true antecedent (P) with a false consequent (*not-Q*) can falsify a conditional rule."

[2] Gigerenzer & Hug (1992) do go some way to dismantling the dichotomy in question at the end of their paper. They admit that "there is no simple and unique division line between structure and content" and that "we can now see that this opposition is not a dichotomy; there is a continuum between these poles". (p. 168 –169)

> universal logic which characterizes *all* valid arguments or the relationships between the meanings of *all* expressions. In practice, different *logical systems* are developed, each with its own particular class of arguments. What class this is depends on the kinds of expressions found in the logical language the logical system uses [i.e. the logical constants of that system]. ... It should be noted, however, that this is not the only way new logical systems can be developed. We can also consider the same set of logical constants under a new interpretation. This too results in a different class of valid argument schemata .... Strictly speaking, then, *a logical system is characterised by its logical constants together with the interpretations placed on them* (my emphasis).

This excerpt is intended to illustrate what "the science of reasoning", i.e. logic, actually involves, so as to disabuse us of the notion that to label something 'logical' is a self-evident thing to do.

Perhaps, as so often seems to be the case when concepts are exported and have a life of their own outside their home disciplines, there is conflation of a folk-theoretic notion with a more scientific understanding of the concept.[3] This 'folk' notion of logic would presumably be what enables us to judge right off that only the first argument given below is valid. For we could surely all agree that

(1) All men are mortal.
Socrates is a man.
Therefore, Socrates is mortal.

should be judged valid in any logical system worth its salt, while

(2) All men are mortal.
All horses are mortal.
Therefore, all men are horses.

should not get that status. There are several remarks to be made about this.

Firstly, the folk-theoretic notion that enables us to make this distinction is always applied to examples for which it is seemingly blindingly obvious. For although the *modus ponens* schema is surely among the most uncontroversial, there are some other less palatable arguments arising from a material implication reading of conditionals. Our intuitions would be much more divided about the following (from Stalnaker, 1975, p. 269):

(3) The butler did it.
Therefore, if he didn't, the gardener did.

[3]Stenning has shown this to be the case with the notion of 'innateness', which has long been replaced with the concept of 'heritability' within biology, but which endures in nature/nurture debates outside biology (Stenning, 2007).

Indeed, Stalnaker judges this to be "intuitively absurd", yet it relies on nothing more than the material implication reading of the conditional in propositional logic to count as valid. Such an argument illustrates what is known as a paradox of material implication. In this case, the paradox is that whenever the antecedent is false, the whole conditional is true. Since making the premise true forces the antecedent of the conditional conclusion to be false, the whole conditional becomes true, and so the argument is valid.[4]

The collision of our intuitions with the functioning of logical connectives such as the material implication do not only occur with these awkward paradox-based arguments, like that given above, but also with more acceptable premises and no obvious incompatibility in the relation to the conclusion. These can nevertheless result in counterintuitive conclusions, as in the following:

> If Jones wins the election, Smith will retire to private life.
> If Smith dies before the election, Jones will win it.
> Therefore, if Smith dies before the election, he will retire to private life.

As Veltman (1986, p. 147) remarks, "What one calls a logically valid argument form with a few pragmatically correct instances is for another a logically invalid argument form with many pragmatically correct instances". (Veltman here makes reference to a conventional role division between pragmatics and semantics, which we will question further on in the chapter.) It seems our folk-theoretic, or intuitive, notions of logic, validity, and good argument, are by no means always conservatively represented in common formal logical systems. Conversely, the consequences of formalisms confront our folk-theoretic notion of what a good argument is. So, in the context of judging reasoning, our folk-theoretic judgements about 'logic' must either be taken to be just that; or they should be retired to make way for specific logical systems. Much of Stenning and van Lambalgen's work can be seen as an effort to discover the specific logical systems at work in our reasoning behaviour; much of formal semantics can also be seen as an effort towards this where reasoning is based on natural language. Indeed, any theory of reasoning needs to be built on semantic theory – that's the bottom line. We now present some arguments which illustrate why this should be so.

[4]The other paradox is that whenever the consequent of a conditional is true, the conditional is true. This is illustrated in the following argument:

(i) If pigs can fly, then there is life on earth.
Pigs can fly.
Therefore, there is life on earth.

4.3 Determining logical form

An important aspect motivating the rejection of a role for logic stems from the assumption that subjects unproblematically extract the intended logical form from the grammatical form of the premises in a reasoning task. As has been amply illustrated in the foregoing chapters, for syllogistic and conditional reasoning tasks, and as we'll see in the next chapter in the context of the Wason selection task, this is most definitely not the case. But much previous reasoning research assumed that logical form could be transparently read off grammatical form. That means assuming, among other things, that 'and' can always be translated as (propositional logical) conjunction $\wedge$, 'or' as disjunction $\vee$, and, perhaps most problematically, 'if ... then' as the material implication $\rightarrow$. This is evident in the very first sentence of Wason's (1968a) article on the selection task (my emphasis):

> This investigation is concerned with the difficulty of making a particular type of inference from *conditional sentences, statements of material implication* of the form, "if P then Q" (P$\supset$Q)."

This is simply a case of broken telephone between disciplines. Semanticists who study the logical structure of natural language make no such assumptions about even the possibility of straightforward and unique translations into logical form. Veltman's (1985, p. 3) comments illustrate:

> As a logician, you can do no more than devise a logic for conditionals and try to persuade your readers to adopt it. ... It cannot be [the best logic for conditionals], not because this actual logic would not be good enough, but simply because there is no such thing.

Now this is quite an extreme opinion; many logicians would beg to differ, but the point is made. For semanticists, it is an open research question what the appropriate logical form is for many natural language constructions, not least conditionals.

Since Wason's 1968 article there has certainly been more attention given to the difference between grammatical and logical form; sometimes merely lip-service – for instance allowing that conditional assertions in natural language may also be read biconditionally (Johnson-Laird, 2001, p. 438)[5] – but sometimes also proper

[5]The difference between grammatical and logical form is related to the more general distinction between a sentence and a proposition. Awareness of the distinction is not always lacking. In Johnson-Laird & Savary (1999), for instance, we find the following cautionary note (p. 193):

> A point to bear in mind, however, is the difference between a sentence or clause and the proposition that it expresses. Most sentences can be used to express many different propositions, ...

So far so good: these theorists seem to be aware that language needs to be processed before a meaning is generated, that is, they seem to be about to take into account 'reasoning to an

consideration, as we have seen in the discussion of Chapter 1 of the similarities and differences between Politzer (2004) and Stenning and van Lambalgen (2008) regarding logical form for the conditional.

Grammar $\neq$ form

Empirically, Fillenbaum's striking studies (1978, as discussed in previous chapter) illustrate precisely subjects' refusal to make exact 'translations' automatically, or even with prompting. Recall that Fillenbaum's subjects converted such sentences as

(4) Clean up the mess or I won't report you

into sentences such as

(5) If you don't clean up the mess I'll report you.

The thing to notice is the switched negation in the second sentence; a 'strict' paraphrase would be the counterintuitive

(6) If you clean up I'll report you.

Fillenbaum labels this phenomenon 'pragmatic normalization'; here we can subsume it into what Stenning and van Lambalgen (2008) termed 'reasoning to an interpretation' and take it as evidence that normal language speakers do not glean logical form directly from grammatical form. The issue of normativity – to what extent the process of pragmatic normalisation exhibited by Fillenbaum's subjects is *warranted* and thus what status such a process should get – is taken up more fully later on in the current chapter. For now, it should be clear that phenomena like normalisation need to be accounted for in a theory of reasoning no matter what normative status they get. In contrast to assuming that reasoning subjects simply read logical form off grammatical form, and as mentioned previously, Stenning and van Lambalgen emphasise that there are two basic stages to reasoning behaviour, namely reasoning *to* an interpretation and reasoning *from* an interpretation. The former stage is concerned with establishing "the domain about which one reasons and its formal properties" – including making decisions about logical form; then one can go on to reason *from* an interpretation, which

---

interpretation', in Stenning and van Lambalgen's terms. This is not the case, as they continue with:

> ... e.g. the disjunction ['Either there wasn't a king in the hand or else there was an ace in the hand'] refers to different hands of cards depending on the circumstances of its utterance. It is laborious to keep writing, 'the proposition expressed by the sentence' and so unless the distinction matters we will use 'assertion' to refer to sentences or the propositions they express.

And that's all that Johnson-Laird & Savary have to say about the difference between sentence and proposition.

is "guided by formal laws". These two stages fit well onto the analysis of subject behaviour given in the previous chapters: both Chapter 1 and Chapter 2 can be seen as investigations of the considerations involved in imposing logical form, that is, reasoning *to* an interpretation. We have explained differences across schooled and unschooled subjects as largely differences at this stage of the reasoning task. Note however that it's not always so much a difference in interpretation which matters, but a difference in the conditions under which the subject imposes one or other interpretation. For example, we saw that with syllogistic materials unschooled subjects appear reluctant to ascribe a (descriptive) law-like reading to the quantified premises on a universal domain; they often looked for further restriction of the domain before allowing the quantification to serve as a basis for predicting about instances. The difference between law-like and contingent versions would not show up in logical form for the proposition but in the semantics accompanying it – i.e. the variability allowed in the domain.

Grammar ≠ form: some implications

Once the space between grammatical form and logical form has been noticed, all sorts of other things start to shift and dissemble. For instance, we can see that the *proposition's* elemental status in propositional logic does not transfer to the *sentence* in natural language.[6] Consider the suppression effect task results, as reported in Byrne (1989). Byrne's main thesis is that "in certain contexts subjects reject instances of the valid modus ponens and modus tollens inference form in conditional arguments" (p. 61). This is demonstrated by what Byrne labelled the "suppression effect" which becomes visible (in her study) only in *between* subject comparisons on different conditions. The different conditions are described in more detail in Chapter 1, here an example suffices. When subjects are presented with the conditional premise 'If she meets her friend then she will go to a play', and the categorical premise 'She meets her friend' then rates of *modus ponens* elicitation are very high – 96%. When subjects are presented with these premises accompanied by an extra conditional such as 'If she meets her family then she will go to a play' the elicitation rates remain high. However, in groups where the materials presented contained an extra premise like 'If she has enough money then she will go to a play', the elicitation rate drops to 38%. Now for this pattern to count as evidence of "suppression" of an inference, you have to assume, among other things, that the grammatical features of the premises – specifically that the two conditions are presented in different sentences – are maintained in the logical form attributed to them. However, it is much more plausible (as even Byrne acknowledges in her formulation of premises as "additional" or "alternative") that subjects roll the two conditional premises into one proposition with a complex

[6]Relevant here is the finding, reported in the previous chapter (Kurvers, 2002), as well as in other studies (e.g. Garton & Pratt, 1998) that concepts such as 'word' and 'sentence' are the products of literacy.

antecedent, as it were, which vary in form according to the type of relation specified in the second conditional premise (see next section for discussion on the role of 'content' in determining form). So for instance, in the case given above, we would expect the combination:

> If she meets her friend then she will go to a play.
> If she meets her family then she will go to a play

to be represented as

> If she meets her friend *or* her family, she will go to a play.

and similarly, for the additional condition 'If she has enough money then she will go to a play':

> If she meets her friend *and* she has enough money, she will go to a play.

Indeed, the only reason not to do this would be if you paid special attention to the sentential presentation of the premises – something which more than a third of the subjects appear to have done. Again, we do not call this so much a logical achievement as a literal stance towards the materials, which some subjects take, and others apparently don't.

4.4 Is it really content *xor* form?

A central argument against the role of logic in human reasoning is based on the content-blindness of logic; the argument is that logical rules are purely formal and thus apply whatever the content is, if they apply at all. This is concisely expressed in Evans (2002, p. 983):

> Such influences [of content and context] are necessarily nonlogical because the deduction paradigm requires people to make inferences based on logical form for *arguments whose particular content and context is irrelevant.*[7]

But, or so the argument goes, human reasoning is highly sensitive to content, and thus cannot involve recourse to logic. This reasoning is often used in defence of domain-specific theories of reasoning – viz. (Fiddick et al, 2000, p. 2):

> A central method used to test domain-specific reasoning theories has been to see whether the inferences people make vary as a function of the content they are asked to reason about.

[7]The equivocation of a specific experimental area 'the deduction paradigm' with logic is itself an interesting equivocation!

The assumption of content-blindness (or topic-neutrality) of logic is closely related to the division of language into logical and non-logical elements. For example, in the sentence "All bears are white", it is assumed there are two sorts of lexical items: logical ones, i.e. "all", and non-logical ones such as "bears", "white". Content is given by the non-logical elements of a sentence. And the logic of the sentence is given by the logical elements (as you may have guessed), so that "All A are B" has the same logical meaning whatever you substitute for A and B.[8]

This is an appealing view on natural language but it is too simplistic. Firstly, as we've discussed above, grammatical form does not always straightforwardly portray logical form, so the logical elements cannot just be read off a natural language sentence. Secondly, natural language does not cleft so neatly into 'logical' and 'non-logical' elements as one might hope. There are many words and constructions which carry logical inferences – not just implicatures – as part of their meaning.[9] Again, we see this articulated in the opening pages of the logic textbook (Gamut, 1991, p. 8), where it is made clear that the 'logical' elements of natural language form a much bigger set than the propositional connectives, and a not easily delimitable set at that:

> Logical constants other than those mentioned so far [i.e. those of propositional and predicate logic] are, for example, modal expressions, like *possibly* and *necessarily*, which are treated in modal logic ..., and temporal expressions and constructions like *it was the case that*, *it will be the case that*, *sometime*, *never*, and the tenses of verbs, which are treated in tense logic .... The set of possible logical constants is an open one. .... A sharp boundary cannot be drawn between purely descriptive terms and the rest .... [Rather,] there seems to be a gradual transition from structural aspects of meaning, which fall within the range of linguistic theories, and descriptive content, which does not.

[8]This point has also been taken up by Stenning and van Lambalgen. For them it is related to "a superficial reading of the classical definition of validity, say for a syllogism such as

> All A are B.
> All B are C.
> Therefore, all A are C.

The validity of the schema is taken to mean something like 'whatever you substitute for A, B and C, if the premises are true for the substitution, then so is the conclusion" (Stenning and van Lambalgen, 2008, p. 28). This is a superficial reading because it ignores the domain-dependence of the applicability of such a schema.

[9]Here again I make reference to a conventional distinction in the semantics-pragmatics literature between inferences and implicatures. The term 'implicature' was coined by Grice (see Grice, 1975, reprinted in Grice, 1989) to explain the difference between speaker and sentence meaning. An implicature is suggested, intended, or meant, by what is said. Later in the chapter we evaluate the tenability of a principled distinction between 'what is said' and 'what is meant'.

The key phrase in the above excerpt is that "[t]he set of possible logical constants is an open one." There is more to logic than that captured with disjunction, conjunction, implication, negation, universal and existential quantification.[10] Terms and phrases such as *only* and 'x knows that p', illustrate.[11] Both of these have logical properties which formal semanticists have studied extensively. Take the case of *only*. It may not strike one as 'logical' immediately. Yet *only* interacts with the focus of the sentence – determined by phonology – to fix the logical form of a stressed sentence.[12] So for example, the following two sentences:

(7) Sarah only WRITES books.

(8) Sarah only writes BOOKS.[13]

express different propositions: the first is true in situations where Sarah reads for instance magazines but no books, the second in situations where Sarah has quit her job as a gossip columnist. This is not mere pragmatic adjustment – we are talking here about the truth-conditions of the expressed propositions. From an inferential point of view, from the first example we can conclude 'Sarah does not read books'; from the second the conclusion 'Sarah doesn't write newspaper columns' follows. A roughshod translation would encode the two sentences as having the same form, perhaps simply into p, since the grammar does not suggest any differing 'logical' aspect – we have to take phonology (or context) into account to see that – thereby losing these essentially logical differences between them. (Second-order) predicate logic could fare better but since most psychology of reasoning studies have dealt with propositional logic, and to illustrate the point, we can stick with it.

For a rather trivial but comparable example, reflect on the fact that the Barbara syllogism comes out in propositional logic as:

p.
q.
Therefore, r.

Evidently, there is continuous discrimination in what needs to be encoded into logical form and what not, from natural language contexts. That discrimination is exercised when we translate quantified sentences into the formal language of predicate logic, while for conditional-based arguments it often suffices to translate them into the much simpler language of propositional logic. It is equally exercised

[10] This view is called the 'First Order thesis' by Barwise (1989, p. 37).

[11] Compare 'x knows that p' with 'x *thinks* that p', which does not have the property of veridicality for the embedded clause p. A variety of epistemic logics have been developed which capture such properties of knowledge and belief. For a nice example of how far you can get with this kind of approach, see Baltag, A., Moss, L.S., Solecki, S. (1998).

[12] Matts Rooth's dissertation gives the logical analysis of 'only' (*Association with focus*, 1985), in which he argues that it has a quantificational structure.

[13] Capitals here indicate a pitch accent.

by our reasoning subjects who are trying to figure out what logical distinctions are important for the task at hand. The distinction between logical and non-logical in natural language is not simply given by the constants of propositional logic.

But surely this independence from specific content is exactly what defines logic? *Within* specific logics, we do achieve topic-neutrality, in the sense that only the logical constants contribute to the logical relations between propositions. In a formal logical language such as that of propositional logic, sentence letters represent 'content' and do not directly contribute to logical structure. One way to characterise logical constants in a language is to see which parts of the language are invariant under permutation of the domain (Tarski, 1936); the idea is that those parts that are not affected by permutations of the objects in the domain under discussion are the parts that are neutral with respect to what the language is representing.[14] This is a technical property of the logic, and to take it as a defining feature of logic needs further justification by appeal to some non-technical property (McFarlane, 2000, p. 19). The non-technical property to which recourse is usually made to is *formality*. But what does it meant to say that logic is formal?

The formality of logic

As we've seen, reasoning researchers are inclined to demarcate logical from non-logical reasoning by making recourse to the formal nature of logic. That is to say, that logic concerns itself only with the "form" of statements, abstracting from content. This is by no means a view peculiar to reasoning researchers. It is a pervasive view of logic. McFarlane (2000, p. 63) goes so far as to say that "This kind of talk is so common as to be nearly invisible." The problem is that we don't know what we *mean* when saying logic is formal. Historically, formal consequence has been differentiated from material consequence, to explain the difference between, for instance (Read, 1994, p. 237), the source of validity in

[14]Historically, the equivocation of syntax with logical form can be seen as deriving (albeit in a distorted fashion) from a particular view of logic: namely that logic is about implication relations between propositions. On this view, logic inheres in the formal structure of language. This can be contrasted with an older understanding of logic as the study of judgements and inferences, where both are mental notions. Logic in this sense is an epistemic tool: inference is an act in which a judgement is made on the basis of other, already made judgements. (The content of a judgement is a proposition.) Sundholm (1994) locates the shift as far back as Bolzano, whom he says makes the two key reductions, namely "(i) that of the correctness of the judgement to that of the truth of the propositional content and (ii) that of the validity of an inference between judgements to a corresponding logical consequence among suitable propositions [namely those which form the content of the judgements]." This is interesting from our point of view because it indicates a change in the locus of logical authority: the internal perspective of an epistemic agent is stripped away. Logic's subject matter becomes the relation between propositions, or even well-formed formulae, which Sundholm calls "meta-mathematical objects of an uninterpreted formal language".

(9) All cats are animals.
All animals have tails.
So cats have tails.

and

(10) Iain is a bachelor.
So Iain is unmarried.

Traditionally only formal consequence falls under the jurisdiction of logic; again, this is difficult to make precise. Read will go on to argue that the distinction between formal and material consequence does not delineate validity correctly, but even if it was correct, in principle, the line is difficult to draw in practice.[15] As we have seen above, the difference between logical and descriptive aspects of natural language is a sliding scale rather than a dividing line.

In reasoning research 'formality' is sometimes even identified in descriptive content. Take for example, the question which opens Norenzayan, Smith, Kim and Nisbett's (2002) article on cross-cultural differences in reasoning behaviour: "Is the Pope a bachelor?" According to the authors, answers to this are representative of one of two reasoning styles. Responding "no" indicates an intuitive, similarity-based approach to the problem (the Pope isn't like other bachelors), while "yes" (the "correct" answer) indicates an analytic rule-based approach (he's an unmarried male). Evidence for a broader style difference thus (Norenzayan et al, 2002, p. 654):

> The 'bachelor' problem illustrates an important theoretical distinction in the psychology of reasoning. According to this distinction, human thinking is guided by two separate classes of cognitive strategies that implement different computational principles. One can be described as intuitive, experience-based, or holistic, whereas the other can be

[15] Going even further than this, McFarlane (*ibid*, p. 51) shows, that the even notion of formality attributed classical logic is a slippery one, being used in at least three common senses:

- To say that logic is **1-formal** is to say that "its norms are *constitutive* of concept use *as such*"
- **2-formal**: "its characteristic notions and laws are indifferent to the particular identities of different objects. 2-formal notions and laws treat each object the same (whether it is a cow, a peach, a shadow, or a number). Mathematically, 2-formality can be spelled out as invariance under all permutations of the domain of objects."
- To say that logic is **3-formal** "is to say that it abstracts entirely from the semantic content or 'matter' of concepts"

McFarlane will go on to argue that logic is best understood as formal in the first sense – perhaps surprisingly. Here it suffices to realise that simply making reference to the formality of logic, and judging an argument to be logical on the grounds that it is formal, will not do, unless it is specified further what formal should be taken to mean!

> described as formal, rule-based or analytic.

The authors go on to relate the two systems hypothesis mentioned here to the widespread research interest into "dual process" theories of reasoning (Wason & Evans, 1975, Evans, 2003). And further on, in more detail (Norenzayan et al, 2002, p. 678):

> Formal reasoning is rule-based, emphasizes logical inference, represents concepts by necessary and sufficient features, and overlooks sense experience when it conflicts with rules or logic. Intuitive reasoning is experience-based, resists decontextualising or separating form from content, relies on sense experience and concrete instances, and overlooks rules and logic when they are at odds with intuition.

They go on to report two studies in which East Asian student subjects relied more heavily than their American counterparts on 'intuitive' processes, leading to less accurate responses – seemingly raising the possibility that East Asian students are less logical. Qualms about such general dichotomies aside, what's interesting to note is that the locus of logic has widened in this study, to cover material inferences. The normative force of the 'logical' answer has also widened accordingly. In other words, for Norenzayan et al, material consequence judgements such as that the Pope is a bachelor count as evidence for logical reasoning, while in other psychology of reasoning studies, asking the question "Is the Pope a bachelor?" would not be seen to have a logical component at all. Here again we see the problems of relying on folk-theoretic concepts in an academic domain: they can be inconsistently applied and thereby make diverse results seem convergent.

One last point on formality: taking permutation invariance to be the hallmark of the formal elements might seem to rule out many natural language counterparts of logical constants – the different interpretations of conditionals across the abstract and thematic variations of the Wason selection task provide a good example of this. But if one allows for the relevant logical differences in the selection task, such as the difference between deontic and descriptive conditionals, this problem disappears. Again, the matter of the status of different interpretations is taken up later on in the chapter.

4.5 Classical first-order logic as basic

> Classical logic has been immensely successful. But this very success has enshrined certain formats and procedures, that also have drawbacks. For instance, many themes suffer from what may be called 'system imprisonment'. We have to discuss the behaviour of [say] negation inside specific formal systems, such as propositional or predicate logic – even though these systems do not correspond to meaningful distinctions in the 'open space' of actual reasoning.

(van Benthem, 2000)

Almost all empirical results on reasoning have been the victims of 'system imprisonment', in the sense that the meaningful distinctions have been presupposed to be those of propositional or at most predicate logic, and not those apprehended by subjects in the tasks. This is because outside – and inside – logic it has generally been assumed that classical logic *is* logic, or at least provides the 'basic case' – a view expressed for instance in Cherniak's influential book (1986), on the thesis that the "acceptance of logic" is a precondition for rationality (p. 76):

> I will deal almost entirely with classical logic. This is not to prejudge the issue of the adequacy of nonstandard logics; the case of classical logic is basic, and the argument should be generalizable to other logics.

Cherniak goes on to describe the complexity of classical logic, and in turn uses this as to argue against the acceptance of logic as a precondition for rationality.

Meanwhile, inside logic we get the following kind of opinions: Barwise, on the proper place of first-order logic in logic in general:

> [First order logic] is just an artificial language constructed to help investigate logic, much as the telescope is a tool constructed to help heavenly bodies. From the perspective of the man in the street, the [first-order] thesis is like the claim that astronomy is the study of the telescope.[16]

I don't want to create the idea that logicians 'know better'; rather I want to point out that logicians study different logics. Classical logic is one of the many exportable products of that study, not the answer; it is a topic of lively discussion in what sense, if any, it is more 'basic' than other logics.

More importantly, there are two senses in which propositional logic can be argued to be basic: basic relative to other logics, or basic relative to the logic of natural language. Classical logic can be basic relative to other logics, in the sense that other logics are built on it (e.g. quantified predicate logic makes use of the same set of logical constants and semantics at the sentence level), without being basic relative to natural language. And even in this sense, classical logic is not basic, for other logics, such as intuitionistic logic, are not embeddable in it.[17] The fear is that letting classical logic go as a norm will bring on a relativist

[16]The first-order thesis is the claim that "logic is the study of the properties of *and, or, not, implies, every, some* and *identity* and that anything that cannot be defined in terms of these is outside the domain of logic" (Barwise, 1989, p. 37).

[17]Intuitionistic logic was developed by the mathematician L.E.J. Brouwer with the aim of modelling the constructive reasoning of intuitionistic mathematics. The basic premise in this style of reasoning is that *reductio ad absurdum* should not be used to prove the existence of something. For example, one should not conclude from the impossibility of, say, no object having the property A, that there is indeed an object with the property A. In this logical system the law of the excluded middle (i.e. that $p \vee \neg p$ is a tautology) no longer holds (Gamut, 1991, p. 140).

free-for-all: if any reasoning (good or bad) can be cast in a custom-made logic, then what can logic possibly tell us about what good reasoning is? But pluralism in logics need not lead to relativism, or at least not a vicious relativism. It may be perfectly appropriate to apply different logics depending on the aims and possibilities of the situation: in co-operative conversation a default logic is most useful; in legal debate or scientific reasoning a more rigorous interpretation process may be necessary, recruiting a more classical-type logic. In both cases, it's not the case that anything goes. A perspectival view does not preclude criteria of rightness in reasoning within a chosen validity definition, and even inherent rightness in choosing this. A normative system can be appropriate relative to parameters without being relativistic.

If classical logic provides the norm, the next question is how the normative interpretations of premises have been supposed to be reached. We critically examine this issue in the following section.

4.6 Just what *are* the given premises?

In the foregoing I have been continuously making reference to 'literal meaning' and claiming that this is the notion that schooled subject access but which unschooled subjects do not have available in reasoning tasks. But what do I mean by 'literal meaning' and how does it relate to other notions of meaning? And what normative status does it have in the context of the reasoning task, if any? Firstly, let us be clear about the role such a notion has in reasoning tasks. Recall the summary offered in Braine (1978, p. 2), epitomising a still dominant view inside (and outside) psychology of reasoning:

> Practical (i.e., everyday) reasoning uses all the information at a person's disposal, whereas formal reasoning is concerned with whether conclusions follow from premises. ... formal reasoning makes two demands not made in everyday reasoning: (a) Reasoners must compartmentalize information (i.e. *restrict the information used to that contained in the premises*) and (b) they must take a special attitude in comprehending the premises, by attempting to discover the *minimum commitments of the premises as they are worded.*

This excerpt sketches a picture in which the logical reasoner, reasoning solely on the basis of the given premises, must make use only of the literal meaning of the premises, or, phrased in Braine's terms, as the "minimum commitments of the premises as they are worded" (*ibid*). Subjects who fail to do this bring in world knowledge and make use of extra assumptions which are not entailed by the premises. Making this distinction, as Braine requires of the logical reasoner, requires being able to draw a sharp line between information contained in the

premises and that merely implied by them. As we will see in the following, this line can only be drawn relative to a theoretical perspective.

4.6.1 Literal meaning in everyday use

The notion of 'literal meaning' has an everyday sense. Often it is used in differentiating from non-literal meanings such as in metaphor or irony. Indeed, the Oxford English Dictionary defines 'literal' as "taking words in their usual or most basic sense without metaphor or exaggeration." Iglesias (2006) has shown that this function hides a multi-faceted notion. This is primarily because literality is always used in contrast with non-literality; Iglesias calls the literal/non-literal distinction "a cluster of productive analogical dichotomies", each of which is applied differently in different settings, including disambiguation, transfer, extension (including conventionalised figurative language, metaphors), indirect speech acts, and implicature.[18] He identifies a common structure in these contexts of use in which the *first* interpretation is non-literal, and a literal meaning is calculated only afterwards, because of some additional information which leads to a cancellation of the original interpretation. This is illustrated in the following exchange, a case of disambiguation (Iglesias, 2006, p. 137):

> A is in Granada sitting in front of the Alhambra and says to B:
>
> (11) This in front of us is a marvel.
>
> B interprets that A is talking about the Alhambra and later reports A's words to C in the following way:
>
> (12) A said that the Alhambra is a marvel.
>
> Then C replies that he finds this implausible, for A is particularly insensitive to this kind of architecture. In this moment B remembers that A is particularly fond of cars and that, when A uttered (11), a Rolls Royce was parked in front of them. B then says:
>
> (13) Well, what A *literally* said is 'This in front of us is a marvel'.

Presumably B would then go on to explain that there was also a Roll-Royce in full-view when A made his utterance. In this case, 'literal meaning' reduces to verbatim citation, and it differs from the original reported speech in that it unfixes the original reference. In other uses of the term 'literally', 'verbatim' would not be an adequate explanatory substitute term; think here of metaphors in which 'literal' means rather 'non-figurative', for example. Furthermore, Iglesias emphasises, the different uses are mutually irreducible, because they operate on different levels of interpretation, and thus that there is no paradigmatic sense of 'literal' from which other uses lend their meaning.

[18] We can understand the difference between the two answers to Norenzayan's 'bachelor problem' as a preference for literal vs figurative/metaphorical readings of the question.

In a similar usage to the everyday sense – i.e. to differentiate from non-literal language usage – literality has also played an important role within philosophy of language discussions about the structure of meaning. But here it is attached to a theory of meaning, which assumes that literal meaning is not only always interpreted *first* – but it also forms the deductive core of the interpretative process. Both processing priority and semantic or informational precedence are adduced to literal meaning in contrast to intended, perhaps non-literal, meaning.

An influential illustration of this view of meaning is found in Searle's paper *Metaphor* (1979) in which the view is propounded that an initial sentence meaning is computed and only then, after a mismatch with the context of use, is a metaphorical meaning generated, on the back of the literal meaning as it were; an early rebuttal of such an account of meaning is Bartsch's (1984) widely cited rejection of the suitability of Searle's account to explain how metaphor works.[19]

4.6.2 Literality in theories of meaning

Recently the role of the notion of literal meaning in theories of meaning has been addressed by François Recanati in his book of the same title, *Literal Meaning* (2004). I will make use of the language-philosophical arguments he lays out in this book, as they are concise yet precisely equipped to elucidate, and challenge, the assumption that subjects should – or even can – reason 'solely on the basis of the given premises'; an assumption which, as we have seen, is evident in most psychology of reasoning research.

In fact, this is a reasonable assumption given the view which has dominated philosophy of language in the last decades, as exemplified in Searle's account of metaphor (see above). According to this view, knowledge of the meanings of words, plus knowledge of compositional rules of the language, allows the hearer to interpret any utterance, at least 'literally'. This gives 'sentence meaning'. But sometimes the speaker wants to communicate something other than what has literally been expressed. At this point pragmatics, in the form of conversational maxims, can be employed, to generate a meaning other than the initial literal meaning – namely, 'speaker meaning'. Take as an example the sentence "I've had breakfast". Under normal circumstances, on the dominant 'literalist' view, the speaker meaning is that the speaker "has had breakfast today", but this is an enriched version of the semantic core of the utterance: the literal or sentence meaning, i.e. the speaker "has had breakfast at some point previously", since, to anticipate slightly, this latter statement already expresses a full proposition.[20] Yet it would be true even if the speaker had only had breakfast "twenty years ago and never since" (Recanati, p. 8). (On top of this there is what the speaker

[19] In fact Bartsch was probably one of the first theorists in the tradition to acknowledge "that there is not always a sentence meaning available to start with." (1984, p. 29).

[20] The matter of the pronoun 'I' is addressed in the next paragraphs.

conveys by uttering "I've had breakfast", namely perhaps an answer "No, thank you" to the question "Would you like a croissant?")

Essential here is that it is assumed that "we can legitimately ascribe determinate contents (such as truth-conditions) to natural language *sentences*, independently of what the speaker actually means" (Recanati, p. i). This entails a basic distinction between 'sentence meaning' and 'speaker meaning'[21], as described above, even when they coincide, and ascription of a certain *status* to sentence meaning, as the deductive core of meaning. Against this, and along similar lines to those described by Iglesias for everyday use of 'literal', Recanati will argue that the basic notion is 'what is said', that this is determined on the basis of content *and* context, as we'll specify further later on; and that 'literal meaning' is at best a post-hoc theoretical construction.

Recanati describes how the difference between sentence meaning and speaker meaning has been seen to cleave parallel to semantic and pragmatic contributions to interpretation. The semantic part of interpretation proceeds deductively, because "knowing a language is like knowing a theory by means of which one can deductively establish the truth-conditions (or the proposition expressed by) sentences"; on the other hand, pragmatic interpretation is characterised by its defeasibility, primarily because "there is no limit to the amount of contextual information that can affect pragmatic interpretation" (p. 54).

Now this picture as stated has obvious gaps, because certain linguistic elements *rely* on contextual information to get their semantic meaning – think of indexicals such as 'here', 'I', or unarticulated constituents such as in 'it's raining' (which we take to mean it's raining *here* rather than just somewhere[22]). In other words, semantic interpretative processes don't always deliver propositions, but rather what Recanati calls 'semantic schemata', and they need to be augmented by contextual input to reach the level of proposition, the proposition expressed by the sentence. This type of (mandatory) process whereby the meaning of the sentence is completed, made propositional, is called saturation.

But the literalist still has a case to make for purely semantic interpretation: she can include the standard contextual inputs needed for saturation by making them semantic rules. For instance, 'I' can be standardly be considered to refer to the speaker. Such contextual input has been distinguished from more unruly *pragmatic* input, involving world knowledge for instance by Ken Bach[23], whom

[21] Actually Recanati goes on to discuss a triadic distinction (Recanati, p. 5):

sentence meaning
vs
what is said
vs
what is implicated.

I won't use this division as it's not needed to make the point at hand.

[22] Example from Perry.

[23] Kaplan's (1978) distinction between *content* and *character* has the same aim. For him, the

Recanati quotes:

> There are two quite different sorts of context, and each plays quite a different role. Wide context concerns any contextual information relevant to determining the speaker's intention and to the successful and felicitous performance of the speech act ... Narrow context concerns information specifically relevant to determining the semantic values of [indexicals] ... Narrow context is semantic, wide context pragmatic.[24]

The problem is that there is no sharp line to draw between these two types of context. Even pure indexicals such as 'here' require Bach's wide context to acquire the appropriate semantic value: "as is well known, 'here' can refer to this room, this building, this city, this country, and so on ..." (Recanati, p. 58) so that although there maybe be a rule for automatically determining the content from the narrow context, *which* narrow context must be used, is determined by pragmatic processes concerning the speaker's intention and the wide context.

Moreover, Recanati argues, most context-sensitive expressions are not indexical but rather simply *semantically underdeterminate*[25]. He gives the example of the possessive phrase 'John's car', which means something like "*the car that bears relation R to John*, where 'R' is a free variable", which is assigned a value, *not* on the basis of of a rule, or as a function of narrow context, but one determined by wide context, what the speaker means for it to intend.[26] Crucially, this process is mandatory, because you cannot reach the level of proposition without settling on a value for 'R'.

I think that a similar argument can be used against the idea that conventionalised conversational maxims can be included in the deductive process to reach an enriched meaning. To take the "I've had breakfast example" example described above, a Gricean could argue that the conventionalised conversational maxim of

---

content of a given expression "is always taken with respect to a given context of use", while the character of an expression "determines *how* the content is determined by the context" (pp. 83–84, emphasis added). In the case of 'I', its character would be represented "by *the function (or rule if you prefer) which assigns to each context that content which is represented by the constant function from possible worlds to the agents of the context*" (p. 84, his emphasis).

[24]This is from a handout of a presentation given by Bach, 'Semantics vs Pragmatics', in 1996. In the published paper 'You don't say?' (2001, p. 21) he reiterates this point: what the sentence says, in the semantic sense, "excludes anything that is determined by [the speaker's] communicative intention (if it included that, then what is said would be partly a pragmatic matter)" – and we can't have that!

[25]Computational semanticists have shown the implausibility of ambiguity as an alternative explanation. Ambiguous expressions are computationally intractable; indeteminacy is much better-behaved computationally. See for example Kamp & Reyle (1993) for a discussion of this.

[26]I imagine candidates for the relation R would be, for instance, ownership; temporary assignment such as on a bumper car ride; design copyright if John is a car designer, etc. One might call these varieties of 'ownership', but the process of deciding which one applies would be a pragmatic process, one that Recanati calls 'loosening', in which the application of the (unarticulated) predicate 'ownership' is widened (see p. 26 of Recanati).

relevance offers at least a quasi-deductive explanation of why we usually interpret the sentence to mean "I've had breakfast today" (and further conversational implicature could deliver us the conveyed sentence). But then the Gricean would have to use a different argument to explain why a sentence like "I've had measles" doesn't get the same treatment, and, again, wide context – i.e. world knowledge – will here be needed to distinguish the two cases. The point is that the use of world knowledge cannot be excluded by attempts to explain the stable nature of certain implicatures.

To return to Recanati: if he is right about the pervasiveness of semantic underdeterminacy, then sentence meaning becomes a problematic notion. In fact, the underdeterminacy does not even have to be pervasive; all we need to show is that in some ordinary 'literal' (i.e. not metaphorical) usage, in Recanati's words, there is "no such thing as 'what the sentence says' in the literalist sense, that is no such thing as a complete proposition autonomously determined by the rules of the language with respect to the context but independent of the speaker's meaning." (p. 59) This is enough to fatally undermine the literalist picture.

There are still possible interpretations of the notion of 'sentence meaning', in terms of a pragmatically-informed notion of 'what is said' (see p. 59). Certainly the first candidate, what Recanati calls the 'minimal proposition', seems to align well with the interpretation that Stenning and van Lambalgen assume the sceptical reasoner is using. The minimal proposition is defined in terms of the fully integrated pragmatic notion of 'what is said', but minus the contextual ingredients – especially the optional ones, as conversational implicatures would be. One abstracts, as it were, over contexts, to distill the semantic content from the fleshed out pragmatic-semantic content. In Recanati's words this proposal is that the minimal proposition forms a kind of 'common denominator', namely "what is asserted in all contexts in which the sentence is uttered and the indexicals are given the same semantic values as in the current context" (p. 61). The problem is, this 'minimal proposition' is not always part of what is said. Recanati offers the example 'The ham sandwich left without paying.' The speaker who utters this does not minimally mean that 'the ham sandwich left without paying', but a *different* proposition altogether, namely, 'The person who ordered the ham sandwich left without paying'[27]. So it's not the case that the minimal proposition can be seen as the semantic 'core' of the sentence.

A second candidate is Perry's notion[28] of 'reflexive proposition', which is the semantic content of a sentence *prior* to saturation. For the utterance u of a sentence 'I am French' the reflexive proposition would be '*the utterer of u is French*'. Again, this is emphatically not to be understood as somehow part of what is said, but is a theoretical construction, where content is specifically solely determined by the rules of the language, and it is not intended to be considered

[27]This relies on the (optional) pragmatic process of 'loosening' again.

[28]Although Recanati notes it stems originally from Reichenbach's *Elements of symbolic logic*.

as constituting a distinct stage in normal language comprehension. Hence this account predicts a 'one-step' model of language processing, a claim to which we will return.

If one accepts that 'what is said' is *not* determined solely by semantic processes which take recourse to context only to assign values to variables via semantic rules, what would be the alternative? Recanati proposes that 'what is said' is characterised by the availability principle, that is, the hearer must be consciously aware of what was said. 'What is said' is equated with "[the semantic content of] the conscious output of the complex train of processing which underlies comprehension" (Recanati, p. 16). Semantic and pragmatic factors have equal share in determining propositional content. This is certainly a more psychological stance and as mentioned already, seems more psychologically plausible than the dominant view in which semantic factors generate an initial core sentence meaning which is then subject to pragmatic adjustment.

4.6.3 Literal meaning dethroned

Taking this perspective on meaning, however, turns psychology of reasoning results on their head. It's a much more democratic sense of meaning and does not automatically grant the experimenter's intended meaning logical superiority because it is based 'solely on the given premises'; or in other words, on 'literal meaning', with no contextual factors, background knowledge, or attribution of speaker intention allowed to interfere with interpretation. If anything, the experimenter's interpretation is here seen to be an artifact of a particular moment in the history of theorising about language. Let us relate in more detail this perspective on meaning to the empirical results already presented. We have the two core notions we need: 'what is said' in Recanati's fully-fledged pragmatic sense of the notion, and the specification of 'sentence meaning' as a post-facto theoretical construction, for example that of the 'reflexive proposition'. We can now explain the difficulty unschooled subjects have with the given premises, not as a lack of ability to reason with abstract material, as Luria would have had, but as a lack of controlled ability to extract a decontextualised 'reflexive proposition' from the premises; we can explain their apparent elaboration on the premises as the normal interpreter's strategy of taking context into account in attributing truth-conditions. In turn, it is precisely such an ability which explains the highly schooled undergraduate's ability to reason with semantically underdetermined premises.[29] A concrete example would be the necessary interpretation for the syllogistic task, as described in Chapter 2: the subject needs to reason with a generalisation on an unspecified domain; we might say the domain selecting pa-

[29]The mechanism by which this occurs could possibly be a result of explicit knowledge of language – although what exactly is important is unclear, because Luria's subjects were able to reason with quantified premises after only a brief period in education, while in Scribner and my studies subjects with even some years of schooling did not always choose this type of response.

rameter remains unsaturated and in this sense the subject needs to reason with a 'minimal proposition'. The key difference is that we would not term the difference styles of answering as more or less logical, but rather as more or less savvy to the theoretical construct that is 'literal meaning'.

In fact, the difference between what was said and how it is to be taken is one which even highly literate individuals apparently struggle to make. Support is provided by Astington and Olson's (1990) finding of undergraduates' insensitivity to the distinction between assertion and implication in interpretation. They presented university undergraduates with the following material:

> It's Adam's birthday tomorrow. Barbara is just sneaking out of the house to buy a present for him when he sees her and asks her where she is going. Barbara says, "We're out of milk. I'm going to the store."
>
> A Barbara means that she is going to buy milk.
>
> B Barbara concedes that she is going to buy milk.
>
> C Barbara asserts that she is going to buy milk.
>
> D Barbara implies that she is going to buy milk.

The majority of the subjects chose "asserts" as the appropriate verb. Notice that this result was garnered even though the story could be argued to prime the correct choice "implies", by suggesting that Barbara wants to mislead Adam. Choosing "asserts" is an indication either that the subject has not differentiated between what was actually said and what they interpreted it to mean, or if they do have this ability, it has not been connected to knowledge of specific metalinguistic concepts such as 'assertion' and 'implication'.

4.6.4 Subjects are savvy *and* confused

Johnson-Laird & Savary (1999) makes explicit a telling assumption in current psychology of reasoning in claiming that "the [mental] model theory of sentential reasoning aims to characterise the deductions of naive individuals, that is, *those with no training in logic*" (p. 193, my emphasis). Implicit in this statement is the assumption that by quizzing only undergraduates who haven't followed logic courses, one gets to see 'natural' cognition, which is presumably also universal. Logic, by contrast, is here seemingly a technology which we can become skilled at using (and which would give us an unfair advantage in reasoning tasks), but which we acquire in a process of explicit learning. This formulation brings out the strangeness of the standard psychology of reasoning paradigm. We want to see if people can do logic; but we don't want to look at people who have actually learnt to do logic.

Moreover, the usual subjects aren't naive. The educational context plays a large role in preparing subjects for the discourse of a logical reasoning task, as the foregoing chapters have highlighted. It is reasonable to conjecture that undergraduate subjects have learnt that divorcing personal experience (to some degree) from cognitive processing is a vital element of success in the academic context – they successfully play the role of the 'universal' cognitive agent. The emphasis here is on role – their minds may not be representative of the universal specification of what a human mind is supposed to be like – but they have some idea of what's peculiar to their cognition and what's not, and bear this in mind when responding. In this sense, perhaps, their response in reasoning tasks can be taken to be universal.

Stenning and van Lambalgen express the differences between subjects in terms of *sceptical* or *credulous* attitudes towards the discourse (Stenning & van Lambalgen, 2008, p. 29). For Stenning and van Lambalgen, a credulous attitude involves constructing a model of the discourse which is the same as the speaker's intended model, while a sceptical attitude means not using any information "save the explicitly stated premises" and entertaining "all possible arrangements of the entities that make these statements true". While the credulous reasoner uses closed world reasoning, the sceptical reasoner would be more inclined to use classical logic to generate conclusions, according to Stenning and van Lambalgen. This is tied up with the sceptical reasoner's aim of "finding only conclusions which are true in *all* interpretations of the premises" (p. 29), and not making use of "whatever general and specific knowledge we have" to narrow interpretative possibilities, as the credulous reasoner would do.[30]

Although certainly the credulous attitude tallies well with how I have described unschooled subjects' concerns, I would disagree that schooled subjects 'logical' performance especially in the syllogistic task is sufficiently explained by Stenning and van Lambalgen's sceptical stance. As Stenning and van Lambalgen elsewhere point out, the premises in a reasoning task "vastly underdetermine the information to be extracted" – i.e. the conclusion (p. 300). In fact, they go on to add, "the psychology of reasoning has suffered as a consequence of its neglect of this fundamental fact." Yet, in supposing that sufficiently sceptical subjects can reach conclusions by using only information given in the premises, they seem to be ignoring this insight with regards to their own research.

As we've seen, it is problematic to suppose that sceptical or classical logical reasoning is simply reasoning on the basis of the given premises, exactly because the givenness of the premises needs further specification. Moreover, speaking in these terms engenders an idea of language processing in which a core semantic meaning is initially computed, and is subsequently subjected to pragmatic ad-

[30]In fact using simply using general knowledge does not guarantee a specific fixed interpretation; interpretation is much more free to roam than this suggests. The opposition is really between what types of general knowledge and how they constrain available interpretations.

justment. This is a view we have seen successfully challenged by Recanati, and one that Stenning and van Lambalgen themselves later disagree with, where they conclude that "Current evidence from neuroimaging seems to point to a one-step model", where in a one-step process "all the available information [semantic and pragmatic] is brought to bear on the computation of the meaning and sentence boundaries do not have a privileged role" (p. 304).

The classical logical reasoner can better be described as interpreting the premises in a highly artful way, to reach what we have termed 'literal meaning'. But, again, we should be clear that literal meaning is not some kind of basic or core (semantic) meaning, to which supplementary (pragmatic) adjustments are made when reaching the 'intended' or 'speaker' meaning. Literal meaning is rather the outcome of a peculiar stance towards language in which certain parameters of interpretation are ignored – the language processor is run on an incomplete input as it were – such as when subjects reason with a quantified statement on an underspecified domain, as we saw in Chapter 2, or with hanging anaphora, for instance. Hence the distinction between 'literal meaning' and 'intended' or 'speaker meaning' should not be understood as a reflecting a difference in processing stages, rather literal meaning should be understood as a very specialised variety of intended meaning. If this is correct, it would mean that sceptical interpretation is a two-step process – which is not to say that credulous interpretation would always have to be one-step[31].

Another difference between what I call 'literal' meaning and the outcome of a sceptical stance to interpretation would be the range of possibilities that a subject considers. Whereas for Stenning and van Lambalgen, a sceptical reasoner considers "all possible arrangements of the entities that make these statements true", a 'literal' reasoner would only do this in a very circumscribed circumstances and would more generally be liable to rely on a pre-given, learned, interpretation of a premise – possibly because of its resemblance to materials used in educational contexts. The 'literal' reasoner is thus considered to be much less independently capable of considering all possible interpretations of the premises and much more reliant on a learned language genre which supplies these 'literal' meanings. This view is supported by Fillenbaum's studies which show that even highly literate subjects find it difficult to retain a 'literal' meaning when it describes an insensible situation, and plenty of examples from my own data, as well as in other studies.[32]

[31]The sentence "I was writing a letter, but then I spilled coffee on it" would be an example of a case where repair processes are employed in a credulous interpretative mode.

[32]When a term is replaced with a semantically similar but incorrect term, people have difficulty in detecting the distortion. This tendency to overlook distortions in statements is known as the Moses Illusion (Park & Reder, 2004), because it has been discovered in studies where subjects are asked "How many animals of each kind did Moses take on the Ark?". Most subjects simply respond "two", even while they know that it was Noah, not Moses, who took the animals on the Ark (Erickson & Mattson, 1981)! The tendency is very robust, even when subjects have been forewarned about possible mistakes in the sentences, read the materials aloud beforehand, and are under no time pressure.

To differentiate more precisely this idea of 'literal meaning' from 'sceptical' meaning, we use Stenning and van Lambalgen's own example. They consider the following discourse:

> Once upon a time there was a butcher, a baker and a candlestick maker. One fine morning, a body was discovered on the village green, a dagger protruding from its chest. The murderer's footprints were clearly registered in the mud. ...

For this discourse, according to Stenning and van Lambalgen, we would adopt a credulous stance and would be inclined to think there are at least three distinct people in the domain, i.e. that the butcher is not the baker is not the candlestick-maker. And it's an "active question" whether any of them are the murderer or even the corpse (Stenning and van Lambalgen, p. 28). This is to be contrasted with the discourse below:

> Some woman is a baker. Some woman is a butcher. Some woman is a candlestick-maker. Some person is a murderer. Some person is a corpse. All women are men.

Now according to them, "cued perhaps by the 'logical puzzle' style of the discourse", the discourse is likely to be understood with a sceptical attitude: subjects would be inclined to entertain many possibilities about how many people there are, considering for instance even the possibility that there is only one person who is all the things mentioned above. I would think that in general subjects would be prompted not so much to be sceptical, as to be confused, by the second discourse, because it has little or no discourse structure or coherence – it's unclear what discourse relations connect the sentences to each other, and thus also whether there is any relation between the five sequentially introduced referents.[33] World knowledge (i.e. that a person usually has only one occupation) would suggest that the first three referents are distinct. On the other hand, world knowledge tells us that it's very unusual to have a female butcher; at the very least we must be talking about modern times – but candle-making is an archaic occupation, so that doesn't mesh well. In short, it's difficult to find a way to fit these sentences to a single discourse model. This lack of discourse cohesion is a means to make subjects aware of interpretational processes and enable them to steer off well-worn interpretational paths but this would be more veering off than consciously exploring the interpretational landscape! I would thus predict that subjects have much more difficulty in reaching the sceptical stance for discourses such as the one given above – and certainly subjects who are not as literate as the university undergraduates from which Stenning and van Lambalgen draw their subjects. Contexts can be thought up for all possible combinations of predicates

[33]Basic discourse relations are considered to include narration, elaboration, background, result, restatement. Notice that world knowledge is often needed to infer which discourse relation holds between two sentences.

and individuals which make them the best fit interpretation-wise – imagine an interrogation in which the respondent is not at liberty to explicitly mention names but can volunteer one the above sentences when presented with an individual's name – but even the sceptical subject would not explicitly consider all of these. Rather the sceptical stance is awareness of interpretational 'scaffolding' (such as inferring of discourse relations) – i.e. automatic but non-lexicalised elements of discourse processing, which can be sequentially examined but cannot always be done without, perhaps even for computational reasons.[34] Witness for example, in the upcoming chapter, undergraduate subjects in the Wason selection task who could be enticed to remove the assumption that the anaphoric element in the rule, i.e. 'one side/other side', reduced to the definite 'front/back' – but only when relying on a biconditional reading of the conditional.

So although it is clear that there has been a detrimental neglect of interpretational processes, as Stenning and van Lambalgen maintain, they do not go far enough in repairing this. They have not yet put enough emphasis on the role of language usage and conventions, conversational or otherwise, and especially at the level of discourse, in shaping reasoning behaviour. This entails that they do not distinguish enough between unintentional interpretational obstacle courses – as leading to 'repair' processes in interpretation (p. 117) – and language genres which cue a perhaps school-based 'symbolic processing' approach to the material. Their description of the sceptical stance covers a host of finer distinctions that can be made about cautious interpreters. Anticipating slightly, this enables us to understand how it is that unschooled subjects can reason sceptically, or rather, classically, without necessarily taking the 'literal' meaning of a sentence, when prompted by the appropriate context (such as in a debate or law-court).

4.7 The normativity issue

Recall our explanation of the difficulty unschooled subjects had with the given premises not as a lack of ability to reason with abstract material, as Luria would have had, but as a lack of controlled ability to extract a decontextualised 'reflexive proposition' from the premises; we can explain their apparent elaboration on the premises as the normal interpreter's strategy of taking context into account in attributing truth-conditions. In turn, it is precisely the ability to extract a 'reflexive proposition' which explains the highly schooled undergraduate's ability to reason with semantically underdetermined premises, as we saw in Chapter 3.[35]

Who's more logical? On this account it would seem that the savvy, schooled

[34]This is indeed much in line with Stenning & van Lambalgen's description of 'system 2' reasoning processes – see Chapter 4 of their book.

[35]In fact, one could even maintain the schooled and unschooled subjects are employing the same strategies but within different contexts, with resulting different possibilities for interpretation.

subject at best is able to take a logical 'stance' towards the premises but is not inherently more logical. A criticism which the reasoning researcher (or ideal language philosopher) might level at this account, is the lack of objectivity it seems to bring with it – indeed the dismissal of logic as irrelevant because of its topic neutrality runs along these very lines. But the 'loss of objectivity' criticism loses bite once one realises that 'sentence meaning', supposedly the source of objectivity, is a theory-driven notion. This enables us to see that any objectivity we had before was a result of theoretical choices made in determining sentence meaning. Here we come very close to describing the theory-grounded sense[36] in which logic has normative status, as defended by Stenning and van Lambalgen (after Husserl). That is, in brief, that any specific logical system has normative status only within a choice of interpretation which supports it; logic only provides "consequences of choices of parameters" (Stenning & van Lambalgen, 2008, p. 301). That makes normativity within any given domain an empirical issue. Only once this has been specified does any specific logic have normative force.

But again, the sentence-oriented theorist might object that meanings are being multiplied beyond necessity. If you grant that context plays a role in determining truth-conditions, then the same sentence, in different contexts, can get different truth conditions. Depending on your theoretical viewpoint this is problematic. Grice argued that sentences can be attributed definite truth conditions even though there is apparent variability, because the variability is external to the truth conditions and part of generalised conversational implicature. The example of the natural language sentence 'p and q' illustrates. Strawson (1952, pp. 80-81, referenced by Recanati) claims that this sentence is not equivalent to $p \wedge q$ because in the former case the order of the clauses can affect the truth conditions ('They had a child and got married' *vs* 'They got married and had a child'); Grice (1989, pp. 47-50, also referenced by Recanati) responds:

> [W]e may consider the temporal implication in 'They got married and had a child' as a conversational implicature, external to what is said, rather than considering it part of the truth-conditions of the utterance in a certain type of context. In this way, we are able to maintain that the truth-conditions of 'p and q' are determined by the truth-table for '$\wedge$', independently of the context of the utterance.[37]

Recanati's response is that Grice is here begging the question – the linguistic ambiguity he must avoid by attributing definite truth conditions is a product of his assumption that variation in propositional content must be accounted for in terms of variation in linguistic meaning. This is not so for the contextualist, who can happily allow variation in propositional content without linguistic meaning varying, precisely because the contextualist is not committed to propositional content being determined solely by traditionally semantic processes.

[36] This can be contrasted with an ontologically- or epistemically-grounded account.

[37] Grice is here relying on the sub-maxim 'Be orderly', part of the maxim of Manner.

The reasoning theorist might counter that reasoning studies are concerned with *normative*, not *descriptive* forms. Unlike formal semanticists, who aim to uncover the *actual* logical form people attribute to natural language sentences, reasoning researchers aim to uncover and understand the gap between actual behaviour and the norm. But this does not go through, because there is no basis – other than perhaps simplicity – for assuming that, for instance, the material implication is the normative form for a conditional construction in natural language. We might wonder why it seems to have acquired this status. Why *did* Wason assume he could presume that 'if p then q' is always equivalent to '$p \supset q$' when any first-year logic course teaches otherwise? In fact many theorists seem committed to the assumption that material implication maintains its status as correctness norm, while allowing subjects other interpretations. The experimenter needs to assume she has (privileged) access to the normative form in hand in order to judge behaviour as divergent; yet the basis for this normative knowledge is unclear. It might seem to have acquired this status because the implication has been represented into a formal system – i.e. formality serves as a basis for normativity. Rather a mixture of historical and instrumental reasons has wrongly been taken to endow normative status on specific interpretations of natural language constructions.

In addition to this, I think that the source of literalist endowment of normative status on certain interpretations stems from the way semantics and pragmatics are seen to cleave on this account; namely, semantic processing is logical because it uses only deductive processes, while pragmatics is considered extra-logical because it involves defeasible processes. This is an assumption rather than an observation. In fact, as we've seen, semantic processing is by no means deductive and in cases of underdetermination involves mandatory recourse to 'wide' contextual factors.

4.8 Summary, conclusions, and outlook

In this chapter the conceptual scaffolding behind the 'deduction paradigm' which has dominated psychology of reasoning studies has been revealed and partially dismantled. This is achieved by confronting assumptions about the construction of meaning and the relation between natural language and logic with the differing perspectives from logicians and philosophers of language on these topics. We saw that logical form is not simply read off grammatical structure, and that what is coded into logical form is determined relative to a goal in an ongoing process of discrimination – this is indeed exactly what Stenning and van Lambalgen have called "reasoning to an interpretation". Next, I challenged the centrality of the notion of literal meaning in theories of reasoning, where reasoning logically requires using 'only information contained in the premises'. Privileging this interpretation, i.e. the 'literal meaning' of a premise, among possible interpretations, is justified only within a theory of meaning which takes it to be the semantic

core of meaning. But this is an untenable thesis because it turns out that this semantic "core" cannot be determined without reference to contextual factors. As such, 'literal meaning' is dethroned in explanations of logical reasoning. Finally, a monolithic view of logic which presumes all logical reasoning can be captured in the classical logical systems of predicate and propositional logic is challenged.

Turning again to Gigerenzer and Hug's question (1992, p. 127):

"What counts as human rationality: reasoning processes that embody content-independent formal theories, such as propositional logic, or reasoning processes that are well-designed for solving adaptive problems?"

it should now be clear that the answer can very well be: both! What we've seen is that subjects are able to focus on 'literal meaning'; we might have called this a 'Logical stance' towards language. These subjects are being Logical with a capital 'l'. On the other hand, much more basic is the logical nature of language use, but this does not associate simply with the logics which have so far made it out into the world. Meanwhile, logicians continue to systematize the logic of natural language and of human interactions. This logic, with a little l, is simply the mechanics of informational exchange. As much as information exchange is considered an adaptive problem, our ability to manipulate symbols in reasoning processes shows we meet Gigerenzer and Hug's second criterion. And in this sense, we are all mostly logical with a little 'l'. That this does not mean we reason logically all the time is demonstrated in the next chapter.

4.8.1 Predictions and experimental suggestions

The current proposal yields specific experimental predictions about unschooled subjects having more difficulty with premises with definite articles, such as "the man had three dogs", than premises with indefinite articles "a man had three dogs", because "the man" cannot get a semantic value; the premise must be evaluated from the 'reflexive proposition' in which saturation has yet to occur. We might expect subjects who cannot easily access this construct to find definite version of premises more difficult than their indefinitely formulated counterparts. Note that it is an open question whether we should expect to find this only in unschooled populations. Given Dabrowska's (1997) results (see previous chapter for details) we might find the difficulties with taking the 'Logical stance' to be much more widespread that previously supposed.

More generally, viewing literal meaning as a theoretical construct throws up many questions about the undergraduate's expectations regarding language use in reasoning tasks, and how these differ from their understanding of language use in other situations. Relevant work in this regard is already being carried out in the field of experimental pragmatics (see Gibbs & Moise, 1997, and Geurts, 2002), but it does not yet look at changes in interpretational strategies across task domains, as this account predicts. Further research here would be very valuable.

Chapter 5

Remining the Wason selection task

5.1 Introduction

Given the conclusions of the previous chapter, it might seem as though existing paradigms used in psychology of reasoning studies are ill-suited to test human reasoning. We saw that the equivocation of logical with grammatical form, and the reliance on a flawed notion of literal meaning as the basis of logical reasoning, have, amongst other things, led experimenters to substantially underdescribe the semantic structure of reasoning materials, and thereby to underestimate the complexity of performance in reasoning tasks. Thus, it can seem as though the tasks themselves are not effectively testing reasoning. This need not be seen so. This chapter is dedicated to showcasing that, even while arguing that the standard testing means have not been penetrating enough, one can still use standard experimental paradigms to investigate reasoning, albeit in a much broader sense, i.e., one that takes the meaning-making processes accompanying and resulting from it seriously.

The experimental paradigm used to achieve this is Wason's (1968) selection task. The selection task has become a veritable cornerstone in the psychology of reasoning. This is evident in the fact that the variation in results generated by different versions of the original task – especially thematic versions (about which more later) – forms the empirical basis for social contract theory proposed in evolutionary psychology (Cosmides, 1989, Cosmides & Tooby, 1989), for the rational analysis theory of reasoning, which posits probabilistic considerations to explain performance (Oaksford & Chater, 1994), and as important input for both the adaptive rationality theory (Gigerenzer & Hug, 1992), and the dual-process theory of reasoning (Evans, 2003).

Yet, as was argued in the previous chapter, much work in reasoning relies on an oversimplistic view of the relation between logic and natural language. In this case, the view is manifested in the assumption that the so-called thematic version of the rule – often formulated as a 'drinking age rule' – has the same logical

form as the rule in Wason's original task. Stenning and van Lambalgen (2001, 2004) have demonstrated why this reflects an inadequate formalisation of the two rules. In other words, the difference in subject performance across these two tasks cannot be explained merely as a 'content-effect', and thereby non-logical, as we saw in the previous chapter. In fact, the original task is much more complex than the thematic versions of it, and the main aim of this chapter is to extend the semantic analysis offered by Stenning and van Lambalgen (2001, 2004) of the original task, which uncovered a nest of semantic parameters which need to be set before the descriptive task can be completed successfully. However, there are clear parallels between the thematic and abstract tasks, and these might facilitate performance in the more difficult 'abstract' original version of the task. This is also investigated in the current chapter, although the focus remains on the 'abstract' versions. In addition to the original task, several experimental variations on Wason's original rule are investigated, replicating and furthering the studies conducted by Stenning and van Lambalgen. As will become evident, university undergraduates exhibit concerns and confusions in reasoning tasks similar to those identified in less schooled groups: they are concerned to establish what the conditions for truth of the premises are; they rely on everyday usage of constructions used in the premises, and yet can under certain circumstances be prompted to take the intended 'normative' interpretation; they struggle to ascertain their role in the task situation. Such high-level similarities lend plausibility to the idea that universal semantic concerns drive reasoning behaviour in illiterate and highly literate subjects, as well as everyone in between.

5.2 Some background to the selection task

Inaugurated in 1968 by Peter Wason, the original selection task is presented to subjects as follows, including card graphic (Wason, 1968a)[1]:

> Below is depicted a set of four cards, of which you can only see the exposed face but not the hidden back. On each card, there is a number on one of its sides and a letter on the other.
>
> Also below there is a rule which applies only to the four cards. Your task is to decide which if any of these four cards you *must* turn in order to decide if the rule is true.
>
> Don't turn unnecessary cards. Tick the cards you want to turn.
>
> **Rule:** *If there is vowel on one side, then there is an even number on the other side.*

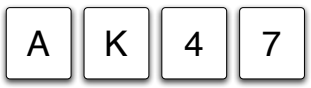

[1]This is the traditional reference although the task is discussed earlier in Wason (1966).

Wason saw the task as a means to investigate the extent to which people reasoned according to the normative theory of scientific reasoning then in vogue: Popper's falsificationism. Popper's philosophy of science concentrated on so-called 'contexts of justification', in which experimental results contributed to scientific knowledge not by confirming existing hypotheses but by falsifying them. Truth is approached by ruling out falsehoods, not by verifying putative truths. For Wason, this mode of scientific thinking was represented in the logic of the selection task, in which the logical subject would seek ways of falsifying the given rule, instead of seeking confirmatory evidence for it.

Student populations overwhelmingly select the A card for turning, and typically more than half also select the 4 card. These results suggested to Wason that people do not reason according to Popper's theory, since only a scant minority select the 7 card, which could potentially falsify the rule by providing a counterexample to it. This is a startling discovery, especially when you consider that the subjects of the tests are intelligent and motivated university students, the scientists of tomorrow! Such an inimical result was not left alone for long; Wason's task has been repeated countless times since then and in myriad variations. The results from the original task have proved extremely robust, but several variations have yielded responses analogous to the normative A, 7 choice – although, as will become clear, there are key logical differences between the original task and these variations.

The variations which yielded normative performance from subjects have been commonly characterised as 'thematic', making use of familiar content in various degrees (Wason & Shapiro, 1971, Johnson-Laird, Legrenzi & Legrenzi, 1972, Griggs & Cox, 1982). An example would be the rule "If a letter has a second class stamp, it is left unsealed".[2] Perhaps the most well-known among these thematic variations is the 'drinking-age' rule, presented in the form "If a person drinks alcohol, they must be over 18 years of age." When presented with such a rule, subjects overwhelmingly choose the response equivalent to A and 7 in the original task (just how equivalent remains to be seen). The results from these tasks contrast with those from the original 'abstract' task and were understood to illustrate 'content-effects' – that is, subjects reason differently with materials which differ in content, but not in logical form. As was described in the previous chapter, such findings have been adduced as evidence against logic-based theories of reasoning. Cheng and Holyoak (1985) had the insight that such rules might be easier for subjects because they describe what *should* be the case, and as such, the task becomes one of detecting violations of the rule, not assessing the rule itself. Cosmides formulated this in terms of social contracts and 'cheater detection'; Stenning and van Lambalgen (2001, 2004) show that this embeds in the

[2]This example also serves to illustrate that familiarity is a context-dependent notion – American students did not achieve ceiling performance on this version, perhaps because the deontic reading was not obvious to them (Stenning and van Lambalgen, 2004, p. 485).

more general notion of *deontic* reasoning. Deontic reasoning is concerned with the logic of obligation and permission; note the use of the modal verb "must" in the drinking-age rule. The rule doesn't describe or purport to describe, what is the case, as a descriptive would, but rather stipulates what *should* be the case. Many of the rules used in variations on the selection task can be interpreted deontically even though they often don't contain modal verbs. Crucially, the original formulation is *not* interpretable deontically, but rather descriptively. As the current chapter will show, the selection task is a much more complex task when dealing with descriptive rather than deontic rules.

5.3 Experiment

The experiment consisted of written questionnaires followed by in-depth tutorial interviews with ten subjects, on a selection of variations on the Wason selection task. The main aims of this study were to gather data on interpretative possibilities, thus extending the work of Stenning and van Lambalgen (2001), and to collate this to subjects' performance across different conditions. A subsidiary aim was to engage the subject with the tasks until stable understanding of the normative competence model for the standard task was achieved, if possible, under the assumption that the transcript thus obtained would provide rich data about what the students need to learn in order to succeed at the standard task.

The written tasks preceding the interviews constitute the standard investigative tool in this area, and served as a base-line control for the rest of the experiment (a 'no learning' condition). The interviews were conducted over two sessions, often with a break overnight. All ten subjects were undergraduate students of Edinburgh University, and first-language English speakers. None of them had previously taken a course in logic.

5.3.1 The written task

The written questionnaire consisted of four versions of the selection task, given in table 5.1.

In addition there was an immediate inference task which will not be discussed on its own here, but will be mentioned when findings from it are relevant in our discussion of the other conditions.[3] Subjects were given as long as they needed

[3]Here is a sample condition of the task with instructions (which were the same for all conditions):

> Assume the sentence in bold is true. Which (if any) of the other sentences below must then also be true? Tick the circle next to the ones you judge must be true.
> **If there is a vowel on one side of a card, then there is an even number on the other side.**

This was followed by up to five sentences, including variously phrased disjunctions, conjunctions,

Condition	Background rule	Rule
original	On each card there is a letter on one of its sides and a number on the other side. The letters are A and K; the numbers are 4 and 7.	*If there is an A on one side of the card, then there is a 4 on the other side.*
conjunctive	as above	*There are As on one side of the cards and 4s on the other side.*
two-rule	Each card has a letter on one side and a number on the other side. The letters are U and I; the numbers are 3 and 8.	**Rule 1:** *If there is a U on one side of the card, then there is an 8 on the other side.* **Rule 2:** *If there is an I on one side of the card, then there is an 8 on the other side.*
arrow rule	Each card has an arrow on one side and a symbol (+ or –) on the other side.	*If there is an upward-pointing arrow on one side of the card, then there is a '+' on the other side.*

Table 5.1: The different experimental conditions in the written task

to complete the written tasks. Reference was made to a subject's written answer only when a discrepancy was observed between that answer and that given in the dialogue. This gave the opportunity to match written answers with subjects' reconstruction of their own reasoning.

The first condition was Wason's original formulation, with almost identical instructions. They were as follows:

> Below is depicted a set of four cards, of which you can see only the exposed face but not the hidden back. On each card, there is a letter on one of its sides and a number on the other side. The letters are A and K; the numbers are 4 and 7.
>
> Also below there is a rule which applies only to the four cards. Your task is to decide which (if any) of these four cards you must turn in order to decide if the rule is true. Don't turn unnecessary cards. Tick the cards you want to turn.

negated conditionals and quantified sentences (using only *every* as in the interview protocol).

Next the subjects were given a conjunctive formulation of the rule. As for all conditions, the conjunctive task used the instructions given above with the pertinent (background) rule. It was included because the difficulty in Wason's selection task is often supposed to lie in its conditional formulation. Conditionals are acquired late (Bowerman, 1986); they are known to be linguistically complex (Comrie, 1986). So why not do the same task without the conditional? That thinking prompted the inclusion of a conjunctive condition in both the written and interview parts of this study. Moreover, Stenning and van Lambalgen (2004) conducted a written test both of a similar condition and subsequently of the same formulation used here. They posit a deontic reading behind the predominant choice for A and 4. By including the conjunctive condition in the dialogues the current study is able to further evaluate this claim and others that Stenning and van Lambalgen make about the conjunctive condition.

In the two-rule task both rules are conditional but here the subject must discern which cards can decide which one of two rules is true. Stenning and van Lambalgen (2001) developed the two-rule task with the aim of investigating the effect of presenting subjects with two rival hypotheses. This provides a means to assess the Bayesian explanation of performance offered by Oaksford and Chater (1994), but Stenning and van Lambalgen (2004) also present the task as a means to reduce non-classical deployment of truth-values which clearly play a role in the original task. Oaksford and Chater (1994) base their 'rational analysis' model of subject reasoning, on this type of behaviour. That is, subjects in Wason's task might see their task as one of providing evidence for the rule – what Wason labelled 'verification bias' – in order to 'show' that it is true, instead of, or as well as, that of seeking potential falsifiers as a means to disprove the rule. In fact, subjects should seek to check whether the rule is false, and conclude that it is true if there is no evidence to the contrary. But subjects often do not conclude truth from non-falsity, as Stenning and van Lambalgen (2001) have also observed.

By presenting subjects with two rules, and told one is true – and in fact, given the current set-up, only one *can* be true, although it's not clear subjects see this immediately – Stenning and van Lambalgen hoped to background the issue of other available truth-values for either rule. The classical normative choice is for the single card 3, which falsifies one or other of the two rules, depending on which letter is on the other side. The idea is that the potential for subjects to view their role as seeking 'verifiers' should be minimised by such a set-up.

In the 'arrow-rule' condition, instead of letters and numbers, cards had arrows and symbols ('+' or '–') on them. The motivation for including this condition was to try to minimise the possibility that subjects would choose cards simply by 'matching'. Evans (see for instance Evans, 1998) has suggested that the A, 4 choice is the result of superficial processing of the rule in which the subject simply chooses cards which 'match' the possibilities mentioned in the rule, and that deeper interaction with the material is hindered because it is so abstract, and there is no thematic link to make between antecedent and consequent. If

matching does indeed lie behind card choices, subjects should be insensitive to, for instance, negations, in the conditional clauses. Although it has been already shown that matching bias cannot explain the full range of responses to negated conditionals (Oaksford & Stenning, 1992), the idea behind matching bias, namely that subjects resort to engaging with the materials at a superficial level, is still an interesting one, especially given the results garnered in thematic versions of the task – where conceptual associations between antecedent and consequent are strong.

The purpose of the arrow rule condition was to see if matching bias would be reduced by a version of the rule in which a thematic link can be thought up about the two sides of the card. In the arrow rule, it is easy to make a conceptual association between upward-pointing arrows and plus signs, both stereotypically signalling some kind of positivity. This kind of natural alignment might aid subjects' processing of the material at a deeper level, and thereby increase the choice for an upward arrow and a minus sign, the equivalent to the A, 7 choice in the original task. In contrast to previously investigated 'thematic' materials, an effort was made here to engender a thematic link without provoking a deontic reading of the rule. However, to anticipate slightly, as the results show (see table 5.4), this did not happen – the only subject to choose the combination of upward arrow and minus sign, was also the only subject who had chosen 'A, 7' already in the original task. Because of this result, the arrow rule condition will not be discussed further in this chapter. It was also not returned to in the dialogues with all students.

The failure to improve performance is nevertheless informative. It could be the result of two things. First, the conceptual association between the upward-pointing arrow and the plus sign was not independently verified, so it might be that these materials did not achieve their aim. This seems unlikely. The second possibility is that matching bias is simply not a good explanation of what motivates the A and 4 choice. If the subjects processing was indeed so superficial because they could not construct a conceptual association between the antecedent and consequent situations, then material which does engender a conceptual association between the antecedent and the consequent should result in different choices to the original material. Since this was not the finding in the current study, it undermines the claim that the choice of A and 4 is the result of this shallow 'matching' behaviour. The finding corroborates Stenning and van Lambalgen's (2004) claim that the reason that other thematic material has aided performance is because the rule is interpreted deontically – not simply because it contains familiar material. Moreover, as we will see in the transcripts, in the overwhelmingly majority of cases the choice for A and 4 reflects full semantic engagement with the material, generating a coherent and stable interpretation of the conditional rule. It is thus not the result of superficial processing.

5.3.2 Interview materials and protocol

Once subjects were finished with their written form, the interview began. The interview protocol was as follows: the subject was presented with a rule, typed up on a sheet of paper, and actual cards, which were laid out on the table in front of the subject and interviewer. The rules were dealt with in the order given below. The subject was asked for their choice of cards, and then the interviewer prompted the subject to consider the possibilities for the back of each card in turn, and the consequences of each of these on the truth of the rule. Finally, the subject was asked to turn each card and asked for their response on turning. The subject was allowed to change their choice of card at any point, but they were always asked to restate their choice at this point. At all times, the interviewer sought to elicit and elucidate the subjects' responses without indicating approval or disapproval at specific choices. Of course this is difficult to achieve in practice, especially with the aim of tutoring the subject to insight. It is however important, because, as the results will show, subjects might achieve the classical normative analysis without, for example, selecting cards on the basis of this – see further discussion in section 5.4 and especially section 5.4.2. The interviews were filmed and transcribed.

The order of conditions in the interviews was as follows:

1. The subject was presented with the original rule. Once the subject had considered each card, and had been given the chance to change their card selection, we moved onto the next condition.

2. The conjunctive rule. Here the same procedure was followed. At this point, depending on what choice the subject had made in the original task, the subject was presented with the following three conditions, interspersed with a return to the original task:

3. The universally quantified condition: same instructions, but with the rule phrased thus:

   Every card that has an A on one side has a 4 on the other side.

4. The 'split anaphor' conditions. The instructions in these conditions were the same as those for the original condition; the rules were

   If there is an A on the face of the card, then there is a 4 on the back of the card.
   If there is an A on the back of the card, then there is a 4 on the face of the card.

5. The two rule task – as described above.

6. The drinking age rule. Here the instructions were:

> Below is a set of cards. These cards represent drinkers in a bar. On one side is written the drinker's drink, and on the other, their age. The bar has the following rule about drinking:
>
> *If a person drinks alcohol, then they must be over 18 years of age.*
>
> Determine for each card whether you have to turn it in order to check whether or not the drinkers in the bar are complying with the rule.

7. The original task with compliance instructions as used in the drinking-age rule (although this was only conducted with two subjects.)

The reasons for this order of conditions are discussed in greater detail in section 5.4.2 on the impact of the tutorial engagement with subjects.

5.4 Results and discussion

A cursory observation of the results, given in tables 5.4 and 5.3[4], already makes it hard to ignore interpretational variability; for many subjects their performance across tasks within the written work is inconsistent, in the sense that no single interpretation of, for instance, a conditional rule, can explain their choices. The only subject to give the intended normative response of A and 7 for the original rule also answers '8' in the two rule task; the only subject to get this latter task right is one of the many who chose A and 4 in the original task. Moreover, three subjects give discrepant answers for the original version and the arrow-rule versions.[5]

Comparing written answers with final answers in the interview shows that performance for the original task tends towards the intended norm as the interview progresses: eight of ten subjects settle on A and 7 by the end of the interview, whereas only one subject had originally answered 'A and 7' in the written task. The same thing happens in the two rule task: seven of the ten subjects settle on the equivalent to A and 7, the 3, by the end of the interview, compared with one subject in the written work.

I will now further analyse the results as follows. Firstly, observe the astonishing range of card choices made by subjects in the written task. In most reasoning studies, effort is made to explain the most common choice, with the rest being more or less ignored. However, the approach here is that we should operate from the assumption that subjects have engaged with the task and have good reasons for their choice; as such, *every* selection should be sought to be explained, preferably by reference to different settings of a limited set of semantic parameters. So for each condition, a range of interpretations is considered which can explain card

[4]Subjects' names have been changed.

[5]As noted, these two conditions most probably have the same underlying logical form, unlike other thematic versions of the task, such as the drinking-age rule.

Condition	original	conjunction	two rule	arrow rule
Subject				
Peter	A, 4	none	8	up, +
Ted	A, 7	none	8	up, −
Stuart	A, 4	A, 4	8	up*
Molly	A	none	8	up
Oona	A, 4	A, 4	8	up, +
Barbara	A, 4	none	3	up, +
Christopher	A, 4	A	8	up*
William	A, 4	none	8	up, +
Rochelle	A, 4	A, 4	8	up*
Philippa	A	K, 7	8	up

* indicates discrepancy between responses for Wason's rule and the novel arrow rule condition

Table 5.2: Responses to the written questionnaire

selections, and the relations between these various interpretations are considered. Evidence for each interpretation is sought in the justifications for the turnings offered by subjects in the tutorial dialogues. At the end of this section, it should be clear what interpretations plausibly lie behind the various responses in each condition and how they are related to each other. This makes it much easier to analyse data across conditions in future quantitative work.

Secondly, I discuss the responses of subjects over the course of the interview. This cleaves into two approaches. In section 5.4.2, an attempt is made to convey the extent of consistency and variety within individual subjects' responses. In the subsequent section, 5.4.3, higher-level semantic issues which play out in all subjects' transcripts over all conditions are described and analysed. In this section I follow up on the taxonomy of semantic concerns identified in Stenning and van Lambalgen's (2001) quantitative study. The focus will be on new findings which bear on our understanding of the original condition.

Finally, the implications of the foregoing analysis for the significance of the selection task in theories of reasoning is briefly considered.

Condition	original		conjunction	split anaphors		two rule	drinking age
	1	2		A on face	A on back		
Subject							
Peter	A, 7	–	A, 7	A	7 → 4,7	8 → U,I,8 → 3	whiskey, 16
Ted	A, 7	–	all	A	4, 7	8 → all → 3	whiskey, 16
Stuart	A, 4	A, 7	A, 7	A	4, 7	3	whiskey, 16
Molly	A	A, 7	none	A	4, 7	3	whiskey, 16
Oona	A, 4	A, 4	A, 4	A	4	8 → 3	whiskey, 16
Barbara	A, 7	-	none	A	4, 7	8 → U,I,8 → U,I,3	whiskey, 16
Christopher	A, 7	–	K, 7	A	4 → 4,7	8 → any → 3	whiskey, 16
William	A, 4	A, 7	none	A	4, 7	8 → all	whiskey, 16
Rochelle	A, 4	A, 4	all	A	4, 7	U,I,3 → all → U,3	whiskey, 16
Philippa	A, 7	-	none	A	4, 7	3	whiskey, 16

Table 5.3: Responses in the interview

5.4.1 Explaining the modal choices in the written task

In the following we look at the range of choices made in the written work, suggest interpretations which support these choices, and then seek evidence for these interpretations in the transcripts of the dialogues. Connecting interpretations to choices does not mean a blanket presumption of logicality on the behalf of the subjects, since the logical forms need still to be motivated, but it helps to connect behaviours across tasks and gives insight into difficulties with the task as intended. The following describes and provides evidence for the various logical forms which would yield modal answers on the original, conjunctive and two rule tasks.

Original task

The normative answer with a classical logical reading of the original task is the choice of the A and 7 cards. Yet in countless studies this is chosen by a very small minority of reasoning subjects. What are the commonly occurring choices, and what interpretation of the rule would justify them? This is what we now consider, before turning to the dialogue data to assess whether such readings are evident in the subjects' considerations.

As can be seen in the above tables, seven of the ten subjects in this study chose **A and 4** for the original rule in the written questionnaire preceding the tutorial dialogues. A fairly typical result is that around half the subjects chose this option (e.g. Wason and Johnson-Laird 1970, Stenning and van Lambalgen, 2001). As Stenning and van Lambalgen (2001, p. 291) have pointed out, the choice of A and 4 can be motivated by decomposing and normalising the conditional.

The rationale behind this is as follows. Conditionals with known consequents are very odd, maybe even downright unacceptable – as evidenced in the examples "If polar bears are difficult to hunt, then polar bears are white". Fillenbaum (1978) found that subjects paraphrase conditionally phrased threats and promises by reversing the clauses, a phenomenon he termed 'pragmatic normalisation'. Might subjects do the same thing with conditionals with consequents which are known to be true, while the antecedent is unknown?

If subjects, in the process of interpreting the anaphor 'one side . . . other side', decompose the anaphor in the original conditional into two separate conditionals, viz:

(1) If there is an A on the (visible) face then there is a 4 on the (invisible) back *and*
if there is a A on the (invisible) back, there is an 4 on the (visible) face.

and normalise it, as Fillenbaum describes, then they would generate:

(2) If there is an A on the face, then there is a 4 on the back, *and*

if there is a 4 on the face, there is an A on the back.

The reading of the original rule, generated by this split and normalise procedure, as given above, will be referred to as a **normalised conditional reading**[6] – and although it is a biconditional of sorts, it is not the same as an anaphoric biconditional "there is an A on one side if and only if there is a 4 on the other side", which would require all cards to be turned. On this interpretation, the K card is irrelevant, and the 4 card can serve as a falsifier of the second condition. Similarly to Stenning and van Lambalgen, in my data many subjects did make this distinction between the K and 4 card. Importantly, the 7 card is also irrelevant on this reading – because the normalised rule specifies nothing about the situation in which there is an A *on the back* of the card. There was direct and indirect evidence that the anaphor decomposition as described above does occur. For instance, in the below transcript, the subject reverses the direction of the conditionality when they come to discuss the 4 card.[7]

Rochelle in the original task:

E: OK. So a K here on the back of the 4 would make the rule not true?
S: Yeah.
E: But you said a K and a 4 here (indicating the K) wouldn't matter.
S: Yes. Because ... yeah no that doesn't seem to make any sense. But. It's because, it's saying if there is an A on the other side then there has to be a 4 on the other. But this is a K, so it doesn't really matter what's on the other side of it. But *if it's a 4 on this side there has to be an A on the other side of it for it to actually work.* So if I had a K there then that would be wrong.

The other possible resolution of the problem of a conditional with an unknown antecedent and known consequent would be to simply reduce the anaphoric ref-

[6]I use this terminology although it is inconsistent with Stenning and van Lambalgen (2004, p. 509), who use the term 'anaphoric' for it. I do this because the intended normative interpretation is better labelled the anaphoric conditional reading.

[7]Further evidence of this occurring, was found in both the two-rule conditon – see further on in the chapter – as well as in the split anaphor condition. All subjects chose to turn the 4 card when presented with the rule "If there's an A on the back of the card, then there's a 4 on the face of the card". The following response was typical:

Oona, in split anaphor condition:

[rule: If there's an A on the back then there's a 4 on the face of the card]
S: In that case, well *this is 4 and it's the face*, I'm assuming this is also a face (pointing to the A card) as in, according to what I am looking at this is a face, so *there should be an A on the back of this one (the 4).*
E: So you choose just the 4?
S: Yup.

And the subject also judges the 7 card *"irrelevant, it doesn't apply"*. In section 5.4.3 we will see in more detail the responses to the split anaphor condition.

erence 'one side ... other side' to the deictic 'this side ... that side', i.e. 'face ... back' and end up with a **deictic conditional reading**[8] :

(3) If there is an A on the face then there is a 4 on the back.

Given this interpretation, the correct choice is **just the A** card. Two of ten subjects chose this in the written questionnaire in the current study, and in other studies the percentage of subjects choosing this option has been only somewhat higher (33% in Wason and Johnson-Laird, 1970, and 24% in Stenning and van Lambalgen, 2001). In the dialogue, one of the subjects, having chosen just A initially, changed their choice to A and 7 in the course of discussing the cards, but not before articulating the anaphor-fixing explicitly:

Stuart, original task:

E: OK. And the 7..
S: (interrupting) is not relevant at all.
E: OK. What could you find on the other side?
S: You could find an A or a K. But *it doesn't say that if there's a 7 on one side, then there must be a K on the other side.*
E: OK. So ...
S: (interrupting) or assuming, *I am assuming, sorry, that we are starting with these cards, as in this (pointing at the 7) is a 7 card, not a letter card, it's a 7 card.* ... I know it sounds strange when I am saying it myself, because it says when there's an A on the other side there must be a 4 on the other, which means this (pointing at the 7) can't be an A, cause there's a 7 on one side, but, *I am assuming this (running his hand over the top of the cards) is the start point of each card.*
E: OK, and *if you take away that assumption?*
S: Then *if there's an A on this side (pointing to the underside of the 7), it would invalidate the rule.*

The subject is thus able to 'unfix' the anaphor, once they have made their own assumption explicit. However, this does not translate to card choice. Here we get a soupçon of Oaksford and Chater's so-called information gain strategy. The exchange continues:

E: OK. *So would that mean that you should turn the 7, or not?*
S: Well *you could turn the 7, but it says don't turn any cards you don't have to, and you only have to turn the A.*
E: OK. So the 7 could have an A on it, which would invalidate the rule, but..

[8]Again here I follow Stenning and Lambalgen's (2004) terminology although it is not clear what happens to the reference when the card is turned over. The fact that many subjects changed their choices after consideration of turning cards might be related to this point – i.e. considering the action of turning prods them out of a deictic reading of the anaphor.

S: (interrupting) It could have, but it could also have a K on it, so if you turned that (the 7) and it had a K, it would make no difference to the rule, and you would have turned a card that was unnecessary, which it says not to do.
E: But what if it had an A on it?
S: But what if it had a K on it?

We will come back to this in excerpt the discussion both of 'interference' effects, and the varying agency of the cards to prove or disprove the rule, later on.

The conjunctive condition

In this task, subjects were presented with the rule

(4) There are As on one side of the cards and 4s on the other side.

Now the intended reading of the statement was as a universal generalisation, which we could rephrase as:

(5) For all cards, they have As on one side and 4s on the other side.

On this reading, the intended normative answer is to turn no cards, since the K and 7 cards already falsify the statement. Only five of the ten subjects chose to turn no cards in the written condition; and only three of the five stuck to that choice in the interview. (The other two changed their responses to 'A and 7' and 'all cards'.) Clearly, the rule was open to interpretation, probably stemming from the rather awkward original phrasing, and specifically the mismatch between the plural 'As'/'4s' and the singular anaphoric 'one side – other side'. In the intended interpretation the anaphoric binding should occur per card, but as we will see this is not always the case. Indeed, it seems most natural to read the anaphor as referring to 'the cards' as a singular set – which would mean that 'one side' is the same side for all the cards – say the face, and 'other side' their backs. In this case the rule would only be true when there is a row of cards with an A facing up on all of them, and a 4 on the back of each card, or, vice versa a row of 4s with As on the back.

A more natural way to express the intended interpretation would be simply to say:

(6) Each card has an A on one side and a 4 on the other side.

. . . but this was judged to be too easy a task for the subject to solve. The fact that this more obvious formulation was *not* chosen might also bias subjects against taking it.

There is another, 'existential' reading of the original rule which also supports the 'no cards' choice – namely, that which takes the anaphor to refer to the

background rule ("On each card, there is a letter on one of its sides and a number on the other side"):

(7) There are cards with As on the 'letter side' and there are cards with 4s on the 'number' side.

Since both an A and a 4 are in plain view, the rule is shown to be true, without turning any cards. So the choice 'no cards' may reflect the subject's belief that the rule is already false, *or* already true. That this is a live option is supported by evidence from the transcripts, viz. the following subjects' 'realisation' that no cards need to be turned:

William in the conjunctive condition:

S: OK. Um I wasn't sure exactly what that was all about. ... *I think that's already true, cause there is an A there (pointing at the A) and there is a 4 there (now at the 4), so I guess that's already proven, just by looking at it.*
E: OK. So you don't need to turn any of the cards?
S: No, but I think I probably ticked that I did. Cause it's quite confusing.

A good illustration of why ticked boxes do not provide enough information! Furthermore, the subject's comment, "it's quite confusing" suggests that the pragmatic expectations afforded by the task clash with the choice of no cards. It is possible that subjects entertain a reading which supports the 'no cards' choice, but then dismiss it because of reluctance to give an answer which does not require turning something. In a sense answering 'no cards' might suffer from the same kind of bias as the 'no valid conclusion' option in multiple choice syllogistic conclusion tasks (described for instance in Newstead et al, 1992). There is evidence for this to be found in the transcript of Ted, who had chosen 'no cards' in the written task but when it comes to the interview is very clear about distancing himself from that option: first he chooses A, 7, then settles on 'all cards'. When the experimenter points out that this is different from his written response, the subject doesn't want to change his answer, but says his original choice of no cards "doesn't make sense in the slightest anymore". The experimenter persists:

Ted in the conjunctive condition:

[at this point subject has chosen to turn all cards]
E: So you read it as *every* card should have a vowel on one side and an even number on the other, is that what you thought?
S: This is really strange. Yeah. *I can't even work out why I would have done that now. Cause you need to turn some of them to prove or disprove it.*

Of the remaining subjects, three answered A and 4 in the written task; one answered K and 7; one answered A. In the interview four of these latter subjects changed their answers and only one stuck to A and 4. What interpretations could lie behind the choices

– A and 4,
– K and 7,
– A and
– A and 7 and
– all cards?

We now discuss each of these card selections in more detail.

The choice of **A and 4** was here made by only three subjects in the written task, and this dropped to one in the dialogue. This is a very different pattern of responses to that garnered by Stenning and van Lambalgen (2004). As mentioned above, they conducted written tests on two different conjunctive formulations. The first was

(8) There is a vowel on one side, and there is an even number on the other side.

and the second was identical to the formulation used here except with "... vowel ... even number ..." instead of "... A ... 4 ...". In their study, with the first formulation 31 of 69 (45%) subjects chose A and 4, and with the second formulation – the same as that used here – 70% chose A and 4! Stenning and van Lambalgen propose that the dominant reading lying behind this choice is a deontic one: "Every card *should* have a vowel on one side and an even number on the other" – which means that the relation between card and rules is very different than that intended. Each card is judged individually against the rule, and the truth of the rule itself is not under interrogation, as is also the case in the drinking-age formulation of the rule. Stenning and van Lambalgen propose that subjects reason that K and 7 already flout the rule, so only A and 4 are worth checking (2004, p. 515). I find it implausible that subjects would ignore the fact that K and 7 flout the rule. Rather subjects judge the K and 7 to be irrelevant, provided they do not have a 4 or an A on the other side, which would mean they plump for the choice **all cards**. This is evidenced in the following transcript:

Ted, conjunctive rule the first time round:

[subject has chosen 'no cards' in the written]
S: Um ... I'd turn... the A, ... and .... I'd probably turn all of them. (pause) I'm not sure.
E: OK let's go through all of them.
S: ... (Pointing at K) I would turn that ... because there might be an even number on the back of that, which would mean that the rule was wrong ... because there's not a vowel on the one side and an even number on the other, it's a consonant and not a vowel, which is not right. Um ...
E: *And if there's an odd number, then what?*
S: *If there's an odd number [on the back of the K], then it's fine.* Um ... (pause, finger on the 4 now) Yeah *and the same applies to the 7, so if there's a vowel on the back of the 7, then the rule's wrong, but if there's a consonant on the back, then it's right.*

Note the subject's description of the certain findings being 'fine' or 'right', not 'making the rule true'. Nevertheless his phrasing suggests he is indeed judging the rule against the cards, by a one-by-one strategy, instead of vice versa.

In the above situation it seems the rule is read as 'pairing' As and 4s. So as long as As and 4s stick together, and Ks with 7s, the rule is fine. It is thus the same as the normalised conditional reading as described in the original condition. In this case, as mentioned above, any card can falsify, but there is a difference in the verifying agency of the A and 4 on the one hand and K and 7 on the other. Turning A to find a 4 would verify the rule, while turning the K to find a 7 would exclude the rule from applying to it. At most the K/7 combination is consistent with the rule, but it doesn't provide any evidence for the truth of the in the sense mentioned in the original condition. But is this necessarily a *deontic* reading of the rule? We could also see it as a restricted-domain reading of the generalisation expressed in the rule. Recall that in Chapter 2 it was argued that unschooled subjects had difficulties reasoning with quantifiers when the context set was unspecified and often attempted to clarify the context set before reasoning with the quantified premise. Could the same behaviour be happening here? That is, the subjects in the current study, when confronted with the rule, which I have labelled conjunctive, for continuity's sake, but which could just as well have been labelled existential, struggle to figure out the intended context set in which the rule applies. Cards, such as the K and 7, which obviously cannot comply with the rule are disqualified from the domain of 'application' of the rule. *Within* this domain, the truth of the rule is assessed, so it is not a straightforward card-by-card checking task, as a deontic reading suggests.

If the deontic reading is indeed employed we should expect to see subjects importing modal verbs into their reformulations of the conjunctive rule, and justifying their choices as means to check the cards, not the rule. As it happened, only one subject chose A and 4 in the dialogue, so there is scant material in which to seek support for Stenning and van Lambalgen's claim that a deontic reading lies behind the choice for A and 4. Moreover, there are other readings available which justify a choice for A and 4 without positing a deontic reading of the rule.

For instance, choice of **A and 4** can also be understood as resulting from a reading of the rule as applying *only* to A or 4 cards, as such paraphrasing as

(9) For all cards, if they have an A or a 4 on one side, then they have an A *and* a 4 on them.

Now this reading strictly requires all cards to be turned – as any card potentially falsifies the rule – unless we again assume the subject exercises Fillenbaum's pragmatic normalisation, which restricts the antecedent reference to known cases, so that the rule becomes

(10) For all cards, if they have an A or a 4 on their face, then they have an A *and* a 4 on them.

This, in turn, is effectively the same as the previously mentioned normalised conditional reading:

(11) For all cards, if there is an A on the face, then there is a 4 on the back, *and* if there is a 4 on the face, there is an A on the back.

Stenning and van Lambalgen (2004, p. 521) acknowledge the link between a conjunctive suppositional interpretation of a conditional – in which one assumes the truth of the antecedent and answers subsequent questions "from within this suppositional context" – and the conjunctive formulation. In the current study, four of the ten subjects in the immediate inference task indicated that a conjunctive formulation is entailed by a conditional.[9] The link the other way is made explicit by the next subject, who, when presented with the conjunctive rule in the interview, having chosen A and 4 in the previous original condition, proclaims:

Oona, conjunctive condition:

S: Exactly the same, I'd turn these two (the A and 4), 'cause there are As on one side and 4s on the other. *It's the same statement, just written in a different way. Isn't it? Because they've missed out the 'if', that's all, that's all they've missed out.*

The subject here makes some comments about doing tasks like this before; she mentions medical stats material, aimed at testing sensitivity to presentation and framing effects. This may be relevant because if she conceives of this task as similar, then it is an exercise in interpretation of the rule, and particularly across conditions, and not so much to do with card selection. As such, her choice of A and 4 is a result of interference from the previous condition, and might be different had the conditions been presented in a different order. The subject makes several comments about "reading too much into [the rule]" and "trick questions"; these suggest she indeed sees the task as a means to test framing effects. She also justifies turning the cards "out of curiousity"; this makes clear she doesn't see card choice as the main focus of the task.

There was one subject who reiterated the choice of A and 4 before changing to 'all cards' and her justification certainly supports a suppositional conditional reading of the rule. Although she imports the modal verb 'should' into her for-

[9]To be more precise: the truth of "If there is a vowel on one side of a card, then there is an even number on the other side" entails "There are vowels on one side of the cards, and even numbers on the other." Interestingly enough, Stuart offered an additional statement which was entailed by the conditional: "If a consonant is on one side of the card there may be an even or odd number on the other side". This actually contradicts his selection of the conjunction, unless we assume an existential reading of the conjunction. It also suggests that the immediate inference task might be interpreted by the subjects not as the experimenter intends it – that is, interpretation of the original statement should be treated as fixed while the other statements are examined. However, subjects might be reinterpreting the original statement anew with each given option.

mulation, note that it is here ambiguous between an epistemic and a deontic reading.

Rochelle, conjunctive condition:

S: I think you need to turn the A and the 4 again.
E: OK. Why?
S: Cause *if there is an A, if there's, that's an A then there should be a 4 on the other side*, and the same with that [4]. But *with those two, the K and the 7, it doesn't matter.* I don't think. Well it would prove that it's untrue, but if we want to prove it's true, then ... (gestures with her hand)

This subjects use of the phrasing 'proving true' and 'proving untrue' are discussed in section 5.4.3; below we gone on to discuss the plausibility of a deontic reading.

Further, a choice for **A and 4** is also consistent with an existential reading of the rule, i.e.

(12) There are cards with As on one side and 4s on the other side.

In this case, the cards are turned in the search for evidence and the proof that the rule is true takes the form of an existence proof. A subject might want to choose cards conditionally here, i.e. it would only be necessary to turn a second card if the first one did not witness the rule. Since only one subject maintained the choice for A and 4 in their interview (see above) we are left without evidence that such a reading justified any subject's choice of A and 4 in the written task. Regarding the formulation used here, Stenning and van Lambalgen state, without giving a reason, that "it is implausible that this rule might be interpreted existentially" (2004, p. 522).

Also the choice of **just A** can be seen as an existential reading but this time coupled with a fixed, or in Stenning and van Lambalgen's (2001) terminology, "asymmetric" interpretation of the anaphor as 'face – back', so that the rule is reformulated:

(13) There are cards with an A *on their face* and a 4 *on their back.*

With this reading of the rule, only the A card is relevant as a potential 'witness' for the rule. It would be interesting to see what happens if a subject should turn the A to find a 7 – would this mean the rule is false straightaway, or would the subject rather unfix the anaphor and then choose to look at the 4 card to see if it could witness the existential? The only subject to choose just A in the written condition changed immediately to a choice of K and 7 in the interview – see the excerpt from Christopher below for details, without any reference to his earlier choice. As such we have no evidence that this formulation lies behind his initial choice for just A. In fact, given the close connection between the conjunctive and

conditional formulations elsewhere, it might be more prudent to suppose he is taking a conditional face reading of the rule, that is

(14) If there is an A on the face then there is a 4 on the back.

The subject did make an analogous choice for just the 'up' arrow in the arrow-rule condition, but did not choose 'just A' in the original formulation of the conditional.

The choice of **K and 7** is at first puzzling. An interpretation of the rule which supports it is

(15) There are As or 4s on one or other side of the cards.

The subject (Philippa) who chose this in the written questionnaire had chosen just A for the original condition, which is interesting because, as we saw above, this might result from a simple reduction of the anaphoric 'one side – other side' to 'face – back'. A disjunctive reading of the conjunction rule also indicates insensitive handling of the anaphor, here by basically treating 'one side – other side' as 'one or other side'. In fact, I think this could be the source of the disjunctive element in the posited interpretation, and the conjunction in the original formulation is able to be rephrased as a disjunction because of a 'free choice' reading of the set {A, 4}, viz:

(16) On one or other side of each card, there is an A or a 4.

This might seem far-fetched, but, as we'll see, the transcripts suggest this is indeed behind the choice for K and 7. In particular, the subject mentioned above changes her answer to 'no cards' in the interview, but recalls a disjunctive reading motivating her written answer.

Philippa, in the conjunctive condition:

S: So that's saying every card ... [rereads rule] but then *K and 7, that's very blatantly obvious that they don't have As and 4s.*
... So I don't need to turn anything, cause that's rubbish. Cause there's a K and there's a 7.

E: [In your written] you put K and 7.

S: *I think I read it as each card either has an A or has a 4*, so that means that this [K] could have had an A or a 4, and this [7] could have had an A or a 4. These two (the A and 4) definitely follow the rule. That's how I read it. But now I read it out loud ... it says 'and', 'As and 4s', so I am going against what I wrote down.

In the dialogues, the following subject switched to a choice of K and 7, with the following justification (the subject ticked A in written task; has just done original task correctly):

Christopher in the conjunctive task:

S: OK. Well this is basically saying that ... so this means that *there is going to be at least an A or a 4 on each card.*
E: What, this rule?
S: *This rule says that there is going to be either an A or a 4.* So which would mean there'd be a 4 here (pointing to the K) and a A (on the 7) here.
[so K and 7 must be turned]
E: And A and 4, we don't have to turn?
S: Um, no, I don't think it would be necessary, because *they [A and 4] already conform with the rule as it were.*

Finally, two subjects chose **A and 7** in the dialogue and their justification reveals a conditional reading. This is plausibly an interference effect from the previous condition, as identified in Oona's responses – see section 5.4.2, or a result of imputing precedence order on the basis of the order of the conjuncts. Conjunctive phrasings often carry an implicit temporal ordering, as in

(17) They got married and had a baby.

is read as

(18) First they got married, and only after that they had a baby.

Similarly, in the case of the conjunctive rule, subjects might assume the order of the conjuncts confers some kind of precedence order on them, meaning that sentence (4) is interpreted as

(19) First there are As on one side, and only 'after that' [when that is the case] there must be 4s on the other side.

This aligns well with the following subjects' justification for the irrelevance of the 4 card.

Peter in the conjunctive task (formulated with vowels and even numbers):

S: OK, well ... I think you would have to turn the A over again, cause if it was an odd number that would disprove the rule. So you need to check that. Umm. You'd need to check the 7 as well.
E: Why?
S: Cause *if there was a vowel on the other side [of the 7] there'd be a vowel and an odd number.* Which again disproves the rule. (pauses) I don't think the 4 matters because, ... if it wasn't a vowel, if it was a consonant, then it wouldn't disprove the rule, *it would just be an even number on the back of that [the K], which doesn't affect the rule.* And (pause) I don't think the K matters either.
...

E: So that combination, the K and the 4, doesn't disprove the rule – is that what you are saying?

S: Yes, because it doesn't say, that, (pause) erm, any even number on one side has to have a vowel on the other. ... *It just says if there is a vowel there has got to be an even number on the other side.*

In the following example the conditional is not explicitly formulated but the emphasised text suggests that subject understands that the rule *does* say that 'if there is an A on one side there must be a 4 on the other':

Stuart in the conjunctive task:

S: You need to turn the A, to check if there is a 4 on the other side, which would validate it. You don't need to turn the 4, because *it doesn't say if there is a 4 on one side there must be an A on the other.*
[subject proceeds to reformulate the consequences of turning the A: a 4 would 'concur with', rather than 'validate' the rule]
... if you turn the 7 (pointing at the 7), initially, there has to be a K. Because if you turn the 7 and there is an A, that invalidates the rule. So *you should turn the 7 as well.*
...

S: Yeah, *if you turned that (the 7) and that (the A) you could make your mind up whether the rule was right or not.*

The amazing variety of readings for the conjunctive formulation of the rule shows that it was by no means a straightforward condition for subjects. The arguments from the previous chapter, namely that grammatical form can belie logical form, and that the 'literal meaning' of a given statement is not always apparent, and even when it is, it is certainly not always available to highly literate subjects in reasoning tasks, are here amply supported by the excerpts. As such, the overall finding here concurs strongly with Stenning and van Lambalgen's declaration that: "[the conjunctive condition illustrates] how unnatural it is for naive subjects to adopt an 'is-this-sentence-literally-true' perspective rather than a 'what-are-the-experimenter's-intentions' perspective" (2004, p. 521). This is, moreover, a verdict highly redolent of our findings in chapters 2 and 3 with less schooled subject groups. In fact, one could easily mistake Stenning and van Lambalgen's (2004, p. 520) general verdict regarding their dialogues with undergraduate students, that "the interpretation of sentence semantics is highly malleable under the forces of task pragmatics", for a verdict from earlier reported studies from Scribner or even Luria. The earlier proposal that 'linguistic' culture supercedes other cultural factors in explaining reasoning behaviour seems to be only further supported by the findings of the current chapter, where we see highly educated subjects' are concerned with the same parameters as the subjects of Chapter 1: what is my role in the task? how should I interpret these sentences in line with this? what adjustments to semantic parameters would achieve this?

The two-rule task

In my study, all except one subject chose only the 8 card in the written version of the task (see table 5.4). These results are very different from those garnered by Stenning and van Lambalgen (2004, p. 516), where the most popular choice was for just the 3 card, and they found as many subjects choosing a letter card in combination with 8, as choosing just the 8 card. As such, the (limited) current findings are unable to provide support for Stenning and van Lambalgen's conclusion that (2004, p. 517) this manipulation of the original task is "substantially easier" than the original task.[10] Note, though, that in the course of the interviews many subjects changed their minds to either 'all of U, I and 8', or 'any one of U, I and 8', and other combinations before settling on just the 3 card. No-one stuck to their original choice of only the 8 card. The implications of this are considered below.

The choice of **just the 8** card is logically consistent with the modal choice of A and 4 in the original task, as is apparent when one considers the normalised conditional interpretation of the original rule, described above. Positing the same reading of the two rules in this new task explains the choice for just the 8 card in the following way.

Suppose the subject assumes a normalised conditional reading for both rules. Their interpretation can be paraphrased and results in card choices as follows:

(20) **Rule 1**:
If there is an U on the face, there is an 8 on the back *and*
if there is an 8 on the face, there is an U on the back
Rule 2:
If there is an I on the face, there is an 8 on the back *and*
if there is an 8 on the face, there is an I on the back

Analogously to the original task, the first rule requires U and 8 to be chosen; the second rule requires I and 8. Now the subject may use the given information that one rule is true to deduce that just the 8 card needs to be turned, since turning it will tell you which of the rules has a false second conjunct. This information rules out one of the rules and allows one to conclude the other is true.

This interpretation of the rule also suggests that subjects attribute the property of *strong falsity* to the conditional – that is, the falsity of "if p, then q" entails that "if p, then not-q" holds. In the immediate inference task of the current study, five subjects indicated a strong falsity reading of a negated conditional. Stenning and van Lambalgen (2004, p. 499) also identified strong falsity in many of their subjects; Fillenbaum (1978) observed the attribution of strong falsity in the majority of his paraphrasing subjects. As is evident in the transcripts below, this

[10] Stenning and van Lambalgen's (2004) conclusion is based on a pen-and-paper task with a larger subject group; my results only provide the suggestion that their findings are less equivocal than might have seemed until now.

leads subjects to 'expect' certain outcomes once having turned the 8 card. In the two-rule task strong falsity is equivalent to an assumption that the cards are split into two *types* – i.e. either Us are always paired exclusively with 8s, or Is are, and vice versa regarding the 3. The subject's task is then to figure out which pairing holds. However, rather puzzlingly, on the type reading *any* card is decisive; it is perhaps why some subjects switched to 'any card' after first choosing just 8. The 8 card might be seen as more informative because it falsifies one rule and *simultaneously* witnesses the other. This is something the other cards don't do, and as we'll see later, subjects are reluctant to call a conditional true without having a witnessing case. Note the similarity here to the interpretations given to the conjunctive rule. This is evidenced in Philippa's transcript below.

How can the **U, I and 8** combination be understood? In my data the switch to U, I and 8 was often made after the subject was confronted by the experimenter with the possibility that the cards are *not* sortable into types – such as, that a letter card could have a 3 on the back, even while the same letter was behind the 8, or by possibility of finding 8s behind both letter cards. It confuses subjects because it undermines the agency of the 8 card to decide between the two rules, thus often leading to conclude that the task set-up is flawed. Subjects then cease to rely on the given information that one rule is true. Even though one of the rules has been dismissed, by the 8 card, the subject might feel that they need to verify that the other rule is true, by checking that there is not a 3 card behind the letter face. If this is indeed the reason for the choice, then this might be expressed as a dependent card choice, and U, I should be seen as a projection of *either* U or I, depending on what's behind 8.

Barbara, two-rule task:

[subject chose 3 in the written questionnaire, has chosen just the 8 card]

E: On the other side of the 8 what could there be?

S: A U or an I. And um, *depending on which letter there was it would tell you which rule applied.*
... *if there was a U on the other side of the 8, then to prove rule 1, you'd expect there to be a 3 on the other side of the I*, then you could prove that rule.

E: OK. And what if there wasn't a 3 on the other side of the I?

S: If there was a 8?

E: Ja, say there was an 8. Or do you think that means there's something wrong with the cards?

S: Yeah, *well if only one of the rules is true, it would suggest that there's something wrong with the cards, as it would mean that both the U and the I have an 8 on them.*

E: ... If you had to decide which ones to turn before you turned any of them, which ones would you choose?

S: *I'd turn over the 8. And I suppose you'd turn those two (the U and I) over as well.*

William in the two-rule task:

[subject has chosen just 8]
E: OK. So let's actually turn these.
[turn 8, U to reveal U, 8 respectively]
E: ... OK. *Now we turn the I and find an 8. What does that tell you?*
S: *It tells you that your statement is wrong.*
E: Which statement's wrong?
S: *The statement that tells you one of the two rules holds true.*
E: Oh right, OK, *so in fact you would want to say that two of the rules hold true?*
S: *Yeah.*

The subject below struggles with issues which recurred in many transcripts. Firstly, she is unsure whether the U and I cards need to be turned. As she points out, an 8 on the back of either of them doesn't prove anything – so finding an 8 behind I "doesn't disprove rule 1, but fits nicely" with rule 2. However – and here we need to posit a normalised conditional to make sense of her reasoning – finding, say an I on the back of the 8 *does* provide definitive proof that rule 1 is false. In other words, the I/8 combination has different agency, depending on which is on the face, clearly indicating she is maintaining the normalised conditional reading as described above.

Philippa, two rule condition:

S: Well *this one here (the U) I'd obviously have to see if there's an 8 on it. If there's an 8 on it ... well that's nice, doesn't actually prove the rule, just fits with the rule. If there is not an 8 then that disproves the rule, so that's crucial.* I'd also have to turn this one (the I), because if there wasn't an 8 on the other side of this then that disproves rule 2. *If there was an 8, it ... doesn't disprove rule 1, but it fits nicely with rule 2. This (8) is also important because i f I turn this over and it was an I, then that actively disproves rule 1 because rule 1 says if there's a U on one of the sides there must be an 8.* So maybe just that ...I think actually maybe just this one (the 8) is crucial. ...So *I only need to do this one (the 8) because it tells me that one of the rules is definitely correct.*
E: OK so say you turned that (8) over and found a U and an 8?
S: Yep. Then I know that rule 1 is correct and rule 2 is wrong.
E: OK. So just the 8 then?
S: *Because if they hadn't told me that one of the rules was definitely true, then I'd have to do all three (U, I, 8). But because I know one of them is true, I only have to do the 8.*

This last turn deserves attention. From the subject's earlier formulation that the 8 card "tells me that one of the rules is definitely correct", one might think

she is exhibiting confirmation bias. But this last turn shows that she realises the 8 card only has this agency because of the task situation, i.e. one rule is "definitely true". This can only be understood if we assume she reasons that the 8 card *disproves* one rule by falsifying the reverse conditional ("if there's an 8 on the face, then ...") for one of them, and it is this, coupled with the information that one rule is true, which allows her to conclude the other one is true. Indeed, this reasoning is consistent with her previous explanation that an 8 on the back of the I "fits with the rule" but doesn't "actually prove it". The continuing irrelevance of the 3 card also strongly supports this normalised reading.

Turning the cards leads to the subject abandon the normalised reading and to reassess the best way to disprove a rule:

> S: So if there's a U then there's an 8 ... and I just said rule one is correct ... [subject now turns the I to reveal an 8] OH! *Oh I didn't think about the fact that they weren't mutually exclusive.* Oh ... Oh but ... [turns 3 to reveal a U] ... OH! aha ... very clever ... so actually rule two is correct and rule one is wrong. Yup, because I thought they were like mutually exclusive, ok.
>
> E: *What do you mean, mutually exclusive? If the 8 was with a U then the 3 would have an I?*
>
> S: *Yeah, that's what I thought.* So actually I should have turned all four of them, I think. Well because ... or maybe I could have just ... if I turned the 3, I didn't even think about the 3 ... If I turned the 3 ... if this had been an I that would have proved that rule two was wrong. And if it had been a U, it would have proved that rule one was wrong. *This one (3) was the crucial one.*

Given Philippa's previous reasoning, finding an 8 behind both U and I needn't undermine her choice of the 8 card; she does however seem nonplussed by the finding. This suggests that either the normalised reading of the conditional attributed to her is inaccurate, which seems unlikely given the above analysis, or, more likely, that she generates the conclusion that the rules are made according to 'type' during the course of the task, from her previous conclusion about which rule is true. In other words, the falsification of one rule, achieved by turning the 8 card, prompts her to a strong falsity reading of the conditionals. Stenning and van Lambalgen (2004, p. 500) reported that when subjects were confronted with the possibility of an 8 behind both the U and I cards, they became very confused. Stenning and van Lambalgen categorise this as inappropriate transferral from 'truth of the card' to 'truth of the rule', which is certainly an accurate description of William's testimony above, but in Philippa's does not seem to be the case. In fact I think this is rather a peculiarity of the two-rule task, where subjects seem to interpret the task as 'which rule provides an accurate description of the cards?'. The information that only one rule is true, coupled with the conclusion as to which one is true, in the course of the task prompts a 'type'

reading of the cards; this is what in turn engenders the strong falsity reading of the conditional.

There is more discussion about subjects' understanding of truth in section 5.4.3.

5.4.2 Tutoring to insight?

In the description of the experiment I mentioned that a subsidiary aim of the interview was to engage with the students until they reached the intended normative answer of the original task, achieved by so pitifully few subjects in pen-and-paper studies. The underlying question was: what would it take to get subjects to do the task as Wason intended, if indeed they can do it at all? The protocol was designed with exactly this in mind; first, by drawing the subjects' attention to each card in turn. Perhaps this would be enough to get them to see the agency of the 7 card. Next, the variations on the original rule might encourage the student pay attention to exact wording for each task and thus be more inclined to see the 'literal meaning' of each rule – without forgetting, as we've seen in the previous chapter, that this is a notion which is theoretically and contextually determined. Additionally here, the universally formulated rule "Every card that has an A on one side has a 4 on the other side" serves as a check whether subjects who have already chosen A and 7 will reiterate the choice in the differently-worded version. Third, the original rule was split into the two anaphoric referent versions, as described in the materials section (see section 5.3.2), and this was presented to subjects to provoke them to consider both possibilities in a return to the original task. Finally, if all this failed, the protocol included the drinking-age rule with the hope that cross-task transfer would occur, i.e. that subjects would see the parallels between the violating agency of the 16-year card and the falsifying nature of the 7 card, and choose the 7 card accordingly.

In fact, many of these interventions proved unnecessary, as five of the ten subjects settled on the combination of A and 7 as their choice in the first round of tutorial dialogue. Only one of these subjects had already chosen this in the written test. Three of the others had chosen A and 4; one had chosen just A. Often, the subject changed their mind simply in the act of discussing what could be on the other side of each card. This is evident in the following transcript, where the subject first retracts his choice of 4 card, and then adds the 7 card as soon as he turns his attention to it

Peter, original task first time around:

[subject has ticked A, 4 in written]

S: (spontaneously) ... I need to turn the 4 over to check there is an A there, cause there could be another letter.

E: And what would it mean if there is another letter [on the back of the 4]?

S: Um.

E: If you did turn the 4 and it was a K, it wasn't an A?
S: If it wasn't an A then the rule mustn't be true (frowns), I don't think, cause, if there's an A on one side, or there's not a *... Hmmm ... Actually it wouldn't disprove the rule, cause, just cause it's not an A, it doesn't mean that all As don't have 4s on the back.* So *even better in fact, you only have to turn one card, the A.*
E: Ok so now let's go through each of these cards.
...
E: OK. So that [K] isn't relevant?
S: No. And I don't think the 7 either... *ah no, actually I do think the 7 is.*
E: OK, why's the 7 relevant?
S: Because, *if there's an A on the other side, then there's not a 4.* It's not, that like ... *it'll be a 7 and an A, rather than like a 4 and an A, which the rule states is the case.* So, *third time lucky, I think the A and the 7*, you've got to turn over the prove or disprove the rule.

What is interesting in this transcript is that the subject comes to 'insight' himself, both with regards to the 4 card and the 7 card. All the experimenter does is prompt the subject to consider each possible combination of front and back explicitly. This seems to suggest that Wason's intended interpretation is particularly unequivocal and stable; and that once subjects have 'seen' it, they will choose it above others and let it guide their card selection. Perhaps surprisingly, however, this is by no means always the case. Rather, the insight seems extremely local. Firstly, the subject above goes on to choose A and 7 in the conjunctive rule condition which is presented directly following this exchange, suggesting that he is still applying this same conditional reading to the now conjunctive rule. The subject then goes on to choose the 8 card in the two-rule task (the equivalent to the 4 card in the original task). This shows that we should not be too quick in attributing to the subject some general insight which then drives interpretational strategies, as might be formulated in terms of 'paying attention to literal meaning' or whatever. In some cases it seems like the subject has to reassert this seemingly rather unnatural choice against a default, i.e. the normalised conditional which supports the choice for A and 4. At any rate, the data suggest that the subject's interpretation of the conditional varies per task.[11]

In fact, these tutorial interventions proved to be instructive because of how much they reveal about what van Lambalgen (2003) labelled "the unbearable lightness of interpretations". Subjects who had chosen A and 7 went on to choose the equivalent of A and 4 in the arrow-rule, described above, and in the quantified 'every card' version of the rule, and the 8 card (the equivalent choice to the 4 card in the original task), in the two-rule task. Even the solitary subject who had

[11]Thompson (2000) reports an analogous finding reported in a more general context – subjects' responses to the same material varies according to the task setting, i.e. suppression-effect task requirements compared to selection task requirements.

chosen A, 7 in the original task right from the written task and reiterated it in the tutorial, chose the 8 card in the two-rule task. Also, in the split-anaphor version of the original task, specifically with the rule "If there is an A on the back of the card, there is a 4 on the front", many subjects chose the 4 card, again, going against earlier 'insight' that this card was not able to give definitive evidence that the rule was either true or false (see more discussion in section 5.4.3 below, on confirming versus falsifying). This is not to say that subjects are just being fickle in their choices; also below we'll see why the 4 choice in the 'A on back' rule makes sense even when you've settled on A and 7 in the original task.

On the other hand, as we have seen above in section 5.4.1, there were several subjects who grasped the agency of the 7 card – namely, it potentially falsifies the rule – but do not think that this means it should be turned. In the given case, the subject's reasoning aligns well with a phenomenon identified in Stenning and van Lambalgen (2004) regarding dependencies between card choices (p. 503):

Stuart, original task:

S: Then *if there is an A on this side (pointing to the underside of the 7), it would invalidate the rule.*
E: OK. So would that mean that you should turn the 7, or not?
S: Well *you could turn the 7, but it says don't turn any cards you don't have to, and you only have to turn the A.*

Stenning and van Lambalgen suggest that this phenomenon might be a way to limit the cognitive demands of the task. This suggestion is supported by the current study, for if we look at the above exchange in context – see below, we see that the subject is first fixing the anaphors, then perfects the conditional to a biconditional when this assumption is dropped. In other words, turning the cards over, turning 'face' into 'back' and vice versa, also changes the direction of the conditionality. It seems the subject needs to set these parameters to reduce the cognitive demands of the task, in its intended form, which require the conditional is read asymmetrically while the anaphor is simultaneously read symmetrically i.e. 'one side' refers to both 'face' and 'back' of the card. Rather, some subject seem to cycle between the choice for 'just A' and 'A and 4'. Exactly the same behaviour is identified in Stenning and van Lambalgen (2001, p. 288 and also 2004, p. 510), labelled 'interference effects'.

Stuart, original condition continued:

[in the course of considering cards he has chosen just A]
E: OK. And the 7 ...
S: (interrupting) is not relevant at all.
E: OK. What could you find on the other side?
S: You could find an A or a K. But *it doesn't say that if there is a 7 on one side, then there must be a K on the other side.*
E: OK. So...

S: (interrupting) or assuming, *I am assuming, sorry, that we are starting with these cards, as in this (pointing at the 7) is a 7 card, not a letter card, it's a 7 card.* And this rule is only saying if there is an A on one side there must be a 4 on the other. *So if you turned this [the 7] and found an A, then the rule isn't saying that if there is a 7 on one side, then there cannot be an A on the other...*

In the last turn we see the subject ably making explicit his assumption that the anaphoric 'one side – other side' is read asymmetrically, that is, it reduces to 'face – back', and when he is asked to reason without it, he accordingly changes his verdict on the agency of the 7 card, but then immediately switches to the assumption that the condition must be read symmetrically, i.e. as the normalised conditional interpretation as described above:

E: OK, and *if you take away that assumption?*
S: Then *if there is an A on this side (pointing to the underside of the 7), it would invalidate the rule.*
E: OK. *So would that mean that you should turn the 7, or not?*
S: Well you could turn the 7, but it says don't turn any cards you don't have to, and you only have to turn the A.
S: *With the assumption, that this [the top of the cards] is the start point, that's (the A) all you need to turn.*
...
But, *if we are not using that assumption, then you would have to turn this, 4, as well, and if this had anything other than an A on it, then it would invalidate the rule.*

This last excerpt encapsulates the 'interference' phenomenon: giving up the asymmetry of the anaphor means giving up the asymmetric (unnormalised) reading of the conditional; conversely, giving up the normalisation means giving up the symmetry of the anaphor and reducing it to a face-back reading. Both subjects who never reached the normative A, 7 choice exhibited this behaviour; that is, they seemed unable to uncouple an asymmetric reading of the conditional – from an asymmetric, 'face-back' reading of the anaphor. So for instance, the subject below has reached the insight that her earlier 'perfecting' of the conditional – note her use of the word "so" – is unwarranted:

Oona, original task, after several attempts:

S: ...it depends how you read it. I'm saying if you read into it, way too much, then you'd say, ah, *if there's an A on the one side then there's a 4 on the other side, so if there's a 4 on one side there should be an A on the other side.* But not necessarily, because where does it say that there has to be a 4? If 4 is facing up why should A be on the back side of it?

Nevertheless, when the experimenter asks the subject to rechoose *without* making the assumption that 'if 4 is facing up there should be an A on the back side', the subject switches to a fixed anaphor reading:

E: So take that insight, well what you just said, 4 doesn't have to have A on the other side of it, then what about the 7?
S: Yeah,... cause it could have an A on the back, couldn't it? But we are not talking about cards which have 7 side up.

She then goes on to reiterate her choice for A, and 4, and *not* 7.

There is another possible source of subjects' 'defaulting' behaviour to the modal A, 4 choice, and the posited normalised conditional reading which motivates it. Above the possibility was raised that constraints on cognitive load might be operating to cause subjects to couple a symmetric (i.e. biconditional) reading of the conditional with a symmetric reading of the anaphor. A complementary suggestion is that this coupling might reflect more typical use of conditional sentences. In the Wason selection task, the directionality of the conditional is at variance with the symmetry of the anaphoric element 'one side ... other side'; when 'one side' refers to the back, the variance is even stronger, because the conditional runs from an unknown antecedent to a known consequent (we can see what is on the visible face of the card, and not what's on the back). It was mentioned that conditionals with known consequents are often marked, or even unacceptable. In seeking to make sense of the materials, subjects may reduce their interpretational range to those in which the antecedent is known, and the consequent unknown, so that the 'face' card is the only one to ever appear in the antecedent of the conditional rule.

Again we are provided with an occasion where the similarities between subjects across the educational range are apparent. Recall the argument of Chapter 2, which was that especially unschooled subjects' difficulties with *all* sentences was to a large extent due to the mismatch between their use in reasoning tasks and the way they occur in everyday language. Here we have seen that when highly literate subjects are presented with a conditional with an anomalous informational structure, they struggle to interpret it 'literally' and seek rather a reading which more closely resembles everyday language usage. The challenge for these subjects, as it was for the subjects of Chapter 1, is to figure out for which key parameters the problematic premise differs from its everyday cousin, and then adjust these to solve the task. As we also saw in the earlier subject group, the achievement of the normative intepretation seems to be very much an individual affair. The following section illustrates this in reverse.

Does the drinking age rule help?

As already explained, all but two subjects had already settled on the choice for A and 7 before they were presented with the drinking rule. So there was negligible opportunity for investigating whether the parallels between the tasks could be used to induce the normative choice in the original task. However, where there *was* the opportunity, it appeared that no amount of parallel-drawing helped the subject, even when compliance instructions (asking "whether the cards comply

with the rule" instead of whether it is true) were used. The difficulty very clearly lies in the need to decouple the anaphor from the conditional, not in the non-compliance of the 7 card.

Oona, in the thematic then abstract tasks:

E: OK. So you think whiskey and 16 years.

S: Those are the most relevant ones, yeah. Because once you are over 18, no-one's going to tell you that you have to drink alcohol and you can't drink lemonade anymore. What happens if you really like lemonade? ...

E: Now the instructions are slightly different. You have to determine for each card whether or not these cards obey this rule (the original rule).
[subject says A needs to be turned, K doesn't]

S: This one (the 4) you do need turn over to make sure the card obeys the rule. It's that if thing again. Actually no, it's not the if thing. If there's an A on the one side then there's a 4 on the other side. It doesn't say if there's a 4 on the one side there'll be an A on the other side.

E: Mhm. So the 4 could K or A.

S: I would expect it to have an A though.

E: But if it doesn't, does it matter?

S: If it doesn't, it doesn't matter, because the rule is saying, if there's an A on the one side then there's a 4 on the other side, so A is the most relevant card here, but 4 still comes into it.

E: And 7?

S: Nope. 7 and K don't matter. Unless.... you turned over 7 and found an A!

E: So say you turned over 7 and found an A.

S: It still doesn't matter, because the thing doesn't say, if there's a 7 on one side, then there'll be an A on the other side.

E: OK. But say, say, [bewildered] um ...

S: If I turned 7 over and found an A, what would I do?

E: Ja. Let's turn it over [turns the 7 to reveal an A]

S: Oh no! There's an A.

E: What does that mean.

S: [long pause] That ... it's not complying to the rule. Because say you reverse the statement, there's a 4 on one side, there should be an A on the other side, but there's not. There's an A on one side and 7 on the other side. What's that 7 doing there? How did it get there?

E: So knowing now, what you know about that 7, would you want to turn it over?

S: Yeah. I would now, cause I know there's an A on the other side.

Only when she sees her own perfection of the conditional does she conclude that the 4 card is unimportant, and she has to reestablish this insight in the very next step, when considering the 7 card. The subject goes on to retract her choice of the 7 card, because she can't be certain that it has an A on the other side. When the experimenter tries to provoke her to think otherwise by comparing the drinking age rule to the original task, she is quite explicit that the two are not the same:

> S: ... this (pointing to the drinking age cards) is different, because *if you are 19 it doesn't matter what you are drinking.* Just like if you are drinking lemonade it doesn't matter what age you are. If you are drinking whiskey it does matter. If you are 16 it does matter. But *in this case (pointing to the AK47 row), if there's an A on one side, and there is a 4 on the other side, then there should be an A on the other side of this one (the 4).*

This excerpt provides irrefutable evidence that the abstract original rule condition is of a different order of complexity to the thematic variations, here in the form of the drinking-age rule, and that, crucially, this complexity stems from the combination of the anaphoric 'one side–other side' with a conditional formulation.

Understanding falsification *vs* verification

Several subjects *did* exhibit increasing grasp of central concepts needed to perform in the task, specifically the asymmetry between confirming or validating a rule and falsifying or invalidating a rule. The subject below initially has acute difficulty in fixing his concepts. Confusion stems from the lack of distinction between confirmation and proof: 4 and 7 have equal status to 'validate' and 'invalidate' the rule respectively. This has been problematic earlier on in his transcript, evidenced by his inconsistent use of the word 'validate'. In the conjunctive condition, he first of all says a 4 on the other side of the A 'validates' the rule, then corrects this to 'concurs with' when the interviewer repeats it back to him. The last turn in the second excerpt below provides clear illustration of his struggle to draw adequate conceptual distinctions; he finally manages it.

Stuart in various conditions:

[conjunctive condition]

E: *So now let's say we turned it (the A) and found a 4. You said that validated the rule.*

S: *It concurs with the rule, yeah.*

...

[original task again: subject has just said all cards need to be turned]

E: OK so you turn the letter A because if the A has a 4 what does that mean?

S: It means the rule is right.
E: Is that enough then?
S: No
E: *OK and if the A has a 7?*
S: *It means the rule is wrong, or it implies the rule is wrong, it doesn't mean it. Actually yes it does. If it (the A) has a 4, it implies the rule is right but it doesn't mean the rule is right. But if it has a 7, it means the rule is wrong.*

This excerpt is interesting also from a linguistic point of view, for the sheer variety of terminology the subject tries out. First he uses the word 'validate' to describe the A/4 combination in relation to the rule. When this is repeated back to him, he switches to 'concur'. Later on, he says the same combination A/4 "means the rule is right". In a subsequent turn he spontaneously adjusts this terminology, after considering the A/7 combination, saying that A/4 "implies the rule is right, it doesn't mean it". Making the distinction between "meaning" and "implying" enables him to distinguish different agency for the A/4 and A/7 combinations. (Here we should understand "implying" not in the language-theoretical sense, i.e. generating an implicature, but in the colloquial sense of "suggesting".) The subject goes on to choose to turn the 7 card. Later on, his insight into the asymmetry of validating and falsifying, and especially the functioning of this asymmetry in the tasks, is again evident in his very self-assured answer in the two-rule task:

Stuart, two rule condition:

E: Which cards?
S: (thinks for a good 30 seconds) *You need to turn the 3* (decisively).
...
Because rule 1 says if there's a U on the fa... this side, one side, there's an 8 on the either. Bear in mind that a card must have a letter and a number. Rule 2 says if there's an I, it must have an 8. *Which means if I turn this 8 (pointing to the 8 card), it could have a U or an I, and that wouldn't tell me either or.* If I turn a U (pointing to the U card) and it has an 8 that would validate it, not an 8 it would invalidate it, I (pointing at the I), same as the 8 (pointing at the U card), but if I turn the 3, then it's definitely got either a U or an I on the other side of the 3, so *if I turn this (the 3) and find out what that is on the other side of this card, then it'll tell me which one of these rules is false.* Like if it's got a U, then rule one is false, which means rule 1 must be right. [..] *You wouldn't actually be checking which rule is right, you would be eliminating which rule is wrong.*

5.4.3 Task-level semantic parameters

The simplicity of the selection task is deceptive. What Wason perhaps envisioned as an exercise in very basic scientific thinking actually involves a whole package of settings or assumptions which the subject must make in the context of the task, and which mostly have little to do the power of falsification, as has hopefully become clear in the course of examining the transcripts. These settings have to do with, apart from the interpretation of the rule itself, the relations between the cards and the rule, the status of the rule in the context of the task, the notion of truth which is called for in the task. All these factors feed into how a subject treats a falsifying card, but are quite separate from understanding what the notion of falsification is, as the following section demonstrates.

'False' *vs* 'wrong'

In the original condition, as we saw in section 5.4.2, subjects were quick to offer correct analysis of the role of the 7 card in disproving the conditional rule, yet commonly failed to connect this with their choice of turnings – either they still didn't think it necessary to turn it, or they included cards which were overridden by it, such as the 4 card. We have seen in subject 3's testimony that potential falsifiers needn't be chosen, and here such a verdict is repeated:

Oona in original condition:

E: *And if there was an A on the other side of the 7?*
S: *Then the rule is false.*
...
E: OK and the 7, you wouldn't want to turn?
S: Um, no, because ... if there is a 4 behind here (the A) then the rule is definitely true, if there is an A behind here (the 4) then, you know, that's further confirmation of the rule being true. *If there is an A behind here (the 7) then the rule's messed up.* (laughs).
E: *But you don't think you need to turn it?*
S: *No.*

... or they may not be considered enough to evaluate the rule:

Rochelle in two rule task:

S: *Oh no if I'd turned the 3 and found an I then the second rule, then that would be untrue. So I'd presume that that (indicating the first rule) was true. If one had to be true. And [I'd] probably not turn anything else.*
E: So, in conclusion, what do you want to turn? [goes through subject's sequence of choices]
S: *Well it would be easier most definitely if you just turned all of them over. But I'm sure there is a way of doing it without actually having to turn them all over.*

E: *Which one do you think are necessary to turn over?*
S: *I think the U and the I are, still, ... but then also the 3, because it disproved, but I don't know if I am just thinking that because it did in this case* ...I think I might just turn all of them.

The last turn is especially puzzling; what could the subject mean when he says that he is unsure about the necessity of turning the 3 card because it disproved the rule only "in this case". Such a rationale suggests the subject sees his role as that of giving a strategy for assessing the truth of the relation between cards and rule more generally; as such the task itself is a sample? As we saw above, this At any rate, the above behaviour indicates confusion about what is needed to establish the truth of a rule, including questions regarding what the exact relation is between the rule and the presented cards. Does it apply to all of them as a set? Or should each card be evaluated individually against the rule? And only these cards? What counts as evidence that the rule is true or false? And what is my job in all this?

This study yielded plentiful evidence that subjects were labouring under similar difficulties to those identified by Stenning and van Lambalgen's in their qualitative study (2001). Stenning and van Lambalgen identified, among others things, the following aspects of the task to be sources of difficulty for subjects: confusion between truth of rule and compliance of cards to rule; 'cards as sample' reading; 'degrees' of truth/lack of brittleness of truth; dependencies between cards leading to contingency in card selection; choice on the basis of possible information gain.

Before we explore subjects' understanding of these various factors, it is instructive to sketch what the required settings are, which the 'scientific' thinker employs in correct performance. Firstly, the scientist must have the correct assumption regarding the **relations between the cards and the rule**. This involves understanding that:

- The cards are not a 'sample' from a larger domain, as Stenning and van Lambalgen (2004, p. 502) have pointed out. The use of conditionals in everyday language to express generalisations operating normally on an open-ended domain (see also Chapter 2), probably makes this a counterintuitive assumption. But, unlike conditional generalisations in everyday language, the subject must here consider the rule to refer to the cards as a complete set in themselves. If, however, the subject does see the cards as sample, it makes the task impossible to solve in its intended form. Enterprising subjects find several means to resolve the quandary: they take a type reading of the cards; or they might be inclined to probabilistic reasoning, giving answers in terms of whether the rule is likely to true in general or not (Stenning and van Lambalgen, 2004, p. 502-3). Another possibility, mentioned above, is that the subject sees the cards as a random selection from a larger sample, but their task as development of a strategy which applies more generally, to any coupling of cards and rule.

- The cards are not made according to type – so that a combination of, say, 4 on the back of an A card does not mean that A will be on the back of the 4 card. However, cards *are* made according to the type specified by the background rule, so that only 4s or 7s, and no other numbers, will occur on the back of a letter card. Especially in the two-rule condition many subjects seemed to assume the cards had been made according to one or other type – more discussion of this follows below. The 'scientific' subject sees that each card must be treated individually.

- However, this is not to say that each card has equal agency. In fact, they all play a different role. The K card is simply irrelevant. The 7 card can falsify, the 4 card can provide verification for the rule, and the A has the potential for both. The task is exactly to specify which cards can affect the truth of the rule, but only in the context of the whole set. So for instance, the A card, if considered on its own, would 'prove' the rule if it had a 4 behind it. However, given that there are *other* cards, the force of this evidence weakens, and creates a dependency within the cards: a 7 behind the A card will falsify the rule immediately, but finding a 4 behind the A requires further turning of the 7 card before conclusions can be drawn.

Next, let's consider what is required regarding the semantic status of the rule in the context of the task and the **concepts of truth and evidence** which are needed for the task.

- Understanding of asymmetric agency of confirmation or verification and falsification. In the intended construal of the task, finding one counterexample to the rule proves that it is false, while finding one instantiation of the conditional isn't enough to conclude that it is true. Most subjects do grasp this asymmetry; if not immediately, then in the course of the task, as illustrated in section 5.4.2.

- In fact, a subject should conclude the rule is true even if there is no witnessing case for it. This is application of the 'classical' notion of truth in the sense that propositions can only take on one of two truth values, so that if there's no evidence to the contrary the subject should conclude that the rule is true. How classical is this though? When does a proposition come to take on these values? As we'll see in the section on positive evidence below, in the split anaphor condition this requirement on truth becomes very counterintuitive for subjects. Additionally, as Stenning and van Lambalgen (2004, p. 498–9) have suggested, if subjects interpret the conditional as a causal relation, then even when there is positive evidence, as in a witnessing case, it might not suffice for concluding the truth of the rule.

- Finally, and this is related to the above discussion regarding the relation between the rule and the cards, the subject should be clear on the status of

the rule. The truth of the rule is determined by the cards and not vice versa – which is not the case in the deontic version of these rules. In the deontic task, the status of the rule remains intact regardless of what is on the cards. The rule is stipulated, and each card complies with it, or not. Each card can thus be judged according to the rule, independently of the other cards. It is also clear that each card represents an individual and is thus unrelated to all others – a 16-year old drinking whiskey carries no information about who's drinking lemonade. In the descriptive task, however, as we've seen above, there is a much more nuanced relationship between rule and cards. The truth of the rule *depends* on the relation of the cards to the rule, but it does not depend on all cards equally: indeed the task, as intended, is to specify which cards of the set can affect the truth of the rule. Many subjects avoid using the terms 'true' and 'false' and rather use terms such as right and wrong, correct and incorrect, even calling the rule "a lie" at some point. What this indicates about their theory of truth is discussed in more detail below.

In the following sections I show how the setting of these high-level task parameters causes difficulty for subjects and suggest how the settings relate to each other.

New finding: the need for positive evidence

The discussion in section 5.4.2 should have made readily apparent that even when subjects have reached the normative A,7 answer in Wason's original rule condition, they often go on to choose cards in other conditions which do not seem to reflect this 'insight'. In other words, just going on the choice of cards, it appears that the adoption of the intended reading for the conditional in that condition does not provoke subjects to take that same reading of a conditional later on – even when it is an almost identical formulation. However, as the discussion below makes clear, there are important distinctions between conditions which change the significance of card selections.

In the current study this phenomenon came to the fore in the split anaphor condition, when a subject was presented with the rule "If there is an A on the back of the card, then there is a 4 on the face of the card". Every single subject in the study wanted to turn over the 4 card in this condition, often when they had already dismissed the significance of the 4 card in the original task. The solitary subject who had adopted the normative interpretation of the conditional in the original rule from the very first, in this condition initially chose only the 7 card but added the 4 card when faced with the clearly uncomfortable conclusion of a true rule in the case of a K behind the 7.

Peter in 'A on back' condition:

S: In this case I think you would need to turn the 7. And... that would be the only one you need to turn. ... if there is an A on the back [of the 4] then it fits the rule (pointing to the 4) and if there is a K on the back then it doesn't apply to the rule, so it doesn't matter, which leaves this one (pointing to the 7) cause if there is an A on the back of here, and it's a 7, then you've disproved the rule.

E: OK and if there is not an A on the back of that (the 7)?

S: If there is not an A on the back ... (thinks) then ... maybe you do need to turn that one (pointing to the 4). *If there is not an A on the back [of the 7], then it doesn't disprove the rule and it doesn't prove it.* So you'd have to turn the 4 I think.

The 4 card was of concern to subjects because it was the only potential source of evidence for the rule; as we see in the next turn the subject goes so far as to say that a K behind both 4 and 7 disproves the rule.

E: *And what if the 4 also didn't have an A on the back – what would that mean for the rule?*

S: *Well then.... for this set of cards, the rule... it would disprove the rule I suppose.* Cause if there is not an A on the back of this card (pointing to the 4), (pause) then... there isn't a 4 on the face of it. OK, *if there is not an A on the back, then none of these cards have an A on the back, and a 4 on the face, which is what the rule states. So for this set of cards it's disproven, it's not true.*

Extremely common in this condition was the specification of the turning order 'first 7, and if there's no A, then 4':

Christopher in 'A on back' condition

S: If there is a A on the back of this card (the 7) then it's finished, you basically don't care anymore. *Whereas if there is a K [behind the 7], all it seems to prove really is that this (the rule) could be true.* [...] *I suppose we do... need to turn this card (the 4), just to affirm the rule.*

Stuart, 'A on back' condition:

[subject has chosen 4 and 7]

S: *... you would turn this [7] one first. Because if you turned this [7] and it had a K, you would still have to turn this (the 4), to make sure that it had an A,* but if you turned (the 7) and it had an A, you would know that the rule is wrong, because it would have an A on the back and a 7 on the front. So you turn this [7] first, and that [4] second. You would only bother to turn this (the 4) if you found a K.

The excerpts show that subjects are fully aware of the asymmetric agency of falsification and verification: an A behind both 4 and 7 unanimously meant that the rule was disproved, showing that the subjects understood that the A/7 combination overrules the verifying contribution of an A/4 combination. However, the excerpts also illustrate that subjects are not only concerned with falsification in this condition.

Is this a different interpretation of the conditional than the normative one in the original case? It is impossible to say. The key factor here is the lack of the A card, and specifically the dual role that the A card plays in the original task, as potential falsifier *and* as potential witness to the rule. Here only the 4 card is a potential witness; only the 7 card can falsify. In a sense this condition provides a more stringent test of falsification than Wason's original task; but it also shows that there is implicit existential import in the conditional. Look again at what our scientific thinker (Peter) concludes, on considering the possibility that neither the 7 nor the 4 have an A on them:

> S: ...OK, *if there is not an A on the back [of the 4], then none of these cards have an A on the back, and a 4 on the face, which is what the rule states. So for this set of cards it's disproven, it's not true.*

In other words, the truth of the conditional depends on the existence of a witnessing case. In the original condition, the A card disguises this requirement – because if it doesn't have a 7 then it not only doesn't falsify the rule, but it also witnesses the rule. As such, there is no possibility to fail both to falsify and to witness the rule. If the subject has considered the consequences of turning the A they will have realised that it failing to falsify – i.e. not finding a 7 – here *coincides* with providing evidence for truth. However, in this split anaphor condition, a K behind both the 4 and 7 card leaves the subject in the uncomfortable position of having no evidence either way. Now according to the material implication reading of the conditional, one should conclude that the rule is true. But subjects do not do this. Does this mean they are thinking unscientifically? No. Rather the original task at this point comes apart from the Popperian paradigm, in which one concludes from a failure to falsify that the unfalsified hypothesis is corroborated, but certainly not that it is true.

Card-checking *vs* rule-checking

How may the complexity of relations between card and rules affect subjects' understanding and performance in the descriptive case? What construal of the task might result? A few options are:

- **card-checking**: choosing A and 4 because they *may* satisfy the rule; K and 7 patently already do not. Plausibly this is related to a generic reading of the cards, which in turn brings on a conjunctive reading. In this case the task would be understood as: what cards should you turn to see if they

satisfy the rule? Under this reading the status of the rule would not be under question; the task would be more one of evaluating the 'truth' of the cards given the rule.

- **rule-forming**: on the basis of what findings would you say this is a true rule about these cards? Choice of A, 4 is reasonable here too; if you substitute 'hypothesis' for 'true rule'. Cards here are chosen on the basis of the potential evidence they supply for the rule; thus increasing the likelihood/probability that it is true. Hence the task here is about evaluating the rule, but more at the stage of 'discovery' than 'justification'.

- **rule-checking**: this is the intended interpretation; the rule is considered true if it is not disproved by the given set of cards.

Subject Oona's testimony from the original condition illustrates the first of these possibilities: A and 4, she says, because there should be an A on the other side of the 4, '*just like if I turn over A there should be a 4*'. The subject also specifies that 4 should be turned over first. I conjecture that these choices arise because she interprets the task as card-checking, as opposed to rule-checking: which cards can be checked against this rule? Her utterances suggest she takes the rule to be true, and her job is to specify which cards are worth checking to see if they satisfy it.

'True' *vs* 'correctly applied'

Confusion regarding the notions of truth and falsity manifest themselves in subjects' language use. The different agencies of specific letter/number combinations to falsify, validate, or prove a certain rule and the extent to which this is grasped might be supposed to be evidenced in the subject's use of the terms. That is, coarse linguistic distinctions suggest that the subject has coarse conceptual distinctions, which hamper performance in the task. However, this is not always the case. Notice for instance the emphasised sections of the following excerpt:

> **Ted in 'every card' condition:**
>
> S: Uh maybe,... yeah no yeah the A and the 7, ... because the A, well naturally you'd need to see what's on the back of it, um.. *if it's a 4 then the rule's correct, if it's a 7, the rule's incorrect.* K ... doesn't matter... because ... we're not concerned with Ks in this case, as I said, doesn't matter whether it's a 4 or a 7.
>
> ...
>
> S: Um ... and with the 7, you need to turn that to check whether or not the rule's correct, cause *if it's ... an A, then the rule's wrong, if there is a K, then the rule's right.*

This subject correctly chooses A and 7 at the first pass at the task, yet keeps using a symmetric 'correct'/'incorrect' distinction until much later on in the interview.

Note that the symmetric 'right'/'wrong' distinction used to justify turning the 7 card (see last turn) captures exactly the reasoning about falsification that Wason sought from his subjects. The emphasised excerpts suggest the subject is using 'correct' as in 'correct description of this card'. The subject is taking a 'case-by-case' reading (see next section) of the rule to work his way through the cards. In a case-by-case reading of the cards, each card is judged according to the rule, and only if all four fit the rule can it be judged true. So we can understand the subject's language as referring to the criteria of 'fit' of card to rule, for each card. The symmetric linguistic distinctions he draws are thus expressing adequate discernment for each card's relation to the rule. The subject correctly choses A, 7 because he realises that the rule is true if precisely *these* cards 'fit' the rule; the K and the 4 are not significant in this sense: they do not affect the 'fit' of the set of cards to the rule so the rule doesn't apply to them.

Regarding the fit of rules to cards, there is interesting testimony to be found in the conjunctive condition:

Rochelle, conjunctive condition:

S: I think you need to turn the A and the 4 again.
E: OK. Why?
S: Cause if there is an A, if there's, that's an A then there should be a 4 on the other side, and the same with that [4]. But *with those two, the K and the 7, it doesn't matter.* I don't think. Well *it would prove that it's untrue,* but if we want to prove it's true, then ... (gestures with her hand) [continued below]

How can this curious distinction of 'proving untrue' and 'proving true' be understood? The subject continues to use it as the dialogue continues, and eventually settles on 'all cards', chosen to fulfil both these criteria. I couldn't help thinking she had some sensible distinction to make, but just wasn't labelling it well. 'Proving untrue' is in the above turn doing double duty as a description of applicability criterion and falsification, as in, the rule doesn't apply to the K and the 7 because they make it untrue. This testifies to a card-checking approach as outline above; as in, which set of cards can satisfy the rule? Being prompted to consider what she has just said seems to make the subject reconsider her choices.

[continued from above]
E: What would prove it's untrue?
S: *If you turned the K and it had a 4 on it, then that would make the rule untrue.*
E: So but you don't think you need to turn the K and the 7, is that right, or do you?
S: Well *you need to turn the A and 4 to prove it, but then ... well I suppose you actually have to turn the K and the 7 as well to .., to make sure that it, that it's not untrue.*
E: To make sure you're not disproving it?

S: Yeah. So actually probably you have to turn them all.

E: But for different reasons though?

S: Yeah. To *turn those two [A, 4] to prove it and to turn those [K, 7] to make sure you're not disproving it.*

Presumably the K and 7 should not have As or 4s on them; then we would be in the restricted-domain biconditional reading of the conjunction.

Another example of confusion about the status of 'truth' in the task:

Oona in the original task:

[she has answered A and 4, but '4 first']

E: So let's think about what could be on the back of the A. On the back of the A there could be a 4 or a 7.

S: *There should be a 4.*

E: There should be a 4?

S: Yes, because *it says if there is an A on one side then there is a 4 on the other side.*

E: Ja, but you have to check whether that's true or not.

S: Ah! OK fine. If it is true that's what it should be. But yeah you're right, it could be a 4 or a 7.

E: So *say you turned it over and there is a 7. What would that mean?*

S: The rule is wrong. Well it's, *it's a lie!*

For the rule to be a lie, it must be purporting to tell the truth – as a characterisation of the cards. Elsewhere she says that a 7-A combination would mean the rule '*is messed up*', yet insists this is no reason to turn the 7. Again, this makes sense if one supposes that she doesn't take the status of the rule to be under scrutiny, and rather takes her task to be one of investigating its applicability to the cards, and, as such, the task is one of card-checking. Her use of the phrase '*messed up*' aligns well with this; rather than simply stating that the rule is falsified, the rule doesn't 'apply' properly to the cards. This is not a straight deontic reading; rather, truth is read as 'applicability'.

5.5 Summary, conclusions and outlook

> "Psychology is in some ways harder once one acknowledges interpretational variety, but given the overwhelming evidence for that variety, responding by eliminating it from psychological theory is truly the response of the drunk beneath the lamp post" (Stenning and van Lambalgen, 2004, p. 491).

In this small qualitative study we found much new and corroborating evidence of interpretational variety, in a range of manipulations of Wason's selection task. In particular, the current findings strongly supported earlier explanations of the modal choice of A and 4 in the original task arising from 'interference effects', i.e. the influence of the anaphor element on the conditional direction, and could now relate this to another popular choice, just A. For the first time subjects were interviewed about their understanding of a conjunctive phrasing. This condition provoked a surprisingly range of responses and an attempt was made to explain the interpretations lying behind these and to relate them to other rule formulations. Different findings in the two-rule manipulation entailed a necessary extension of the analysis provided by Stenning and van Lambalgen (2004) on this task.

The posited range of interpretations for the various conditions was then placed into the context of the broader semantic settings of the task. The tutorial dialogue setting clearly helped subjects to reach the intended interpretation of the task, with almost all reaching the normative selection by the end of the interview. However, there was also great variability in subjects' interpretations of the different conditions, suggesting that the normative interpretation of the original task serves as a highly scaffolded interpretation for the subjects. Slight variations in the conditions led subjects to adopt different interpretations of the conditional. In the split anaphor condition we saw that a seemingly different selection could represent the same reading of the conditional, and the task demands interact with this to produce a different selection.

Briefly the various possible relations of cards to rule and the ways these affected the subject's construal of their task and the conditions for truth and falsehood in the task were discussed. Novel findings here are the existential import assumed by subjects which surfaces only in the split anaphor version of the task; the assumption of strong falsity in the two-rule task and the card-per-card reading interacting with a descriptive understanding of the rule to produce a notion of truth which resembles 'applicability'.

I hope to have convinced readers that the Wason selection task is by no means just a simple test of falsification, but that the basic set-up can be fruitfully used to investigate a wealth of semantic issues that subjects are concerned with in reasoning about the truth of a descriptive rule. Most earlier research has been purely quantitative and has thus failed to reveal these factors, but, as Stenning and van Lambalgen suggest, it was time to start looking further than in the pool of light around the lamppost.

5.5.1 Experimental suggestions

The natural next step would be a quantitative study on the basis of the findings here. However, the suggestion would be to collate data from subjects *across* conditions, to investigate the extent to which there is stability of interpretation. In

section 5.4.2 it became clear that individual subjects adjust their interpretation of a rule and according truth conditions surprisingly easily. However, the small size of this data set might be magnifying the seemingly endless interpretational caprice. With a bigger data-set one would be more able to identify tendencies towards one or other readings, across conditions but within subjects. Furthermore, the tutorial data have shown how certain choices are highly connected for some subjects, such as the choice for 'just A' and 'A and 4' in the original task, and this kind of information can now be used to identify patterns of related but non-equivalent responses in a big data set.

An issue arising from this study which deserves further attention is the connection between conditional and conjunctive phrasings such as in the original and conjunctive tasks. This connection is implicit in Athanasiadou & Dirven's (1995, 1997) analysis of course-of-event conditionals – in particular the commitment to both the antecedent and consequent in such conditionals – as Chapter 2 described. There remain many open questions about the circumstances under which conditional phrasings have conjunctive force, as well as those in which conjunctive formulations are interpreted conditionally, one being the seeming conversion of the conjunct order into a kind of suppositional precedence ordering, perhaps similarly to or by means of an imposition of temporal order.

Conclusions

Summary

The dissertation opened in Chapter 1 with a fresh analysis of reasoning behaviour in subjects with varying but low educational levels. We saw that Luria's blanket negative conclusions about such reasoners can be replaced by a more informative account which takes into account the semantic and pragmatic constraints in the task which shape subjects' responses, including matters such as the anomalous epistemic structure of the question-answer interaction, the various readings possible for the generalisations expressed in the major premises, and the relation of the protagonist to this generalisation. The analysis indicated that the commonalities with schooled subjects are much greater than previously reported. In Chapter 2 the remaining differences between groups, specifically the observed differences between *all* and *if* premises, were explained by an original and exploratory synthesis of corpus analysis with semantic theory.

Chapter 3 contextualised the results from these chapters in the broader debate about the cognitive consequences of literacy, and considered some of the manifold social and cultural dimensions which feature in a full account of literacy's effects. This prompted a 'local' approach to theorising, in the form of the evaluation of a very specific proposal from David Olson, namely that the results from reasoning research signify ability to distinguish 'literal' meaning, in the form of 'what is said', from 'what is meant'. Although this was found to be an attractive suggestion, it was argued that the notion of 'literal meaning' has been understood too simplistically, in its role in reasoning tasks.

A more qualified understanding of literality was outlined in Chapter 4, and the resultant decoupling of literal meaning from logicality in reasoning tasks was described. This formed a key part of the chapter's more general critical examination of the theoretical underpinnings of mainstream reasoning research. Assumptions such as that logical form is given by grammatical form, that there is a clear and unchanging division between the logical and non-logical, or contentful, parts of

natural language, and the identification of logical reasoning with classical logical norms, were interrogated and repudiated. It was proposed that the schooled stance towards reasoning tasks, described as reasoning Logically, has served as an overly narrow description of a more basic ability to generate and manage meaning, an ability in which interpretation and reasoning are intrinsically connected, and that this is more aptly seen as the locus of logic in reasoning.

Finally, in Chapter 5 this approach was applied to an analysis of qualitative data from a small study of Wason's selection task and a number of variations on it. Wason's selection task has been an important source of evidence for negative conclusions about reasoning. The last chapter challenged these, by describing plausible interpretations of the linguistic materials as well as broader matters of task construal, in an effort to show that the range of responses observed reflects full engagement with the task and a similar concern for semantic parameters observed in the subjects reported in the opening chapters. In this way, the common basis to behaviour across different subject groups became more apparent.

Finding logic

The introduction sketched a situation in reasoning research for which Krueger and Funder's (2004, p. 318) contention that "only irrationality is newsworthy" seemed pretty accurate. In the course of the current study we have seen that the focus on negative findings is accompanied by fluctuations in classifications of logical ability. When Luria and Scribner investigated the reasoning behaviour of illiterate farmers, they contrasted it with their educated peers, who were found to be exemplary reasoners. But researchers who have concentrated their investigations within highly educated subject populations have documented failings judged at least as serious as Luria's findings. In the introduction we read of Wason's despair at his undergraduate subjects: "What makes people so narrow minded and so cognitively prejudiced? Why did they find these trivial games so difficult?" (Wason, 1968b, p. 172) – it feels as if this exclamation could just as well have come from Luria's writings. It can seem that reasoning researchers have directed their energies towards wherever they find the errors, and thus seem to uncover more and more irrationality in human reasoning.

In this dissertation I sought to provide a counterbalance to the catalogue of reasoning errors. The aim was not to show that everyone is logical all the time, but to provide a more constructive and nuanced account of the role of logic in reasoning behaviour. But I would go so far as to say that as researchers of human reasoning, it is our job to make rationality newsworthy. The burden is on us to find the sense in what our subjects do, not to try to trick them out. This approach changes what counts as an interesting result. When we look for logic in semantically-driven interactions, the details of a particular interpretation, the justification given by the subject for it, and the support that can be found

for it in everyday language use, become essential means to assess the adequacy of a subjects' response and are no longer viewed as unconnected or detracting from the logical aspects of it. There is moreover plenty of logic to be found, as I hope to have convinced you over the course of this work. Subjects engage fully with presented premises and set semantic parameters appropriate both to their linguistic experience and to their construal of the task. The range of choices available to subjects sometimes reflect their educational background, for instance, in the varying ease with which subjects accepted an *all* premise as expressing a law-like generalisation – as described in Chapter 2 – but more often reflect their individual ability to set and reset semantic parameters in a self-controlled fashion, demonstrated in both Chapter 1 and the last chapter, Chapter 5.

As we saw in Chapter 4, the context-dependence of meaning, and the dynamic character of interpretation that it entails, has implications for the status of classical norms for reasoning, since they are assumed to be determined by the literal meaning of the premises. Once literal meaning loses its sovereignty as bearer of logicality, the criteria for an account of good reasoning change. On this account reasoning skills are to be judged on skill in discerning differences between interpretations and fitting them to the task at hand in the best way possible, including ability in matching conceptual distinctions within an interpretation with those which the situation requires, creativity in articulating and constructing possible interpretations, in particular other than the default which may be derived from typical language use, and insight into the consequences of adopting specific interpretations. This is not the same as saying that reasoning is domain-dependent; it is to say it is done so on the basis of parameters which are set according to the task at hand, in the process of constructing meaning. It also entails that for many other behaviours which have been labelled as systematic 'reasoning errors', this has been done so in haste, and they might also be rehabilitated when semantic analysis is applied to them.

There are broader implications of this account of reasoning and its relation to language and literate practices. Firstly, we should do away with the idea that language is stable and monolithic, and take seriously the dynamic nature of meaning. This impacts on the way that languages are studied, as well as the way they are taught, in both cases by removing the search for, or teaching of, some kind of a linguistic bedrock, the same for speakers across contexts, which, as we have seen, massively underestimates the interpretational flexibility language users have at their command. Secondly, the findings here suggest that although education is largely a matter of learning new ways of using symbolic systems, this learning is something which occurs quite locally and sporadically. I think that this is so because these new ways with language are taught with only a vague sense of why such an enterprise is valued, and how it connects with what we do when using language. A better understanding of our meaning-making behaviour will make education more able to connect up with what we do anyway, and thereby more able to effectively contribute to cognitive development.

Where to next

Apart from the specific experimental suggestions to be found at the end of each chapter, there are some natural avenues for further research suggested by the dissertation as a whole. Firstly, studies which extend the analysis of language used in cognitive tasks and relate it to general language use would be very valuable. In other words, further investigation of the relation between language used for thinking and language used for talking. Factors underlying subjects' behaviour do not lie within neat disciplinary boundaries: processing constraints presumably affect choices just as much as conscious deployment of learnt standards of reasoning, as well as, as we have seen, automatic attribution of linguistic meanings, and socio-pragmatic norms. Hence, further research can probably best be done in collaboration with linguistic researchers, primarily semanticists, but psycholinguists and sociolinguists could also help bridge the gap between language use for thinking and language use for communication purposes.

Secondly, the qualitative data presented here yield many ideas for further quantitative studies, and these should be followed up on, although I think that a combination of both qualitative and quantitative data collection remains optimal for this area of research. With the less educated subject populations quantitative research is difficult both because it is difficult to find suitable populations and because of the necessarily verbal nature of the testing, but as we saw in Chapter 3 simply broadening the subject profile beyond the student undergraduate already produces a range of linguistic behaviours which have not been fully investigated. The suggestion would be to conduct a broader range of quantitative tasks – including for instance, interpretation-oriented tests – with a broader range of subjects.

Finally, it would be very constructive to provide appropriate formalism for the analyses given here of subject behaviour. Partial suggestions in this direction were made in Chapter 1, but these should be greatly extended to a more complete formal model of reasoning behaviours. Apart from the benefits in terms of precision and perspicuity regarding claims and predictions, this is important because it allows the common features of diverse cognitive tasks to be identified, especially tasks classified under problem-solving, decision-making and such like. These are traditionally separate areas of study, while it is fairly obvious that they share many features of reasoning tasks. Stenning and van Lambalgen (2008) have already, by the use of a specific formal apparatus, been able to model features shared by diverse tasks such as the false belief task and the suppression effect task. Continuing in this vein is essential if we are to develop a more unified theory of human cognition.

Bibliography

Akinnaso, F. N. (1992). On the differences between spoken and written language. *Language and Speech 25*(2), 97–125.

Allal, L. (1999). Metacognitive regulation of writing in the classroom. In A. Camps and M. Milian (Eds.), *Metalinguistic Activity in Learning to Write*. Amsterdam University Press.

Astington., J. W. and D. R. Olson (1990). Metacognitive and metalinguistic language: learning to talk about thought. *Applied Psychology: An International Review 39*, 77 – 87.

Athanasiadou, A. and R. Dirven (1995). Typology of *if*-clauses. In E. Casad (Ed.), *Cognitive linguistics in the redwoods*, pp. 609–654. Mouton De Gruyter.

Athanasiadou, A. and R. Dirven (1997). Conditionality, hypotheticality, counterfactuality. In A. Athanasiadou and R. Dirven (Eds.), *On conditionals again*, pp. 61–96. John Benjamins.

Bach, K. (2001). You don't say? *Synthese 128*, 15–44.

Baltag, A., L. S. Moss, and S. Solecki (1998). The logic of public announcements and common knowledge and private suspicions. In *Proceedings of TARK 1998*, pp. 43 – 56.

Bartsch, R. (1984). The structure of word meanings: polysemy, metaphor, metonymy. In F. Landman and F. Veltman (Eds.), *Varieties of formal semantics*. Foris Publications, Dordrecht.

Barwise, J. (1985). Model-theoretic logics: background and aims. In J. Barwise and S. Feferman (Eds.), *Model-Theoretic Logics*. New York: Springer Verlag.

Barwise, J. (1989). *The Situation in Logic*, Volume 17. Stanford, California: CSLI. Distributed by Cambridge University Press.

Beal, C. R. and J. H. Flavell (1984). Development of the ability to distinguish communicative intention and literal message meaning. *Child Development 55*(3), 920–928.

Berry, J. (1971). Ecological and cultural factors in spatial perceptual development. *Canadian Journal of Behavioral Science 3*(324-336).

Berry, J. and P. Dasen (1974). *Culture and cognition: readings in cross-cultural psychology*. Methuen, London.

Berry, J. W., Y. H. Poortinga, M. H. Segall, and P. S. Dasen (Eds.) (2002). *Cross-Cultural Psychology: Research and Applications*. Cambridge University Press.

Bitbol, M. (2002). Science as if situation mattered. *Phenomenology and the Cognitive Sciences 1*, 181–224.

Bloom, B. S. and L. J. Broder (1950). *Problem-solving processes of college students: an exploratory explanation*. University of Chicago Press, Illinois.

Bobaljik, J. D. (2002). Floating quantifiers: Handle with care. In L. L. S. Cheng and R. P. E. Sybesma (Eds.), *The Second Glot International State-Of-The-Article Book: The Latest in Linguistics*. Mouton de Gruyter.

Bonitatibus, G. (1988). Comprehension monitoring and the apprehension of literal meaning. *Child Development 59*, 60–70.

Bonnefon, J. F. and D. J. Hilton (2005). Getting the point of conditionals: An argumentative approach to the psychological interpretation of conditional premises. In G. Kern-Isberner, W. Rödde, , and F. Kulmann (Eds.), *Conditionals, Information, and Inference: International Workshop (WCII 2002)*. Berlin: Springer Verlag.

Bourdieu, P. and translated from the french by J.B. Thompson (1991). *Language and symbolic power*. Polity Press Cambridge, Mass.

Bowerman, M. (1986). First steps in acquiring conditionals. In E. C. Traugott, A. ter Meulen, J. S. Reilly, and C. A. Ferguson (Eds.), *On Conditionals*. Cambridge University Press, Cambridge.

Bradford, B. (1997). Upspeak in british english. *English Today. 51, 13:3*, 29–36.

Braine, M. (1978). On the relation between the natural logic of reasoning and standard logic. *Psychological Review 85*, 1–21.

Braine, M. (1990). The" natural logic" approach to reasoning. In W. Overton (Ed.), *Reasoning, Necessity, and Logic: Developmental Perspectives*. Lawrence Erlbaum.

Brisson, C. (2003). Plurals, all and the nonuniformity of collective predication. *Linguistics and philosophy 26(2)*, 129–184.

Brubaker, R. (1984). *The limits of rationality: an essay on the social and moral thought of Max Weber*. George Allen & Unwin, London.

Bruner, J. (1957). Going beyond the information given. In H. G. et al (Ed.), *Contemporary Approaches to Cognition*. Cambridge University Press, Cambridge.

Bruner, J. (1990). *Acts of Meaning*. Harvard University Press, Cambridge, Massachusetts.

Bruner, J. (1991). The narrative construction of reality. *Critical Inquiry 18:1*, 1–21.

Bruner, J. (1996). *The culture of education*. Harvard University Press, Cambridge, Massachusetts.

Byrne, R. (1989). Suppressing valid inferences with conditionals. *Cognition 31(1)*, 61–83.

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

Camps, A. and M. Milian (Eds.) (1999). *Metalinguistic Activity in Learning to write*. Amsterdam University Press, Amsterdam.

Carruthers, M. J. (1990). *The Book of Memory: A Study of Memory in Medieval Culture*. Cambridge University Press, Cambridge.

Ceci, S. and A. Roazzi (1994). The effects of context on cognition: postcards from brazil. In R. Sternberg (Ed.), *Mind in Context: Interactionist Perspectives on Human Intelligence*. Cambridge University Press, Cambridge.

Ceci, S. J., T. B. Rosenbum, and E. DeBruyn (1998). Laboratory versus field approaches to cognition. In R. J. Sternberg (Ed.), *The Nature of Cognition*. MIT Press, Cambridge, Massachusetts.

Celce-Murcia, M. and D. Larsen-Freeman (1983). *The Grammar Book: An ESL/EFL teacher's course*. Newbury House, Rowley, Massachusetts.

Chase, V., R. Hertwig, and G. Gigerenzer (1998). Visions of rationality. *Trends in Cognitive Science 2(6)*, 206–214.

Chater, N., M. Oaksford, R. Nakisa, and M. Redington (2003). Fast, frugal and rational: how rational norms explain behaviour. *Organizational behavior and human decision processes 90*, 63–86.

Cheng, P. W. and K. J. Holyoak (1985). Pragmatic reasoning schemas. *Cognitive Psychology 17(4)*, 391–416.

Cherniak, C. (1986). *Minimal Rationality*. MIT Press, Cambridge, Massachusetts.

Clark, H. H. and A. Bangerter (2004). Changing ideas about reference. In I. A. Noveck and D. Sperber (Eds.), *Experimental Pragmatics*. Palgrave Macmillan.

Cole, M., J. Gay, J. A. Glick, and D. W. Sharp (1971). *The cultural context of learning and thinking: An exploration in experimental anthropology*. New York: Basic Books.

Cole, M. and S. Scribner (1974). *Culture and thought*. Wiley, New York.

Cole, M., D. Sharp, and C. Lave (1976). The cognitive consequences of education: Some empirical evidence and theoretical misgivings. *The Urban Review 9*(4), 218–233.

Comrie, B. (1986). A typology of conditionals. In E. C. Traugott, A. ter Meulen, J. S. Reilly, and C. A. Ferguson (Eds.), *On Conditionals*. Cambridge University Press, Cambridge.

Cosmides, L. (1989). The logic of social exchange: Has natural selection shaped how humans reason? studies with the Wason selection task. *Cognition 31*, 187–276.

Cosmides, L. and J. Tooby (1989). Evolutionary psychology and the generation of culture, part ii. case study: A computational theory of social exchange. *Ethology and Sociobiology 10*, 51–97.

Cosmides, L. and J. Tooby (1992). Cognitive adaptations for social exchange. In *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*. Oxford: Oxford University Press.

Cummins, D., T. Lubart, O. Alksnis, and R. Rist (1991). Conditional reasoning and causation. *Memory and Cognition 19*(3), 274–282.

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

Dabrowska, E. (1997). The LAD goes to school: a cautionary tale for nativists. *Lingistics 35*, 735–766.

Dasen, P. (1977). *Piagetian Psychology: Cross Cultural Contributions*. Halsted Press.

Declerck, R. and S. Reed (2001). *Conditionals: A Comprehensive Empirical Analysis*. Mouton de Gruyter.

Dieussaert, K., W. Schaeken, W. Schroyens, and G. d'Ydewalle (2000). Strategies during complex conditional inferences. *Thinking and Reasoning 6*(2), 125–161.

Dolz, J. and S. Erard (1999). Metaverbal activities as an approach to teach spoken and written genres. In A. Camps and M. Milian (Eds.), *Metalinguistic Activity in Learning to Write*. Amsterdam University Press.

Donaldson, M. (1978). *Children's Minds*. Fontana Collins, Glasgow.

Ehri, L. (2000). Development of the ability to read words. In R. Barr, M. L. Kamil, and P. B. Mosenthal (Eds.), *Handbook of reading research*. Lawrence Erlbaum Associates.

Erickson, T. D. and M. E. Mattson (1981). From words to meaning: A semantic illusion. *Journal of Verbal Learning & Verbal Behavior 20*(5), 540–551.

Evans, J. (1972). On the problems of interpreting reasoning data: Logical and psychological approaches. *Cognition 1*, 373–384.

Evans, J. (1998). Matching bias in conditional reasoning: Do we understand it after 25 years? *Thinking and Reasoning 4*, 45–82.

Evans, J. (2002). Logic and human reasoning: An assessment of the deduction paradigm. *Psychological Bulletin 128*, 978–996.

Evans, J. (2003). In two minds: Dual-process accounts of reasoning. *Trends in Cognitive Sciences 7*(10), 454–459.

Evans, J., S. Newstead, and R. Byrne (1993). *Human Reasoning: The Pychology of Deduction*. Hove, Sussex, UK: Lawrence Erlbaum.

Evans, J. and D. Over (1996). *Rationality and Reasoning*. Psychology Press, Hove.

Evans, J. S. B. T. (1989). *Bias in human reasoning: Causes and consequences.* Hillsdale, NJ: Erlbaum.

Evans-Pritchard, E. E. (1937). *Witchcraft, Oracles and Magic Among the Azande*. Clarendon Press, Oxford.

Evans-Pritchard, E. E. (1956). *Nuer Religion*. Oxford University Press.

Feldman, C. (1991). Oral metalanguage. In D. R. Olson and N. Torrance (Eds.), *Literacy and orality*. Cambridge University Press, Cambridge.

Fiddick, L., L. Cosmides, and J. Tooby (2000). No interpretation without representation:the role of domain-specific representations and inferences in the wason selection task. *Cognition 77*, 1–79.

Fillenbaum, S. (1978). How to do some things with if. In D. Rumelhart, D. Norman, J. Cotton, and R. Klatzky (Eds.), *Semantic factors in cognition*. Erlbaum, Hillsdale, NJ.

Finnegan, R. (1973). Literacy versus non-literacy: the great divide? In R. Horton and R. Finnegan (Eds.), *Modes of Thought, Essays on Thinking in Western and Non-Western Societies*. Faber and Faber, London.

Ford, C. and S. A. Thompson (1986). Conditionals in discourse: A text-based study from english. In E. C. Traugott, A. ter Meulen, J. S. Reilly, and C. A. Ferguson (Eds.), *On Conditionals*. Cambridge University Press, Cambridge.

Ford, M. (1995). Two modes of mental representation and problem solution in syllogistic reasoning. *Cognition 54*, 1–71.

Fordham, S. and J. Ogbu (1986). Black students' school success: Coping with the "burden of 'acting white"'. *The Urban Review 18*(3), 176–206.

Gamut, L. T. F. (1991). *Logic, Language and Meaning Volume 1*. University of Chicago Press, Chicago.

Garnham, A. and J. Oakhill (1987). Interpreting elliptical verb phrases. *Quarterly Journal of Experimental Psychology 39*(4), 611–627.

Garton, A. and C. Pratt (1998). *Learning to be Literate: The Development of Spoken and Written Language*. Blackwell.

Gebauer, G. and D. Laming (1997). Rational choices in wason's selection task. *Psychological Research 60*, 284–293.

Gellner, E. (1988). *Plough, Sword, and Book: The Structure of Human History*. Collins Harvill, London.

Geurts, B. (2002). Donkey business. *Linguistics and philosophy 25*(2), 129–156.

Gibbs, R. and J. Moise (1997). Pragmatics in understanding what is said. *Cognition 62*, 51–74.

Gigerenzer, G. (1999). The adaptive toolbox. In G. Gigerenzer and R. Selten (Eds.), *Bounded Rationality: The Adaptive Toolbox*. MIT Press, Cambridge, Massachusetts.

Gigerenzer, G. and D. Goldstein (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review 103*, 650—669.

Gigerenzer, G. and K. Hug (1992). Domain-specific reasoning, social contracts, cheating, and perspective change. *Cognition 43*, 127—171.

Gigerenzer, G. and R. Selten (Eds.) (1999a). *Bounded Rationality: The Adaptive Toolbox*. MIT Press, Cambridge, Massachusetts.

Gigerenzer, G. and R. Selten (Eds.) (1999b). *Bounded Rationality: The Adaptive Toolbox*. MIT Press, Cambridge, Massachusetts.

Gilovich, T. (1991). *How We Know What Isn't So: The Fallibility of Human Reason in Every- day Life*. The Free Press: New York.

Glick, J. (1975). Cognitive development in cross-cultural perspective. In J. Horowitz (Ed.), *Review of child development research*, Volume 4. Chicago University Press.

Goodman, N. (1947). The problem of counterfactual conditionals. *The Journal of Philosophy 44(5)*, 113–128.

Goodnow, J. (1990). The socialization of cognition: what's involved? In J. Stigler, R. Shweder, and G. Herdt (Eds.), *Cultural Psychology: Essays on Comparative Human Development*. Cambridge University Press.

Goody, J. and I. Watt (1963). The consequences of literacy. *Comparative studies in society and history V*(3), 304–345.

Gray, P. (1991). *Psychology: An introductory psychology textbook*. Worth Publishers, New York.

Greenfield, P. (1966). On culture and conservation. In J. Bruner, P. Greenfield, and R. Olver (Eds.), *Studies in Cognitive Growth*. New York: Wiley.

Greenfield, P. (1972). Oral or written language: the consequences for cognitive development in africa, the united states and england. *Language and Speech 15*(2), 169–78.

Grice, H. P. (1989). Logic and conversation (originally published 1975). In *Studies in the Ways of Words*. Cambridge, Massachusetts: Harvard University Press, Cambridge, Massachusetts.

Griggs, Richard A.; Cox, J. R. (1982). The elusive thematic-materials effect in wason's selection task. *British Journal of Psychology 73(3)*, 407–420.

Haan, G. (2007). How illiterates interpret syllogistic problems. Master's thesis, University of Amsterdam.

Haegeman, L. (2003). Conditional clauses: External and internal syntax. *Mind & Language 18 (4)*, 317–339.

Haiman, J. (1978). Conditionals are topics. *Language 54*, 564–589.

Hamill, J. (1990). *Ethno-logic: the anthropology of human reasoning*. University of Illinois Press.

Harman, G. (1989). *Change in view: principles of reasoning* (2nd ed ed.). MIT Press, Massachusetts.

Harman, G. (1999). *Reasoning, meaning and mind*. Clarendon Press, Oxford.

Harris, R. (1986). *The Origin of Writing*. Duckworth Press.

Harris, R. (2000). *Rethinking Writing*. Athlone Press.

Havelock, E. A. (1982). *The literate revolution in Greece and its cultural consequences*. Princeton University Press.

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

Hill, C. and L. Anderson (1993). The interview as a research tool. *New Ideas in Psychology 11*(1), 111–125.

Hilton, D. J. (1995). The social context of reasoning: Conversational inference and rational judgement. *Psychological Bulletin 18(2)*, 248–271.

hl, D. W. (1985). Determiners and context sets. In J. van Benthem and A. ter Meulen (Eds.), *Generalized Quantifiers in Natural Language*. Foris Publications, Dordrecht.

Husserl, E. (2001). *The Shorter Logical Invesigations. Translated by J.N. Findlay*. Routledge, London.

Hutchins, E. (1980). *Culture and Inference: A Trobriand Case Study*. Harvard University Press, Cambridge, Massachusetts.

Hwang, M. (1979). A semantic and syntactic analysis of if-conditionals. Master's thesis, University of California, Los Angeles.

Iglesias, M. H. (2006). Meaning 'literal'. In M. Frapolli (Ed.), *Saying, meaning, referring: essays on the philosophy of Francois Recanati*. Palgrave: London.

Ingold, T. (2000). *The perception of the environment: essays in livelihood, dwelling and skill*. Routledge, London.

Inhelder, B. and J. Piaget (1958). *The growth of logical thinking from childhood to adolescence (transl. Anne Parsons and Stanley Milgram)*. Routledge & Kegan Paul, London.

Inhelder, B. and J. Piaget (1969). *The psychology of the child (transl. Helen Weaver)*. Routledge & Kegan Paul, London.

Jahandarie, K. (1999). *Spoken and Written Discourse: A Multi-Disciplinary Perspective*. Stamford: Abex Publishing Corporation.

Johnson-Laird, P. (1983). *Mental models*. Cambridge University Press.

Johnson-Laird, P., P. Legrenzi, and M. S. Legrenzi (1972). Reasoning and a sense of reality. *British Journal of Psychology 63*(3), 395–400.

Johnson-Laird, P. N. (2001). Mental models and deduction. *Trends in Cognitive Sciences 5*(10), 434–442.

Johnson-Laird, P. N. and F. Savary (1999). Illusory inferences: a novel class of erroneous deductions. *Cognition 71*(3), 191–229.

Johnson-Laird, P. N. and E. Shafir (1993). The interaction between reasoning and decision making: an introduction. *Cognition 49*, 1–9.

Kamp, H. and U. Reyle (1993). *From discourse to logic: Introduction to modeltheoretic semantics of natural language, formal logic and discourse representation theory*. Kluwer Academic Dordrecht; Boston.

Kaplan, D. (1978). Dthat. In P. Cole (Ed.), *Syntax and Semantics 9*. Academic Press.

Karanth, P., A. Kudva, and A. Vijayan (1995). Literacy and linguistic awareness. In B. de Gelder and J. Morais (Eds.), *Speech and Reading: A Comparative Approach*. Psychology Press.

Karanth, P. and M. G. Suchitra (1993). Literacy acquisition and grammaticality judgments in children. In R. J. Scholes (Ed.), *Literacy and Language Analysis*. Lawrence Erlbaum Associates.

Karmiloff-Smith, A. (1994). Precis of: Beyond modularity: a developmental perspective on cognitive science. *Behavioural and Brain Sciences 17*(4), 693–745.

Kearney, M. (1975). World view theory and study. *Annual Review of Anthropology 4*, 247–270.

Keenan, E. L. and D. Westerståhl (1997). Generalized quantifiers in linguistics and logic. In J. v. Benthem and A. ter Meulen (Eds.), *Handbook of Linguistics and Logic*, Chapter 15. Amsterdam: Elsevier.

Kida, T. (2006). *Don't Believe Everything You Think: The 6 Basic Mistakes We Make in Thinking*. Prometheus.

Krueger, J. I. and D. C. Funder (2004). Towards a balanced social psychology: Causes, consequences and cures for the problem-solving approach to social behaviour and cognition. *Behavioural and Brain Sciences 27*, 313–376.

Kurvers, J. (2002). *Met ongeletterde ogen. Kennis van taal en schrift van analfabeten.* Amsterdam: Aksant Academic Publishers.

Langacker, R. W. (1991). *Foundations of Cognitive Grammar: Volume II*. Standford University Press.

Lechler, A. (2004). Interpretation of conditionals in the suppression task. Master's thesis, Edinburgh University.

Lee, E., N. Torrance, and D. R. Olson (2001). Young children and the say/mean distinction: verbatim and paraphrase recognition in narrative and nursery rhyme contexts. *Journal of Child Language 28*, 531–543.

Levi-Strauss, C. (1966). *The Savage Mind*. University of Chicago Press, Chicago.

Levy-Bruhl, L. (1926/1910). *How natives think*. George Allen & Unwin, London.

Lukes, S. (2000). Different cultures, different rationalities? *History of the human sciences 13*(1), 3–18.

Luria, A. (1976). *Cognitive development: its cultural and social foundations*. Harvard University Press, Cambridge, Massachusetts.

Luria, A. and L. Vygotsky (1992). *Ape, primitive man and child: essays in the history of behaviour*. Harvester Wheatsheaf, Hertfordshire.

Mackie, J. (1974). *The Cement of the Universe*. Oxford University Press.

MacWhinney, B., E. Bates, and R. Kliegl (1984). Cue validity and sentence interpretation in english, german and italian. *Journal of Verbal Learning and Verbal Behavior 23*(2), 127–150.

McCarthy, M. and R. Carter (2001). Ten criteria for a spoken grammar. In E. Hinkel and S. Fotos (Eds.), *New perspectives on grammar teaching in second language classrooms.* Mahwah, New Jersey: Laurence Erlbaum Associates.

McFarlane, J. (2000). *What does it mean to say that logic is formal?* Ph. D. thesis, University of Pittsburgh.

McLuhan, M. (1962). *The Gutenberg Galaxy: The Making of Typographic Man.* University of Toronto Press.

McNeil, N. M. and M. W. Alibali (2004). You'll see what you mean: students encode equations based on their knowledge of arithmetic. *Cognitive Science 28*, 451–466.

Miller, K., M. Kelly, and X. Zhou (2005). Learning mathematics in china and the united states:cross-cultural insights into the nature and course of preschool mathematical development. In *Handbook of Mathematical Cognition.* Psychology Press.

Moore, C., K. Pure, and D. Furrow (1990). Children's understanding of the modal expression of speaker certainty and uncertainty and its relation to the development of a representational theory of mind. *Child Development 61*(3), 722–730.

Morais, J., L. Cary, J. Alegria, and P. Bertelson (1979). Does awareness of speech as a series of phones arise spontaneously? *Cognition 7*, 323–331.

Myerhoff, B. G. (1974). *Peyote Hunt: The Sacred Journey of the Huichol Indians.* Cornell University Press.

Narasimhan, R. (1991). Literacy: its characterization and implications. In D. R. Olson and N. Torrance (Eds.), *Literacy and orality.* Cambridge University Press.

Newstead, S. E., P. Pollard, J. S. B. T. Evans, and J. L. Allen (1992). The source of belief bias effects in syllogistic reasoning. *Cognition 45*, 257–284.

Nickerson, R. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology 2*(2), 175–220.

Nisbett, R. E. and T. D. Wilson (1977). Telling more than we can know: verbal reports on mental processes. *Psychological Review 84(3)*, 231–259.

Norenzayan, A. (2006). Cultural variation in reasoning. In *Biological and Cultural Bases of Human Inference.* Routledge, London.

Norenzayan, A., E. E. Smith, B. Kim, and R. E. Nisbett (2002). Cultural preferences for formal versus intuitive reasoning. *Cognitive Science. 26*, 653–684.

Noveck, I. and D. Sperber (Eds.) (2004). *Experimental pragmatics*. Palgrave MacMillan, New York.

Oakhill, J., P. Johnson-Laird, and A. Garnham (1989). Believability and syllogistic reasoning. *Cognition 31(2)*, 117–40.

Oaksford, M. and N. Chater (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review 101*, 608–631.

Oaksford, M. and K. Stenning (1992). Reasoning with conditionals containing negated constituents. *Journal of Child Language: Learning, Memory and Cognition 18*, 835–854.

Oberlander, J., K. Stenning, and R. Cox (1999). Hyperproof: Abstraction, visual preference and multimodality. In M. d. R. Lawrence S. Moss, Jonathan Ginzburg (Ed.), *Logic, Language and Computation*. Stanford, California.

Ogbu, J. (1990). Cultural model, identity, and literacy. In J. S. . R. S. . G. Herdt (Ed.), *Cultural psychology: Essays on comparative human development*. Cambridge University Pres.

O'Hear, A. (1997). *Beyond Evolution: Human Nature and the limits of evolutionary explanation*. Clarendon Press, Oxford.

Olson, D. (1993). Writing, literal meaning and logical proof. In R. J. Scholes (Ed.), *Literacy and Language Analysis*. Lawrence Erlbaum Associates.

Olson, D. R. (1977). From utterance to text: The bias of language in speech and writing. *Harvard Educational Review 47*, 257–81.

Olson, D. R. (1994). *The World on Paper: The Conceptual and Cognitive Implications of Writing and Reading*. Cambridge University Press, Cambridge.

Olson, D. R. and J. Astington (1990). Talking about text: how literacy contributes to thought. *Journal of pragmatics 14*, 705–721.

Ong, W. J. (1982). *Orality and Literacy: The Technologizing of the Word*. Methuen, New York.

Park, H. and L. Reder (2004). Moses illusion: Implication for human cognition. In *Cognitive illusions.* Hove: Psychology Press.

Partee, B. (1985). Situations, worlds and contexts. *Linguistics and Philosophy 8*, 53–58.

Pelletier, F. and N. Asher (1997). Generics and defaults. In J. van Benthem and A. ter Meulen (Eds.), *Handbook of Logic and Language*. Elsevier.

Piattelli-Palmarini, M. (1994). *Inevitable Illusions: how mistakes of reason rule our minds*. John Wiley, New York.

Pinnock, P. (1994). *Xhosa: a cultural grammar for beginners*. African Sun Press, Cape Town.

Politzer, G. (1981). Differences in interpretation of implication. *American Journal of Psychology 94*, 461–477.

Politzer, G. (1986). Laws of language use and formal logic. *Journal of Psycholinguistic Research 15*, 47–92.

Politzer, G. (2004). Reasoning, judgement and pragmatics. In I. A. Noveck and D. Sperber (Eds.), *Experimental pragmatics*. Palgrave MacMillan, New York.

Politzer, G. and L. Macchi (2002). Reasoning and pragmatics. *Mind and Society 1(1)*, 73–93.

Politzer, G. and I. Noveck (1991). Are conjunction rule violations the result of conversational rule violations? *Journal of Psycholinguistic Research 20*, 83–103.

Raven, D. (1996). The enculturation of logical practice. *Configurations 4*(3), 381–425.

Read, C., Y. Zhang, H. Nie, and B. Ding (1986). The ability to manipulate speech sounds depends on knowing alphabetic reading. *Cognition 24*, 31–44.

Read, S. (1994). Formal and material consequence. *Journal of Philosophical Logic 23: 3*, 247–265.

Recanati, F. (1996). Domains of discourse. *Linguistics and philosophy 19(5)*, 445–475.

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

Robinson, E., H. Goelman, and D. Olson (1983). Children's understanding of the relation between expressions (what was said) and intentions (what was meant). *British Journal of Devevelopmental Psychology 1*, 75–86.

Rogoff, B. (2003). *The Cultural Nature of Human Development*. Oxford University Press.

Rogoff, B. and J. Lave (1984). *Everyday cognition: its development in social context*. Harvard University Press, Cambridge, Massachusetts.

Rooth, M. (1985). *Association with focus*. Ph. D. thesis, University of Massachusetts.

Saxe, G. (1988). Candy selling and math learning. *Educational Researcher 17*(6), 14–21.

Schluchter, W. (1981). *The rise of Western rationalism*. University of California Press, Berkeley.

Schwarzschild, R. (1996). *Pluralities*. Springer.

Scribner, S. (1997). *Mind and Social Practice: selected writings of Sylvia Scribner*. Cambridge University Press, Cambridge.

Scribner, S. and M. Cole (1981). *The psychology of literacy*. Harvard University Press, Cambridge, Massachusetts.

Searle, J. (1979). Metaphor. In J. Searle and A. Ortony (Eds.), *Metaphor and thought*. Cambridge University Press.

Siddiqui, S., R. F. West, and K. E. Stanovich (1998). The influence of print exposure on syllogistic reasoning and knowledge of mental-state verbs. *Scientific Studies of Readingp 2*(1), 81–96.

Simons, P. (1995). Meaning and language. In B. Smith and D. W. Smith (Eds.), *The Cambridge Companion to Husserl*. Cambridge University Press, Cambridge.

Sloman, S. (1998). Categorical inference is not a tree: the myth of inheritance hierarchies. *Cognitive Psychology 35*, 1–33.

Smith, B. (1989). Logic and formal ontology. In J. Mohanty and W. McKenna (Eds.), *Husserl's Phenomenology: A textbook*, pp. 29–67. Lanham: University Press of America.

Smith, B. and D. W. Smith (Eds.) (1995). *The Cambridge Companion to Husserl*. Cambridge University Press, Cambridge.

Smith, E. R. and F. D. Miller (1978). Limits on perception of cognitive processes: a reply to nisbett and wilson. *Psychological Review 85(4)*, 355–362.

Sober, E. (1979). Psychologism. *Journal of the Theory of Social Behaviour 8 (2)*, 165–191.

Sportiche, D. (1988). A theory of floating quantifiers and its corollaries for constituent structure. *Linguistic inquiry 19*(3), 425–449.

Stalnaker, R. (1975). Indicative conditionals. *Philosophia 5*(3), 269–286.

Stanovich, K. (1993). Does reading make you smarter? literacy and the development of verbal intelligence. *Advances in child development and behavior 24*, 133–80.

Stanovich, K. (1999). *Who is rational? Studies of individual differences in reasoning*. Lawrence Erlbaum Associates.

Stein, E. (1996). *Without Good Reason: The Rationality Debate in Philosophy and Cognitive Science*. Oxford University Press.

Stenning, K. (2002). *Seeing reason. Image and language in learning to think.* Oxford: Oxford University Press.

Stenning, K. (2007, July). The origin of specialties,. Presentation at ESPP 2007, Geneva.

Stenning, K. and R. Cox (2006). Reconnecting interpretation to reasoning through individual differences. *Quarterly Journal of Experimental Psychology 59*(8), 1454 – 1483.

Stenning, K. and M. van Lambalgen (2001). Semantics as a foundation for psychology: A case study of wason's selection task. *Journal of Logic, Language and Information 10*, 273–317.

Stenning, K. and M. van Lambalgen (2004). A little logic goes a long way: basing experiment on semantic theory in the cognitive science of conditional reasoning. *Cognitive Science 28(4)*, 481–530.

Stenning, K. and M. van Lambalgen (2008). *Human reasoning and cognitive science*. Cambridge, MA.: MIT Press.

Stenning, K., P. Yule, and R. Cox (July 1996). Quantifier interpretation and syllogistic reasoning. *Cognitive Science Conference, UCSD, La Holla, Ca.*, 678–683.

Stich, S. P. (1990). *The fragmentation of reason: preface to a pragmatic theory of cognitive evaluation*. MIT Press. Cambridge, Massachusetts.

Stock, B. (1990). *Listening for the Text: On the Uses of the Past*. Baltimore: Johns Hopkins University Press.

Strawson, P. (1952). *Introduction to logical theory*. Methuen: London.

Street, B. (1984). *Literacy in Theory and Practice*. Cambridge Studies in Oral and Literate Culture, Cambridge.

Sundholm, G. (1994). Ontologic versus epistemologic: some strands in the development of logic: 1837 - 1957. In *Logic and Philosophy of Science in Uppsala*.

Sweetser, E. E. (1984). *Semantic structure and semantic change: a cognitive linguistic study of modality, perception, speech acts, and logical relations*. Ph. D. thesis, Berkeley.

Tarski, A. (1936/1956). *The Concept of Truth in Formalized Languages*. Oxford University Press.

Thompson, V. A. (2000). The task-specific nature of domain-general reasoning. *Cognition 76*, 209–268.

Torrance, N. and D. Olson (1987). Development of the metalanguage and the acquisition of literacy: A progress report. *Interchange 18*, 136–146.

Traugott, E., A. ter Meulen, J. Reilly, and C. Ferguson (Eds.) (1986). *On Conditionals*. Cambridge University Press, Cambridge.

Tsolwana, H. (1996). *New Xhosa Phrase Book and Vocabulary*. J L van Schaik.

Tulviste, P. (1991). *The cultural-historical development of verbal thinking*. Nova Science Publishers.

van Benthem, J. (2006). Where is logic going, and should it? *Topoi 25*, 117–122.

van Lambalgen, M. (to appear). Evolutionary considerations on logical reasoning. In D. Westerstahl (Ed.), *Proceedings of Twelfth International Congress on Logic, Methodololgy and Philosophy of Science, Oviedo, Spain, 2003.*

Veltman, F. (1986a). Data semantics and the pragmatics of indicative conditionals. In E. C. Traugott, A. ter Meulen, J. S. Reilly, and C. A. Ferguson (Eds.), *On Conditionals*. Cambridge University Press, Cambridge.

Veltman, F. (1986b). *Logics for Conditionals. PhD thesis*. University of Amsterdam.

Wason, P. (1966). Reasoning. In B. M. Foss (Ed.), *New Horizons in Psychology*, pp. 135–151. Penguin.

Wason, P. (1968a). Reasoning about a rule. *Quarterly Journal of Experimental Psychology 20*(3), 273–281.

Wason, P. (1968b). 'on the failure to eliminate hypotheses ...' – a second look. In P. Wason and P. Johnson-Laird (Eds.), *Thinking and Reasoning*. Penguin.

Wason, P. and J. Evans (1975). Dual processes in reasoning? *Cognition 3*(2), 141–154.

Wason, P. and P. Johnson-Laird (1970). A conflict betweeng selecting and evaluating information in an inferential task. *British Journal of Psychology 61*(4), 509–515.

Wason, P. and D. Shapiro (1971). Natural and contrived experience in a reasoning problem. *Quarterly Journal of Experimental Psychology 23*, 63–71.

Wason, P. C. and P. N. Johnson-Laird (1972). *Psychology of Reasoning: Structure and Content*. Boston: Harvard University Press.

Weber, M. (1978). *Economy and society, Vol. 1 and 2*. University of California Press, Berkeley.

Wilkins, M. (1928). The effect of changed material on the ability to do formal syllogisms. *Archives of Psychology 102*.

Willemsen, M. (2001). *Logical Reasoning and Culture*. Ph. D. thesis, Tilburg University.

Wilson, B. R. (Ed.) (1970). *Rationality*. Basil Blackwell, Oxford.

Winch, P. (1964). Understanding a primitive society. *American Philosophical Quarterly 1*, 307–324.

Woodworth, R. and S. B. Sells (1935). An atmosphere effect in syllogistic reasoning. *Journal of Experimental Psychology 18*, 451 – 460.

Samenvatting

Dit proefschrift gaat over ons vermogen om logisch te redeneren, en over hoe dat tot nu toe onderzocht is. Het doel is om de bestaande benadering van logisch redeneren kritisch te belichten en te vernieuwen, op basis van een combinatie van empirisch werk en theoretische analyse. Hoe dit bereikt is, wordt duidelijk door de hieronder gegeven korte beschrijvingen van elk hoofdstuk.

In het eerste hoofdstuk presenteer ik de bevindingen van mijn interview-studie met geletterde, minder geletterde en ongeletterde proefpersonen. Deze studie, uitgevoerd in Zuid Afrika, is vergelijkbaar met de studies van Luria (1976) en Scribner (1977). Proefpersonen krijgen een zogenaamd logisch probleem voorgeschoteld; dit is in de vorm van een tweetal zinnen in de rol van premissen, en een daaropvolgende vraag over conclusies die uit de zinnen zouden kunnen volgen. Deze vraag moeten de proefpersonen goed (weten te) beantwoorden. Maar dat doen ze niet, zo constateerden zowel Luria als Scribner. Deze eerdere studies kwamen uit op voornamelijk negatieve conclusies over de vermogens van ongeletterde mensen om logisch te redeneren. Het doel van mijn hoofdstuk is om de basis van deze negatieve conclusies kritisch te onderzoeken, om te zien in hoeverre ze gerechtvaardigd zijn. We vinden dat de negatieve conclusies inderdaad te snel getrokken waren. Dit komt vooral doordat er niet genoeg aandacht geschonken was aan de semantische en pragmatische context waarin de logische problemen voorgesteld waren. Daarmee bedoel ik dat er verschillende interpretaties mogelijk zijn, van beide de premissen en van de vraag, maar dat het eerdere werk geen rekening hield met die variëteit in interpretaties. Hedendaags linguistisch onderzoek heeft ons geleerd dat veel factoren bijdragen aan wat voor interpretatie een persoon toeschrijft aan een zin. Ons toenemend inzicht in deze factoren kan gebruikt worden om beter in te zien wat proefpersonen doen met de gepresenteerde zinnen. We vinden dat zowel geletterde als ongeletterde proefpersonen veel werk moeten doen om bij een interpretatie van de zinnen te komen, en dat ze inderdaad vaak uitkomen bij een interpretatie die anders is dan de experimentator zelf in gedachten had. Dit is echter geen reden om te denken dat proefpersonen niet

logisch redeneren; eerder laat het de gebreken van vroegere studies zien.

Hoofdstuk 2 gaat verder in op één van de bevindingen van het eerste hoofdstuk, namelijk dat de minder geletterde proefpersonen 'beter' presteerden in logische taken met conditionele zinnen ('als p dan q') dan in taken met universele gekwantificeerde zinnen ('alle x zijn y'). Ik zoek een verklaring hiervoor in het dagelijkse gebruik van zulke zinnen. Met het oog daarop is er een verkennende analyse gemaakt van het Engelse woord 'all' aan de hand van zijn voorkomen in een corpus van gesproken taal. Hetzelfde is gedaan met de vertaling in het Xhosa, '-onke', want de studie van Hoofdstuk 1 was in Xhosa uitgevoerd. De resultaten van deze analyse suggereren dat 'all' en vergelijkbare vertalingen vooral gebruikt worden voor contingente generalisaties, die heel gevoelig zijn voor context, en dat zij minder worden gebruikt voor de uitdrukking van wetmatige generalisaties, die ongevoelig zijn voor context. Denk hierbij aan de verschil tussen 'alle meisjes zijn gaan zwemmen' (alleen dan waar wanneer de bewering betrekking heeft op een specifieke groep op een specifiek moment) en 'alle mensen hebben twee benen' (in het algemeen waar, juist voor zover het niet verwijst naar een specifieke subgroep van de mensheid). Het gebruik van zinnen van de vorm 'alle x zijn y' in logische taken leunt eerder op deze tweede, universele betekenis. Gegeven dat dit in alledaags taalgebruik minder vaak voorkomt, is het niet verbazingwekkend dat ongeletterde proefpersonen sneller problemen met zulke zinnen hebben; dat beweer ik althans. In contrast daarmee, levert een vergelijkbare analyse van zinnen van de vorm 'als p dan q' de bevinding op dat er een grote overeenkomst bestaat tussen hoe zulke zinnen gebruikt worden in dagelijkse taal, en hoe ze gebruikt worden in studies naar logische redeneringen. Dit verklaart waarom 'als ... dan' zinnen makkelijker zijn voor proefpersonen in deze studies.

In Hoofdstuk 3 wordt het onderzoek geplaatst in de bredere context van het debat over de cognitieve gevolgen van geletterdheid. Een kritisch overzicht van empirische resultaten en theoretische standpunten in dit gebied onthult dat veel werk niet nauwkeurig of zelf-kritisch genoeg is geweest; gepaard hiermee geef ik suggesties voor verbeteringen in toekomstig werk. Het werk van theoreticus David Olson is in nader detail beschreven en beoordeeld. In het bijzonder heb ik zijn hypothese over letterlijke betekenis en diens rol in onder andere redeneertaken onderzocht. Een aangepaste versie van deze hypothese is voorgesteld om het verschil tussen proefpersonen in redeneertaken te verklaren.

In nauwe aansluiting daarop ga ik in Hoofdstuk 4 verder met de notie van letterlijke betekenis, 'literal meaning'. Dit hoofdstuk vormt op een bepaalde manier de analytische ruggengraat van het hele proefschrift, omdat de onderliggende aannames van het empirische onderzoek naar menselijk redeneren boven tafel worden gehaald om ze daarna te kunnen toetsen. Welk begrip van taal, betekenis, logica, en welke relaties ertussen, worden gebruikt in de psychologie van het redeneren? Zijn ze redelijk? Is het gerechtvaardigd om redeneervermogens op basis hiervan te beoordelen? Mijn hoofdstelling hier is dat de notie van letterlijke betekenis te onkritisch in gebruik is genomen als de basis van logisch

redeneren. In plaats van te onderzoeken welke mogelijke interpretaties er van premissen in een redeneertaak zijn, en uit te zoeken waarom een proefpersoon tot een bepaalde interpretatie komt, hebben onderzoekers geleund op het idee dat logisch redeneren alleen gebaseerd is op de letterlijke betekenis van premissen, en dat aandacht voor interpretatieprocessen daarom niet nodig is. Gebruikmakend van argumenten van François Recanati, laat ik zien dat de letterlijke betekenis vaak een mythe is, en dat die zeker geen equivalent is van een notie van 'natuurlijke' of 'logische' betekenis. In tegenstelling hiertoe beargumenteer ik dat de notie van letterlijke betekenis een theoretisch construct is, dat is overgebleven uit de vroege taalfilosofie, maar waarvan de houdbaarheidsdatum nu reeds verstreken is. We kunnen niet ontkomen aan de noodzaak om interpretatieprocessen op te nemen in een theorie van logisch redeneren.

Tenslotte doet Hoofdstuk 5 verslag van een interview-study met geletterde proefpersonen (universiteitsstudenten) over een belangrijke redeneertaak in het onderzoeksveld, de selectietaak van Wason. Hierin bouw ik op eerder werk van Stenning en van Lambalgen over het paradigma van Wasons taak, waarin ze op zoek zijn gegaan naar de mogelijke interpretaties die proefpersonen hanteren in de taak, met het oog op de rol daarvan in het verklaren van de latere keuzes van proefpersonen. Veel onderzoek naar deze taak heeft alleen gekeken naar de keuzes van proefpersonen, maar de interpretaties waarop ze gebaseerd zijn, bleven onderbelicht. Dit hoofdstuk is bedoeld om te laten zien dat, zodra we aandacht voor interpretatieprocessen koppelen aan expliciet gedrag, de bestaande paradigma's van het veld nog steeds een zeer vruchtbaar terrein voor onderzoek naar menselijk redeneren vormen.

Abstract

This dissertation provides a critical assessment of investigations into logical reasoning ability as reported in the psychology of reasoning literature. This is achieved through a combination of original empirical research with theoretical analysis, using insights from formal semantics and philosophy of language. In the first chapter, the findings of an interview study with subjects with varying but low education levels are presented. The aim of the chapter is to engage with and challenge the negative conclusions reached in previous studies with illiterate subjects. This is achieved by reanalysing the typical responses with heightened attention to semantic and pragmatic factors which shape subjects' responses. Chapter 2 provides an exploratory analysis of the everyday use of some terms used in reasoning studies. Comparing everyday usage with typical usage of the terms in reasoning studes enables us to explain why certain terms are more difficult for some subjects than others. In Chapter 3 these findings are placed in the context of the broader debate on the cognitive consequences of literacy. The work of David Olson is handled in detail, and an adapted version of his 'literal meaning' hypothesis is proposed to explain the difference between schooled and unschooled subjects' performance in reasoning tasks. Consideration of the notion of 'literal meaning' feeds directly into the next chapter, Chapter 4. This is the analytic backbone for the dissertation, as it explains how much previous work in the psychology of reasoning has relied on an oversimplified picture of the relations between natural language, logic, and reasoning, giving rise to confusions about logical form, its relation to meaning in general and to literal meaning in particular. Finally, Chapter 5 reports on an interview study on Wason's original selection task, building on the work of Stenning and van Lambalgen in this area. This chapter is intended to show that, despite the criticisms of the previous chapter, standard reasoning experiments can provide rich data for a positive theory of human reasoning. The dissertation ends with a summary and concluding remarks.

Titles in the ILLC Dissertation Series:

ILLC DS-2001-01: **Maria Aloni**
Quantification under Conceptual Covers

ILLC DS-2001-02: **Alexander van den Bosch**
Rationality in Discovery - a study of Logic, Cognition, Computation and Neuropharmacology

ILLC DS-2001-03: **Erik de Haas**
Logics For OO Information Systems: a Semantic Study of Object Orientation from a Categorial Substructural Perspective

ILLC DS-2001-04: **Rosalie Iemhoff**
Provability Logic and Admissible Rules

ILLC DS-2001-05: **Eva Hoogland**
Definability and Interpolation: Model-theoretic investigations

ILLC DS-2001-06: **Ronald de Wolf**
Quantum Computing and Communication Complexity

ILLC DS-2001-07: **Katsumi Sasaki**
Logics and Provability

ILLC DS-2001-08: **Allard Tamminga**
Belief Dynamics. (Epistemo)logical Investigations

ILLC DS-2001-09: **Gwen Kerdiles**
Saying It with Pictures: a Logical Landscape of Conceptual Graphs

ILLC DS-2001-10: **Marc Pauly**
Logic for Social Software

ILLC DS-2002-01: **Nikos Massios**
Decision-Theoretic Robotic Surveillance

ILLC DS-2002-02: **Marco Aiello**
Spatial Reasoning: Theory and Practice

ILLC DS-2002-03: **Yuri Engelhardt**
The Language of Graphics

ILLC DS-2002-04: **Willem Klaas van Dam**
On Quantum Computation Theory

ILLC DS-2002-05: **Rosella Gennari**
Mapping Inferences: Constraint Propagation and Diamond Satisfaction

ILLC DS-2002-06: **Ivar Vermeulen**
A Logical Approach to Competition in Industries

ILLC DS-2003-01: **Barteld Kooi**
Knowledge, chance, and change

ILLC DS-2003-02: **Elisabeth Catherine Brouwer**
Imagining Metaphors: Cognitive Representation in Interpretation and Understanding

ILLC DS-2003-03: **Juan Heguiabehere**
Building Logic Toolboxes

ILLC DS-2003-04: **Christof Monz**
From Document Retrieval to Question Answering

ILLC DS-2004-01: **Hein Philipp Röhrig**
Quantum Query Complexity and Distributed Computing

ILLC DS-2004-02: **Sebastian Brand**
Rule-based Constraint Propagation: Theory and Applications

ILLC DS-2004-03: **Boudewijn de Bruin**
Explaining Games. On the Logic of Game Theoretic Explanations

ILLC DS-2005-01: **Balder David ten Cate**
Model theory for extended modal languages

ILLC DS-2005-02: **Willem-Jan van Hoeve**
Operations Research Techniques in Constraint Programming

ILLC DS-2005-03: **Rosja Mastop**
What can you do? Imperative mood in Semantic Theory

ILLC DS-2005-04: **Anna Pilatova**
A User's Guide to Proper names: Their Pragmatics and Semanics

ILLC DS-2005-05: **Sieuwert van Otterloo**
A Strategic Analysis of Multi-agent Protocols

ILLC DS-2006-01: **Troy Lee**
Kolmogorov complexity and formula size lower bounds

ILLC DS-2006-02: **Nick Bezhanishvili**
Lattices of intermediate and cylindric modal logics

ILLC DS-2006-03: **Clemens Kupke**
Finitary coalgebraic logics

ILLC DS-2006-04: **Robert Špalek**
Quantum Algorithms, Lower Bounds, and Time-Space Tradeoffs

ILLC DS-2006-05: **Aline Honingh**
The Origin and Well-Formedness of Tonal Pitch Structures

ILLC DS-2006-06: **Merlijn Sevenster**
Branches of imperfect information: logic, games, and computation

ILLC DS-2006-07: **Marie Nilsenova**
Rises and Falls. Studies in the Semantics and Pragmatics of Intonation

ILLC DS-2006-08: **Darko Sarenac**
Products of Topological Modal Logics

ILLC DS-2007-01: **Rudi Cilibrasi**
Statistical Inference Through Data Compression

ILLC DS-2007-02: **Neta Spiro**
What contributes to the perception of musical phrases in western classical music?

ILLC DS-2007-03: **Darrin Hindsill**
It's a Process and an Event: Perspectives in Event Semantics

ILLC DS-2007-04: **Katrin Schulz**
Minimal Models in Semantics and Pragmatics: Free Choice, Exhaustivity, and Conditionals

ILLC DS-2007-05: **Yoav Seginer**
Learning Syntactic Structure

ILLC DS-2008-01: **Stephanie Wehner**
Cryptography in a Quantum World

ILLC DS-2008-02: **Fenrong Liu**
Changing for the Better: Preference Dynamics and Agent Diversity

ILLC DS-2008-03: **Olivier Roy**
Thinking before Acting: Intentions, Logic, Rational Choice

ILLC DS-2008-04: **Patrick Girard**
Modal Logic for Belief and Preference Change

ILLC DS-2008-05: **Erik Rietveld**
Unreflective Action: A Philosophical Contribution to Integrative Neuroscience

ILLC DS-2008-06: **Falk Unger**
Noise in Quantum and Classical Computation and Non-locality

ILLC DS-2008-07: **Steven de Rooij**
Minimum Description Length Model Selection: Problems and Extensions

ILLC DS-2008-08: **Fabrice Nauze**
Modality in Typological Perspective

ILLC DS-2008-09: **Floris Roelofsen**
Anaphora Resolved