

Personnel

Michiel van Lambalgen (co-director)

Frank Veltman (co-director)

Marian Counihan (PhD student)

Darrin Hindsill (PhD student)

Katrin Schulz (PhD student)

Description of the proposed research

The problem: the role of nonmonotonicity in cognition.

The proposed programme investigates the interplay between language, reasoning and cognition, with particular emphasis on the role of *nonmonotonicity*. This phenomenon can be illustrated by the following example. When presented with the premisses (1) 'If Fred has to write an essay, he is in the library', and (2) 'Fred has to write an essay', most people would draw the conclusion that Fred is in the library. However, when the additional premiss (3) 'The library is closed', is supplied, people tend to withdraw this conclusion and to infer instead that Fred is not in the library. This simple example has been studied in cognitive psychology under the name 'suppression effect' (Byrne 1989), which reflects the fact that psychologists often regard this phenomenon as a failure to derive classical logical consequences. We think it is better described as an ability in nonmonotonic reasoning.

Generally, we say that a relation of logical consequence is *nonmonotonic* if the addition of a new premiss to a given set of premisses may make conclusions previously obtained invalid. Here is another, more recent, example from the psychological literature. Klauer et al. 2000 adduce experimental evidence for the claim that logical inference tasks (i.e. does conclusion C follow from premisses A,B?) are typically not solved using the definition of validity of classical logic ('if the premisses are true, so is the conclusion'), but by the construction of *a single model*. Which model is constructed depends on whether the conclusion is believed or not. If the conclusion C is believed, then a person will construct a model of A,B,C; if this succeeds, then C follows, otherwise nothing follows. (If C is not believed, the person will construct a model of A,B, not-C.) This procedure is clearly nonmonotonic: adding another premiss D may make the set A,B,D,C inconsistent, so that C would, according to the above criterion, no longer follow from A,B,D.

Nonmonotonic reasoning represents a natural way of thinking, and people are good at it. They *have* to be good at it, because planning in an uncertain environment (i.e. reasoning about change, events and actions) requires this type of reasoning. Speaking more formally, both monotonic and nonmonotonic reasoning start from the maxim:

M) 'assume only what is given in the premisses'

but they implement (M) in different ways. Nonmonotonic reasoning takes (M)

to mean: all assumptions beyond those required by the premisses are false; instead, monotonic reasoning interprets (M) as: suspend judgement on statements which do not follow (and whose negations do not follow) from the premisses. In the interesting cases, these two interpretations of (M) can be reformulated as follows. In nonmonotonic reasoning, people construct a *minimal model* of the premisses (which is often unique); in monotonic reasoning, they must consider *all* models of the premisses. Experiments seem to indicate that people are better in nonmonotonic reasoning than in classical monotonic reasoning. For example, in controlled reasoning experiments, it is consistently found that subjects are better at reasoning with normative statements (which is nonmonotonic) than at reasoning with indicatives. The intriguing study Deglin and Kinsbourne 1996 has indicated that the monotonic interpretation of (M) is a feat accomplished by the left hemisphere of the brain, whereas the right hemisphere is not able to make the

required abstraction from truth values. The anthropological study Scribner 1997 indicated that the ability to interpret (M) monotonically is related to literacy.

Nonmonotonicity also pervades natural language. To mention a few examples: the progressive form, perceptual verbs, and generic sentences all induce a strong nonmonotonic inference potential (Miller and Johnson-Laird 1976, Jackendoff 1990, Veltman 1996, van der Does and van Lambalgen 2000, Hamm and van Lambalgen 2000). If the above considerations indicating that nonmonotonicity is inherently linked to the construction of a single model of a set of sentences, are correct, the question arises how models are represented in human memory. (The weakness of the popular 'mental models' theory due to Johnson-Laird 1983 is precisely that it fails to address this question). Stenning (1978,1986,1992) has much to say on this issue, and we intend to use this work as a benchmark against which to gauge the complexity of our logical proposals.

Generally speaking, there exists a large gap between logical and psychological work on reasoning, which has led logicians to neglect investigating the possible psychological relevance of their formalisms, and psychologists to take a rather cavalier attitude toward logical and semantic subtleties. We therefore see the scientific challenge posed by nonmonotonicity as the development of a framework in which to study the interaction of cognition and semantics. In more detail, we may summarise the scientific problems posed by nonmonotonicity as follows.

1. Why (and to what extent) are humans apparently better at nonmonotonic, than at monotonic reasoning? How is this related to the fact that from a formal-computational point of view, nonmonotonic reasoning is generally much more complex than monotonic reasoning?
2. How is nonmonotonic reasoning implemented and learned?
3. In what forms does nonmonotonicity occur in natural language? Is its manifestation the same across languages?
4. Does nonmonotonicity necessitate a new theory of linguistic meaning?
5. Are current approaches to formal semantics adequate for dealing with nonmonotonicity in reasoning and natural language generally?

Methodology

We will study these problems in three subprojects (which are not independent):

1. psychology of reasoning and nonmonotonicity
2. nonmonotonicity and *Aktionsart*
3. reasoning with false assumptions: the counterfactual mood

1. We investigate experimentally the relative difficulties of reasoning with indicative and deontic conditionals. An important finding of earlier experiments, reported in Stenning and van Lambalgen 2001 is that a compositional view of meaning in which interpretation of a statement precedes reasoning with that statement, cannot be upheld; instead the reasoning task partly determines the interpretation chosen. We will further investigate this phenomenon, and try to devise mathematical models which take account of this experimental fact. A rather ambitious goal, motivated by educational concerns, is to try to relate conditions in which a subject learns logical expressions to the actual use of these expressions made by that subject. In this project we would like to employ a postdoc

2. Reasoning about change, actions, and events permeates ordinary life, and several features of natural language, in particular nominalisation and *Aktionsart*, relate to this type of nonmonotonic reasoning. We believe this fact necessitates a drastic change in the concept of linguistic meaning, along the lines suggested by Moschovakis 1990 and worked out in Hamm and van Lambalgen 2000: the meaning of an expression is an *algorithm* which computes the denotation of that expression in a minimal model of someone's world-knowledge. The main thrust of the project will be to develop this concept of meaning and to apply it to nominalisation and aspectual phenomena. There is an intimate relation to the research outlined in the previous project; we believe that the concept of meaning indicated here explains both good performance in deontic reasoning tasks (Cheng and Holyoak 1985) and the 'perspective change' observed there (Gigerenzer and Hug 1992). Also, as mentioned above, recent psychological evidence suggests that construction of a single model for a set of premisses is of great importance in reasoning. In this project we would like to employ a PhD student.

3. Reasoning from contrary-to-fact hypotheses is strongly nonmonotonic. This form of reasoning cannot be understood without reference to the human representation of causal relationships. In this project we will be concerned with the psychology of this type of reasoning and the way it is represented across languages. Therefore the project will consist of psychological experiments determining the precise relationship between causal and contrary-to-fact reasoning, and the construction of a semantics which takes account of the findings of these experiments, and which does justice to the different ways in which this form of reasoning is expressed in different languages. In this project we would like to employ a PhD student. Slightly more detailed descriptions of the three projects can be found in section 'Subprojects' below.

This research proposal is guided by the idea that in the area of nonmonotonic reasoning logic and psychology should go hand in hand. We study nonmonotonicity simultaneously at three levels:

1. mathematical models
2. psychological experiments
3. collection of typological and diachronic linguistic data.

(1) At the level of mathematical modelling, it is our intention to compare, evaluate and perhaps integrate two different approaches to semantics with cognitive ambitions, the dynamic approach, and the computational approach. We subscribe to the basic tenet of cognitive and conceptual semantics (e.g. Jackendoff 1990) that the primary objects with which semantics should be concerned are mental representations. However, we also think that the practitioners of cognitive semantics construct a spurious contrast between logical methods and their own and that cognitive semantics could profit from logical and mathematical techniques. Furthermore, increased application of formal techniques is necessary to bring phenomena involving nonmonotonicity within the scope of natural language processing.

(2) It is our intention to test the resulting formal theories in controlled experiments. This pertains in particular to studying inferences. These experiments will be performed in collaboration with professor Keith Stenning of the Human Communication Research Center (Edinburgh), an expert on the experimental investigation of reasoning. Previous collaboration has led to the paper Stenning and van Lambalgen 2001. At the time of writing, Stenning and van Lambalgen have finished another round of experiments, whose results are now being written up.

(3) The proposed semantics claim to be cognitively relevant. It follows that their predictions should hold across languages, and should also remain valid through language change. We thus require both typological and diachronic studies of specific aspects of language, in particular the various types of conditionals and the way language conceptualizes events. This work will be done in collaboration with the Seminar für Sprachwissenschaft in Tübingen, which has relevant expertise.

It will be clear from the above description that the proposed research requires the collaboration of several disciplines; a 'programme' would provide the ideal opportunity.

Subprojects.

1. Psychology of reasoning and nonmonotonicity.

We mentioned in the introduction that people appear to be better at reasoning involving a nonmonotonic consequence relation, than they are at reasoning with a monotonic consequence relation. However, this way to put things is controversial. Many psychologists would detect a contrast between reasoning based on 'logical formal' and reasoning based on 'content' here.

Experiments in the psychology of reasoning typically concern tasks involving propositional connectives, mainly implication, and tasks involving simple inferences from predicate logic, mainly (although not exclusively) syllogisms. The most famous experimental paradigm for investigating implications is Wason's 4-card task, which comes in two varieties, one involving indicative conditionals (if p then q) and one involving normative or deontic conditionals (if p then *must* q). The traditional interpretation of the outcomes of these experiments is that subjects do not grasp the *truth* condition for an indicative conditional, but do understand when a deontic conditional is *violated*, because in contrast to the indicative conditionals, the deontic conditionals typically involve an easily understood real-life statement such as 'if you want to drink alcohol, you must be over 18'. This is the argument behind the 'evolutionary'

explanation of Cosmides 1989, which claims that humans have evolved 'cheating detectors' for the purpose of detecting violation of norms. Such 'cheating detectors' would not be triggered by the purely indicative conditionals of Wason's original experiment.

Based on experimental and methodological work reported in Stenning and van Lambalgen 2001, we claim that this is an oversimplification. For example, unlike an indicative implication, which can be true or false, a deontic rule has to be taken as given and can only be violated, which implies that for each card it can be decided independently whether that card violates the rule. This greatly reduces the demand on processing. We claim that, more generally, psychology has an unsatisfactory concept of 'logical form' which vitiates any attempt to distinguish between the form and content of an assertion.

It seems to us that more investigation into the precise logical structure of the various tasks is required, and that it is of some importance to make psychologists aware of the logical and semantic issues involved. This logical investigation should provide new directives for cognitive experimentation, based on a more realistic view of semantics. The traditional picture of a subject engaged in a reasoning task is this: the subject first interprets the statements involved, and then reasons on the basis of those statements. However, observations show a furious interaction between interpretation and reasoning (see Stenning and van Lambalgen 2001). It follows that the interaction itself has to be subject to mathematical modelling and that cognitive experimentation should take account of the interaction. This is most naturally done by means of conducting tutorial interviews with subjects, instead of the more traditional 'filling--in--a--form' format. (An additional argument in favour of this technique is that subjects may give the same answer for very different reasons.) From a semantical point of view, these observations show that in the case of the conditional, natural language understanding is not compositional. It is a moot point whether it is then profitable to attempt to give a compositional semantics for the conditional. (A noncompositional approach to semantics will be explored in the subproject *Nonmonotonicity and Aktionsart*)

It is also of importance, not least for educational purposes, to establish a connection between the way the various meanings of the conditional are learned by a subject, and the way in which that subject subsequently uses the conditional in reasoning (both monotonic and nonmonotonic). We have evidence that these meanings differ widely among subjects.

What was said above concerning psychological experiments on reasoning holds *a fortiori* for brain-imaging experiments such as those of Goel et al. 2000, Osherson et al. 1998, which attempt to find brain loci which correlate with specific types of reasoning (syllogistic and probabilistic). In our opinion, these experiments are not as informative as they could be due to insufficient attention to the semantic issues involved. A subsidiary goal of this project is therefore to establish a dialogue with neuroscientists in order design a methodology to investigate brain correlates of reasoning, which is sensitive to semantic issues. The last mentioned research will be ILLC's contribution to the programme of the nascent *Cognitive Science Center Amsterdam*.

II. Nonmonotonicity and Aktionsart.

The claim to be investigated in this subproject is that a concept of meaning based on nonmonotonicity is cognitively plausible, and moreover, that this concept of meaning can be given a precise mathematical form. Cognitive linguists have rightly emphasised the pervasiveness of nonmonotonicity in natural language. For instance Jackendoff argues that lexical semantics cannot be formulated in terms of necessary and sufficient conditions, but needs default conditions. Consider the following example.

- a. Harry climbed the mountain
- b. Harry climbed down the tree
- c. The snake climbed up the tree
- d. ?The snake climbed down the tree

The prototypical meaning of 'climb' is an upward, scrambling movement; that is, by default both conditions are satisfied. However, contextual information may cancel one or the other default implications. The last sentence shows that one cannot cancel both defaults.

Cognitive linguists also correctly point to the importance of *coercion*, the fact that in a certain context an expression may take on a meaning which it does not ordinarily have (cf. Langacker 1998).

However, the conclusion often drawn from these observations, that formal semantics is incapable of doing justice to the intricacies of language, is not justified. One simply has to look for a subtler formalism. In semantic formalisms which use a monotonic logic, the meaning of an expression is defined

as a function which assigns to every possible world a denotation (of the expression) in that world. Meaning is thus evaluated with respect to a fixed possible world structure. By contrast, in a nonmonotonic logic the denotation of an expression is not given once and for all, but determined by a *scenario*, that is, a set of statements containing all available knowledge concerning the situation at hand. The denotation is then considered only in 'preferred' or 'minimal' models of that scenario (often these models are even uniquely determined). As a consequence, new information, encoded as statements added to the scenario, may change the denotation, because the set of preferred or minimal models of the new scenario may differ from that corresponding to the old scenario.

This concept of meaning bears a strong resemblance to the notion of meaning introduced by Moschovakis 1990 in the context of computability theory. Moschovakis clothes Frege's (1892) distinction between sense and reference in a modern garb by equating the sense of an expression with the algorithm that corresponds with the expression, and the reference of the expression with the result of the computation. Moschovakis then goes on to propose a formal criterion for intensional identity, a problem unresolved since Frege.

We believe that a promising formalism for modelling nonmonotonic features of language use is obtained by combining Moschovakis' ideas with the rich formalisms for talking about events, processes and plans developed in Artificial Intelligence. Implicitly this was done already in Hamm and van Lambalgen 2000, where it is shown that the complexities of the English gerund can be formalized very well in such a framework. ('Implicitly', because at the time we were not aware of the over-arching framework provided by Moschovakis.)

The phenomena studied in Hamm and van Lambalgen 2000 include nominalisation, aspect (in particular the progressive), and causality. The pivotal notions here are 'event' and 'nonmonotonicity'. Of course there exists a voluminous literature on aspect, e.g. Comrie 1976 on the typology of aspect, Dowty 1979 on the rudiments of an 'aspectual calculus', or Verkuyl 1993 for a comparison of different approaches. Although currently much attention is given to the role of events in natural language semantics (cf. e.g. Pustejovsky's collection *The grammar of events*, 2000), the resulting theories are never axiomatised, so that it is unclear what does, or does not, follow from the theory. In addition, the needs of natural language processing call for a higher degree of formalisation than has been customary, which is what the present project aims to provide.

We claim that it is possible to develop a complete theory of *Aktionsart* in the framework outlined above, taking full account of such phenomena as causality, agency and telicity. Further interesting challenges are provided by aspectual phenomena in other languages than English. If, as we claim, the formalism reflects deep-seated properties of human cognitive processing of information about objects, events and states-of-affairs, then the formalism should be able to account for the data from a variety of languages. One focus of attention here is the work of Koptjevskaja-Tamm, which claims that in essence all nominalizations found in the languages investigated fall either under the perfect or under the imperfect rubric. This claim needs further semantic research.

We are currently investigating the representation of events in English diachronically, and try to determine whether the manifest syntactic changes correspond to a semantic change.

III. The counterfactual mood.

If everything we ever thought and said would be purely descriptive in nature, nonmonotonicity would be the exception rather than the rule. But in much we say we don't restrict ourselves to a description of the facts, we ponder possibilities that may, or might, or might have been the case, and as soon as these modalities are brought into play nonmonotonic inferences abound. 'If this match had been struck, it would have lit'. Yes, but if it had been wet or if there hadn't been any oxygen around, it wouldn't have.

In this project we are particularly interested in the 'counterfactual' modalities, which in English are exemplified in the sequences *might + have + past participle* and *would + have + past participle*. The reason for choosing the counterfactual mood as one of our topics is not only because they give rise to nonmonotonic inference patterns, but also because we think that our methodology will prove to be very fruitful here, since both language comparative studies and psychological experiments will help to inform the theoretical development.

A question that lends itself very well to empirical investigation is the question whether the modalities at issue express an epistemic or an alethic modality. Are they used to express something about the speaker's epistemic state or are they used to state a fact that is independent of what the speaker does or

does not know? A simple puzzle, devised by Hansson 1989, may illustrate the difference at stake here: You enter a town that you believe to have just two snack bars, A and B. There you see a man walking along the street with a hamburger, and you form the belief that at least one of the snack bars is open. Now as you approach one of the snack bars, it happens to be A, you see that the lights are on there. As a result you form the belief that that A is open. Question 1: Do you also believe 'If A had not been open, B would have been open'? Question 2: What if somewhat later you find out that snack bar B is in fact closed. Would you still believe 'If A had not been open, B would have been open'?

Most semanticists agree that the answer to the second question should be 'no', but the research community is really divided about the answer to the first question. The ones favouring an epistemic interpretation (for example Morreau 1992) say 'yes', the ones favouring an alethic interpretation (for example Rott 1999) say 'no'.

Experiments performed by Fillenbaum 1974 suggest that in understanding counterfactual conditionals, the alethic concept 'causality' is the primary notion, but more experimentation is needed. If future experiments corroborate these findings, the event-based semantics of Hamm and van Lambalgen 2000 has a role to play here. In fact, we may very well end up with a theory that gives a dynamic twist to the theory developed in project 2.

Sentences of the form 'It would have been the case that..' make no sense if they are presented without a context. As the examples above may have indicated, they typically occur in the context of a clause of the form 'If it had been the case that...!', and so far, all semantic investigations -- which started with Stalnaker 1968 -- have been restricted to their occurrence in that particular context.

However, conditionals are not the only place where we can find the counterfactual mood. Consider:

- (i) He did not drink any wine. He would have got sick.

One of the questions we want to pursue is in exactly which contexts the modal perfect can occur. That this is by no means a trivial question is perhaps best illustrated by comparing (i) with (ii):

- (ii) * He drank too much wine. He would not have got sick.

Why is it that (ii) makes no sense, unless one inserts something like 'otherwise' in front of the second sentence? The framework of dynamic semantics (Kamp and Reyle 1993, Groenendijk and Stokhof 1991, Veltman 1996), offers a way to get to grips with questions like this, as it is designed to study 'meaning in context'. Here, a start has been made in the work of Kibble (1994), Kaufmann (1997) and von Stechow (2001). There is still a lot of work to be done, however, as the theories developed so far leave so many questions open that they lack predictive power in concrete cases.

Linguistic comparison of the counterfactual construction in different languages comes in with a cluster of questions having to do with the interplay of tense and mood in the English *would+ have* construction. Again, let us illustrate the problems by an example. For an indicative conditional to make sense it is not necessary that the event described in the antecedent precedes the event in the consequent. There is nothing wrong with a sentence like:

If he left the interview smiling, it went well.

However, in the counterfactual mood, this cannot be done.

- * If he had left the interview smiling, it would have gone well.

One wonders if this phenomenon is due to the peculiar way in which of tense and mood are combined in the English modal perfect, or if there is a deeper, semantic or cognitive reason for it, which also affects the counterfactual mood in other languages.

References

- Beller, S. (2001) A model theory of deontic reasoning about social norms. *Proceedings CogSci 2001*, 91-96.
- Byrne, R.M.J. (1989) Suppressing valid inferences with conditionals. *Cognition* **31**, 61-83.
- Cheng, P. and K. Holyoak (1985) Pragmatic reasoning schema. *Cognitive Psychology* **17**, 391-416.
- Comrie, B. *Aspect: an introduction to the study of verbal aspect and related problems*. Cambridge University Press.
- Cosmides, L. (1989), The logic of social exchange; has natural selection shaped how humans reason? *Cognition* **31**, 187-276.
- Deglin, V.L. and M. Kinsbourne (1996). Divergent thinking styles in the hemispheres: how syllogisms are solved during transitory hemisphere suppression. *Brain and cognition* **31**, 285-307.
- van der Does, J. and M. van Lambalgen (2000), A logic of vision. *Linguistics and Philosophy* **23**, 1-92
- Dowty, D. (1979) *Word meaning and Montague grammar*. Reidel.
- Fillenbaum, S., 1974, Pragmatic normalization: Further results for some conjunctive and disjunctive sentences. *Journal of Experimental Psychology*, 102, 574-578.
- von Fintel, K. (2001)., Counterfactuals in a Dynamic Context. In: Michael Kenstowicz (ed.), *Ken Hale: a Life in Language*, MIT Press.
- Gigerenzer, G. and K. Hug (1992), Domain specific reasoning: social contracts, cheating, and perspective change. *Cognition* **43**, 127-171.
- Goel, V. et al , 2000, Dissociation of mechanism underlying syllogistic reasoning. *Neuroimage***12**, 504-514.
- Groenendijk, J. and M. Stokhof, 1991, 'Dynamic Predicate Logic', *Linguistics and Philosophy* **14**, 39-100.
- Hamm , F. and M. van Lambalgen. (2000) Event calculus, nominalization and the progressive. Submitted to *Linguistics and Philosophy*. Available from <http://www.semanticsarchivenet>.
- Hansson, S. O., 1989, A new semantical approach to the logic of preference. *Erkenntnis* **31**, 1--42.
- Jackendoff , R.S. (1990) *Semantic structures*. MIT Press.
- Johnson-Laird, P.N. (1983) *Mental models: towards a cognitive science of language, inference and consciousness*. Cambridge University Press.
- Johnson-Laird , P.N. and R.M.J. Byrne (in press) Conditionals: a theory of meaning, pragmatics and inference. To appear in *Psychological Review*.
- Kamp , H. and U. Reyle (1993). *From discourse to logic*. Kluwer.
- Kaufmann, Stefan, (1998), Context Inheritance in Discourse Processing. Presented at the Annual Meeting of the Linguistic Society of America, New York, January 09, 1998.
- Kibble, R. (1994), Dynamics of epistemic modality and anaphora. In Bunt, H., Muskens, R., and Rentier, G.(editors), *International Workshop on Computational Semantics*, Tilburg, NL. ITK.
- Klauer, K.C. et al (2000) On belief bias in syllogistic reasoning. *Psychological Review* **107**, 852-884.
- Koptjevskaja-Tamm, M. (1993) *Nominalizations*. Routledge.
- Langacker , R. W. (1998), Conceptualization, symbolization, and grammar. In. M. Tomasello (ed.), *The new psychology of language: cognitive and functional approaches to language structure*. Erlbaum.
- Miller, G.A. and P.N. Johnson-Laird (1976). *Language and Perception*. MIT Press.
- Morreau, M. 1992, 'Epistemic Semantics for Counterfactuals', *Journal of Philosophical Logic* **21**, 33-62.
- Moschovakis, Y.N. (1990) Sense and denotation as algorithm and value. In: J. Väänänen (ed.), *Logic Colloquium '90. Lecture Notes in Logic*. A.K. Peters.
- Osherson, D. et al (1998) Distinct brain loci in deductive versus probabilistic reasoning. *Neuropsychologia* **36(4)**, 369-376.
- Pustejovsky , J. and C. Tenny (2000) *Events as grammatical objects*. CSLI Press.
- Rott, H. , 1999, 'Moody Conditionals: Hamburgers, Switches, and the Tragic Death of an American President', in J. Gerbrandy et. al. (eds.) , *JFAK. Essays Dedicated to Johan van Benthem on the Occasion of his 50th Birthday*, Vossiuspers AUP, Amsterdam.
- Scribner, S. (1997), *Mind and social practice*. Cambridge University Press.

- Stalnaker, 1968, 'A Theory of Conditionals', in *Studies in Logical Theory, American Philosophical Quarterly* monograph 2, Blackwell, Cambridge, pp. 98--112.
- Stenning, K.D. (1978) Anaphora as an approach to pragmatics. In M. Halle, J. Bresnan, G.A. Miller (eds.) *Linguistic Theory and Psychological Reality*. MIT Press.
- Stenning, K.D. (1986) On making models: a study of constructive memory. In T. Myers, E.K. Brown and B.O. McGonigle (eds.) *Discourse and Reasoning*. Academic Press.
- Stenning, K.D et al (1992) Representations of individuals and the processing of reference change. *Quart. J. Exp. Psych.* **45A(3)**, 373-398.
- Stenning, K.D. and M. van Lambalgen 2001, Semantics as a foundation for psychology: a case study of Wason's selection task. *J. Logic, Language and Information* **10(3)**, 273-317.
- Veltman, F. 1996, 'Defaults in Update Semantics', *Journal of Philosophical Logic* **25**, 221--261.
- Veltman, F., (in preparation), Counterfactual Subordination. Invited talk, to be presented at the First International Workshop on Semantics, Pragmatics and Rhetoric. Donostia, San Sebastian, November 2001.
- Verkuyl, H.J. (1993), *A theory of aspectuality*. Cambridge University Press.
- Yarlett, D and Ramscar, M (2001). Structural Determinants of Counterfactual Reasoning. Proceedings CogSci 2001, 1182-1187.