Title
Cognitive U ncertainty in Syllogistic Reasoning: An Alternative M ental M odels Theory

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 23(23)

ISSN
1069-7977

Authors
Quayle, Jeremy D.
Ball, Linden J.

Publication Date
2001

Peer reviewed

Cognitive Uncertainty in Syllogistic Reasoning: An Alternative Mental Models Theory

Jeremy D. Quayle (j.d.quayle@derby.ac.uk)
Cognitive and Behavioral Sciences Research Group,
Institute of Behavioural Sciences, University of Derby,
Mickleover, Derby, DE3 5GX, UK

Linden J. Ball (l.ball@lancaster.ac.uk)
Psychology Department, Lancaster University,
Lancaster, LA1 4YF, UK

Abstract

In this paper we propose a mental models theory of syllogistic reasoning which does not incorporate a falsification procedure and clearly specifies which conclusions will be generated and in what order of preference. It is assumed the models constructed vary in terms of the number of uncertain representations of end terms, and the directness of the representation of the subjects of valid conclusions. These key factors determine which quantified conclusion will be generated, as well as the varying tendency to respond that "nothing follows". The theory is shown to provide a close fit to meta-analysis data derived from past experiments.

Introduction

The categorical syllogism is a deductive reasoning problem comprising two premises and a conclusion (see example below).

Some artists are beekeepers
All beekeepers are carpenters
Therefore, some artists are carpenters

The premises feature three terms which refer to classes of items or individuals: an end term in each premise, and a middle term which appears in both premises. The formal structure of a syllogism is determined by its mood and its figure. The term mood refers to the different combinations of quantifer that can be featured in the premises and conclusion. Four standard quantifiers are used in English language syllogisms: All, Some, No/None, or Some... are not. The term figure refers to the four possible arrangements of the terms within the premises: A-B, B-C; A-B, C-B; B-A, C-B; and B-A, B-C (where A refers to the end-term in the first premise, C refers to the end-term in the second premise, and B refers to the middle term). As each premise can contain one of four quantifiers, and there are four figures, 64 standard premise pairs are possible. The logically valid conclusion to a syllogism is a statement which describes the relationship between the two end terms in a way that is necessarily true, given that the premises are true.

The principal challenges for a theory of syllogistic reasoning are: (1) to explain how people are able to reach the right conclusion for the right reason (i.e., logical competence), and (2) to explain the systematic variations in difficulty and responding between different forms of syllogism (i.e., performance). Responses to syllogisms vary both in terms of the forms of quantified conclusion generated, and in the tendency to respond that there is no valid conclusion. One theory that provides a good fit for the data, and that has received considerable support and attention in the literature is the mental models theory (e.g., Johnson-Laird & Byrne, 1991).

This theory, which can be said to have its routes in early Euler circles, set-based accounts (e.g., Erickson, 1974, 1978), is one of the most comprehensive theories of syllogistic reasoning competence and performance. It assumes that the reasoning process begins with the construction of a mental model of the premises within working memory that makes explicit the minimum amount of information. From this initial model a parsimonious conclusions is formulated, the validity of which is tested in a search for counter-examples– a process which may involve 'fleshing out' the initial model. If a falsifying model cannot be constructed, then the conclusion is generated (since it must be valid), otherwise it is rejected. Reasoners who experience difficulty whilst reasoning will be inclined either to respond with an incorrect quantified response that is consistent with current models, or to say that there is no valid conclusion.

A central assumption of the mental models theory is that syllogisms vary in terms of the number of mental models it is necessary to construct in order to test putative conclusions. With some syllogisms a single model is sufficient for a valid conclusion to be generated (termed one-model syllogisms), with others it is necessary to construct two or three models (termed multiple-model syllogisms). Multiple-model syllogisms place greater, and less manageable demands on limited working memory resources than one-model syllogisms, and consequently, yield the smallest proportion of valid conclusions and the largest proportion of erroneous no valid conclusion (NVC) responses.

Although good support has been found for mental models accounts of performance (e.g., Johnson-Laird & Bara, 1984), the theory does have some notable weaknesses. First, the theory does not clearly specify how conclusions are formulated, and so cannot adequately explain the clear preference for some conclusions over others. For example, according to Johnson-Laird and Steedman (1978), the theory predicts an average of 3.3 different responses per syllogism. Second, there is increasing evidence to suggest that reasoners construct only a single mental model and do not engage in a falsification process at all (e.g., Polk & Newell, 1995; Bucciarelli & Johnson-Laird, 1999; Newstead, Handley & Buck, 1999).

In spite of these weaknesses, we find many assumptions of the mental models theory to be intuitively plausible, and believe that there is scope for a new Euler-circle inspired mental models account which retains some of these assumptions, specifies the manner in which conclusions are generated more precisely, and incorporates the idea that conclusions are generated after the construction of a single minimal mental representation.

A Cognitive Uncertainty Theory

In accordance with the mental models theory, we propose a theory that assumes people reason with syllogisms by constructing abstract analogical models of the logical relationships between the terms described in the premises within working memory. We would argue that the goal of reasoners when constructing models of syllogisms is to represent mentally both the semantic meaning of each premise and the order of the terms within each premise in the simplest and most probable form. Simple models are those which do not explicitly represent all of the different possible relationships between the end terms. These are constructed for reasons of cognitive economy, since they should not place such high loads on limited working memory capacity as would the consideration of more complex alternative models. Probable models are those which represent the most likely or 'available' situation where a number of alternatives are possible.

Just as the written and spoken forms of terms within premises are read or heard in a particular serial order, so the mental representations constructed of terms within premises are intended to be scanned mentally in a direction which corresponds with these forms of presentation. For ease of explanation we shall refer to this intended scanning direction as 'left to right'.

The Representation of Quantifiers

If a 'universal' quantifier (all or no/none) precedes a term, then the complete class of items or individuals to which the term refers will be represented in the model. For example, in our notation, "All A" would be represented as: [A]

When a premise has the universal, affirmative quantifier all (e.g., "All A are B"), we suggest that reasoners represent the premise as: [A B]

This representation (and the representations constructed for all other forms of premise) is intended to represent both the meaning of the premise and the order of the terms within the premise. It features a description of the class A (the subject of the premise) in terms of the class B. This is the simplest most economical way of representing this premise, since it avoids the need to represent members of the B class who are not also members of the A class. The representation is equivalent to the conditional statement "If it is an A, then it is a B". Just as the statement "All A are B" is intended to be read from left to right (or spoken in a corresponding serial order), so this representation is intended to be scanned from left to right. When scanned from left to right, this representation is unambiguous. However, when scanned from right to left, the representation is ambiguous, suggesting a fallacious identity interpretation of the premise (i.e., ambiguity in the representation leads to the assumption that "All A are B" also means "All B are A").

Twenty eight syllogisms have at least one premise featuring the All quantifier. Eighteen of these are determinate syllogisms. For fourteen of these, the ambiguity of this representation does not affect valid conclusion generation. However, the assumption that participants construct ambiguous representations of All premises can account for all preferred conclusions for the remaining fourteen syllogisms.

The premise "No A are B" (featuring the universal, negative quantifier) would be represented as: [A] [B]

If an 'existential' quantifier precedes a term, then an incomplete class of items or individuals to which the term refers will be represented in the model. For example, "Some A" would be represented as: A) or (A

Grice's (1975) maxim of quantity states that speakers should be as informative as possible, and should not deliberately withhold information which they know to be true. It follows from this notion that it would be wrong for speakers to use the word "some" when they know "all" to be true (see also Begg, 1987; Newstead & Griggs, 1983). We argue, therefore, that although a logician's definition of the quantifier some is "at least one and possibly all", complete classes of items or individuals are not represented when a term is preceded by "some". Hence, we suggest that the premise "Some

A are B" (featuring the existential, affirmative quantifier) would be represented as: [A) B]

With this premise, the possibility that there could be A's that are not B's is not represented in the model, although this may be understood implicitly.

The premise "Some A are not B" (featuring the existential, negative quantifier) would be represented as: (A [B]

In this instance, the possibility that there could be A's that are B's is not represented in the model, but again, this may be understood implicitly.

Model Construction

It is suggested that the construction of models occurs as follows:

1) One of the two premises is picked to be the first premise represented in the model.
2) A model of the first premise is constructed in which the first term in the premise is described in relation to the second term.
3) The model of the first premise is augmented so that it features a representation of the first term in the second premise described in relation to the second term in the second premise. When a universal set is represented in the model of the first premise, and the term referring to that set is preceded by some in the second premise, the universal representation is reduced to an existential representation in the combined model (e.g., syllogisms 3 and 5). As with the first premise, the model is constructed such that the semantic meaning of the quantifier and the serial order of terms in the second premise are retained within the model (see examples below).

1. Some A are B All B are C	2. Some A are B No C are B	3. No A are B Some B are C
[A) B C] Some A are C (valid)	[c] [A) B] No C are A (invalid)	[a] [B) c] No A are C (invalid)
4. Some A are not B All C are B	5. No B are A Some B are C	6. All B are A No B are C
(A [c B] Some A are not C (valid)	[B) c] [a] Some C are not A (valid)	[B A] [C] No A are C (invalid)

Syllogism 1 exemplifies a situation in which it is necessary to represent members of one end-term class who are also members of the other end-term class in order to construct a model in which the semantic meaning of each premise and the order of the terms within each premise is represented. In contrast, syllogisms 2 to 6 exemplify situations in which it is not necessary to represent members of one end-term class who are also members of the other end-term class.

We acknowledge that the model shown for syllogism 5 is not the only possible model that could be constructed. For example, the possible intersection between the A term and C term, or the possible containment of the A term within the C term could be represented in a model. We would suggest, however, that if it is not necessary to represent an overlap in class membership in a model in order to represent the meaning of each premise and the order of the terms, then no overlap will be represented. The model for syllogism 5 shows no overlap between the A term and the C term, and so, the relationship between these terms shown in the model is equivalent to "No C are A" or "No A are C". This may be considered a rational approach to model construction, since the situation "No X are Y" (where X and Y are two properties picked at random) is almost always true (cf. Chater & Oaksford, 1999). Hence, the model shown for syllogism 5 represents the most probable situation out of a number of possible alternatives.

Certainty and Uncertainty within Models

In some instances, the representation of one end term in relation to the other end term in a model is 'certain'. By this we mean that the class represented by one end term (e.g., the A's in syllogism 4) cannot incorporate members of the class represented by the other end term (e.g., the C's in syllogism 4) unless members of one end-term class (e.g., the A's in syllogism 1) are already members of the other end-term class (e.g., the C's in syllogism 1). In other instances, however (e.g., in syllogisms 2, 3 4 and 5), the representation of one end term is 'uncertain' in relation to the other end term. That is, members of one end-term class who are not represented as being members of the other end-term class could possibly be members of that other end-term class – in our notation, lowercase letters denote uncertain representation of a term. For example, in the model for syllogism 2, as the possibility that there could be some A's that are not B's is not explicitly represented, it is possible that some or all of the C's represented could be A's – hence, the C term is represented by a lowercase letter. In the models for syllogisms 3 and 5, the only certain representations are of the B's that are C's. As some or all of the A's that are represented could be C's, and some or all of the C's that are not B's could be A's, both the A's and the C's are shown in lowercase.

Suggested Conclusion Generation

The conclusion that is initially 'suggested' by a model is the one that follows the serial order in which the terms are represented (left-to-right in our notation). The subject of this conclusion will be the end term on the left. The quantifier that is chosen is determined by the representation of that term. If an existential set is represented, then the conclusion will have an existential quantifier (either some or some… are not), otherwise the conclusion will have a universal quantifier (either All or No). If the first end term in the model is represented outside the second end term, then the conclusion will be negative (either No or Some… are not), otherwise it will be affirmative (either All or Some).

Although conclusions that are suggested by a model will be the preferred responses, the initial conclusion suggested by left-to-right scanning is not always the one that is generated. The initial conclusions suggested by left-to-right scanning of the models constructed for the six example syllogisms are shown beneath the models.

With some syllogisms the representation of the subject of the conclusion suggested by left-to-right scanning is uncertain (e.g., the C term in syllogism 2 and the A term in syllogism 3). Uncertainty of this nature in a model should cause reasoners to lack confidence or certainty over the validity of an initial conclusion, and should motivate them to scan the model from right-to-left in search of an alternative conclusion. If the subject of the conclusion suggested by right-to-left scanning has a certain representation, then this conclusion will be favoured over the conclusion initially suggested (when scanned from right to left the model for syllogism 2 suggests the valid conclusion "Some A are not C" and the model for syllogism 3 suggests the valid conclusion "Some C are not A"). Since two stages of conclusion generation are required with these problems, they load cognitive resources more heavily than those with models where the subject of the conclusion suggested by left-to-right scanning is certain (e.g., syllogism 1). Consequently, they should yield more logical errors – in particular, the generation of invalid conclusions that are consistent with left-to-right scanning (which should be the second most common quantified response) – lower feelings of certainty in participants, and high levels of fallacious NVC responding.

With some indeterminate syllogisms, the subjects of conclusions suggested by both left-to-right and right-to-left scanning have uncertain representations (see syllogisms 8 to 12). With these problems we suggest that the conclusion suggested by left-to-right scanning will be generated more frequently than the one suggested by right-to-left scanning. This is because some reasoners will fail to generate a conclusion through right-to-left scanning due to the extra cognitive demand involved, while most reasoners will be able to generate an initial conclusion through left-to-right scanning. These syllogisms should yield lower feelings of certainty in participants and more NVC responses than syllogisms with models containing one or no uncertain representations.

Direct and indirect representation of the subject

With most determinate syllogisms the subject of the valid conclusion is represented 'directly' in the model by members of the class referred to by the subject of this conclusion. For example, the subject of the valid conclusion to syllogism 2 is "Some A", and the representation of this in the model is also "Some A". With some determinate syllogisms, however, the subject of the valid conclusion is represented 'indirectly' in the model by members of the middle-term class who are also members of the class referred to by the subject of this conclusion. For example, the subject of the valid conclusion to syllogism 3 is "Some C", although the representation of this in the model is "Some <u>B</u>". In this instance it is necessary to convert mentally the representation of the subject of the conclusion from "Some B" to "Some C" before a valid conclusion can be generated. As this additional step in the reasoning process is necessary, we suggest that these problems are more difficult than those where the subject is represented directly. Hence, syllogism 3 should yield more logical errors (including a greater proportion of NVC responding) and lower feelings of certainty in participants than syllogism 2.

Implied Conclusion Generation

Not all individuals respond with a conclusion that is suggested directly by a model. A small proportion will respond with a conclusion that is implied by a model. We suggest that implication affects responding in the following way:

- As some and some… are not are given similar interpretations (e.g., see Begg & Harris, 1982), whenever a some conclusion is suggested by a model, a some… are not conclusion will often be generated, and whenever a some… are not conclusion is suggested by a model, a some conclusion will often be generated.
- Conclusions with universal quantifiers (all and no) imply that existential alternatives are also true (some and some… are not respectively).

Categorising Syllogisms

When models are constructed and conclusions generated according to the assumptions we have outlined, five types of determinate syllogism (assigned labels D1 to D5 in Table 1) and four types of indeterminate syllogism (assigned labels I1 to I4 in Table 1) can be identified. Example syllogisms 1 and 4 are classified as D1 problems, syllogism 6 is a D2 problem, syllogism 2 is a D3 problem, syllogism 5 is a D4 problem, and syllogism 3 is a D5 problem.

Examples of indeterminate syllogisms together with models and suggested conclusions are given below.

7.	8.	9.
All A are B	Some B are not A	Some A are not B
Some B are C	No C are B	Some C are B
[A B) C]	[c] (B [a]	(a [c) B]
Some A are C	NVC	NVC
	No C are A	Some A are not C

10.	11.	12.
No B are A	Some B are not A	Some B are A
No B are C	Some B are C	Some C are B
[B] [a] [c]	(B [a] [B) c]	[B) a] [c) B]
[a] [B] [c]	(a [B) c]	[c) B) a]
or	or	NVC
[c] [B] [a]	((cB [a]	Some C are A
NVC	NVC	
No A are C	Some A are not C	
or	or	
No C are A	Some C are not A	

Example syllogism 7 is an I4 problem, syllogisms 8 and 9 are I1 problems, syllogisms 10 and 11 are I3 problems, and syllogism 12 is an I2 problem. With two of these (types I4 and I1), it is possible to construct a model in which the semantic meaning of each premise and the order of terms in the premises are represented with little difficulty. With the remaining two, however, the construction of a 'unified' model is less straightforward, as either: (1) the premises suggest models with a split representation of the middle term (type I2 and some I3), or (2) the premises do not dictate the order in which the end terms should be represented in a model (type I3). With type I2 and I3 problems, reasoners may feel highly confident that there is no necessary relationship between the end terms, and so, respond that there is NVC. However, as reasoners display a bias towards the generation of quantified responses where NVC responses are appropriate (e.g., Johnson-Laird & Steedman, 1978; Johnson-Laird & Bara, 1984), with I2 and I3 syllogisms we have considered strategies reasoners may adopt in order that models might still be constructed and quantified conclusions generated. With I2 syllogisms a quantified conclusion may be generated if reasoners construct a model in which a representation of the end term which forms the subject of a some premise (e.g., "Some <u>C</u> are B" in syllogism 12) is contained within the class referred to by the middle term in the other premise (see second model for syllogism 12). There are good reasons for assuming that models of this nature are constructed: (1) they are able to represent the semantic meaning of each premise and the order of the terms within each premise, and (2) they feature a representation of a highly probable relationship between the end terms.

With I3 syllogisms, a unified model may be constructed after applying a simple conversion procedure to one of the premises (i.e., by switching the two terms in a premise around without changing the quantifier). Once conversion has taken place, it is possible for reasoners to construct models like those constructed for I1 or I2 syllogisms (depending upon mood) and to generate quantified conclusions (see second models for syllogisms 10 and 11).

Table 1: Nine types of syllogism as a function of figure, together with NVC rankings.

Type	AB BC	BA CB	AB CB	BA BC	Total	NVC Rank
D1	4	4	4	4	16	1
D2	1	1	0	3	5	2
D3	0	0	2	0	2	3
D4	0	0	0	2	2	3
D5	1	1	0	0	2	4
I1	5	5	2	0	12	5
I2	2	2	0	0	4	6
I3	0	0	5	7	12	6
I4	3	3	3	0	9	2.5

Rank values have been assigned to each type of syllogism in Table 1 according to the level of NVC responding that should be associated with them according to the theory (a 1 ranking denotes a low level of NVC responding). With the exception of I2 and I3 syllogisms (where high levels of NVC responding should be associated with strong feelings of certainty), these rankings should correlate negatively with the levels of certainty associated with each type of syllogism.

Predicting performance

We have used Chater and Oaksford's (1999) meta-analysis data (on conclusion quantification and NVC responding) together with Johnson-Laird and Steedman (1978, Experiment 1 and 2), and Johnson-Laird and Bara's (1983, Experiment 3) data (on conclusion term order) in order to test how well the cognitive uncertainty theory explains syllogistic reasoning performance.

Quantification

The theory makes ranking predictions for most common conclusions, as well as second, third and fourth most common conclusions. Predicted rankings match data rankings in 175 out of 256 cases. When the proportions of quantified responses associated with correct matches are summed, the theory accounts for nearly 91% of quantified responses in exact order of commonality.

Term Order in Preferred Conclusions

The theory makes a definite prediction about the order of terms in preferred conclusions for 50 of the 64 premise pairs. It matches the term order for preferred conclusions in 47 of these, and therefore, directly predicts the term order for over 73% of preferred quantified responses in the data—although the theory can accommodate over 95% of these.

NVC responding

Rank values have been assigned to each syllogism based on predicted feelings of certainty in participants, and how much NVC responding would be expected according to the theory (see Table 1). There is a significant correlation between these rankings and the percentages of NVC responses in the meta-analysis data (Spearman's Rho = .885, p< .001). The theory, therefore, accounts for 78% of the variance in the NVC data.

Discussion

A new theory of syllogistic reasoning, inspired by early Euler circles, set-based, explanations has been proposed. The theory has been shown to provide a good fit to meta-analysis data derived from past experiments, in terms of: (1) predicting order of preference for quantified responses, (2) predicting term-order in preferred conclusions, and (3) accounting for the varying tendency to give NVC responses.

The new theory was developed in response to an increasing awareness that key assumptions of the mental models account proposed by Johnson-Laird and colleagues were failing to receive support in the literature. In particular, serious doubts have been raised concerning the conclusion falsification and fleshing out processes outlined in the mental models theory. The new account overcomes this problem by suggesting that only a single model is constructed based on an interpretation of the semantic meanings of the premises. Models may be logically accurate, or suggest a fallacious or ambiguous interpretation of All quantifiers. Difficulty is said not to be caused by cognitive loads associated with model construction, but instead by the cognitive loads associated with 'searching' a model in order to identify a conclusion that does not feature an uncertain representation of the subject.

The mental models account may also be criticised for not adequately specifying the manner in which conclusions are formulated, and hence, not explaining why clear orders of preference for different quantified conclusions are evident in the data (e.g., see Chater & Oaksford's meta-analysis of past experiments). The new theory clearly specifies which conclusion will be identified first from a model, whether other conclusions will be identified that are also suggested by a model, and which conclusions will be generated to a lesser degree after being implied by a model. Importantly, clear psychological justifications are given for the predicted orders of preference for these quantified responses.

References

Begg, I. (1987). Some. Canadian Journal of Psychology, 41, 62–73.

Begg, I., & Harris, G. (1982). On the interpretation of syllogisms. Journal of Verbal Learning and Verbal Behaviour, 21, 595–620.

Bucciarelli, M., & Johnson-Laird, P.N. (1998). Strategies in syllogistic reasoning. Cognitive Science, 23, 247–303.

Chater, N., & Oaksford, M. (1999). The probability heuristics model of syllogistic reasoning. Cognitive Psychology, 38, 191–258.

Erickson, J.R. (1974). A set analysis theory of behaviour in formal syllogistic reasoning tasks. In R.L. Solso (Ed.), Theories of cognitive psychology: The Loyola Symposium. Hillsdale, N.J.: Lawrence Erlbaum Associates Inc.

Erickson, J.R. (1978). Research on syllogistic reasoning. In R. Revlin & R.E. Mayer (Eds.), Human reasoning. Washington, DC: Winston.

Grice, P. (1975). Logic and converstion. In P. Cole & J.L. Morgan (Eds.), Studies in syntax, Vol. 3: Speech acts. New York: Academic Press.

Johnson-Laird, P.N., & Bara, B.G. (1984). Syllogistic inference. Cognition, 16, 1–62.

Johnson-Laird, P.N., & Byrne, R.M.J. (1991). Deduction. Hove: Lawrence Erlbaum Associates Ltd.

Johnson-Laird, P.N., & Steedman, M. (1978). The psychology of syllogisms. Cognitive Psychology, 10, 64–99.

Newstead, S.E., & Griggs, R.A. (1983). Drawing inferences from quantified statements: A study of the square of opposition. Journal of Verbal Learning and Verbal Behaviour, 22, 535–546.

Newstead, S.E., Handley, S.J., & Buck, E. (1999). Falsifying mental models: Testing the predictions of theories of syllogistic reasoning. Memory and Cognition, 27, 344–354.

Polk, T.A., & Newell, A. (1995). Deduction as verbal reasoning. Psychological Review, 102, 533–566.