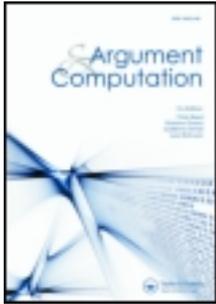


This article was downloaded by: [Princeton University]

On: 21 June 2012, At: 09:22

Publisher: Taylor & Francis

Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



Argument & Computation

Publication details, including instructions for authors and subscription information:

<http://www.tandfonline.com/loi/tarc20>

The processes of inference

Sangeet Khemlani ^a & P. N. Johnson-Laird ^b

^a Navy Center for Applied Research in Artificial Intelligence, Naval Research Lab, Washington, DC, 20375, USA

^b Department of Psychology, Princeton University, Princeton, NJ, 08544, USA

Available online: 21 Jun 2012

To cite this article: Sangeet Khemlani & P. N. Johnson-Laird (2012): The processes of inference, *Argument & Computation*, DOI:10.1080/19462166.2012.674060

To link to this article: <http://dx.doi.org/10.1080/19462166.2012.674060>



PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: <http://www.tandfonline.com/page/terms-and-conditions>

This article may be used for research, teaching, and private study purposes. Any substantial or systematic reproduction, redistribution, reselling, loan, sub-licensing, systematic supply, or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The accuracy of any instructions, formulae, and drug doses should be independently verified with primary sources. The publisher shall not be liable for any loss, actions, claims, proceedings, demand, or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

The processes of inference

Sangeet Khemlani^{a*} and P.N. Johnson-Laird^b

^aNavy Center for Applied Research in Artificial Intelligence, Naval Research Lab, Washington, DC 20375, USA; ^bDepartment of Psychology, Princeton University, Princeton, NJ 08544, USA

(Received 2 December 2011; final version received 20 February 2012)

Mental models represent possibilities, and the theory of mental models postulates three systems of mental processes underlying inference: (0) the construction of an intensional representation of a premise's meaning – a process guided by a parser; (1) the building of an initial mental model from the intension, and the drawing of a conclusion based on heuristics and the model; and (2) on some occasions, the search for alternative models, such as a counterexample in which the conclusion is false. System 0 is linguistic, and it may be autonomous. System 1 is rapid and prone to systematic errors, because it makes no use of a working memory for intermediate results. System 2 has access to working memory, and so it can carry out recursive processes, such as the construction of alternative models. However, it too is fallible when the limited processing capacity of working memory becomes overburdened. The three systems are embodied in a unified computational implementation of the model theory, called *mReasoner*, which is a recent departure in the theory. We review its three systems as they apply to reasoning about the properties of sets of individuals, and we explore how these systems can be extended to other domains of reasoning.

Keywords: mental models < Interdisciplinary links with Computational Argument>; rationality < Interdisciplinary links with Computational Argument>; cognitive science; *mReasoner*; computational modelling; inferences; deduction; reasoning

Individuals with no training in logic can make certain deductions with ease. For instance, if they learn that

Senator Smith is on the Appropriations Committee

Everyone on the Appropriations Committee is Texan

they can infer that Senator Smith is Texan. Indeed, not to make this inference could be a mistake. It is *valid*, because its conclusion holds in every case in which its premises hold, i.e. the conclusion must be true if the premises are true (Jeffrey 1981, p. 1). Yet, despite the simplicity of the deduction, its underlying mental processes are not accessible to introspection, and so they remain a challenge to cognitive science.

Early psychological accounts postulated that reasoning depends on a mental implementation of formal logic (Inhelder and Piaget 1958, p. 305). Individuals were supposed to extract the logical form of premises to derive a conclusion using formal rules of inference in a process akin to a logical proof, and then to restore the appropriate contents to the form of the conclusion (Johnson-Laird 1975; Osherson 1975; Braine 1978; Rips 1983). But, these theories had difficulty in explaining a robust empirical result: the contents of premises can affect inferences in ways that logic alone cannot predict (Wason and Shapiro 1971). The discovery of these effects led to the theory of mental models, which takes meaning rather than logical form to be central to reasoning. In this article, we focus on its account of deductive reasoning, which postulates that reasoners use the contents of

*Corresponding author. Email: skhemlani@gmail.com

premises to simulate the world under description. In the past, the theory has been applied piecemeal to various sorts of reasoning. We describe how we have begun to unify it in a single computational implementation called *mReasoner* (for “model-based Reasoner”). To illustrate how it works, we focus on monadic assertions, which are about the properties of sets of individuals. However, to set the scene for this new unification, we begin with the general principles of the theory.

The model theory of reasoning

The mental model theory – the “model” theory, for short – postulates that when individuals understand discourse, they construct a simulation of the possibilities consistent with what the discourse describes (Johnson-Laird 2006). The theory accordingly depends on three main principles:

- (1) Individuals use a representation of the meaning of a premise, an *intension*, and their knowledge to construct mental models of the various possibilities to which the premises refer.
- (2) The structure of a model corresponds to the structure of what it represents, i.e. the model is *iconic* as far as possible (Peirce 1931–1958, Vol. 4). In some cases, such as the representation of negation (Khemlani, Orenes, and Johnson-Laird in press), models have to include abstract symbols.
- (3) To minimise the load on working memory, mental models represent what is true and not what is false. Only *fully explicit* models represent what is false, and reasoners should have difficulty constructing such models.

On this account, a conclusion is *probable* if it holds in most of the models of the premises, *necessary* if it holds in all of the models of the premises, and *possible* if it holds in one or more models of the premises. A conclusion is *consistent* with the premises if it holds in at least one model of the premises. Also, it is invalid if it fails to hold in at least one model of the premises, namely, a model that is a *counterexample* to the conclusion. The model theory is an alternative to theories based on mental rules of inference or probabilistic calculations. Table 1 summarises many of the key theoretical differences.

The model theory predicts several phenomena that have been corroborated in experiments. Inferences are faster and more accurate when they depend on only a single model than when they depend on several models (Bauer and Johnson-Laird 1993; Evans, Handley, Harper, and Johnson-Laird 1999). Semantics and general knowledge can block the construction of models of possibilities and add various sorts of relation among the entities in a model (Johnson-Laird and Byrne 2002; Quelhas, Johnson-Laird, and Juhas 2010). Individuals use counterexamples to refute invalid conclusions, especially if these conclusions are consistent with the premises but do not follow from them (Bucciarelli and Johnson-Laird 1999; Johnson-Laird and Hasson 2003). And reasoners err predictably when an inference requires them to consider what is false (Johnson-Laird and Savary 1999; Khemlani and Johnson-Laird 2009; Kunze, Khemlani, Lotstein, and Johnson-Laird 2010).

The model theory applies to reasoning of many sorts, including inferences based on quantifiers such as *all artists* and *some bohemians* (Johnson-Laird 1983; Khemlani, Lotstein, and Johnson-Laird under review a, under review b), and sentential connectives such as *and*, *or*, and *if* (Johnson-Laird and Byrne 1991). Table 2 summarises the domains to which it applies. Each of these extensions calls for novel assumptions, and many of them have been implemented in separate computer programs (for a review see Johnson-Laird and Yang 2008). For instance, the extension to temporal reasoning postulates a separate iconic “time line” to represent temporal assertions, such as, “A happens *before* B” and “A happens *during* B”, and it has been implemented in a

Table 1. Distinguishing characteristics of psychological theories of reasoning based on mental rules, probabilities, and mental models.

Distinguishing characteristic	Psychological theories of reasoning		
	Mental rules	Probability heuristics model	Mental models
Domains to which the theory has been applied.	All of first-order logic.	Monadic assertions including quantifiers, such as “most”, which are outside first-order logic; and reasoning based on conditionals (<i>if-then</i>).	See Table 2
Input to the proposed mechanism	The logical forms of premises	Sentences	Sentences
Linguistic analysis	Not applicable	Not specified, because the theory is at the computational level, not the algorithmic level	Parser recovers a representation of their meaning: an intensional representation, which is used in the creation and modification of mental models.
Heuristics	None	“Fast and frugal” heuristics approximate rational, context dependent probability calculations, e.g. Bayes’s theorem	Heuristics derived from patterns of valid deductions use intensions to yield putative conclusions and check them in initial models.
Deliberations in inference	Formal rules of inference, such as modus ponens, are applied to the logical forms of premises.	None	Search for alternative models that may serve as counterexamples or validate initial conclusions.
Effect of context and content	They can affect logical form, but the theory applies only thereafter.	They establish prior probabilities.	They can modulate the interpretation of quantifiers, connectives, and other terms, blocking the construction of models and adding relations to models.
Output	A proper subset of the valid conclusions in first-order logic, because there are valid inferences that the theory cannot capture.	Probabilistic conclusions.	Mental models, psychologically plausible conclusions, and valid conclusions that are beyond the competence of naïve reasoners.

computer program (Schaeken, Johnson-Laird, and d’Ydewalle 1996). With each new extension, the danger is that the theory splits into separate fragments. Each fragmentary theory may give a satisfactory account of its domain, but the fragments may no longer fit together or rest on common background assumptions. This problem is hardly unique to the model theory, and part of the appeal of unitary architectures is that they obviate it (Newell 1990; Anderson 1993). In our view, the time has come for a unification of accounts of reasoning: one theory, one architecture, and one computational model.

As a consequence, we have begun to unify the model theory and to implement it in mReasoner, which we describe in the following sections of the article. The theory postulates an architecture in which there are three main systems. System 0, as we refer to it, is linguistic. It parses each premise in order to create an *intensional* representation of its meaning. System 1 uses the intension to build an *extensional* representation, i.e. a mental model of a possibility, and its heuristics rely on this

Table 2. The mental model theory and its extensions.

Domain of inference	Example premises and/or queries	Source(s)	Implemented in mReasoner version 0.8?
Syllogistic reasoning	All As are Bs	Bucciarelli and Johnson-Laird (1999)	Yes
	Some As are Bs		
	Some As are not Bs	Khemlani et al. (under review a)	
	No A is a B		
Reasoning about consistency	Can both A and B be true at the same time?	Johnson-Laird, Legrenzi, Girotto, and Legrenzi (2000) Johnson-Laird, Girotto, and Legrenzi (2004)	Yes
Set membership inferences	A is a B.	Khemlani et al. (under review c)	Yes
	A is not a B.		
The interpretation of negation	It is not the case that ...	Khemlani et al. (in press)	Yes
Sentential reasoning	A and B	Johnson-Laird, Byrne, and Schaeken (1992)	In progress
	A or B or both		
	A or else B	Johnson-Laird and Byrne (2002)	
	If A then B		
	If and only if A then B		
Modal reasoning	A is possible	Bell and Johnson-Laird (1998)	In progress
	A is not possible		
	A is necessary	Goldvarg and Johnson-Laird (2000)	
	A is not necessary		
Multiple quantification	Some As are not in the same place as all Bs	Johnson-Laird, Byrne, and Tabossi (1989)	In progress
Numerical quantifiers and quantifiers outside first-order logic	More than three of the A are B	Kroger, Nystrom, Cohen, and Johnson-Laird (2008)	In progress
	More than half the A are B	Neth and Johnson-Laird (1999)	
Extensional probabilistic reasoning	The probability of A is ...	Johnson-Laird (1994)	No
	A is more likely than B.		
	What is the probability that A?	Johnson-Laird, Legrenzi, Girotto, Legrenzi, and Caverni (1999)	
	Which is more likely, A or B?		
Spatial reasoning	A is on the right of B	Byrne and Johnson-Laird (1989)	No
	A is in front of B	Jahn et al. (2007) Mackiewicz and Johnson-Laird (2012)	
Temporal reasoning	A happens before B	Schaeken et al. (1996)	No
	A happens while C		
	A happens after B	Juhos et al. (2012)	
	A happens during B		
Causal reasoning	A will cause B	Goldvarg and Johnson-Laird (2001)	No
	A causes B		
	A caused B		
	A prevents B	Frosch and Johnson-Laird (2011)	
	A allows B		
	A and only A will cause B		
Deontic reasoning	A permits B	Bucciarelli and Johnson-Laird (2005)	No
	A obligates B		
	A prohibits B		
	A permits not B		
Relational reasoning	A is taller than B.	Goodwin and Johnson-Laird (2005)	No
	A is taller than B to a greater extent than C is taller than D.	Goodwin and Johnson-Laird (2006)	
Counterfactual reasoning	If A had not occurred then B would not have occurred.	Byrne and Tasso (1999) Byrne (2005)	No

model and the intension to draw a rapid initial conclusion. System 2 carries out more powerful processes, and it searches for alternative models, including counterexamples in which the initial conclusion fails to hold. This system can evaluate, supplement, and even correct initial inferences. The search for alternatives uses various operations, which manipulate the initial model by adding, rearranging, or removing properties of individuals. The search can also flesh out initial models and add relations among the entities in the models in order to yield additional inferences. The two systems often work in concert. For example, system 1 can operate implicitly to determine the tense of a spontaneous conclusion while system 2 operates more explicitly to determine its contents (Juhos, Quelhas, and Johnson-Laird 2012).

The linguistic processes in the first stage may be autonomous – we make no strong claims about the matter. The other two stages, however, correspond to the familiar distinction between the rapid intuitions of system 1 and the slower deliberations of system 2 in dual-process theories of cognition (Johnson-Laird 1983, Chap. 6; Sloman 1996; Stanovich 1999; see, e.g. Evans 2003, 2007, 2008; Verschueren, Schaeken, and d’Ydewalle 2005; Kahneman 2011). However, the model theory goes beyond dual-process accounts of reasoning, because it is the first to be implemented computationally, and it replaces descriptive labels such as “associative” and “deliberative” with an architectural distinction in computational power. It makes the strong assumption that system 1 has no access to working memory for intermediate results, and so the high-level computations it can carry out are computationally constrained to those that can be performed by a finite-state automaton. It can, therefore, work only with a single model at a time, and is unable to carry out any sort of recursive processes, such as counting beyond a small finite number. In contrast, system 2 has access to working memory, and is therefore more powerful computationally: it can search for alternative models, and it can count and carry out other arithmetical operations, at least until they overload its processing capacity. Because almost all reasoning is computationally intractable, no finite system can cope as problems increase in complexity, e.g. with the addition of more premises.

We now turn to a description of how mReasoner works for inferences based on monadic assertions, which are about the properties of sets of individuals. For instance, the assertion, “Everyone on the Appropriations Committee is Texan” is monadic, because it assigns the property, *Texan*, to all members of a set of individuals. The noun phrases in monadic assertions typically refer to sets of individuals, usually as a result of a determiner, such as “all”, “most”, or “some”, which in combination with a noun or nominal yield a quantifier, such as “all artists,” “most of the bohemians”, and “some of the cadgers”.

mReasoner: A unified computational model of reasoning

mReasoner is a computational implementation of the model theory of reasoning. Its architecture is based on three main systems, which, as we mentioned earlier, construct intensional representations (system 0), build an initial model and use heuristics to formulate a putative conclusion (system 1), and search for alternative models of the intensions (system 2). We outline each of these three systems in turn.

System 0 processes: parsing premises to compose intensional representations

Models are built by consulting the representation of the meaning of each premise, i.e. an *intensional* representation, which is composed out of the meanings of words and the grammatical relations among them. Accordingly, the first process in mReasoner is a shift-and-reduce parse that makes use of a context-free grammar and a lexicon (Hopcroft and Ullman 1979). It uses the meanings of words in the lexicon to compose an intensional representation that depends on the grammatical relations among the words. The lexical entries consist of a word (such as “all”), its part of speech

(“determiner”), and a specification of its semantics. The grammatical rules specify how a string of syntactic constituents can be reduced to a higher order grammatical constituent, culminating in a well-formed sentence, and each grammatical rule is paired with an appropriate semantic rule. The parser applies the matching semantic rule as it uses a grammatical rule to reduce a string of words to a single constituent, such as a “noun phrase”. The uses of a standard parser, a context-free grammar, and a rule-by-rule compositional semantics, are neither novel nor empirical claims of the theory. The point instead is to illustrate how the meanings of assertions containing determiners, both orthodox ones, such as “all” and “some”, and unorthodox ones, e.g. “most” and “few”, can be captured using the values of parameters. Orthodox determiners are those that can be represented in the first-order predicate calculus, which is the standard version of logic in which variables range over individuals. The unorthodox determiners cannot be captured in this calculus, but call for one in which variables range over predicates, i.e. sets of individuals. The model theory accordingly postulates that quantified assertions express relations between sets (for a summary of this view, see Cohen and Nagel 1934, pp. 124–125; Khemlani and Johnson-Laird in press). Other psychological theories likewise treat quantifiers as relations between sets; some of them make use of diagrammatic representations to handle relations (Ceraso and Provitera 1971; Erickson 1974; Ford 1995), and others rely on formal rules of inference (Stenning and Yule 1997; Guyote and Sternberg 1981; Geurts 2003; Politzer, van der Henst, Luche, and Noveck 2006). What distinguishes the model theory is that it relies on models of individual entities and properties to represent sets.

The program captures the meanings of quantifiers in the values of six parameters. They constrain various operations, including building models, which we describe in the next section of the article. As an illustration of the parameters, consider the assertion:

Some artists are bohemians.

The first parameter in the intension is the cardinality of the entities in a model representing the set in the initial noun phrase, e.g. the number of tokens for the set of artists in the example above. This value is set by default to 4 in the lexical entry of the quantifier, and so the initial model contains four such tokens. The system generates the same predictions for syllogistic reasoning regardless of whether the default value is 3, 4, or more, but future studies may allow researchers to determine the parameter empirically. The value is mutable, and it can be changed at a later stage of processing (see the account below of searching for counterexamples). This parameter also includes the boundary conditions on the cardinality, e.g. it must be greater than or equal to 1 in the case of “some”. Unlike Aristotle, modern logic treats universally quantified assertions, such as “all artists are bohemians”, as making no claims about the existence of artists. Likewise, in daily life, an assertion such as “all trespassers are prosecuted”, can be true even if there are none. At present, mReasoner finesses this problem, but in principle it can be handled in the parameters.

The second parameter in the intension is the cardinality of the set referred to by the quantified phrase as a whole, such as “some artists”, which for this determiner has a default setting of 2, which is less than the value of the first parameter. However, in the case of “all artists”, the second parameter is the same as the cardinality of “artists,” i.e. the default value of the first parameter. Of course, the second parameter changes if a change is made to the first parameter. The third parameter states the constraints on the relation between the two cardinalities, e.g. “some artists” is represented with fewer tokens than the set of artists as a whole, but the number must be greater than zero in order to capture the existential force of the determiner. The fourth parameter states the polarity of the determiner, that is, whether it is affirmative or negative. The fifth parameter states whether the determiner is universal (e.g. “all,” “no”) or existential (e.g. “some,” “most”), which affects the strategies used to search for counterexamples. Also, the sixth parameter states the relation between the sets referred to in the subject and in the predicate of the assertion. In the

case of monadic assertions, the relation is usually set-inclusion or its negation, e.g. a subset of artists is included, or not included, in the set of bohemians. However, other relations, such as set membership do also occur, e.g. “artists are of varying abilities”, which means that the set of artists is a member of the set of those sets of individuals who vary in ability from one to the other.

The set of parameters may seem complicated, but readers should bear in mind that the parameters can capture the meaning of other sorts of determiner such as: “most”, “many”, “at least three”, “more than half”, which include determiners that cannot be expressed in the first-order predicate calculus. The present set of parameters is illustrated in Table 3 for the assertions that occur in Aristotelian syllogisms: “All As are Bs”, “Some As are Bs”, “No As are Bs”, “Some As are not Bs”, and three representative examples of assertions that occur outside syllogisms, “Most As are Bs”, “Neither A is a B”, and “Exactly five of the As are Bs”. The set of parameters is incomplete, because an extension of the theory to deal with quantified relations, such as, “All philosophers have read some books”, calls for further parameters to represent the respective scopes of the quantifiers. The likely interpretation of the preceding example contrasts in scope with an assertion in the passive voice: “Some books have been read by all philosophers” (see, e.g. Johnson-Laird and Byrne 1991). Unlike the first assertion, the second implies that philosophers have read the same books.

In summary, intensions are collections of parameters that, together with the semantic content of the open-class words in a sentence, such as “artists” and “bohemians”, capture the meaning of the sentence. They provide the data needed to build and to modify models. The order of the parameters has no bearing on how the intension is used in these processes. The compositionality of intensions follows the tradition of formal semantics, as does the assignment of truth values to assertions if their intensions can be mapped into independent models of the world (Partee 1996). From this perspective, a mental model captures what is common to a set of possibilities (Barwise 1993). But, the relation between sentences and their intensions is also compatible with cognitive and constructionist approaches to grammar (see e.g. Goldberg 2003; Langacker 2008). The system can be extended to deal with subtle distinctions in meaning among determiners, such as, “all”, “each”, “every”, and “any” (Langacker 2008, p. 292). Likewise, it can be extended to deal with numerical, proportional, and scalar quantifiers. However, our focus in the next section is on the monadic assertions that occur in Aristotelian syllogisms. These inferences are from two premises, in which each has one of the four initial forms in Table 3.

System 1 processes: the construction and interpretation of initial models

System 1 uses the intension of the first premise in a syllogism to build an initial model, and it updates this model given the subsequent premise. As an illustration, consider the assertion:

Some artists are bohemians.

The intension’s parameter values are shown for “Some As are Bs” in Table 3. The first parameter stipulates that the number of artists (As) by default in the initial model is four:

Artist
Artist
Artist
Artist

In this diagram, each row represents an individual as an artist, and so the model has a small finite number of mental tokens. Of course, mental models represent real individuals, not words, which we use here for convenience. The inferential system in mReasoner is able to treat the model above as representing “all artists” and not, say, “exactly four artists”, because it has access to the intension of the premise, which constrains the possible interpretations of models and their possible modifications.

Table 3. A summary of mReasoner's six parameters in the intensions of six sorts of monadic assertions, where '?' signifies a default value that can be modified provided it meets the boundary conditions.

Assertion	The six parameters in a monadic intension					
	i. Cardinality of overall set of As and its boundary conditions	ii. Cardinality of set referred to by the quantifier	iii. Constraints on ii.	iv. Polarity of the determiner	v. Universal quantifier	vi. Set-theoretic relation between subject and predicate
All As are Bs	?4 ≥ 1	?4	= cardinality in i.	Positive	True	Include
Some As are Bs	?4 ≥ 1	?2	≤ cardinality > 0	Positive	False	Include
No As are Bs	?4 ≥ 1	?4	= cardinality	Negative	True	Include
Some As are not Bs	?4 ≥ 1	?2	≤ cardinality > 0	Positive	False	Not-include
Most As are not Bs	?4 ≥ 2	?3	< cardinality cardinality > 1/2	Positive	False	Not-include
Neither A is a B	2	2	= cardinality	Negative	True	Include
Exactly five As are Bs	5	5	= cardinality	Positive	True	Include

Note: When the polarity of the determiner of an intension (parameter iv.) is negative, it is treated as equivalent to the set-theoretic relation of exclusion (i.e. a value of parameter vi. set to "not-include").

The second parameter specifies that the number of artists who are also bohemians is two by default. Because the fourth parameter states that the polarity of the determiner is positive, two of the artists are updated as bohemians. (Had the parameter been negative, the artists would have been updated with the property of not being bohemians.) The model is accordingly updated to:

Artist Bohemian
 Artist Bohemian
 Artist
 Artist

The boundary constraints in the third parameter are satisfied at this point, i.e. the number of artists who are bohemians is smaller than the number of artists and it is greater than zero.

The fifth parameter states that the assertion is not universal, and so it is possible that there are artists that are not bohemians, and bohemians who are not artists. The model is updated with an extra individual bohemian who is not an artist:

Artist Bohemian
 Artist Bohemian
 Artist
 Artist

Bohemian

mReasoner similarly constructs initial models for other sorts of monadic assertions.

System 1 also accommodates a second premise, such as:

All bohemians are cadgers.

Its intension leads to the updating of the model to:

Artist Bohemian Cadger
 Artist Bohemian Cadger
 Artist
 Artist

Bohemian Cadger

Once system 1 has an initial model of this sort, it can draw a conclusion establishing a new set-theoretic relation, that is, a relation that is not asserted in the premises. Researchers often place heuristics at the forefront of theories of reasoning (see also Ford 1995; Stenning and Yule 1997; Chater and Oaksford 1999; Politzer et al. 2006), but until now proponents of the model theory have downplayed their use. In an effort to bridge the two approaches, mReasoner embodies heuristics in its system 1 processes. However, the system abides by the constraint that any conclusion that heuristics generate must hold in the initial model. We have described the specific heuristics in detail elsewhere (Khemlani et al. under review a), and so here we only outline their general principles. For the inference above, heuristics need to deliver both the quantifier in the conclusion (*its mood*) and the order of the terms that occur in it: “artists” and “cadgers” (*its figure*). Previous heuristics, such as the atmosphere effect (Revlis 1975), have been based on superficial aspects of sentences, such as the determiners that occur in them, or on the informativeness of premises with a view to yielding probabilistic conclusions (Chater and Oaksford 1999). The heuristics in system 1 depend on a very different idea: individuals use their knowledge of the meaning of premises, i.e. their intensions, to guide them to initial conclusions. If a negative premise occurs in a syllogism, any valid conclusion is bound to be negative too. If a premise containing the determiner “some” occurs in a syllogism, any valid conclusion is bound to include it too. So, given that individuals are sensitive to the nature of potentially valid conclusions, they should have acquired the knowledge that as soon as a premise contains a negation, or an existential determiner, the conclusion must be in a negative and existential mood. The premises in our example above are:

Some artists are bohemians.

All bohemians are cadgers.

They accordingly yield a conclusion with the existential determiner, “some”. It follows that only if both premises are based on “all”, can a valid conclusion use this determiner. In general, when two premises differ in mood, one of them dominates the other in determining the mood and figure of the initial conclusion according to the following rank order:

- Some_are not_
- None_are_
- Some_are_
- All_are_

This order is based on a generalisation from valid conclusions to syllogisms, and so it differs from the heuristics invoked in the atmosphere effect, and in other analogous theories. However, at first sight, it parallels the probabilistic heuristics of Chater and Oaksford (1999), which are based on informativeness. Where the two theories diverge is that Chater and Oaksford add a further assumption that individuals avoid drawing conclusions in the mood, Some _ are not _, because they are so uninformative.

The premise in the dominant mood also determines the order of the terms in the conclusion (its figure). Given the preceding premises, the two end terms, which are those that occur in only one premise, are “artists” and “cadgers”. The first premise is the dominant one, and so it determines the figure of the conclusion. System 1 uses the grammatical role of the end term in the dominant premise to play the same role the conclusion, and so it yields the initial conclusion:

Some artists are cadgers.

This conclusion holds in the initial model of the premises, and so it is the output from system 1. Analogous principles apply to other sorts of premises. They account for the well-known figural effect that occurs in syllogistic reasoning, e.g. the tendency to infer the conclusion above rather than its converse, “Some cadgers are artists”. Indeed, conclusions in the predicted figure occurred 82% of the time in six different experiments (see the meta-analysis in Khemlani and Johnson-Laird in press). A similar heuristic determining the figure of conclusions is due to Chater and Oaksford (1999), i.e. the “attachment” heuristic according to which if the least informative premise has an end-term as its subject, it is also the subject of the conclusion; otherwise, the end-term in the other premise is the subject of the conclusion. Once again, however, the theories diverge, because of Chater and Oaksford’s assumption that individuals avoid inferring conclusions of the form, Some _ are not _.

The heuristics in mReasoner rely on the intensions of the premises and the initial mental model. They operate without storing any information in working memory, and so they are rapid. However, the heuristics are fallible. Consider the following inference:

Some artists are bohemians.
 Some bohemians are cadgers.
 The premises yield the initial model:
 Artist Bohemian Cadger
 Artist Bohemian Cadger
 Artist
 Artist
 Bohemian
 Cadger

The heuristics yield the initial conclusion:

Some artists are cadgers.

It holds in the initial model, and so mReasoner yields it as a putative conclusion. In fact, it corresponds to the most frequent conclusion that reasoners tend to draw from premises of this sort (i.e. it occurred on 61% of inferences in the meta-analysis of Khemlani and Johnson-Laird in press). But, it is invalid, and most of the other responses in the meta-analysis were that nothing

Table 4. The three operations that system 2 of mReasoner uses to search for counterexamples to monadic conclusions, their description, an example in which the operation is invoked, a putative conclusion that holds in the initial model, the initial model itself, and the model that the operation yields as a counterexample to the conclusion.

Operation	Description	Example	Putative conclusion	Models	
				Initial model	Modified model
Adding	An individual is added to the set of models.	All Bs are As. All Bs are Cs.	All As are Cs.	B A C	B A C A
Breaking	An individual with multiple properties is broken into two separate individuals.	Some As are Bs. All Cs are Bs.	Some As are Cs. Some Cs are As.	A B C A B C	A B B C A B C
Moving	A property is moved from one individual to another.	No As are Bs. No Bs are Cs.	No As are Cs. No Cs are As.	A \neg B B \neg C C	A \neg B C B \neg C

Note: We list only the different sorts of individuals in each model, and “ \neg ” denotes negation.

follows from the premises, which is correct because no definite conclusion interrelating the end terms does follow from the premises. How can a theory explain that reasoners make both valid and invalid inferences? System 2, to which we now turn, provides the answer.

System 2 processes: the search for counterexamples

In the preceding section, we focused on how mReasoner embodies system 1 and uses heuristics to draw conclusions that are true in the initial model of the premises. However, these conclusions are often not true in other models of the premises. For that reason, the program embodies this second system, which makes a recursive search for alternative models that might falsify a heuristic conclusion. When it finds a counterexample, it also formulates a new conclusion if one is possible, or else declares that no definite conclusion follows about the relation between the end terms. It searches for counterexamples using three operations: *adding*, *breaking*, and *moving* properties in a model. These operations were embodied in an earlier program that dealt solely with syllogisms, and subsequent research in which participants manipulated external models of premises showed that they used these three operations too (Bucciarelli and Johnson-Laird 1999). We describe and illustrate each of the operations in Table 4.

Reasoners may adopt additional strategies for searching for counterexamples. Indeed, when they reason from premises that concern spatial relations (e.g. *in front of*, *to the left of*, *above*) individuals often make minimal, systematic changes to their initial models by “chunking” multiple entities within a model and operating upon those entities as though they were a single unit (Jahn, Knauff, and Johnson-Laird 2007). At present, system 2 implements only the search operations for models based on monadic assertions, but in principle, it can support other sorts of operation. The difficulty is to find empirical methods that reveal the nature of the operations underlying a search for counterexamples.

When system 2 succeeds in finding a counterexample to a conclusion, it attempts to formulate a weaker conclusion in the same figure by adjusting the parameters in the intension of the conclusion (Table 3). For instance, if it finds a counterexample to the conclusion, “All artists are cadgers”, it reduces the value of the parameter specifying the default number of artists who are cadgers from 4 to 3. The result is an intension for the assertion, “Some of the artists are cadgers”. The weaker

conclusion is then checked against both the initial model and the counterexample. If it does not hold in both these models, system 2 weakens it still further. Also, if it is ultimately weakened until it expresses no information, the program responds that no valid conclusion exists. In this way, mReasoner accounts for how individuals infer that no valid conclusion follows from some premises, which many theories of monadic reasoning cannot do. In contrast, if the search fails to find a counterexample, the system asserts that the conclusion is valid, i.e. it holds in all models of the premises.

mReasoner predicts that certain valid inferences should be more difficult than others, and it even predicts that certain valid inferences are beyond the ability of logically naïve individuals. Consider, for instance, the following problem:

No atheists are believers.

All believers are credulous.

What follows?

System 1 predicts the conclusion that no atheists are credulous. System 2 finds a counterexample and weakens this conclusion to some atheists are not credulous. However, this conclusion is invalid, and the valid conclusion – Some credulous individuals are not atheists – cannot be obtained by any further weakening. Hence, mReasoner predicts that naïve reasoners should fail to draw the valid conclusion reliably. Indeed, across six studies, individuals drew the valid conclusion only 3% of the time, which was not reliably different from 0% (Khemlani and Johnson-Laird in press). mReasoner infers this conclusion by forming the converse of the heuristic conclusion, i.e.: “No credulous individuals are atheists”, and then searching for a counterexample. It finds such a model in which all the atheists are credulous and all the believers are credulous too, and these latter credulous individuals are not atheists. No further weakening is possible, and so the valid conclusion is indeed that some credulous individuals are not atheists. mReasoner marks any conclusion that makes use of this conversion operation as beyond the capacity of naïve reasoners.

We have assessed the theory as it applies to syllogistic reasoning. In our meta-analysis (Khemlani and Johnson-Laird in press), we examined seven extant theories: the atmosphere hypothesis (Begg and Denny 1969), an analogous hypothesis in which reasoners are supposed to draw conclusions matching the mood of a premise (Wetherick and Gilhooly 1990), the hypothesis that reasoners make illicit conversions of premises (Revlis 1975), the probability heuristics model (Chater and Oaksford 1999), a theory based on rules of inference (Rips 1994), a program implementing an earlier mental model theory (Johnson-Laird and Byrne 1991), and another model-based program in which verbal formulations are central and no search for alternative models occurs (Polk and Newell 1995). Other theories of syllogisms exist, but their proponents did not consider them complete enough to be entered into the meta-analysis, and the predictions of still another theory have never been published (Khemlani and Johnson-Laird in press). In the meta-analysis, we compared the predictions of each of the seven theories to the conclusions that the participants had drawn in six studies of syllogistic reasoning. The accuracy of a theory depends on the extent to which its predicted responses occur in the data, and the extent to which the responses that it does not predict do not occur in the data. Figure 1 combines these two measures into a single measure of “prediction accuracy”. We have recently examined the performance of mReasoner, and, as Figure 1 shows, it outperforms in accuracy all seven of the theories in our previous meta-analysis.

Other inferential tasks in mReasoner

We have described how mReasoner makes inferences from monadic premises. However, the general procedure of a rapid heuristic inference followed by an attempted falsification is general. It applies to many sorts of valid inference in many sorts of reasoning. Of course, there are limits on the kinds of heuristics the system can implement. At present, the system does not use heuristics

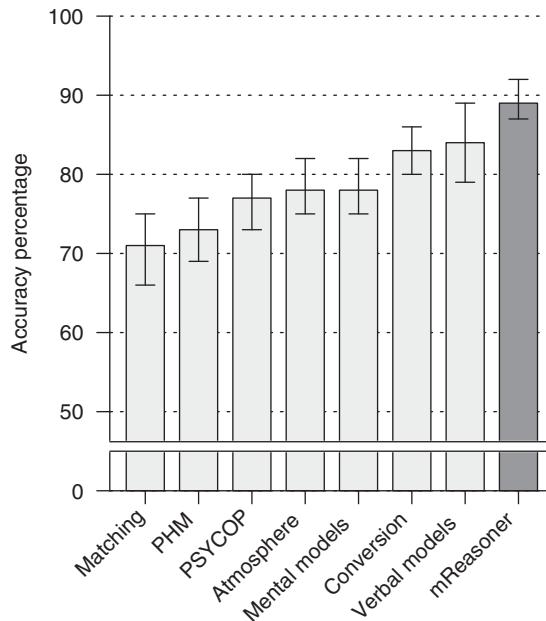


Figure 1. The percentages of accurate predictions of seven different theories of syllogistic reasoning for six experiments (in a meta-analysis in Khemlani and Johnson-Laird in press) to which we have added the predictions of mReasoner. The theories in the original meta-analysis are the matching hypothesis, the probability heuristics model: probability heuristics model PSYCOP, the atmosphere hypothesis, the original mental models theory, the illicit conversion model, and verbal models theory (see the text for references to these theories).

for models of three or more premises. The constraint reflects the intuition that individuals use heuristics only for limited sets of premises, and tend to be lost about what to conclude from a large set of premises.

Psychological experiments on reasoning typically call for the participants to draw a valid conclusion from premises, or else to evaluate the validity of a given conclusion. Sometimes the conclusion is about what is necessarily the case, and sometimes it is about what is possibly the case, though this latter task occurs less often, and some theories of reasoning offer no account of how reasoners carry it out, e.g. Rips's (1994) formal rule theory. It is straightforward in mReasoner: a conclusion about what is possible is valid if there is a model of the premises in which it holds. In daily life, many other sorts of reasoning occur. Individuals may need to infer a likely conclusion, to create an explanation that resolves an inconsistency, and even to detect the inconsistency in the first place. mReasoner is already able to carry out many of these tasks. They include:

- (1) The evaluation of a stated conclusion to determine whether, given the premises, it is *necessarily* the case.
- (2) The similar task of assessing whether a stated conclusion is *possibly* the case.
- (3) The spontaneous formulation of such conclusions from premises.
- (4) The assessment of whether or not a set of assertions is consistent, i.e. whether they could all be true at the same time.
- (5) Given a putative but invalid inference, the formulation of a counterexample that refutes it.

Each of these tasks depends on a high-level procedure, and each of them is carried out using dual-processes: system 1 builds initial models and, where necessary, formulates heuristic conclusions;

and system 2 searches for alternative models and, where necessary, reformulates or rejects conclusions. In this way, mReasoner makes predictions about various inferential tasks at both naïve and expert levels of performance. It provides the beginnings of a unified computational account of reasoning.

General discussion

Inferences come in many flavors: modal inferences about what's necessary or possible, spatial inferences about relations among entities, causal inferences about agents and enablers, and many other inferences besides. In the past, many theories of reasoning concerned only a particular domain or a particular task. Perhaps this compartmentalisation was necessary for researchers to begin to investigate reasoning. It was also successful in that it produced three powerful frameworks for theories of reasoning: formal rules of inference, the probability calculus, and mental models. So far, however, none of them gives a unified account of reasoning covering all domains and all inferential tasks. The mind does not compartmentalise reasoning: conclusions often depend on combinations of different sorts of inference – modal, spatial, causal, temporal – and reasoners reach them without difficulty. Consider the following problem:

When Samantha stands to the right of Fred, she makes him nervous.

Samantha is standing next to Fred.

Is it possible that he's nervous?

Is it necessary that he's nervous?

It is trivial to infer that it's possible but not necessary that Fred is nervous. The modal conclusion depends on a spatial relation causing a particular state of mind. However, the ease of the inference contrasts with the difficulty of accounting for it any current computational model of reasoning. The problem is easy for reasoners, but hard for cognitive scientists. It demands a unified system of reasoning.

This article has described both an architectural organisation that unifies the mental model theory and its computational implementation in mReasoner. Both rest upon three core systems:

- (0) The production of an intensional representation of the meaning of a premise under the control of a parser.
- (1) The construction of an initial model and the use of heuristics to derive an intuitive response or conclusion.
- (2) The search for alternative models, which may invalidate a conclusion, or, say, show that a set of assertions is consistent.

For inferences from monadic assertions, such as “some artists are bohemians”, system 0 constructs intensions that include parameters that capture set-theoretic relations, system 1 draws heuristic conclusions based on intensions and their corresponding initial models, and system 2 searches for alternative models, which may either validate the initial conclusion or lead to its weakening or rejection. The system is unified in its computational structure by incorporating both rapid heuristics and more deliberate, memory intensive, searches for models. Also, the theory is unified by building upon the successes of heuristic-based accounts (Chater and Oaksford 1999) and relying on the representations and high-level procedures of model-based accounts.

Syllogisms are just one domain of monadic reasoning, but we have begun to assess the theory's predictions in other domains, including immediate inferences from one premise to a conclusion (Khemlani et al. under review b), judgments of consistency, inferences about set-membership, and systematic fallacies in reasoning with quantifiers (Kunze et al. 2010). Likewise, we are expanding the theory to handle still other domains of reasoning, including sentential reasoning and probabilistic reasoning (see Table 1 for what is implemented in mReasoner, version 0.8).

Yet, the theory and its implementation are far from a unified account of reasoning. They have several major shortcomings. As Table 1 provides, they have yet to be extended to many domains of reasoning. They embody no general procedures that translate the instructions for different sorts of reasoning tasks into procedures that carry out these tasks. They make no numerical predictions about either accuracy or latency. They embody no principles of learning, and so they cannot learn heuristics. Finally, they offer no account of differences in ability or strategy from one individual to another.

Acknowledgements

This research was supported by a National Science Foundation Graduate Research Fellowship to the first author, and by National Science Foundation grant no. SES 0844851 to the second author to study deductive and probabilistic reasoning. The authors are grateful to Max Lotstein for his help in all aspects of the research, including the computational modelling. The authors thank Ruth Byrne, Vittorio Girotto, Sam Glucksberg, Adele Goldberg, Hua Gao, Catrinel Haught, Niklas Kunze, Greg Trafton, and Marco Ragni, for providing their helpful criticisms.

References

- Anderson, J.R. (1993), *Rules of the Mind*, Hillsdale, NJ: Erlbaum.
- Barwise, J. (1993), 'Everyday Reasoning and Logical Inference', *Behavioral and Brain Sciences*, 16, 337–338.
- Bauer, M.I., and Johnson-Laird, P.N. (1993), 'How Diagrams can Improve Reasoning', *Psychological Science*, 4, 372–378.
- Begg, I., and Denny, J. (1969), 'Empirical Reconciliation of Atmosphere and Conversion Interpretations of Syllogistic Reasoning', *Journal of Experimental Psychology*, 81, 351–354.
- Bell, V., and Johnson-Laird, P.N. (1998), 'A Model Theory of Modal Reasoning', *Cognitive Science*, 22, 25–51.
- Braine, M. (1978), 'On the Relation between the Natural Logic of Reasoning and Standard Logic', *Psychological Review*, 85, 1–21.
- Bucciarelli, M., and Johnson-Laird, P.N. (1999), 'Strategies in Syllogistic Reasoning', *Cognitive Science*, 23, 247–303.
- Bucciarelli, M., and Johnson-Laird, P.N. (2005), 'Naïve Deontics: A Theory of Meaning, Representation, and Reasoning', *Cognitive Psychology*, 50, 159–193.
- Byrne, R.M.J. (2005), *The Rational Imagination: How People Create Alternatives to Reality*, Cambridge, MA: MIT press.
- Byrne, R.M.J., and Johnson-Laird, P.N. (1989), 'Spatial Reasoning', *Journal of Memory and Language*, 28, 564–575.
- Byrne, R.M.J., and Tasso, A. (1999), 'Deductive Reasoning with Factual, Possible, and Counterfactual Conditionals', *Memory & Cognition*, 27, 726–740.
- Ceraso, J., and Provitera, A. (1971), 'Sources of Error in Syllogistic Reasoning', *Cognitive Psychology*, 2, 400–410.
- Chater, N., and Oaksford, M. (1999), 'The Probability Heuristics Model of Syllogistic Reasoning', *Cognitive Psychology*, 38, 191–258.
- Cohen, M.R., and Nagel, E. (1934), *An Introduction to Logic and Scientific Method*, London: Routledge & Kegan Paul.
- Erickson, J.R. (1974), 'A set Analysis Theory of Behavior in Formal Syllogistic Reasoning Tasks', in *Loyola Symposium on Cognition*, Vol. 2, ed. R. Solso, Hillsdale, NJ: Lawrence Erlbaum Associates, pp. 305–330.
- Evans, J.St.B.T. (2003), 'In Two Minds: Dual Process Accounts of Reasoning', *Trends in Cognitive Sciences*, 7, 454–459.
- Evans, J.St.B.T. (2007), *Hypothetical Thinking: Dual Processes in Reasoning and Judgement*, Hove: Psychology Press.

- Evans, J.St.B.T. (2008), 'Dual-Processing Accounts of Reasoning, Judgment and Social Cognition', *Annual Review of Psychology*, 59, 255–278.
- Evans, J.St.B.T., Handley, S.J., Harper, C.N.J., and Johnson-Laird, P.N. (1999), 'Reasoning about Necessity and Possibility: A Test of the Mental Model Theory of Deduction', *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25, 1495–1513.
- Ford, M. (1995), 'Two Modes of Mental Representation and Problem Solution in Syllogistic Reasoning', *Cognition*, 54, 1–71.
- Frosch, C.A., and Johnson-Laird, P.N. (2011), 'Is Everyday Causation Deterministic or Probabilistic?', *Acta Psychologica*, 137, 280–291.
- Geurts, B. (2003), 'Reasoning with Quantifiers', *Cognition*, 86, 223–251.
- Goldberg, A. (2003), 'Constructions: A New Theoretical Approach to Language', *Trends in Cognitive Science*, 7, 219–224.
- Goldvarg, E., and Johnson-Laird, P.N. (2001), 'Naive Causality: A Mental Model Theory of Causal Meaning and Reasoning', *Cognitive Science*, 25, 565–610.
- Goodwin, G., and Johnson-Laird, P.N. (2005), 'Reasoning about Relations', *Psychological Review*, 112, 468–493.
- Goodwin, G., and Johnson-Laird, P.N. (2006), 'Reasoning About the Relations Between Relations', *Quarterly Journal of Experimental Psychology*, 59, 1047–1069.
- Guyote, M.J., and Sternberg, R.J. (1981), 'A Transitive-Chain Theory of Syllogistic Reasoning', *Cognitive Psychology*, 13, 461–525.
- Hopcroft, J.E., and Ullman, J.D. (1979), *Formal Languages and Their Relation to Automata*, Reading, MA: Addison-Wesley.
- Inhelder, B., and Piaget, J. (1958), *The Growth of Logical Thinking from Childhood to Adolescence*, London: Routledge & Kegan Paul.
- Jahn, G., Knauff, M., and Johnson-Laird, P.N. (2007), 'Preferred Mental Models in Reasoning about Spatial Relations', *Memory & Cognition*, 35, 2075–2087.
- Jeffrey, R. (1981), *Formal Logic: Its Scope and Limits* (2nd ed.), New York, NY: McGraw-Hill.
- Johnson-Laird, P.N. (1975), 'Models of Deduction', in *Reasoning: Representation and Process in Children and Adults*, ed. R.J. Falmagne, Hillsdale, NJ: Erlbaum, pp. 7–54.
- Johnson-Laird, P.N. (1983), *Mental Models*, Cambridge, MA: Harvard University Press.
- Johnson-Laird, P.N. (1994), 'Mental Models and Probabilistic Thinking', *Cognition*, 50, 189–209.
- Johnson-Laird, P.N. (2006), *How We Reason*, Oxford: Oxford University Press.
- Johnson-Laird, P.N., and Byrne, R.M.J. (1991), *Deduction*, Hillsdale, NJ: Erlbaum.
- Johnson-Laird, P.N., and Byrne, R.M.J. (2002), 'Conditionals: A Theory of Meaning, Pragmatics, and Inference', *Psychological Review*, 109, 646–678.
- Johnson-Laird, P.N., Byrne, R.M.J., and Schaeken, W.S. (1992), 'Propositional Reasoning by Model', *Psychological Review*, 99, 418–439.
- Johnson-Laird, P.N., Byrne, R.M.J., and Tabossi, P. (1989), 'Reasoning by Model: The Case of Multiple Quantification', *Psychological Review*, 96, 658–673.
- Johnson-Laird, P.N., Girotto, V., and Legrenzi, P. (2004), 'Reasoning from Inconsistency to Consistency', *Psychological Review*, 111, 640–661.
- Johnson-Laird, P.N., and Hasson, U. (2003), 'Counterexamples in Sentential Reasoning', *Memory & Cognition*, 31, 1105–1113.
- Johnson-Laird, P.N., Legrenzi, P., Girotto, V., Legrenzi, M.S., and Caverni, J. (1999), 'Naïve Probability: A Mental Model Theory of Extensional Reasoning', *Psychological Review*, 106, 62–88.
- Johnson-Laird, P.N., and Savary, F. (1999), 'Illusory Inferences: A Novel Class of Erroneous Deductions', *Cognition*, 71, 191–229.
- Johnson-Laird, P.N., and Yang, Y. (2008), 'Mental Logic, Mental Models, and Computer Simulations of Human Reasoning', in *Cambridge Handbook of Computational Psychology*, ed. R. Sun, Cambridge, MA: Cambridge University Press, pp. 339–358.
- Juhos, C., Quelhas, A.C., and Johnson-Laird, P.N. (2012), 'Temporal and Spatial Relations in Sentential Reasoning', *Cognition*, 122, 393–404.
- Kahneman, D. (2011), *Thinking, Fast and Slow*, New York, NY: Farrar, Strauss, Giroux.

- Khemlani, S., and Johnson-Laird, P.N. (2009), 'Disjunctive Illusory Inferences and How to Eliminate Them', *Memory & Cognition*, 37, 615–623.
- Khemlani, S., and Johnson-Laird, P.N. (in press), 'Theories of the Syllogism: A Meta-analysis', *Psychological Bulletin*, in press.
- Khemlani, S., Lotstein, M., and Johnson-Laird, P.N. (under review a), A unified theory of syllogistic reasoning. Manuscript under submission.
- Khemlani, S., Lotstein, M., and Johnson-Laird, P.N. (under review b), Immediate inferences in quantified assertions. Manuscript under submission.
- Khemlani, S., Lotstein, M., and Johnson-Laird, P.N. (under review c), The psychology of set membership.
- Khemlani, S., Orenes, I., and Johnson-Laird, P.N. (in press), 'Negation: A Theory of its Meaning, Representation, and Use', *Journal of Cognitive Psychology*, in press.
- Kroger, J.K., Nystrom, L.E., Cohen, J.D., and Johnson-Laird, P.N. (2008), 'Distinct Neural Substrates for Deductive and Mathematical Processing', *Brain Research*, 1243, 86–103.
- Kunze, N., Khemlani, S., Lotstein, M., and Johnson-Laird, P.N. (2010), 'Illusions of Consistency in Quantified Assertions', in *Proceedings of the 32nd Annual Conference of the Cognitive Science Society*, eds. S. Ohlsson and R. Catrambone, Austin, TX: Cognitive Science Society, pp. 2028–2032.
- Langacker, R. (2008), *Cognitive Grammar: A Basic Introduction*, New York, NY: Oxford University Press.
- Mackiewicz, R., and Johnson-Laird, P.N. (2012), 'Reasoning from Connectives and Relations between Entities', *Memory & Cognition*, 40, 266–279.
- Neth, H., and Johnson-Laird, P.N. (1999), 'The Search for Counterexamples in Human Reasoning', in *Proceedings of the Twenty First Annual Conference of the Cognitive Science Society*, p. 806.
- Newell, A. (1990), *Unified Theories of Cognition*, Cambridge, MA: Harvard University Press.
- Osherson, D. (1975), 'Logical and Models of Logical Thinking', in *Reasoning: Representation and Process in Children and Adults*, ed. R.J. Falmagne, Hillsdale, NJ: Erlbaum, pp. 81–92.
- Partee, B.H. (1996), 'The Development of Formal Semantics in Linguistic Theory', in *The Handbook of Contemporary Semantic Theory*, ed. S. Lappin, Oxford: Blackwell, pp. 11–38.
- Peirce, C.S. (1931–1958), *Collected Papers of Charles Sanders Peirce* (8 vols), in eds. C. Hartshorne, P. Weiss, and A. Burks, Cambridge, MA: Harvard University Press.
- Politzer, G., van der Henst, J.B., Luche, C.D., and Noveck, I.A. (2006), 'The Interpretation of Classically Quantified Sentences: A Set-Theoretic Approach', *Cognitive Science*, 30, 691–723.
- Polk, T.A., and Newell, A. (1995), 'Deduction as Verbal Reasoning', *Psychological Review*, 102, 533–566.
- Quelhas, A.C., Johnson-Laird, P.N., and Juhos, C. (2010), 'The Modulation of Conditional Assertions and its Effects on Reasoning', *Quarterly Journal of Experimental Psychology*, 63, 1716–1739.
- Revlis, R. (1975), 'Two Models of Syllogistic Reasoning: Feature Selection and Conversion', *Journal of Verbal Learning and Verbal Behavior*, 14, 180–195.
- Rips, L.J. (1983), 'Cognitive Processes in Propositional Reasoning', *Psychological Review*, 90, 38–71.
- Rips, L.J. (1994), *The Psychology of Proof*, Cambridge, MA: MIT Press.
- Schaeken, W.S., Johnson-Laird, P.N., d'Ydewalle, G. (1996), 'Mental Models and Temporal Reasoning', *Cognition*, 60, 205–234.
- Slooman, S.A. (1996), 'The Empirical Case for Two Systems of Reasoning', *Psychological Bulletin*, 119, 3–22.
- Stanovich, K.E. (1999), *Who is Rational? Studies of Individual Differences in Reasoning*. Mahwah, NJ: Erlbaum.
- Stenning, K., and Yule, P. (1997), 'Image and Language in Human Reasoning: A Syllogistic Illustration', *Cognitive Psychology*, 34, 109–159.
- Verschueren, N., Schaeken, W., and d'Ydewalle, G. (2005), 'A Dual-Process Specification of Causal Conditional Reasoning', *Thinking and Reasoning*, 11, 278–293.
- Wason, P.C., and Shapiro, D. (1971), 'Natural and Contrived Experience in a Reasoning Problem', *Quarterly Journal of Experimental Psychology*, 23, 63–71.
- Wetherick, N.E., and Gilhooly, K.J. (1995), "'Atmosphere", matching, and logic in syllogistic reasoning', *Current Psychology*, 14, 169–178.