13
MODELS AND RATIONAL DEDUCTIONS

Phil N. Johnson-Laird

Introduction

In daily life, we reason to try to reach conclusions that are true. But, since our premises are often uncertain, a sensible goal is at least to make what logicians refer to as *valid* deductions. These are inferences whose conclusions hold in every case in which their premises hold (see Jeffrey, 1981, p. 1), and so if the premises are true their conclusions are true too, e.g.:

> If I'm in Britain then I drive on the left.
> I'm in Britain.
> Therefore, I drive on the left.

Whatever *rational* might mean – and it is open to many interpretations, rational deductions should at least be valid. They have no counterexamples, which are cases in which their premises are true but their conclusions false.

The inference above is so simple that you might think that all valid deductions in daily life can follow the formal rules of inference for a logical calculus. For instance, one such rule, appropriate for the inference above, is:

> If A then B; A; Therefore, B

where A and B are variables that range over propositions. This rule is from the calculus that concerns the analogs in logic of *if*, *or*, *and*, and *not* (Jeffrey, 1981). Everyday reasoning would be straightforward if it could follow rules that were complete in that they captured all and only valid deductions, that yielded a decision about validity or invalidity in a finite number of steps, and that did so in a tractable way, that is, depending on only reasonable amounts of time and memory. Alas, some inferences in life violate these desirable properties: their logic is incomplete, and they lack a decision procedure or a tractable one. An instructive example (from Johnson-Laird, 1983, p. 140) is:

> More than half the musicians were classically trained.
> More than half the musicians were in rock groups.
> So, at least one of the musicians was both classically trained and in a rock group.

The concept of "more than half" calls for a logic in which variables range over sets, and this logic is not complete, and has no decision procedure (Jeffrey, 1981, Chapter 7). Likewise, reasoning about the domain of two-dimensional spatial relations, such as "The cup is on the left of the saucer", to which we return later, is intractable for all but the simplest deductions (Johnson-Laird, 1983, p. 409; Ragni, 2003). So, even before we consider the limitations of human reasoning, deduction is bounded for any finite device (cf. Simon 2000). And yet it is to a finite device, the human brain, that we owe the proofs of incompleteness, undecidability, and intractability.

Validity is a notion that can be applied to any system of inference provided that we know the conditions in which its assertions are true. What is invalid in one logic of possibilities, for instance, is valid in another logic of possibilities, because the truth conditions for "possible" differ between the two logics (e.g., Girle, 2009). You might therefore think that, given the truth conditions for everyday assertions, rational deduction is nothing more than valid deduction. Whatever you understand *rationality* to mean, this idea is wrong. The reason why is illustrated in these three deductions, which are each valid and each silly:

1. I'm in Britain.
Therefore, it is raining or it isn't raining.

2. I'm in Britain, and if so then I drive on the left.
Therefore, if I'm in Britain, which I am, then I drive on the left.

3. I'm in Britain, and if so then I drive on the right.
Therefore, I'm in Britain or I drive on the right, or both.

In (1), the conclusion is bound to be true given the premises, because it is a tautology. But it is silly, because its conclusion has no relation to the premises. A sensible deduction needs to depend on the premises. In (2), the conclusion is bound to be true given the premises, and it does depend on them. But it is silly, because it merely repeats their contents. In (3), the conclusion is bound to be true given the premises. But it is silly, because it throws information away: the premise implies the conjunction of the two propositions. A sensible deduction is in contrast:

I'm in Britain, and if so then I drive on the right.
Therefore, I drive on the right.

Sensible deductions are accordingly inferences that maintain the information in the premises (and are thereby valid), that are parsimonious, and that yield a conclusion that was not explicit in the premises (see Johnson-Laird, 1983, p. 37; Johnson-Laird & Byrne, 1991, p. 22). If no conclusion meets these requirements then people tend to say that nothing follows from the premises – a judgement that stands in stark contrast to logic, in which infinitely many conclusions follow from any premises whatsoever. As the three examples above illustrate, most of them are silly. So, let us say: *To make a rational deduction is at least to maintain the information in the premises, to simplify, and to reach a new conclusion.*

Can naïve individuals – those who are ignorant of logic or its cognate disciplines – make rational deductions? Yes, of course. Without this ability, Aristotle and his intellectual descendants would have been unable to develop logic. You might suppose, as some theorists argue, that human beings would not have evolved or survived as a species had not at least some of them been capable of some rational deductions. This way of framing the matter recognizes two robust phenomena: people differ in deductive ability (Stanovich, 1999), and some rational deductions pertinent to daily life defeat almost everyone. Perhaps that is why logic exists.

How do individuals – naïve ones should henceforth be taken for granted – make rational deductions? The question has been under investigation for over a century, but only in recent years have cognitive scientists begun to answer it. This chapter outlines one such answer – the theory of mental models. It rests on five principles corresponding to sections in the chapter. Mental models are not uncontroversial, but the final section of the chapter shows that they explain phenomena beyond other theories of human deduction.

Possibilities underlie reasoning

The theory of mental models postulates that all inferences – even inductions, which are outside the scope of the present chapter – depend on envisaging what is possible, given premises such as assertions, diagrams, or direct observations (Johnson-Laird, 2006; Johnson-Laird & Ragni, 2019). A possibility is a simple sort of uncertainty, but a conclusion that holds in all the possibilities to which the premises refer must be valid. The major principle of human reasoning is accordingly that inferences are good only if they have no *counterexamples*, that is, possibilities in which the premises hold, but the conclusion does not.

> Suppose you know about a particular electrical circuit:
> If there was a short in the circuit then the fuse blew.

The salient possibility to which this assertion refers is:

Short in circuit	fuse blew

But, other cases are possible too, namely:

Not a short	fuse blew
Not a short	fuse did not blow

The first of these latter two possibilities occurs when the circuit is overloaded, and the second of them is the norm. It is hard to keep three distinct possibilities in mind, and so people tend to focus on the first one, and to make a mental note that the other cases (in which the *if*-clause of the conditional assertion does not hold) are possible. Of course, people don't use words and phrases to represent possibilities, but actual models of the world akin to those that the perceptual system constructs.

The *mental* models of the premise above are represented in the following diagram:

Short in circuit	fuse blew
. . .	

The first model is the salient one in which the *if*-clause holds. The second model (shown as an ellipsis) is an implicit one representing the other possibilities in which there wasn't a short in the circuit. Mental models underlie intuitive reasoning. You learn that there was a short in the circuit. It picks out only the first model, from which it follows at once:

> The fuse blew.

Possibilities and necessities can be defined in terms of one another: if an event is possible then it is not necessarily not the case, and if an event is necessary, then it is not possibly not the case. So, what evidence shows that people model possibilities rather than necessities? If possibilities are fundamental, then reasoners should be more accurate in inferring what's possible than in inferring what's necessary. Indeed, they are faster and more accurate (Bell & Johnson-Laird, 1998). A typical experiment concerned games of one-on-one basketball, in which only two can play. So, how would you answer the question about this game:

If Alan is in the game then Betty is in the game.
If Cheryl is in the game then David is not in the game.
Can Betty be in the game?

The correct answer is "yes", because both premises hold in case Alan and Betty are in the game. But, now suppose that the question is instead about a necessity:

Must Betty be in the game?

You should grasp at once that this question is harder. In fact, if you deliberate about the matter, you will discover that there are only three possible games: Alan plays Betty, Betty plays Cheryl, or Betty plays David. So, Betty must be in any game. When the correct answer is "no", the question about whether a player must be in the game is easier to answer than the question about whether a player can be in a game. A single model that is a counterexample suffices to justify the negative answer to the first question, whereas all models of possibilities are needed to justify the negative answer to the second question. These differences in difficulty between the two "yes" answers, and the switch between the two "no" answers, show that models represent possibilities rather than necessities.

Deductions are in default of information to the contrary

An assertion, such as:

The fault is in the software or in the cable, or both

implies three possibilities:

It is possible that the fault is in the software.
It is possible that the fault is in the cable.
It is possible that the fault is in the software and in the cable.

People make these three deductions, and reject only the deduction that it is possible that the fault is in neither the software nor the cable (Hinterecker, Knauff, & Johnson-Laird, 2016). Yet, none of these inferences is valid in modal logics, which concern possibilities and necessities. To see why, consider the case in which it is impossible for the fault to be in the software but true that it is in the cable. The premise is true, but the first conclusion above is false. So, the inference is invalid. Analogous cases show that the other two inferences are also invalid. Why, then, do individuals make these inferences? The answer is that they treat the disjunction as referring to the set of possibilities in default of information to the contrary. Conversely, if they know that it is possible that the fault is in the software, possible that it is in the cable, and possible that it is in both, they can use the disjunction above to summarize their knowledge. One reason for the assumption of defaults is their ubiquity in daily life.

Facts often reveal that a valid conclusion that you drew is false. It may be a surprise, but you cope. You no longer believe your conclusion. Yet, in logic, no need exists for you to withdraw it.

Logic means never having to be sorry about any valid conclusion that you've drawn. The reason is that a self-contradiction, such as one between a conclusion and a fact, implies that any conclusion whatsoever follows validly in logic. The jargon is that logic is "monotonic", whereas everyday reasoning is "non-monotonic". The divergence is just one of many between logic and life.

When the facts contradict a valid deduction, you should change your mind, and one view of a *rational* change is that it should be minimal. As William James (1907, p. 59) wrote, "[The new fact] preserves the older stock of truths with a minimum of modification, stretching them just enough to make them admit the novelty." This parsimony seems sensible, and many cognitive scientists advocate *minimalism* of this sort (e.g., Harman, 1986; Elio & Pelletier, 1997). It implies, of course, that the rational step in dealing with a contradiction to a valid inference is to make a minimal amendment of the premises. That's not what happens in daily life. Consider this example:

> If a person is bitten by a cobra then the person dies.
> Viv was bitten by a cobra. But, Viv did not die.
> What would you infer?

Minimalism predicts that you should make a minimal change to the premises, e.g.:

> Not everyone dies if bitten by a cobra.

In fact, most individuals respond in a different way. They create an explanation that resolves the inconsistency, e.g.:

> Someone sucked the venom from the bite so it did not get into Viv's bloodstream.

Hardly a minimal change. Yet, people tend to create such causal explanations, and to judge them to be more probable than minimal amendments to the premises (Johnson-Laird, Girotto, & Legrenzi, 2004). After they have created such explanations, they even find inconsistencies harder to detect (Khemlani & Johnson-Laird, 2012). Just as individuals assume possibilities by default, so they draw conclusions by default. They give them up in the light of evidence to the contrary.

Knowledge modulates the meanings of logical terms

In logic, the analogs of such words as *if*, *or*, and *and*, have constant meanings. Earlier we saw that conditionals – assertions based on "if" – usually refer to three distinct possibilities. But, consider the conditional:

> If God exists then atheism is wrong.

It ought to refer to the possibility in which God does not exist and atheism is wrong. But it doesn't. Knowledge of the meaning of "atheism", as disbelief in God, blocks the construction of the corresponding model. So, the preceding assertion is in reality a *biconditional* equivalent to:

> If, and only if, God exists then atheism is wrong.

Our knowledge modulates our interpretations of words (e.g., Quelhas & Johnson-Laird, 2017). Yet, we are unaware for the most part of these *modulations*. In one experiment (Juhos, Quelhas, & Johnson-Laird, 2012), the participants drew their own conclusion from premises containing a biconditional:

If the client makes an order then the goods are shipped. The goods are shipped. What follows?

They tended to infer:

The client made an order.

They also drew their own conclusion from the similar premises:

If the client makes an order then the goods are shipped. The client makes the order. What follows?

They tended to infer:

The goods are shipped.

They were unaware of a subtle difference between the two sorts of conclusion: their first conclusion was in the past tense, whereas their second conclusion was in the present tense (or in the future tense in Portuguese, which was the participants' native tongue). The model theory predicts this difference. The participants use their knowledge to infer the sequence in which the two events occur, i.e., the ordering the goods comes before their shipping. The conclusion to their first inference refers to the ordering, which comes before the shipping. So, the conclusion calls for the past tense. In contrast, the conclusion to their second inference refers to the shipping, which comes after the ordering. So, the conclusion calls for the present tense, which can refer to future events in both English and Portuguese, or to the Portuguese future tense. (As its native speakers often don't realize, English has no future tense.)

The implementation of modulation in a computer program revealed that some assertions should be judged to be true *a priori*, i.e., without the need for evidence, and that other assertions should be judged to be false *a priori*. An experiment corroborated the predictions (Quelhas, Rasga, & Johnson-Laird, 2017). Individuals judged an assertion, such as:

If Mary has flu then she is ill

to be true. And they judged an assertion such as:

If Mary has flu then she is healthy

to be false. Such judgements run contrary to an influential view in philosophy that the difference between *a priori* assertions and those contingent on evidence is "an unempirical dogma of empiricism" (Quine, 1953, p. 23). Not any more.

Reasoning depends on intuitions or on deliberations, or both

Human reasoners are equipped with two systems for reasoning, which interact with one another. Intuitions depend on mental models, which represent only what is true. As we saw earlier, the assertion:

If there was a short in the circuit then the fuse blew has these mental models:

a short in circuit	fuse blew
. . .	

Suppose you learn:

The fuse didn't blow.

It eliminates the first of the two mental models above, and so only the implicit model remains. It has no explicit content, and so it seems that nothing follows from the two premises. Many people have this intuition. If they think harder, however, their deliberations may lead them to fully explicit models of the possibilities:

a short in circuit	fuse blew
not a short in circuit	fuse blew
not a short in circuit	fuse did not blow

The premise that the fuse did not blow rules out the first two of these models. So, only the third model remains, and it yields the conclusion:

There was not a short in the circuit.

Some people are able to make this valid deduction, but, as the theory predicts, it is harder and takes longer than the inference described earlier. It depends on deliberation and fully explicit models whereas the earlier inference depends on intuition and mental models.

Consider the following inference, where each premise refers to two alternative possibilities:

Either the pie is on the table or else the cake is on the table.
Either the pie isn't on the table or else the cake is on the table.
Could both of these assertions be true at the same time?

Most people say, "yes". Their inference depends on mental models. Each premise refers to the possibility that the cake is on the table, and so it seems that both assertions could be true. But, let's spell out the fully explicit possibilities. The first premise refers to these two possibilities for what's on the table:

the pie	not the cake
not the pie	the cake

The second premise refers to these two possibilities:

not the pie	not the cake
the pie	the cake

A careful examination reveals that the two assertions have no possibility in common. So, the correct conclusion is that they cannot both be true at the same time. The intuitive inference is an illusion, just one of many different sorts (Khemlani & Johnson-Laird, 2017). What they have in common is that they follow from mental models of premises, but not from their fully explicit models. The systematic occurrence of illusory inferences is by far the most unexpected prediction of the model theory. It is a litmus test to show that reasoning is relying on mental models rather than on fully explicit models.

Simulations based on spatial, temporal, and kinematic models

So far, you may have the impression that models are just words and phrases rather than actual models of the world. In fact, models can be three-dimensional structures underlying your grasp, say, of how to get from an office on one floor of a building to its main entrance. If someone asks you to draw a path in the air of your route you can do so: the ability is a prerequisite for navigation, especially the sort that Micronesian islanders carry out without instruments, not even a compass (see, e.g., Gladwin, 1970).

When individuals make two-dimensional spatial inferences, they rely on models. Here is one such problem:

> The cup is on the right of the plate.
> The spoon is on the left of the plate.
> The knife is in front of the spoon.
> The saucer is in front of the cup.
> What is the relation between the knife and the saucer?

The premises describe a layout of the sort in this diagram of a table-top:

spoon	plate	cup
knife		saucer

where the items at the bottom of the diagram are in front of those at the top. It is quite easy to infer:

> The knife is on the left of the saucer.

If instead you used logical rules to derive your answer, your first step would be to deduce the relation between the spoon and cup (using the logical properties of *on the left* and *on the right*):

> The spoon is on the left of the cup.

You would then use axioms for the two-dimensional relations in the remaining premises to deduce from this intermediate conclusion that:

> The knife is on the left of the saucer.

A similar problem is most revealing:

> The plate is on the right of the cup.
> The spoon is on the left of the plate.
> The knife is in front of the spoon.
> The saucer is in front of the plate.
> What is the relation between the knife and the saucer?

The premises are consistent with two distinct layouts, because of the uncertainty of the relation between the spoon and the cup:

spoon	cup	plate		cup	spoon	plate
knife		saucer			knife	saucer

Yet, the relation between the two items in the question is the same in both layouts, and so it follows that:

The knife is on the left of the saucer.

If you used models to infer the conclusion, the deduction should be harder than the previous one, because you have to construct two models. But, if you used logical rules to infer the conclusion, then the deduction should be easier than the previous one. You no longer have to deduce the relation between the spoon and the plate, because it is stated in the second premise. Hence, the experiment is crucial in that mental models and logical rules make opposite predictions. The results of experiments corroborated the model theory, and they did so even when a later study corrected for the fact that the first premise is irrelevant in the second inference (see Byrne & Johnson-Laird, 1989; Schaeken, Girotto, & Johnson-Laird, 1998). Analogous results occur when inferences concern temporal rather than spatial relations (Schaeken, Johnson-Laird, & d'Ydewalle, 1996). And a corollary is that diagrams that make it easy to envisage alternative possibilities both speed up and enhance the accuracy of deductions (Bauer & Johnson-Laird, 1993).

The description of a sequence of events can elicit a kinematic model that itself unfolds in time. Imagine a railway track that runs from left to right. It has a siding onto which cars enter from the left, and exit to the left. Five cars are at the left end of the track: A B C D E. Figure 13.1 is a diagram of the situation. Here is a problem for you to solve:

> All but one of the cars enter the siding. The remaining car at the left end of the track moves over to the right end of the track. Now each car on the siding, one at a time, moves back to the left end of the track and immediately over to the right end. What is the resulting order of the cars at the right end of the track?

Adults, and even 10-year-old children, can deduce the final order of the cars: E D C B A. They do so by envisaging the effect of each move until the sequence ends. In other words, they simulate the sequence in a kinematic mental model (Khemlani et al., 2013). Children accompany their simulations with gestures indicating the moves they would make if they were allowed to move the cars – an outward sign of inward simulation. And they are less accurate if they are prevented from gesturing (Bucciarelli et al., 2016).

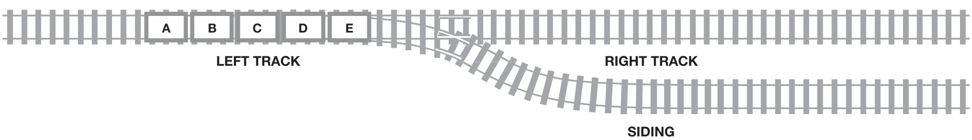

Figure 13.1 A diagram of a small railway track with five cars at the left-hand side of the track

Conclusion

Rational deductions depend on mental models. The theory is controversial in psychology, but it has led to evidence that challenges its rivals. The earliest theories of human reasoning proposed that the mind is equipped with formal rules of inference akin to those of logic (e.g., Rips, 1994). These theories cannot explain the illusory inferences described earlier, and, as we saw, when models were pitted against logic, number of models rather than number of logical steps predicted the difficulty of deductions. Because models represent possibilities, modal logic might be relevant because it deals with possibilities and necessities. But, even though there are many different modal logics (see, e.g., Girle, 2009), human reasoning diverges from all of them. One such divergence is in inferences such as:

> The fault is in the software or in the cable, or both.
> Therefore, it is possible that the fault is in the software.

It is invalid in all modal logics (for the reasons described on p. 220). Yet, as we saw, most people make such inferences (Hinterecker et al., 2016).

A radical alternative is that reasoning relies, not on logic, but on probabilities (e.g., Oaksford & Chater, 2007). This proposal has a more limited purview than the model theory. It offers no explanation of illusory or kinematic inferences. It applies to Aristotelian syllogisms, but the model theory fits the results of experiments better (Khemlani & Johnson-Laird, 2013). Possibilities with appropriate numbers are probabilities. So, the model theory offers an account of probabilistic reasoning (Johnson-Laird et al., 1999; Khemlani, Lotstein, & Johnson-Laird, 2015). In contrast, theories based on probabilities alone cannot distinguish between certainty and necessity, so they are unable to explain inferences about who *must* play in one-on-one games of basketball, as discussed above.

Our reasoning is bounded. We make deductions in domains that are undecidable or intractable in their demands on time and memory. Yet, we are rational in principle. We grasp the force of counterexamples, and spontaneously use them to refute invalid inferences. Our deductions from fully explicit models yield valid conclusions based on all the information in the premises. Indeed, models of possibilities are enough to establish principles of rationality, and the model theory itself.

References

Bauer, M. I., and Johnson-Laird, P. N. 1993. How diagrams can improve reasoning. *Psychological Science* 4, 372–378.

Bell, V., and Johnson-Laird, P. N. 1998. A model theory of modal reasoning. *Cognitive Science* 22, 25–51.

Boolos, G. S., and Jeffrey, R. C. 1980. *Computability and Logic* (3rd ed.). Cambridge: Cambridge University Press.

Bucciarelli, M., Mackiewicz, R., Khemlani, S. S., and Johnson-Laird, P. N. 2016. Children's creation of algorithms: Simulations and gestures. *Journal of Cognitive Psychology* 28, 297–318.

Byrne, R. M. J., and Johnson-Laird, P. N. 1989. Spatial reasoning. *Journal of Memory and Language* 28, 564–575.

Elio, R., and Pelletier, F.J. 1997. Belief change as prepositional update. *Cognitive Science* 21, 419–460.

Girle, R. 2009. *Modal Logics and Philosophy* (2nd ed.). London: Routledge.

Gladwin, T. 1970. *East Is a Big Bird*. Cambridge, MA: Harvard University Press.

Harman, G. 1986. *Change in View: Principles of Reasoning*. Cambridge, MA: MIT Press.

Hinterecker, T., Knauff, M., and Johnson-Laird, P. N. 2016. Modality, probability, and mental models. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 42, 1606–1620.

James, W. 1907. *Pragmatism*. New York: Longmans, Green.

Jeffrey, R. 1981. *Formal Logic: Its Scope and Limits* (2nd ed.). New York: McGraw-Hill.

Johnson-Laird, P. N. 1983. *Mental Models*. Cambridge, MA: Harvard University Press.

Johnson-Laird, P. N. 2006. *How We Reason*. New York: Oxford University Press.

Johnson-Laird, P. N., and Byrne, R. M. J. 1991. *Deduction*. Hillsdale, NJ: Erlbaum.

Johnson-Laird, P. N., Girotto, V., and Legrenzi, P. 2004. Reasoning from inconsistency to consistency. *Psychological Review* 111, 640–661.

Johnson-Laird, P. N., Legrenzi, P., Girotto, V., Legrenzi, M., and Caverni, J-P. 1999. Naive probability: A mental model theory of extensional reasoning. *Psychological Review* 106, 62–88.

Johnson-Laird, P. N., and Ragni, M. 2019. Possibilities as the foundation of reasoning. *Cognition* 193, 130950.

Juhos, C. A., Quelhas, C., and Johnson-Laird, P. N. 2012. Temporal and spatial relations in sentential reasoning. *Cognition* 122, 393–404.

Khemlani, S. S., and Johnson-Laird, P. N. 2012. Hidden conflicts: Explanations make inconsistencies harder to detect. *Acta Psychologica* 139, 486–491.

Khemlani, S. S., and Johnson-Laird, P. N. 2013. The processes of inference. *Argument and Computation* 4, 4–20.

Khemlani, S. S., and Johnson-Laird, P. N. 2017. Illusions in reasoning. *Minds and Machines* 27, 11–35.

Khemlani, S. S., Lotstein, M., and Johnson-Laird, P. N. 2015. Naive probability: Model-based estimates of unique events. *Cognitive Science* 39, 1216–1258.

Khemlani, S. S., Mackiewicz, R., Bucciarelli, M., and Johnson-Laird, P. N. 2013. Kinematic mental simulations in abduction and deduction. *Proceedings of the National Academy of Sciences* 110, 16766–16771.

Oaksford, M., and Chater, N. 2007. *Bayesian Rationality: The Probabilistic Approach to Human Reasoning*. Oxford: Oxford University Press.

Quelhas, A. C., and Johnson-Laird, P. N. 2017. The modulation of disjunctive assertions. *Quarterly Journal of Experimental Psychology* 70, 703–717.

Quelhas, A. C., Rasga, C., and Johnson-Laird, P. N. 2017. A priori true and false conditionals. *Cognitive Science* 41, 1003–1030.

Quine, W. V. O. 1953. *From a Logical Point of View*. Cambridge, MA: Harvard University, Press.

Ragni, M. 2003. An arrangement calculus, its complexity and algorithmic properties. In A. Günter, R. Kruse, and B. Neumann (Eds.), *Advances in Artificial Intelligence: Proceedings of 26th Annual German Conference on AI* (pp. 580–590). Berlin: Springer.

Rips, L. 1994. *The Psychology of Proof*. Cambridge, MA: MIT Press.

Schaeken, W. S., Girotto, V., and Johnson-Laird, P. N. 1998. The effect of an irrelevant remise on temporal and spatial reasoning. *Kognitionswisschenschaft* 7, 27–32.

Schaeken, W. S., Johnson-Laird, P. N., and d'Ydewalle, G. 1996. Mental models and temporal reasoning. *Cognition* 60, 205–234.

Simon, H. A. 2000. Bounded rationality in social science: Today and tomorrow. *Mind & Society* 1, 25–39.

Stanovich, K. E. 1999. *Who Is Rational? Studies of Individual Differences in Reasoning*. Hillsdale, NJ: Erlbaum,