

This is a section of [doi:10.7551/mitpress/11252.001.0001](https://doi.org/10.7551/mitpress/11252.001.0001)

The Handbook of Rationality

Edited by: Markus Knauff, Wolfgang Spohn

Citation:

The Handbook of Rationality

Edited by: Markus Knauff, Wolfgang Spohn


DOI: 10.7551/mitpress/11252.001.0001

ISBN (electronic): 9780262366175

Publisher: The MIT Press

Published: 2021

Funding for the open access edition was provided by the MIT Libraries Open Monograph Fund.



The MIT Press

3.3 Psychological Theories of Syllogistic Reasoning

Sangeet Khemlani

Summary

Psychologists have studied syllogistic inferences for more than a century, because they can serve as a microcosm of human rationality. “Syllogisms” is a term that refers to a set of 64 reasoning arguments, each of which comprises two premises, such as, “All of the designers are women. Some of the women are not employees. What, if anything, follows?” People make systematic mistakes on such problems, and they appear to reason using different strategies. A meta-analysis showed that many existing theories fail to explain such patterns. To address the limitations of previous accounts, two recent theories synthesized both heuristic and deliberative processing. This chapter reviews both accounts and addresses their strengths. It concludes by arguing that if syllogistic reasoning serves as a sensible microcosm of rationality, the synthesized theories may provide directions on how to resolve broader conflicts that vex psychologists of reasoning and human thinking.

1. A Microcosm of Rationality

In 1908, the German scholar Gustav Störring published a 130-page manuscript detailing the results of the first known experiments on human reasoning. His main purpose in conducting them was to develop solutions to long-standing debates between logicians and philosophers, such as what people imagine when they reason. The studies worked like this: volunteers entered a dark room alone, sat down, and received a battery of deductive reasoning problems called “syllogisms,” one after another. Störring recorded his observations of their verbal responses, reaction times, eye movements, gestures, and even their breathing patterns (Störring, 1908).

The research would likely be rejected were it to be submitted to any contemporary psychology journal. For one thing, Störring investigated only four participants. For another, he used an arbitrary experimental design,

and he failed to present any quantitative analysis of their behaviors except for a single table that listed averaged reaction times. But what Störring learned from his research was remarkable (see Clark, 1922; Knauff, 2013; Politzer, 2004). He noticed, for instance, that his volunteers were biased by the structure of the different reasoning problems: for some syllogisms, volunteers chose conclusions immediately, as though they could observe the answer directly. For other problems, they reported the sensation of *Nachdenken*, that is, “a feeling of deliberation.” His volunteers appeared to adopt certain strategies as they carried out the task, and they seemed aware of their strategies well enough that they could articulate them. They reported that they used their imagination and mental imagery on many problems, and when probed further, they were able to depict that experience by sketching out corresponding diagrams. Perhaps the most important discovery was that Störring’s volunteers had a limited ability to select logically *valid* conclusions, that is, conclusions that were true in all the situations in which the premises were true (cf. Jeffrey, 1981, p. 1)—for some problems, they produced correct answers, and for others, they made mistakes.

Scientists often make use of microcosms as a way of understanding broader phenomena. For instance, the geneticist Gregor Mendel examined pea plants to understand genetic inheritance; the entomologist Agostino Bassi studied silkworms to understand bacterial diseases. In his experiments on syllogisms, Störring had analyzed a feasible microcosm of human rationality. In the years that followed, syllogisms played an outsized role in educating contemporary researchers on the processes of thinking and reasoning. Many experiments investigated syllogistic reasoning in isolation, and many more used syllogistic reasoning as a stand-in for reasoning behavior more generally. For instance, Goel and colleagues ran a neuroimaging experiment in which they gave participants syllogistic reasoning problems with and without meaningful contents to discover that certain

brain regions—such as temporal and frontal regions—systematically respond to semantic information (Goel, Buchel, Frith, & Dolan, 2000).

Perhaps syllogisms serve as an attractive microcosm of thinking behavior because of their simplicity and that there is only a finite number of them. Classical syllogisms (i.e., those investigated by Aristotle and Scholastic logicians) are reasoning arguments comprising multiple premises, such as

(1) All of the women are designers.

Some of the employees are not women.

What, if anything, follows?

These syllogisms contain a quantified noun phrase, such as “all of the women,” and these quantifiers can be in one of four separate *moods*, that is, expressions comprising quantifiers and negations, as shown below:

All of the <i>a</i> are <i>b</i> .	(<i>Aab</i>)
Some of the <i>a</i> are <i>b</i> .	(<i>Iab</i>)
Some of the <i>a</i> are not <i>b</i> .	(<i>Oab</i>)
None of the <i>a</i> is <i>b</i> .	(<i>Eab</i>)

The parentheses indicate the abbreviation conventions adopted by Scholastic logicians (i.e., the 12th-century university scholars who gained access to Aristotle’s works). Contemporary psychologists adopted those conventions, and we retain them here. Since syllogisms consist of two premises, the *terms* in the premises (e.g., “women,” “designers,” “employees”) can occur in four different arrangements. These different arrangements are known as *figures*:

Figure 1	Figure 2	Figure 3	Figure 4
<i>a – b</i>	<i>b – a</i>	<i>a – b</i>	<i>b – a</i>
<i>b – c</i>	<i>c – b</i>	<i>c – b</i>	<i>b – c</i>

Other psychologists use different numbering systems for figures—and they sometimes include the conclusion as part of their numbering systems. Here we state the figures in terms of the premises only.

In sum, syllogisms concern 64 separate reasoning problems (4 moods of the first premise \times 4 moods of the second premise \times 4 separate figures). Many experiments on syllogisms focus on only these 64 problems, that is, they provide participants with the pairs of premises and then ask them to infer what follows from them. Typically, reasoners do not consider all the possible valid and invalid responses when they generate conclusions; they tend to describe just one or two. But across the problems overall, their conclusions can be classified into nine different structural patterns, as follows: “All of the *As* are

Cs” (abbreviated as *Aac*); “All of the *Cs* are *As*” (abbreviated as *Aca*); the responses that correspond to *Iac*, *Ica*, *Oac*, *Oca*, *Eac*, and *Eca*; and the response that no valid conclusion follows.

As Störing recognized, some syllogisms are quite easy, and others can be difficult. Consider this problem:

(2) Some of the designers are animators.

All of the animators are undergraduates.

What, if anything, follows?

Before reading further, how might you respond to the problem? If you inferred that some of the designers are undergraduates, you’d be correct. But now consider this problem:

(3) None of the designers are animators.

All of the animators are undergraduates.

What, if anything, follows?

This problem is significantly harder: the correct response is that some of the undergraduates are not designers. Why is (2) easy but (3) difficult? To answer the question, psychologists have run many studies on which inferences people conclude from the 64 syllogisms. Khemlani and Johnson-Laird (2012) compiled six of them together in a meta-analysis, which shows that the most common response to (2) is “Some of the designers are undergraduates” (see figure 3.3.1, left panel, row *Iab Abc*), and the most common response to (3) is “None of the designers are undergraduates” (see figure 3.3.1, left panel, row *Eab Abc*)—the response is an error, since (3) does not rule out the possibility that some of the designers are undergraduates.

In a typical study, responses to a syllogism vary from one individual to the next. Consider how people respond to (1), that is, figure 3.3.1, right panel, row *Aba Ocb*: some draw an (erroneous) conclusion of the form *Oca* about half of the time; others make a different error and conclude *Ica*. Only about a fifth of university students (who typically serve as participants in studies on syllogisms) respond correctly that there is no valid conclusion. The psychologist’s task is to explain the robust patterns of inference across the set of 64 problems. The difficulty for theorists is that reasoners approach the problems with different abilities and appear to develop different strategies (see table 3.3.1). The variability in reasoners’ responses was enough to convince some theorists that the only way to understand how people reason syllogistically is to examine their individual differences (Stenning & Cox, 2006)—some reasoners generate correct

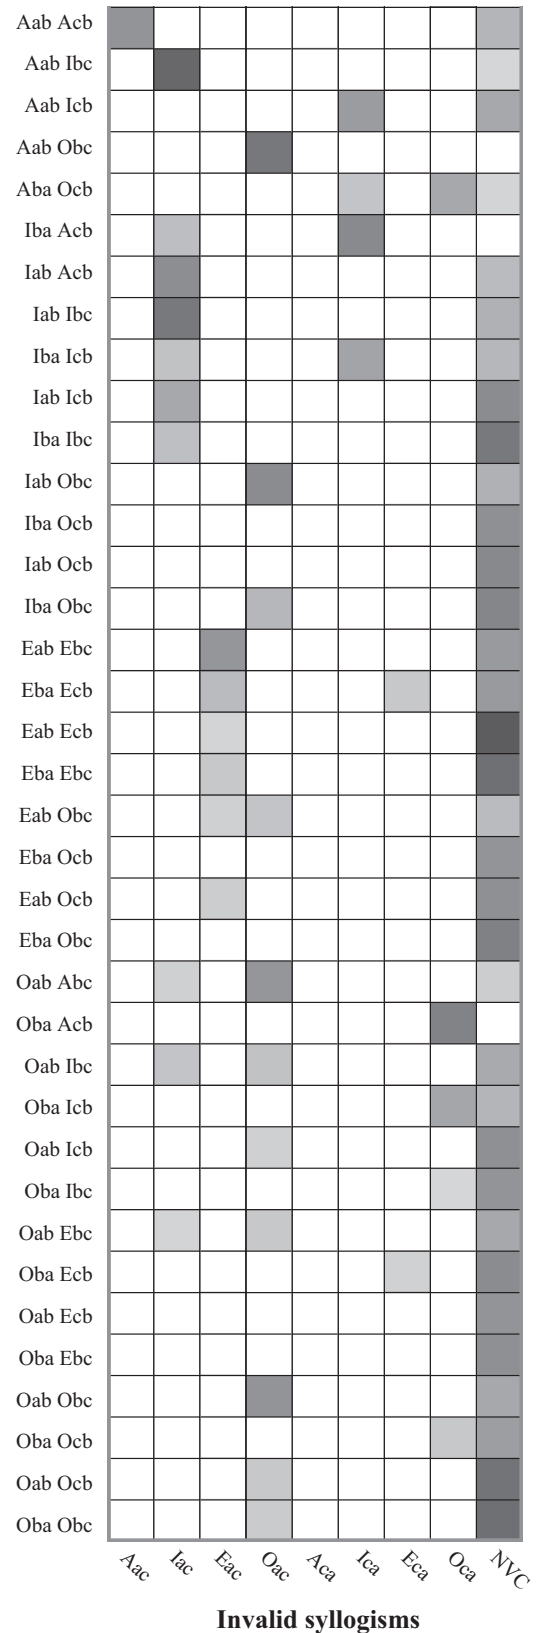



Figure 3.3.1

The percentages of responses to 64 syllogisms in the meta-analysis in Khemlani and Johnson-Laird (2012). Each of the 64 pairs of premises occurs in a row, and each of the possible responses occurs in a column. Abbreviations for premises are as follows: *Aac* = All of the *A* are *C*, *Iac* = Some of the *A* are *C*, *Eac* = None of the *A* is a *C*, *Oac* = Some of the *A* are not *C*, and *NVC* = No valid conclusion. The left panel denotes the 27 syllogisms with a valid definite conclusion, and the right panel denotes the 37 syllogisms without a valid definite conclusion. The grayscale in each cell indicates the proportion of corresponding conclusions (black = 100% and white = 16% or below). Hence, for the top-most valid syllogism, *Aab Abc*, nearly 100% of participants in the meta-analysis responded that *Aac* follows.

Table 3.3.1

Summary of robust sources of differences in syllogistic reasoning performance

Phenomenon	Description
Differences between problems	Some syllogisms are extremely difficult (1% correct solutions), and others are easy (90% correct conclusions). Differences in difficulty are reliable across studies.
Differences among individuals	Accuracy varies greatly: from one reasoner to another, between certain groups of university students, and as a function of intelligence and training.
The figural effect	The figure of a syllogism (i.e., the order in which terms are organized) affects the frequency and type of conclusions reasoners tend to draw.
The effect of contents	Individuals are more likely to accept believable conclusions and less likely to accept unbelievable ones, especially for invalid syllogisms. This is also referred to as the “belief bias” effect.
Developmental differences	Children are unable to draw sensible inferences from syllogisms until they can understand and produce quantifiers. The ability develops through adolescence up to adulthood.
Strategic differences	Reasoners acquire various strategies to solve syllogisms (e.g., which premise is interpreted first, how premises are interpreted, and how counterexamples are constructed).

answers for 85% of syllogisms, and others generate correct answers for only 15% of them (Johnson-Laird, 1983, pp. 118–119).

The study of syllogisms arguably resulted in a major challenge to human rationality. Early work by Woodworth and Sells (1935) suggested that, instead of reasoning their way through syllogisms, people were biased by the “atmosphere” created by the premises, which yielded a predisposition to accept a certain sort of conclusion. The atmosphere effect suggested wholesale irrationality: humans could diverge from normative reasoning behavior predictably and systematically (see chapter 2.3 by Johnson-Laird, this handbook). The result sparked a fascination with the extent to which human reasoning could be characterized as rational (see chapter 1.2 by Evans and chapter 3.1 by Steinberger, both in this handbook), and theorists began to devise accounts of the phenomena underlying syllogistic reasoning. The use of computational and formal tools helped some researchers implement psychological theories of the syllogism and test them against human data. As a result, after decades of research, nearly a dozen theories of the phenomenon had been proposed, and there existed a dire need to sort out the different theoretical proposals. Khemlani and Johnson-Laird (2012) surveyed existing psychological accounts of syllogistic reasoning to discover broad trends between them. The survey suggested that theories tended to fall into one of three groups: one group of theories explained syllogistic reasoning by appealing to sets of heuristics in how quantified statements were processed (e.g., Begg & Denny, 1969; Chater & Oaksford, 1999; Revlis, 1975; Wetherick & Gilhooly, 1995). For example, the so-called matching strategy

(Wetherick & Gilhooly, 1995) posited that for syllogisms such as

(4) Some of the designers are women.

Some of the women are employees.

What, if anything, follows?

people should conclude—erroneously—that “some of the designers are employees.” The reason is because the conclusion matches the mood of the most “conservative” premise, that is, the premise that presupposes the existence of the fewest entities. And indeed, reasoners draw the predicted conclusion 61% of the time (see figure 3.3.1, right panel, row *Iab Ibc*). But about a third of the time, they also accurately infer that “No valid conclusion” follows, and accounts based on heuristics have difficulty explaining the deliberative processes by which reasoners correct their mistakes (see also Ragni, Dames, Brand, & Rieusterer, 2019). In order to account for deliberative reasoning, another group of psychological theories proposed that reasoners mentally simulate the situation described in the premises when they reason about syllogisms (Bucciarelli & Johnson-Laird, 1999; Guyote & Sternberg, 1981; Johnson-Laird & Steedman, 1978; Polk & Newell, 1995). The theories posited that mental simulations help explain both errors and correct responses: reasoners construct, and can make inferences from, initial simulations, but difficult syllogisms demand reasoners to consider alternative simulations (Johnson-Laird, 1983). A third group of theories assumed that syllogistic reasoning depends on mental proofs and rules of inference akin to those in formal logic (see, e.g., Braine & Romain, 1983; Geurts, 2003; Politzer, 2007; Rips, 1994)—but such theories have systematic difficulty explaining how reasoners draw the

conclusion that nothing follows from a set of premises, and so we presently address only the first two groups of theories.

Theories based on heuristics and theories based on deliberation both failed to explain many systematic patterns of syllogistic reasoning (Khemlani & Johnson-Laird, 2012). In retrospect, the debate between the two types of processing presents a false dichotomy. A robust theory of syllogistic inference needs to explain both heuristic and deliberative processing (see Evans & Stanovich, 2013; Johnson-Laird & Steedman, 1978). The most recent psychological accounts of syllogistic reasoning have sought to unify the two kinds of reasoning processes. This chapter reviews these recent theories and their computational implementations, and it summarizes their strengths and weaknesses. It also provides a broader perspective on how the theories address ongoing debates in the psychology of reasoning.

2. Unified Accounts of Syllogistic Reasoning

An ancient idea is that human thinking relies on two different systems: one fast, one slow. Peter Wason, with his students Philip Johnson-Laird and Jonathan Evans, proposed that reasoning processes should be construed in terms of two distinct, interreliant processes (Johnson-Laird & Wason, 1970; Wason & Evans, 1974). As Evans (2008, p. 263) notes, the dichotomy between heuristics and deliberation is closely related to dual processes because heuristics are thought to be a fast, shallow form of processing and deliberation is thought to be a slower, deeper form of processing. In practice, heuristics and deliberative thinking often occur sequentially, that is, a heuristic response is proposed and a deliberative process validates or falsifies it. More general accounts of dual processing are not committed to sequential processing—they permit that fast processes and slow processes can operate in parallel and interact with one another. The introduction noted that previous theories of syllogistic reasoning tended to account for one type of process over the other. It may be that previous theories were easier to formulate because it is difficult to anticipate the interactive effects of two interdependent processes. Yet, if it is indeed the case that human thinking depends on two inextricable processes, those theories were doomed to fail.

Two recent theories of syllogistic reasoning are unique in that they seek to model interactive processing. Both theories are built around the integrative idea that fast, heuristic processing is the result of a biased sampling procedure that can be formalized using probabilistic constraints and that slower, deliberative processing suggests

that reasoning depends on representations referred to as “mental models.” The theory that people construct mental models when they reason originates from Johnson-Laird (1983; see also chapter 2.3 by Johnson-Laird, this handbook), who computationally developed earlier proposals that people build “small-scale models” of reality to anticipate events (Craik, 1943). Johnson-Laird’s “model theory” posits that each mental model represents a distinct possibility or situation in the world. In other words, when reasoners draw inferences from syllogisms, they mentally simulate the situation referred to by the premises. The model theory predicts that problems that require reasoners to consider multiple mental models should be more difficult relative to those that require fewer models. As a result, models help explain reasoning difficulty in many domains (see, e.g., Johnson-Laird & Khemlani, 2013; Khemlani & Johnson-Laird, 2017). But, as Khemlani and Johnson-Laird (2012) show, previous implementations of the model theory tend to make overly liberal predictions of the kinds of syllogistic inferences people are likely to draw. Hence, the two latest theories of syllogistic reasoning add additional constraints that explain why reasoners are reticent to draw overly liberal conclusions from model-based representations. We describe each theory in turn.

2.1 The Probability Sampling Model

Masaki Hattori developed a recent account of syllogistic reasoning called the “probability sampling model” (PSM; Hattori, 2016). The account holds that reasoners interpret a set of syllogistic premises by constructing a prototypical representation of them (referred to as a “probability prototype model”). The probability prototype model uses circles to represent set-membership relations, and so it is closely related to, for example, Euler circles and Venn diagrams (see the meta-analysis in Khemlani & Johnson-Laird, 2012, which reviews other theories based on such diagrammatic systems, and also chapter 13.1 by Jamnik, this handbook). The various intersections of the circles in the prototype model denote different kinds of individuals. For example, a diagram of “All of the designers are women” would include a circle representing the set of designers embedded inside a circle representing the set of women. The intersections of the circles represent the different kinds of individuals consistent with the premise, for example, that one area would represent the designers who are also women, and another would represent women who are not designers. Hattori posits that reasoners annotate the different areas of the prototype model with information about the probability of their occurrence.

Once the probabilities are established, reasoners draw a finite set of random samples to construct a mental model, that is, a small set of tokens that denote the entities referred to in the premises. Hattori's adaptation of mental model theory is more restricted than previous accounts. For example, provided that the two kinds of areas established by a prototype model have equal probability, a sample mental model of "All of the designers are women" can be represented in the following diagram:

designer	woman
¬ designer	woman
designer	woman
¬ designer	woman

Each row of the diagram depicts the results of a random draw from the prototype model, and so the first row depicts a designer who is also a woman. The "¬" denotes the symbol for negation, and so the second row depicts a woman who is not a designer. The additional rows depict additional random draws from the prototype model. In the PSM, the establishment of a sample mental model is the central representation on which a unitary reasoning process operates. The algorithm works by applying a series of tests, one after another, to the sample mental model in order to generate a conclusion. Figure 3.3.2 provides a schematic of how the full theory works.

To test the theory's predictions, Hattori implemented the theory computationally and then ran simulations that compared the theory's predictions against eight separate data sets on syllogistic reasoning. He also compared the PSM's predictions against two other theories: Chater and Oaksford's probability heuristics model (Chater & Oaksford, 1999) and a parameterized version of mental model theory (Hattori's implementation of Johnson-Laird & Bara, 1984). His analyses show that the PSM matches the performance of both theories (Hattori, 2016, p. 308).

One major strength of the PSM is that it can explain how contents affect the kinds of inferences reasoners draw. For instance, consider the following problem:

(5) Some of the Frenchmen are wine-drinkers.

Some of the wine-drinkers are Italians.

What, if anything, follows?

A robust result from studies of syllogistic reasoning is that reasoners should be less apt to infer "Some of the Frenchmen are Italians," since they likely consider it either implausible or rare that a person is both French and Italian at the same time. Hattori's PSM can account for the effect as follows: the meanings of the premises

bias the way people construct the probability prototype model such that implausible areas of the representation are assigned low probabilities. Hattori implemented the process, and he described two studies that corroborate the predictions of his implementation (Hattori, 2016).

However, Hattori's theory is not without its limitations. One limitation is that the theory cannot explain how reasoners might draw inferences from more than two premises: in these cases, the two-dimensional probability sampling model may be unable to represent all of the possible different individuals. Another limitation is that under the PSM, deliberation occurs by applying a set of logical tests, one after another, to a sample mental model. The process is akin to the way heuristic processing operates in previous theories (e.g., Chater & Oaksford, 1999). No evidence at present suggests that reasoners apply such tests in a fixed, systematic order for each and every syllogism; indeed, the evidence suggests that reasoners tend to develop strategies over the course of a study on syllogisms, which would seem to conflict with the PSM. Moreover, running such tests in a fixed order is bound to produce a conclusion for any sample mental model—and so a clear consequence of the tests is that the PSM cannot account for why reasoners often spontaneously respond that "No valid conclusion" follows (Ragni et al., 2019). Other accounts based on sampling and constructing mental models suffer from similar issues (see, e.g., Tessler & Goodman, 2014). In general, it is too taxing for people to constantly and repeatedly apply the same set of tests to the representations they construct, and so the algorithm, although tractable, is limited and not cognitively plausible. In essence, the PSM does not explain the psychological processes that underlie how reasoners deliberate on syllogistic inferences.

The primary theoretical contribution of Hattori's (2016) probabilistic sampling model is that it integrates probabilistic machinery and sampling procedures with the construction of mental models. Another recent computational theory—mReasoner—likewise integrates probabilistic sampling with procedures with the construction of mental models.

2.2 mReasoner: A Unified Computational Implementation of Mental Model Theory

mReasoner (Khemlani & Johnson-Laird, 2021) is a unified computational implementation of mental model theory, which posits that reasoning depends on the construction and manipulation of mental models, that is, iconic simulations of possibilities (Bucciarelli & Johnson-Laird, 1999; Johnson-Laird, 1983; Johnson-Laird, Khemlani, &




Figure 3.3.2

A schematic diagram of how Hattori’s (2016) probability sampling model draws syllogistic conclusions. (Used with permission from Hattori, 2016.)

Goodwin, 2015a). The theory and its implementation are based on three fundamental assumptions:

- **Mental models represent possibilities:** a given assertion refers to a set of discrete possibilities that are observed or imagined (Johnson-Laird, 1983).
- **Iconicity and discreteness:** mental models are iconic, that is, their structures mirror the structures of what they represent (see Peirce, 1931–1958, vol. 4). Models

can also include abstract symbols, for example, the symbol for negation (Khemlani, Orenes, & Johnson-Laird, 2012). They are discrete in that they do not consist of continuous spaces, areas, and regions.

- **Dual processes:** reasoning, including syllogisms, is based on two interacting sets of processes. Rapid, heuristic reasoning occurs as a consequence of building and scanning a single model. Deliberative reasoning,

by contrast, occurs as a result of revising the initial model to search for alternative models of heuristic conclusions (Khemlani & Johnson-Laird, 2013, 2016; Khemlani, Lotstein, Trafton, & Johnson-Laird, 2015).

The computational model makes syllogistic inferences by stochastically constructing a mental model directly from the premises. Hence, unlike the PSM, mReasoner eschews any intermediary representations (e.g., the probability prototype model). Two factors dictate how initial models are constructed: the size of a model (i.e., the maximum number of entities it represents) and the contents of a model. One parameter in the system controls the size of a mental model. It does so by basing the size on a sample drawn from a Poisson distribution. Another parameter governs the model's contents, which are drawn from the most common set of possibilities corresponding to a particular assertion (the canonical set) or else the complete set of possibilities consistent with the assertion. For example, in the case of "All of the designers are women," reasoners tend to consider only one canonical possibility: female designers. But the complete set of possibilities allows for women who are not designers, and a parameter in the system sets the probability of drawing from the complete set.

To illustrate how the system makes heuristic inferences, consider this syllogism:

- (6) All of the designers are women.
 Some of the women are not employees.
 What, if anything, follows?

Suppose that the premises for (6) are input into mReasoner. The system may construct the following initial model:

designer	woman		
designer	woman		
designer	woman	employee	
			employee

The diagram is similar to that of the sample mental model illustrated in the previous section: its rows denote separate individuals. To generate a heuristic conclusion, the system scans the model in the direction in which it was built. In the model above, for instance, the system builds tokens for "designers" first, "women" second, and "employees" third. Hence, the program draws an initial conclusion that interrelates "designers" to "employees" (e.g., some of the designers are not employees). This conclusion matches the preponderance of conclusions that reasoners spontaneously generate. For other sorts of syllogisms, the system draws initial intuitive

conclusions that interrelate "employees" to "designers," again depending on how the model was constructed.

Because the heuristic conclusion depends on just a single model, the system generates it quickly. But, as the example illustrates, the conclusion may be invalid. To correct the error, the program can call on a deliberative component to search for counterexamples to conclusions (Khemlani & Johnson-Laird, 2013). It operates by modifying the initial model, using a finite set of search strategies (Bucciarelli & Johnson-Laird, 1999). When the deliberative system is engaged, it can find a counterexample to the conclusion that some of the designers are not employees:

designer	woman	employee
designer	woman	employee
designer	woman	employee
	woman	

This model is one that represents the premises in (6) but falsifies the heuristic conclusion, and so the program responds that no valid conclusion holds. A separate parameter controls the probability that the deliberative system is engaged. Figure 3.3.3 provides a schematic diagram of the system's architecture.

mReasoner provides a close fit to the results from the data presented in the meta-analysis on syllogistic reasoning compared to alternative theories (Khemlani & Johnson-Laird, 2013). A major strength of the theory is that it is flexible enough to simulate reasoning about problems that include any number of premises, for example, one-premise "immediate" inferences (Khemlani et al., 2015), two-premise syllogisms (Khemlani & Johnson-Laird, 2016) and set-membership inferences (Khemlani, Lotstein, & Johnson-Laird, 2014), and more complex problems composed of three premises (Ragni, Khemlani, & Johnson-Laird, 2014). It does so by relying on a limited set of parameters that provide motivated constraints on how the theory's predictions vary (see figure 3.3.3). The flexibility of the system allows it to model individual differences and strategies in syllogistic reasoning (Bacon, Handley, & Newstead, 2003; Khemlani & Johnson-Laird, 2016; Ragni, Riesterer, Khemlani, & Johnson-Laird, 2018).

mReasoner has several limitations. One limitation is that at present, unlike the probability sampling model, mReasoner does not have machinery capable of explaining effects of content (i.e., the "belief bias" effect). In principle, the machinery is feasible, and a prototype system implements "semantic modulation" for sentential reasoning (see Khemlani, Byrne, & Johnson-Laird,

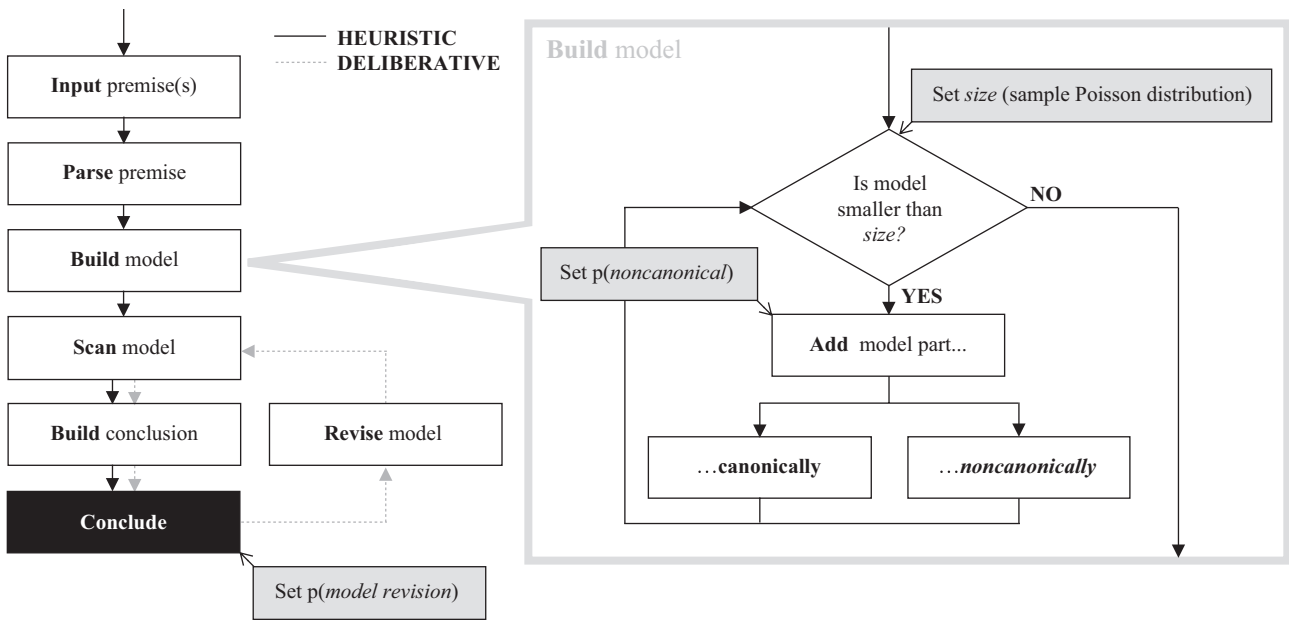


Figure 3.3.3

A schematic diagram of the program implemented in mReasoner for reasoning with mental models of quantified assertions such as “All the designers are women.” White boxes denote procedures, gray boxes denote parameters, and the black box denotes the point at which the program generates a conclusion. The program integrates two separate reasoning pipelines: a fast, heuristic process (represented by solid arrows) that does not revise models and a slow, recursive, deliberative process (represented by dotted arrows) that revises initial models and draws conclusions from the set of alternative models. The diagram also highlights the stochastic algorithm that controls how models are built.

2018). But the system has not been applied to simulate belief bias effects. A second limitation of the theory is that, unlike theories based on Bayesian updating mechanisms (e.g., Chater & Oaksford, 1999; Oaksford & Chater, 2009), mReasoner does not explain how learning and reasoning interact with one another.

2.3 Integrated Models of Reasoning

The two newest theories of syllogistic reasoning described here were developed to achieve parsimony relative to previous theories. They share many commonalities: they both use sampling procedures to build up representations. They both result in the construction of mental models (i.e., mental simulations of possibilities). They’re both implemented computationally, and so their limitations are clearly defined. Because of the clarity afforded by their computational implementations, both have been validated against numerous data sets—indeed, both accounts manage to explain far more data than any single previous account of syllogistic reasoning to date. And they predict many empirical phenomena beyond the narrow scope of the 64 syllogistic reasoning problems. As such, the theories may serve as promising foundations for future research.

3. General Discussion

Why have psychologists spent so much time and effort on understanding syllogistic reasoning? One reason is because Aristotle had made it central to his logic; logic was practically synonymous with syllogistic logic until Frege developed the predicate calculus, when logicians realized that syllogisms were but a small corner of the space of logical analysis. Nevertheless, syllogisms presented a feasible microcosm for study. By the middle of the 20th century, researchers began to recognize that people’s syllogistic inferences could serve as a measure of their general ability to think rationally. Syllogisms were included in verbal reasoning tasks, entrance exams, and aptitude tests (e.g., Ekstrom, French, Harman, & Dermen, 1976; Nester & Colberg, 1984). Researchers of psychiatric disorders even began to use erroneous responses in syllogisms as a way to measure irrational thinking, for example, in patients with schizophrenia (Gottesman & Chapman, 1960; Williams, 1964), a practice that carries on into modern psychiatry (Mujica-Parodi, Malaspina, & Sackeim, 2000). Yet, no psychological account of syllogisms existed when Störring conducted his initial investigations into people’s logical reasoning abilities. His

primary goal was to resolve theoretical debates among philosophers and logicians; hence, it was necessary to develop a theory of syllogisms.

The generation of researchers who followed Störing developed theories that operated under the consensus that errors in reasoning were the result of fallible strategies. New theories continued to flourish late into the century, such that by its end, theorists had proposed a dozen different psychological accounts of syllogistic reasoning. As theories flourished, consensus among researchers dissipated. Khemlani and Johnson-Laird (2012) noted that the failure to arrive at a single, unified account was a catastrophe for the science of human thinking: “skeptics may well conclude that cognitive science has failed [as it] yields no consensus about a small, empirically tractable domain of reasoning.”

The dozen theories of syllogistic reasoning cannot all be correct: they make inconsistent predictions and commitments. A trivial path forward might be to stitch together all of the theories’ predictions by disregarding their inconsistencies; after all, people themselves can be inconsistent. Different people develop different reasoning strategies, and even the same person may use different strategies on different occasions (see, e.g., Ragni et al., 2018). Alas, an amalgamated account is unlikely to explain the robust patterns of responses depicted in figure 3.3.1. Worse still, it is unlikely to concern anything beyond the scope of the 64 syllogistic reasoning problems, and so it cannot provide much of a guide for how humans reason in general. A more productive way to resolve the impasse is to develop a new, unified theory of syllogistic reasoning, one that supersedes the old ones.

Hattori (2016) and Khemlani and Johnson-Laird (2013, 2016) took such an approach in their most recent proposals of reasoning by syllogism. Both proposals provide a framework for explaining (a) the interpretation and mental representation of syllogistic premises, (b) what the mind computes and how it carries out inferential tasks with such assertions, (c) the differences in difficulty from one inference to another and common errors, (d) how contents affect performance, and (e) individual differences in performance from one person to another. And their theories have been implemented in computer programs whose operations were used to simulate empirical data.

Perhaps the most fierce debate between theorists who study human reasoning is the challenge of integrating probabilistic and deductive inference (see, e.g., Baratgin et al., 2015; Evans & Over, 2013; Johnson-Laird, Khemlani, & Goodwin, 2015a, 2015b; Khemlani, 2018). Syllogistic reasoning stands at the nexus of the debate: people

have the ability to draw valid syllogistic inferences, that is, they have the ability to reason deductively. The result is evident in Störing’s seminal studies and in numerous studies that followed. But reasoners’ deductive abilities are limited; they often draw conclusions that are likely—but not guaranteed—to be true, and their own background knowledge biases their tendency to accept valid deductions. Some theorists sought to resolve the discrepancy by developing new probabilistic interpretations of validity (Chater & Oaksford, 1999; see also Evans, 2012; Over, 2009), by emphasizing those aspects of reasoning that are easiest to formulate with probabilistic machinery (e.g., Evans & Over, 2004), and by discounting the relative importance of deduction in daily thinking (e.g., Oaksford & Chater, 2002). But such treatments are controversial, even among theorists who argue that human rationality is fundamentally probabilistic (see Johnson-Laird et al., 2015a). As a result, arguments on how to integrate probability and deduction often result in deadlock. Just as the two theories of syllogistic reasoning reviewed in this chapter may help adjudicate between heuristic and deliberative processes, they may also represent ways to overcome larger deadlocks in debates on rationality: at their core, the theories advocate limited and biased sampling processes best characterized and formalized with probabilistic constraints. But the outputs of the sampling processes are discrete mental simulations—mental models—that help explain deductions, inductions, and errors in reasoning.

Acknowledgments

This research was supported by a grant from the Naval Research Laboratory to study deductive reasoning. I am grateful to my advisor and long-time collaborator, Phil Johnson-Laird, and I also thank Paul Bello, Gordon Briggs, Monica Bucciarelli, Ruth Byrne, Sam Glucksberg, Adele Goldberg, Geoff Goodwin, Hillary Harner, Tony Harrison, Laura Hiatt, Zach Horne, Laura Kelly, Markus Knauff, Joanna Korman, Max Lotstein, Marco Ragni, Nicolas Riesterer, and Greg Trafton for their intuitions and deliberations.

References

- Bacon, A., Handley, S., & Newstead, S. (2003). Individual differences in strategies for syllogistic reasoning. *Thinking & Reasoning*, *9*, 133–168.
- Baratgin, J., Douven, I., Evans, J. St. B. T., Oaksford, M., Over, D., & Politzer, G. (2015). The new paradigm and mental models. *Trends in Cognitive Sciences*, *19*, 547–548.

- Begg, I., & Denny, J. P. (1969). Empirical reconciliation of atmosphere and conversion interpretations of syllogistic reasoning errors. *Journal of Experimental Psychology*, *81*, 351–354.
- Braine, M. D. S., & Rumain, B. (1983). Logical reasoning. In P. H. Mussen (Ed.), *Handbook of child psychology: Vol. 3. Cognitive development* (pp. 263–340). New York, NY: Wiley.
- Bucciarelli, M., & Johnson-Laird, P. N. (1999). 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.
- Clark, R. S. (1922). An experimental study of silent thinking. *Archives of Psychology*, *7*, 5–13.
- Craik, K. J. W. (1943). *The nature of explanation*. Cambridge, England: Cambridge University Press.
- Ekstrom, R. B., French, J. W., Harman, H. H., & Dermen, D. (1976). *Manual for kit of factor-referenced cognitive tests* (Vol. 102). Princeton, NJ: Educational Testing Service.
- 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. (2012). Questions and challenges for the new psychology of reasoning. *Thinking & Reasoning*, *18*, 5–31.
- Evans, J. St. B. T., & Over, D. E. (2004). *If*. Oxford, England: Oxford University Press.
- Evans, J. St. B. T., & Over, D. E. (2013). Reasoning to and from belief: Deduction and induction are still distinct. *Thinking & Reasoning*, *19*, 267–283.
- Evans, J. St. B. T., & Stanovich, K. E. (2013). Dual-process theories of higher cognition: Advancing the debate. *Perspectives on Psychological Science*, *8*, 223–241.
- Geurts, B. (2003). Reasoning with quantifiers. *Cognition*, *86*(3), 223–251.
- Goel, V., Buchel, C., Frith, C., & Dolan, R. J. (2000). Dissociation of mechanisms underlying syllogistic reasoning. *Neuroimage*, *12*, 504–514.
- Gottesman, L., & Chapman, L. J. (1960). Syllogistic reasoning errors in schizophrenia. *Journal of Consulting Psychology*, *24*, 250–255.
- Guyote, M. J., & Sternberg, R. J. (1981). A transitive-chain theory of syllogistic reasoning. *Cognitive Psychology*, *13*(4), 461–525.
- Hattori, M. (2016). Probabilistic representation in syllogistic reasoning: A theory to integrate mental models and heuristics. *Cognition*, *157*, 296–320.
- Jeffrey, R. (1981). *Formal logic: Its scope and limits* (2nd ed.). New York, NY: McGraw-Hill.
- Johnson-Laird, P. N. (1983). *Mental models: Towards a cognitive science of language, inference, and consciousness*. Cambridge, MA: Harvard University Press.
- Johnson-Laird, P. N., & Bara, B. G. (1984). Syllogistic inference. *Cognition*, *16*, 1–61.
- Johnson-Laird, P. N., & Khemlani, S. S. (2013). Toward a unified theory of reasoning. In B. Ross (Ed.), *The psychology of learning and motivation* (Vol. 59, pp. 1–42). Cambridge, MA: Academic Press.
- Johnson-Laird, P. N., Khemlani, S., & Goodwin, G. P. (2015a). Logic, probability, and human reasoning. *Trends in Cognitive Sciences*, *19*, 201–214.
- Johnson-Laird, P. N., Khemlani, S., & Goodwin, G. P. (2015b). Response to Baratgin et al.: Mental models integrate probability and deduction. *Trends in Cognitive Sciences*, *19*, 548–549.
- Johnson-Laird, P. N., & Steedman, M. J. (1978). The psychology of syllogisms. *Cognitive Psychology*, *10*, 64–99.
- Johnson-Laird, P. N., & Wason, P. C. (1970). A theoretical analysis of insight into a reasoning task. *Cognitive Psychology*, *1*, 134–148.
- Khemlani, S. S. (2018). Reasoning. In S. L. Thompson-Schill (Ed.), *Stevens' handbook of experimental psychology and cognitive neuroscience* (4th ed., Vol. 3, pp. 385–427). Hoboken, NJ: Wiley.
- Khemlani, S., Byrne, R. M. J., & Johnson-Laird, P. N. (2018). Facts and possibilities: A model-based theory of sentential reasoning. *Cognitive Science*, *42*, 1887–1924.
- Khemlani, S., & Johnson-Laird, P. N. (2012). Theories of the syllogism: A meta-analysis. *Psychological Bulletin*, *138*, 427–457.
- Khemlani, S., & Johnson-Laird, P. N. (2016). How people differ in syllogistic reasoning. In A. Papafragou, D. Grodner, D. Mirman, & J. Trueswell (Eds.), *Proceedings of the 38th Annual Conference of the Cognitive Science Society* (pp. 2165–2170). Austin, TX: Cognitive Science Society.
- Khemlani, S., & Johnson-Laird, P. N. (2017). Illusions in reasoning. *Minds & Machines*, *27*, 11–35.
- Khemlani, S., & Johnson-Laird, P. N. (2021). Reasoning about properties: A computational theory. Manuscript in press at *Psychological Review*.
- Khemlani, S., Lotstein, M., & Johnson-Laird, P. N. (2014). A mental model theory of set membership. In P. Bello, M. Guarini, M. McShane, & B. Scassellati (Eds.), *Proceedings of the 36th Annual Conference of the Cognitive Science Society* (pp. 2489–2494). Austin, TX: Cognitive Science Society.
- Khemlani, S., Lotstein, M., Trafton, J. G., & Johnson-Laird, P. N. (2015). Immediate inferences from quantified assertions. *Quarterly Journal of Experimental Psychology*, *68*, 2073–2096.
- Khemlani, S., Orenes, I., & Johnson-Laird, P. N. (2012). Negation: A theory of its meaning, representation, and use. *Journal of Cognitive Psychology*, *24*, 541–559.

- Knauff, M. (2013). *Space to reason: A spatial theory of human thought*. Cambridge, MA: MIT Press.
- Mujica-Parodi, L. R., Malaspina, D., & Sackeim, H. A. (2000). Logical processing, affect, and delusional thought in schizophrenia. *Harvard Review of Psychiatry, 8*, 73–83.
- Nester, M. A., & Colberg, M. (1984). The effects of negation mode, syllogistic invalidity, and linguistic medium on the psychometric properties of deductive reasoning tests. *Applied Psychological Measurement, 8*, 71–79.
- Oaksford, M., & Chater, N. (2002). Commonsense reasoning, logic and human rationality. In R. Elio (Ed.), *Common sense, reasoning and rationality* (pp. 174–214). Oxford, England: Oxford University Press.
- Oaksford, M., & Chater, N. (2009). Précis of *Bayesian rationality: The probabilistic approach to human reasoning*. *Behavioral and Brain Sciences, 32*, 69–120.
- Over, D. E. (2009). New paradigm psychology of reasoning. *Thinking & Reasoning, 15*, 431–438.
- Peirce, C. S. (1931–1958). *Collected papers of Charles Sanders Peirce* (C. Hartshorne, P. Weiss, & A. Burks, Eds., 8 vols.). Cambridge, MA: Harvard University Press.
- Politzer, G. (2004). Some precursors of current theories of syllogistic reasoning. In K. Manktelow & M. C. Chung (Eds.), *Psychology of reasoning: Theoretical and historical perspectives* (pp. 223–250). New York, NY: Psychology Press.
- Politzer, G. (2007). The psychological reality of classical quantifier entailment properties. *Journal of Semantics, 24*(4), 331–343.
- Polk, T. A., & Newell, A. (1995). Deduction as verbal reasoning. *Psychological Review, 102*, 533–566.
- Ragni, M., Dames, H., Brand, D., & Riesterer, N. (2019). When does a reasoner respond: Nothing follows? In A. K. Goel, C. M. Seifert, & C. Freksa (Eds.), *Proceedings of the 41st Annual Conference of the Cognitive Science Society* (pp. 2640–2646). Montreal, Canada: Cognitive Science Society.
- Ragni, M., Khemlani, S., & Johnson-Laird, P. N. (2014). The evaluation of the consistency of quantified assertions. *Memory & Cognition, 42*, 53–66.
- Ragni, M., Riesterer, N., Khemlani, S., & Johnson-Laird, P. N. (2018). Individuals become more logical without feedback. In C. Kalish, M. Rau, T. Rogers, & J. Zhu (Eds.), *Proceedings of the 40th Annual Conference of the Cognitive Science Society* (pp. 2315–2320). Austin, TX: Cognitive Science Society.
- 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. (1994). *The psychology of proof: Deductive reasoning in human thinking*. Cambridge, MA: MIT Press.
- Stenning, K., & Cox, R. (2006). Reconnecting interpretation to reasoning through individual differences. *Quarterly Journal of Experimental Psychology, 59*, 1454–1483.
- Störring, G. (1908). Experimentelle Untersuchungen über einfache Schlussprozesse [Experimental investigations of simple inference processes]. *Archiv für die gesamte Psychologie, 11*, 1–27.
- Tessler, M. H., & Goodman, N. D. (2014). Some arguments are probably valid: Syllogistic reasoning as communication. In *Proceedings of the 36th Annual Meeting of the Cognitive Science Society* (pp. 1574–1579). Austin, TX: Cognitive Science Society.
- Wason, P. C., & Evans, J. St. B. T. (1974). Dual processes in reasoning? *Cognition, 3*, 141–154.
- Wetherick, N. E., & Gilhooly, K. J. (1995). ‘Atmosphere’, matching, and logic in syllogistic reasoning. *Current Psychology, 14*, 169–178.
- Williams, E. B. (1964). Deductive reasoning in schizophrenia. *Journal of Abnormal and Social Psychology, 69*, 47–61.
- Woodworth, R. S., & Sells, S. B. (1935). An atmosphere effect in formal syllogistic reasoning. *Journal of Experimental Psychology, 18*, 451–460.

© 2021 The Massachusetts Institute of Technology

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Stone Serif and Stone Sans by Westchester Publishing Services.

Library of Congress Cataloging-in-Publication Data

Names: Knauff, Markus, editor. | Spohn, Wolfgang, editor.

Title: The handbook of rationality / edited by Markus Knauff and Wolfgang Spohn.

Description: Cambridge : The MIT Press, 2021. | Includes bibliographical references and index.

Identifiers: LCCN 2020048455 | ISBN 9780262045070 (hardcover)

Subjects: LCSH: Reasoning (Psychology) | Reason. | Cognitive psychology. | Logic. | Philosophy of mind.

Classification: LCC BF442 .H36 2021 | DDC 153.4/3—dc23

LC record available at <https://lcn.loc.gov/2020048455>