THINKING & REASONING, 2008, 14 (4), 342–364

Psychology Press
Taylor & Francis Group

The psychological puzzle of Sudoku

N. Y. Louis Lee
The Chinese University of Hong Kong, Shatin, Hong Kong SAR of China

Geoffrey P. Goodwin
University of Pennsylvania, Philadelphia, PA, USA

P. N. Johnson-Laird
Princeton University, NJ, USA

Sudoku puzzles, which are popular worldwide, require individuals to infer the missing digits in a 9×9 array according to the general rule that every digit from 1 to 9 must occur once in each row, in each column, and in each of the 3-by-3 boxes in the array. We present a theory of how individuals solve these puzzles. It postulates that they rely solely on pure deductions, and that they spontaneously acquire various deductive tactics, which differ in their difficulty depending on their "relational complexity", i.e., the number of constraints on which they depend. A major strategic shift is necessary to acquire tactics for more difficult puzzles: solvers have to keep track of possible digits in a cell. We report three experiments corroborating this theory. We also discuss their implications for theories of reasoning that downplay the role of deduction in everyday reasoning.

Keywords: Deduction; Problem solving; Reasoning; Inferential strategies.

Correspondence should be addressed to N. Y. Louis Lee, Department of Educational Psychology, The Chinese University of Hong Kong, Shatin, New Territories, Hong Kong SAR of China. E-mail: louis@cuhk.edu.hk

We thank the editor for useful advice about the framing of the paper; Vittorio Girotto and two anonymous reviewers for helpful criticisms; Sam Glucksberg, Adele Goldberg, and Fabien Savary for helpful suggestions; and Winton Au and Fanny Cheung for their assistance in Hong Kong. We also thank Edwin Williams for a helpful discussion and for sending us his program that devises Sudoku puzzles.

http://www.psypress.com/tar DOI: 10.1080/13546780802236308

Psychologists have studied deductive reasoning for over a century and they have proposed various theories about its underlying mental processes in naïve individuals; that is, those who have had no training in logic. Earlier theories followed the lead of Piaget and were based on formal rules of inference (e.g., Inhelder & Piaget, 1958; for an assessment of the deductive paradigm, see Evans, 2002). However, about 25 years ago an alternative theory proposed that deduction was based on mental models constructed from the meaning of the premises and general knowledge: a valid conclusion was one that held in all the models of the possibilities consistent with the premises (e.g., Johnson-Laird, 1983; Johnson-Laird & Byrne, 1991).

Recently psychologists have realised that pragmatic processes have a major impact on reasoning (e.g., Johnson-Laird, 2006; Johnson-Laird & Byrne, 2002), and they have proposed dual-process theories of reasoning in which one system, System 1, exploits automatic processes sensitive to pragmatic factors, and a subsequent system, System 2, uses logic and calculation (e.g., Evans, 2007; Sloman, 1996; Stanovich, 2008). Some theorists have even made the extreme claim that there is no need to postulate deductive ability as part of thinking. In the case of syllogisms, which are simple deductions based on quantifiers such as "all" and "some", naïve reasoners are said to respond solely on the basis of the atmosphere of the premises (e.g., Wetherick & Gilhooly, 1995). Similarly, Oaksford and Chater (e.g., 2002, p. 349) argue that "everyday rationality does not depend on formal systems like logic and only formal rationality is constrained and error prone . . . everyday reasoning is probabilistic and people make errors in so-called logical tasks because they generalize these strategies to the laboratory". Hertwig, Ortmann, and Gigerenzer (1997, pp. 105–106) likewise write: ". . . those who study first-order logic or variants thereof, such as mental rules and mental models, ignore the ecological and social structure of environments". Still others argue that deductive reasoning depends on pragmatic schemas for specific contents (e.g., Cheng & Holyoak, 1985), or on innate modules adapted to deal with specific contents, such as checking for cheaters (e.g., Cosmides, Tooby, Fiddick, & Bryant, 2005). So, just how much ability in *pure* deductive reasoning do naïve individuals really possess?

One way to answer this question is to consider Sudoku puzzles. They are digit placement puzzles that have become popular worldwide. Their popularity shows that people can solve them. But they cannot be solved using pragmatic schemas or innate modules tuned to specific contents, or fast and frugal heuristics, or the atmosphere of premises, or probabilistic reasoning. They are, as we shall argue, puzzles of pure deduction, and their solution depends ultimately on the ability to make valid deductive inferences; that is, to draw conclusions that must be true given the truth of their premises. In terms of dual processes, the rapid intuitions of System 1 may play a part in their solution but the slower deliberations of System 2 are decisive.

Sudoku puzzles derive from Latin squares (Ball & Coxeter, 1987; Euler, 1849). Figure 1 presents three typical puzzles, and the task is to fill in the empty cells in an array with the correct digits. One general rule governs the puzzles: Every digit from 1 to 9 must occur exactly once in each row, in each column, and in each of the nine 3-by-3 boxes into which the array is divided (shown by the bold lines in Figure 1).

This rule contains three quantifiers: *every*, *exactly once*, and *each*, and so its logical analysis calls for first-order logic (Jeffrey, 1981). A proper Sudoku puzzle has one unique solution, but larger puzzles rapidly become intractable (Yato & Seta, 2002).

In order to explain how individuals master Sudoku, we need to draw a distinction between inferential strategies and tactics. We propose the following working definition following Van der Henst, Yang, and Johnson-Laird (2002):

> A *strategy* in reasoning is a systematic sequence of elementary mental steps that an individual follows in making an inference.

Mild

		1		9		2	7	
		9			2		5	
2					3			
3				1	4			2
	8						4	
1			2	8				5
			9					7
	1		3			9		
	4	6		7		5		

Difficult

					8			7
8	1	7			2			
				5			1	4
7		2		9		1		5
5		8		7		3		2
3	7			8				
			4			9	7	3
4			5					

Fiendish

7	5			9			4	6
9		1				3		2
2			6		1			7
	8						2	
1			3		8			5
3		9				2		4
8	4			3			7	9

Figure 1. Three Sudoku puzzles used in Experiment 1: mild, difficult, and fiendish (from Gould, 2005).

We refer to each of these mental steps as a *tactic*, and so a strategy is a sequence of tactics that an individual uses to make an inference. We can illustrate this terminology with the following inference:

If there is a 5 in this cell, then there is not a 5 elsewhere in the row.
There is a 5 elsewhere in the row.
Therefore, there is not a 5 in this cell.

You can make this inference using various strategies. For example, you can make the supposition that there is a 5 in this cell, infer from the conditional that there is not a 5 elsewhere in the row, detect the inconsistency between this conclusion and the second premise, and so reject your supposition. Each step in this strategy—the supposition, the simple inference, the detection of the inconsistency, and so on—is a separate tactic. An alternative strategy is to envisage that the one possibility consistent with both premises is:

not 5 in this cell 5 elsewhere in the row

and to use this possibility to accept the conclusion. In this case, the tactics are to envisage a possibility, and to draw a conclusion from it.

Theorists of deductive reasoning disagree about whether inferential tactics rely on formal rules of inference (Braine & O'Brien, 1998; Rips, 1994) or mental models (Johnson-Laird & Byrne, 1991). They also disagree on whether there is a single deterministic inferential strategy (formal rule theories) or multiple non-deterministic strategies (the mental model theory; see, e.g., Van der Henst et al., 2002). Sudoku puzzles have a lesson to teach both sorts of theory. When participants make deductions in the psychological laboratory, they already know the basic inferential tactics that they use. They know, for example, how to make the simple "modus ponens" deduction from a conditional illustrated above. But, in contrast, they do not know the basic inferential tactics for Sudoku, and so they have to learn them. We now turn to our theory of the psychology of these puzzles.

THE PSYCHOLOGY OF SUDOKU PUZZLES

The theory presented in this section postulates that when individuals first encounter a Sudoku puzzle, without having had any experience with them or instruction on how to solve them, they do not know how to proceed at first. They may try to guess a digit in an empty cell, or think about the set of possible digits that could occur in it. But the first step in their mastery of the puzzles is to acquire a repertoire of simple tactics that can be used in a basic strategy.

A *simple* tactic, by our definition, is one that starts with definite digits in several cells and that enables individuals to deduce from them a definite digit that occurs in another cell. For example, if a row in a puzzle contains the digits 1, 2, 3, 4, 5, 6, 7, 8, and there is one empty cell in the row, it follows that the digit in the empty cell must be 9, because the rule governing all puzzles states that every digit from 1 to 9 must occur exactly once in each row (see above). The inference is a valid deduction, and its premises are the state of the array in the puzzle, and the general rule. This tactic is an instance of one that applies to any *set*, i.e., row, column, or 3-by-3 box: if any *set* contains eight digits, then the empty cell in the set has the missing ninth digit. The tactic uses just a single constraint, i.e., individuals need to consider only one set to deduce the value of the target cell. The tactic is simple and obvious, but alas it cannot be used at the start of any real Sudoku puzzle, because the initial state of the puzzle never has eight digits in the same set (see, e.g., Figure 1). An *advanced* tactic, by our definition, is one that eliminates digits from the set of prior possibilities for one or more cells, and we return to the nature of these tactics later in the paper.

The present theory postulates that there are seven distinct sorts of simple tactic that individuals acquire, but we emphasise here that they are not forced to use any of these tactics. We have devised a computer program that solves Sudoku puzzles without using any of them, and we describe its method of solution later. Hence, the postulate of seven tactics is a psychological claim about how naïve individuals first begin to cope with the puzzles. The tactics are all valid deductions, and they differ on two principal dimensions. The first dimension is binary and concerns the nature of the underlying deduction. Consider the tactic in which eight digits in a row exclude these digits from a target cell, and so the remaining ninth digit must occur there. The tactic is a case of *exclusion* (of the eight digits). In other words, in an exclusion tactic digits already in the array are excluded from a particular cell, which must therefore have the remaining digit, for example:

1 through 8 are already in a set containing one empty cell.
Therefore, they can be excluded from the one empty cell in the set.
Therefore, the digit in the empty cell is the remaining digit of 9

In contrast, consider a tactic in which, say, a 3-by-3 box does not contain a 9, but 9s do occur in two rows intersecting the box. The remaining row in the box must *include* a 9, and so if there is only one empty cell in this row, it must be 9. The tactic is a case of *inclusion* (of the 9). In other words, in an inclusion tactic, a digit that already occurs one or more times in the array must be included in a particular cell, for example:

> 9 occurs in two rows intersecting a box.
> Only one cell in the box is empty in the remaining row.
> Therefore, the digit in the empty cell is 9.

The second dimension on which tactics differ is the number of constraints on the unknown target digit. The simple tactic that we illustrated first uses only a single constraint: the digits in a single set imply the value of the target digit. Another simple tactic uses the distribution of eight digits over two different but intersecting sets, such as a row and a column, to determine the digit in the cell at the intersection. This tactic accordingly uses two distinct constraints: the row and the column. The number of constraints in a relation determining the value of variable is known as its ''relational complexity'' (Halford, Wilson, & Phillips, 1998). Studies have shown that relational complexity affects the difficulty of deductions (Birney, Halford, & Andrews, 2006; Goodwin & Johnson-Laird, 2005). Hence, the greater the relational complexity of a simple tactic, the harder it should be to use: it should take longer, and it should be more likely to lead to error.

Figure 2 illustrates the seven sorts of simple tactics, both exclusion and inclusion tactics, and they are laid out according to their relational complexity. Each tactic is an instance of a valid deduction, but naïve individuals who have never encountered a Sudoku puzzle before will have to discover these tactics for themselves if they are to master the puzzles.

The present theory postulates that relational complexity depends on the number of *sets* that have to be taken into account in order to deduce the digit in a cell. An alternative hypothesis defines relational complexity in terms of the number of *digits* that have to be taken into account in using a tactic. For inclusion tactics, the two hypotheses make opposite predictions because, as Figure 2 shows, the numbers of digits to be taken into account in an inclusion tactic (including those within a box) decline with an increase in relational complexity based instead on sets. We examine this divergence in Experiment 2. The minimum relational complexity in terms of sets is one, as when eight digits in a set can be excluded as digits for the ninth cell. The maximum relational complexity is five, as when a digit's occurrence in two rows and two columns, and the contents of the 3-by-3 box itself, imply that the digit is included in a particular empty cell in the box.

Anyone who is experienced in solving Sudoku puzzles knows that they differ in difficulty, and that they are usually classified at four levels: *easy*, *mild*, *difficult*, and *fiendish* (see, e.g., the compilations of puzzles edited by Gould, 2005). For skilled solvers, easy puzzles take about 10 minutes to complete, whereas fiendish puzzles take over an hour. The present theory

Exclusion tactics

Sufficient digits are distributed over one, two, or three *sets* to exclude all but a definite digit from the target cell. A set is a row, column, or the box, and the number of sets equals the relational complexity of the tactic.

Inclusion tactics

All the other empty cells in the box containing the target cell are in at least one of a number of sets (either rows or columns) in which one possible digit for the target cell occurs. The box itself must contain the possible digit somewhere (which functions as a separate constraint), and so it can be inferred that this digit must occur in the target cell. The relational complexity of the tactic is thus equal to the number of sets (rows or columns) that functions as constraints, plus the box constraint. An X in a cell indicates that the cell is occupied by another number. The value of this number is inconsequential, so the X's do not count as a separate constraint.

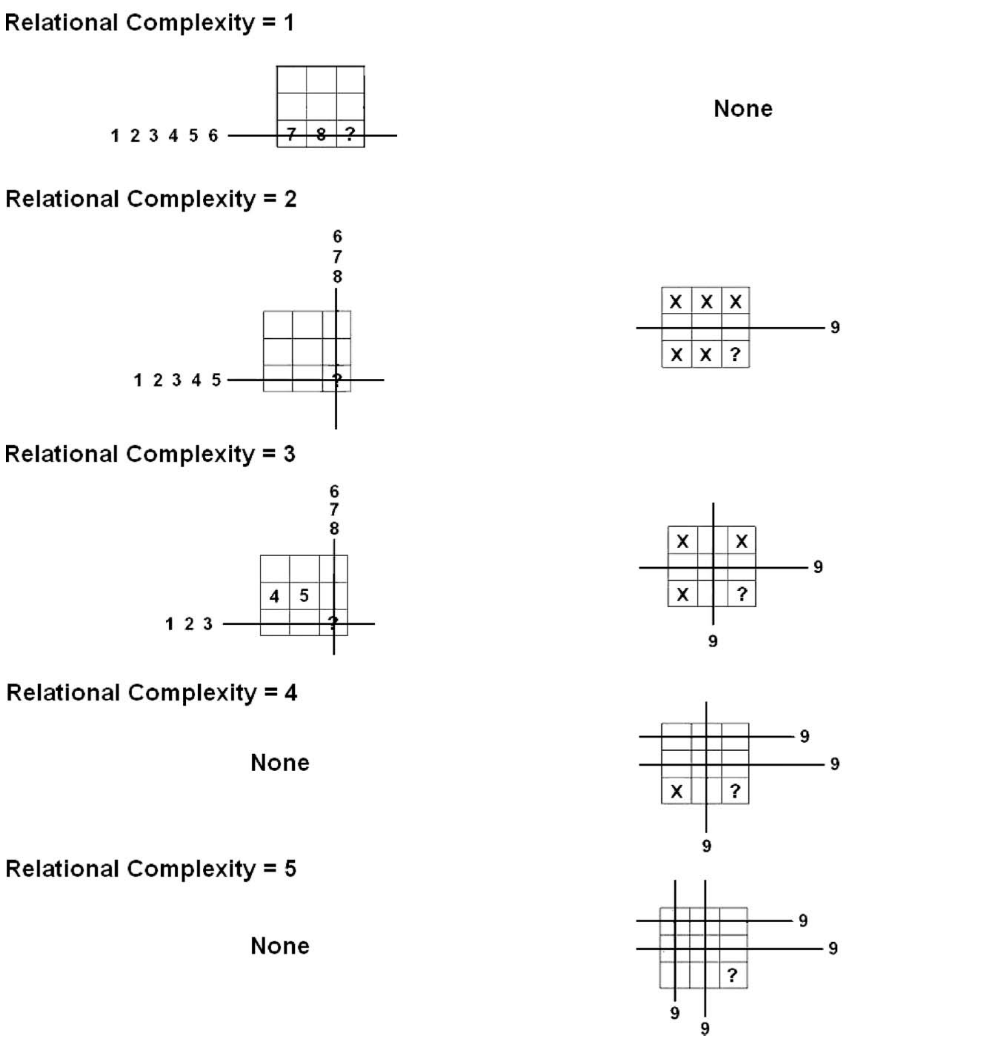

Figure 2. Illustrations of the seven simple tactics for solving easy Sudoku puzzles. The value of "?" is always 9 for each tactic. Exclusion tactics directly exclude possible digits from a target cell so that it can only contain one specific digit. Inclusion tactics use the occurrence of a digit in other cells in a set to infer that it must be included in the target cell. Each line in a diagram with digits at its end signifies ruling out the occurrence of these digits in the empty cells through which the line runs, because these numbers are already in the relevant set.

accordingly needs to explain, first, what aspects of the puzzles contribute to their difficulty and, second, how experience enables skilled individuals to solve the more difficult puzzles. Readers might suppose that the difference in difficulty over the four sorts of puzzle is merely a matter of how many digits are missing from the initial array of a puzzle. In fact, this number does not vary much. The three puzzles in Figure 1, for example, differ vastly in difficulty but have the same number of missing digits. Hence, the number of empty cells in the initial array is not crucial. Likewise, as we will show, the initial steps in the four sorts of puzzle also do not differ reliably in difficulty. So, what is responsible for difficulty?

The answer according to the present theory depends on strategy and tactics: simple tactics suffice for the easier puzzles, but fail utterly with more difficult puzzles. The claim is true, but too coarse to elucidate four levels of difficulty. A more refined analysis calls for the concept of a *stage* in the solution of a puzzle. The initial stage is the array presenting the puzzle. At this stage it is possible to deduce the particular digits in a certain number of cells in the array depending solely on the digits in the initial array, and it is also possible to eliminate possible digits for other cells using advanced tactics, again depending solely on the digits in the initial array. When all of these deductions, which are independent of one another, have updated the array, the puzzle is at its second stage: some cells contain definite digits, and the remainder contain sets of possible digits. Once again, a new set of independent deductions can be made to infer definite digits for cells or to reduce the possible digits for cells. They yield the next stage of the puzzle, and so on . . . all the way to the solution of the puzzle in which all the cells have definite digits. The concept of a stage in a puzzle therefore has a precise logical definition, and the computer program that we devised allows us to examine each stage in the solution of a puzzle.

Individuals are most unlikely to use a strategy that proceeds in an orderly fashion, moving from one complete stage to the next. Indeed, the present theory postulates that because tyros start with easy or mild puzzles, they should soon acquire some simple tactics. Their basic strategy should be to infer a definite digit, and then to check whether it in turn enables them to use another simple tactic to infer a further definite digit, and so on. As soon as they can go no further in this way, they scan the array looking for other configurations of digits that allow them to use a simple tactic. If this hypothesis is correct, then easier puzzles should have on average a higher number of definite digits that can be deduced at any stage in their solution. We used our computer program to test this prediction. The program uses a recursive strategy based on a single advanced tactic, i.e., it works through an array, recursively using each definite digit in a cell to eliminate possibilities for all the cells in the same sets as this cell. We took samples of 10 puzzles at each of the four levels of difficulty from Gould (2005). The program revealed the following mean numbers of definite digits that could be inferred at each stage in the solution of the puzzles: 4.3 for the easy puzzles, 2.3 for the mild puzzles, 1.8 for the difficult puzzles, and 1.3 for the fiendish puzzles (Jonckheere trend test, $z = 6.51$, $p \ll .00001$). A corollary is that the easier a puzzle is, the fewer the number of stages in its solution. The psychological implication of our computational result is that human solvers can infer a definite digit at any point more readily in the easier puzzles than in the harder puzzles, and this factor provides an explanation of the cause of difficulty.

Easy puzzles can be solved using only simple tactics, as can many mild puzzles. But simple tactics alone fail to cope with difficult and fiendish puzzles. The mean of 1.3 definite digits per stage for fiendish puzzles, which we cited above, is just a mean, and there are many stages in the solution of such a puzzle where all that can be inferred are eliminations from the *possible* digits in cells. A crucial shift in strategy is therefore necessary to solve the harder Sudoku (the difficult and fiendish ones): individuals have to keep a record of the possible digits in cells, and to use *advanced* tactics that enable them to eliminate possible digits. The progressive elimination of possibilities ultimately leads to only a single possibility for a cell, but this final inference may not be feasible until individuals have passed through a large number of stages. We reiterate that the use of advanced tactics become feasible only when individuals realise that they need to record the possible digits in cells. We return to the nature of advanced tactics later in the paper, but first we need to consider the evidence supporting the theory's account of simple tactics.

EXPERIMENT 1

Our first experiment was designed to examine the initial performance of naïve individuals encountering their very first Sudoku puzzles. It tested whether they were able to discover simple tactics, whether they relied on them to a greater extent than advanced tactics, and whether any difference in difficulty occurred at the outset amongst mild, difficult, and fiendish puzzles. The theory predicts that naïve individuals should develop simple tactics before they discover the need to develop advanced tactics, which depend on maintaining a record of the possible digits in cells. It also predicts that once individuals have inferred a definite digit they should follow up its consequences, if possible, in another inference.

Method and participants

Ten Chinese University of Hong Kong students (mean age = 20.4 years) acted as their own controls; on their own account, none of them had any prior experience with Sudoku puzzles. They tackled three puzzles: a mild, a difficult, and a fiendish one (selected from Gould, 2005), spending 15 minutes on each of them. Figure 1 above presents examples of the puzzles used in the experiment. They were presented in random orders, and the participants' task was to fill in as many missing digits as possible. After the participants filled in a digit, they had to write a justification explaining why the digit was correct for that cell. In this way we were able to assess whether the participants had become aware of a tactic, and the sequence of justifications revealed the order in which they had filled in the digits while tackling each puzzle. The participants were told the general rule for Sudoku (see the Introduction), but otherwise received no instructions whatsoever on how to go about solving them.

Results and discussion

The participants correctly solved a mean of 2.2 digits for the mild puzzles, 2.1 for the difficult puzzles, and 2.0 for the fiendish puzzles, but this slight trend was not reliable (Page's $L = 123.0$, $p = .25$, *ns*). The main differences in difficulty among the three sorts of puzzles evidently occur after their early stages. To solve only two digits in 15 minutes shows that the puzzles were difficult for the naïve participants, because experienced solvers can solve the whole of an easy puzzle in that time.

As the theory predicts, the participants overwhelmingly used simple tactics (83% of their inferences; Wilcoxon test, $z = 2.97$, $p < .005$). Although we had not predicted the phenomenon, the participants were also more likely to use exclusion than inclusion tactics (on 85% of occasions; Wilcoxon test, $z = 1.89$, $p < .05$). One reason for this bias may be that exclusion tactics are easier than inclusion tactics. In an exclusion tactic, reasoners can start with a putative target cell, and check what digits occur in the same row, column, and box as this cell's. But with an inclusion tactic reasoners have to look for occurrences of the same digit in rows or columns that are appropriately aligned with a box which just happens to have only one free cell that can include the digit. This task is likely to be demanding, at least until the array has four instances of the same digit aligned on a box.

The mean relational complexity of the participants' simple tactical steps was 2.8, but their numbers at different levels of complexity were insufficient for a test of its effects on their performance. As the theory predicts, their strategy followed the principle of following up an inference of a digit with a subsequent tactic that exploited it. When they had a choice between using this strategy and making an independent inference, they tended to use this strategy (87% of occasions; Wilcoxon test, $z = 2.37$, $p < .01$).

The participants' written justifications for their inferences of digits contained cases of six of the seven simple tactics in Figure 2: no one used the inclusion tactic with a relational complexity of five. The following protocols illustrate the use of three different simple tactics:

1. Exclusion tactic with relational complexity of 3 (Participant 5, fiendish puzzle): "1, 2, 3, 9 have already occurred in the row [of the target cell], 4, 5, 8 have already occurred in the column [of the target cell], 7 has already occurred in the 3-by-3 box [of the target cell], only 6 can be filled in [the target cell]."
2. Inclusion tactic with relational complexity of 3 (Participant 9, mild puzzle): "The 3-by-3 box does not contain 3, and the other four empty cells cannot contain 3 [because 3 has already occurred in the same two rows as theirs]". The participant then put a 3 in the appropriate target cell.
3. Inclusion tactic with relational complexity of 4 (Participant 2, mild puzzle): "4 has already occurred in the two rows above, and also cannot be placed in the cell to the left [of the target cell], because 4 has already occurred [in the same column] in the 3-by-3 box below [the cell to the left of the target cell]." The participant then put a 4 in the appropriate target cell.

These protocols are typical, and they show that the participants were well aware of the tactics that they used to infer correct digits for empty cells. Sceptics might suppose that individuals are bound to acquire the simple tactics, because they are necessary to solve the puzzles. In fact, none of the simple tactics postulated in the theory is necessary for the solution of puzzles. Our computer program uses none of them, but instead a general tactic of eliminating possibilities (see the earlier account).

The participants differed considerably in their accuracy. One participant made no errors whereas another, the least competent, made errors on three-quarters of the digits that he inferred. Overall 28% of the assignments of digits to cells were erroneous. Some errors may have been guesses constrained by the possible values for a cell. Others were unconstrained in this way and included mistakes, such as a failure to bear in mind a digit that was already in a set. No reliable difference occurred between the two sorts of errors. We conclude that with experience the proportion of errors declines, along with any tendency to guess.

EXPERIMENT 2

The relational complexity of a simple tactic should directly influence its difficulty. Previous studies of the effects of relational complexity on deduction have examined only small values, because in everyday language relations seldom hold over two or three arguments (see, e.g., Goodwin &

Johnson-Laird, 2005; Halford et al., 1998). In contrast, the simple tactics for Sudoku vary in relational complexity from one to five, and so they allow us to investigate a much greater range of values of relational complexity than has typically been studied. Experiment 2 tested the prediction that tactics should increase in difficulty with their relational complexity, and it focused on inclusion tactics so that we could contrast predictions based on number of digits with those based on number of sets. As Figure 2 shows, the numbers of digits to be taken into account in an inclusion tactic decline with an increase in relational complexity defined instead on sets.

Method and participants

We tested 18 Princeton students (mean age = 26.5 years), who acted as their own controls; none of them reported any prior experience with Sudoku puzzles. Their task was to infer the value of the digit in a given cell in each of nine problems. There were two problems depending on inclusion tactics with relational complexity values of two, three, four, and five, and one trial for an exclusion tactic with a relational complexity value of two. Figure 3 illustrates examples of the five sorts of problems. Each participant received the total of nine problems in a different random order, and they had 4 minutes to solve each problem. The problems were presented on separate print-out sheets, and the participants filled in their answers on the sheets.

Results and discussion

Table 1 presents the percentages of problems that the participants solved, and the overall mean latencies for both correct and incorrect solutions (because of the small percentages of correct solutions in some conditions). Those participants who failed to give any answer within the time limit (24% of responses) were assigned the latency of the time limit. The trend for the inclusion tactics across the four relational complexity values was reliable for latency (Page's $L = 478.5$, $z = 2.3$, $p < .01$) but not for accuracy, although it was in the predicted direction (Page's $L = 455.0$, $z = 0.4$, $p = .34$, *ns*). The departure from the effects of relational complexity when its value was five is probably attributable to the ease of noticing four instances of the same digit in columns and rows intersecting the box containing the target cell. Participants performed better on the exclusion tactic problem, which had a relational complexity of two, than the two inclusion tactics problems with a relational complexity of two, and the difference was reliable for latency (Wilcoxon test, $z = 3.2$, $p < .05$) but not for accuracy (Wilcoxon test, $z = 1.3$, $p = .18$, *ns*). The participants' performance also improved over the trials, as reflected by the

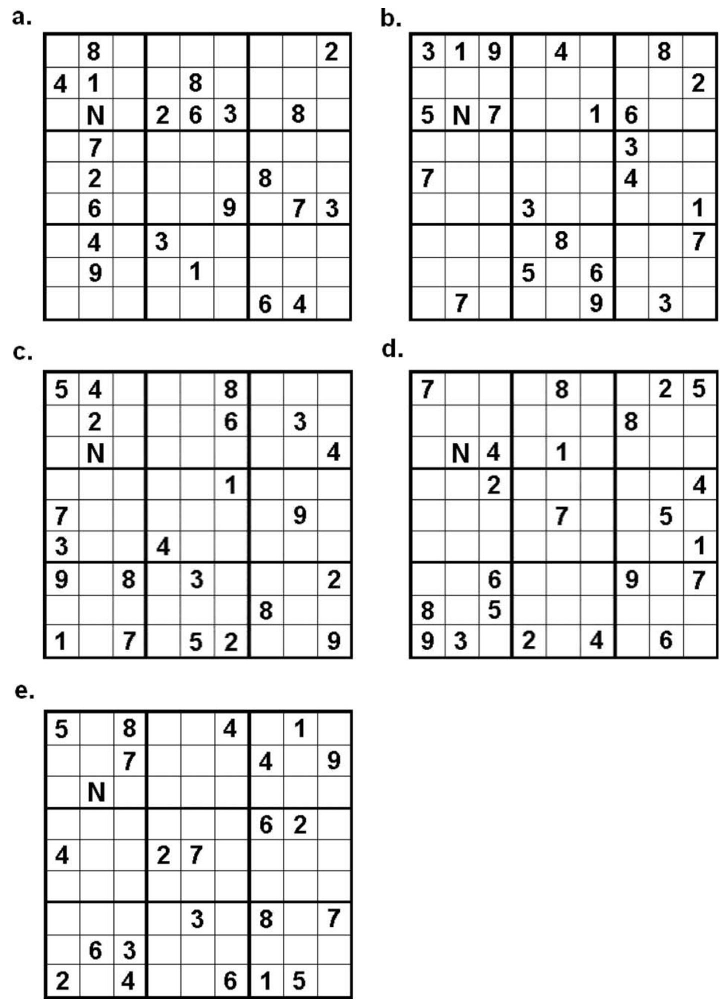

Figure 3. Examples of problems in Experiment 2 in which the task was to infer the digit for the cell marked "N". (a) N = 5, relational complexity of two, (b) N = 2, relational complexity of two, (c) N = 7, relational complexity of three, (d) N = 8, relational complexity of four, (e) N = 4, relational complexity of five.

TABLE 1
Percentages of correct responses and mean latencies of all responses, correct and incorrect, Experiment 2

	Exclusion tactic; relational complexity of two	*Inclusion tactic; relational complexity of two*	*Inclusion tactic; relational complexity of three*	*Inclusion tactic; relational complexity of four*	*Inclusion tactic; relational complexity of five*
Percentages correct	83	67	64	50	58
Overall latencies (s)	47.4	122.1	158.6	170.5	165.0

increase in both latency and accuracy (Page's L $= 4338.0$, $z = 3.2$, $p \ll .001$, and Page's L $= 4330.0$, $z = 3.1$, $p \ll .001$ respectively). Finally, the participants failed to solve 38% of the problems, but in most cases (84% of these responses) they did not infer any digit rather than inferring an erroneous digit (Wilcoxon test, $z = 2.35$, $p < .05$), thereby corroborating Experiment 1's finding that individuals were not inclined to make guesses.

THE NATURE OF ADVANCED TACTICS

Simple tactics alone cannot solve difficult or fiendish puzzles, which require advanced tactics embedded in an advanced strategy in which participants keep a record of the sets of possible digits for cells. Hence, when individuals move on to more difficult puzzles they have to learn to use this strategy. Their impetus to make the shift in strategy is likely to come from their discovery that they can make no further progress in a difficult problem by relying solely on simple tactics. However we have no evidence on this point and, as readers should realise, its collection would call for an empirical study that would make unrealistic demands on participants' time.

Advanced tactics can be analysed as two-step processes. The first step is to infer a set of digits as the only possibilities for certain cells, and the second step is to use these possibilities to eliminate possibilities from other cells. Simple exclusion and inclusion tactics can be extended into advanced tactics that yield not a definite digit, but a set of possibilities. For example, sufficient digits might be distributed over sets to exclude all but *two* digits for a target cell. Likewise, the configuration of possibilities may call for two target cells to each include a pair of possible digits. These extensions of exclusion and inclusion tactics to cases yielding possibilities rather than definite digits provide the first step in advanced tactics. The second step is to use these possibilities to winnow those in other cells.

Figure 4 illustrates two advanced tactics, which we investigated empirically. They are probably among the earliest advanced tactics that individuals are likely to discover. Their first steps do indeed depend on the extensions to exclusion tactics, which we mentioned earlier, or on the prior use of advanced tactics themselves (in a recursive way). Advanced tactic 1 in Figure 4 has a first step yielding the same n digits as the only possibilities for n cells in a set, where $n = 2$ in this case, and its second step is to eliminate these same possibilities from any other cells in the set. For instance, as the figure shows, if two cells in the same row of a box can contain only 5 or 6, then these two digits can be eliminated from any other cell in the box. Advanced tactic 2 in Figure 4 consists of a first step yielding a particular digit that is possible in only two cells in a row or column within a box, and nowhere else in the box, and its second step is to eliminate this digit from any cell in the same row or column outside the box. For instance, as

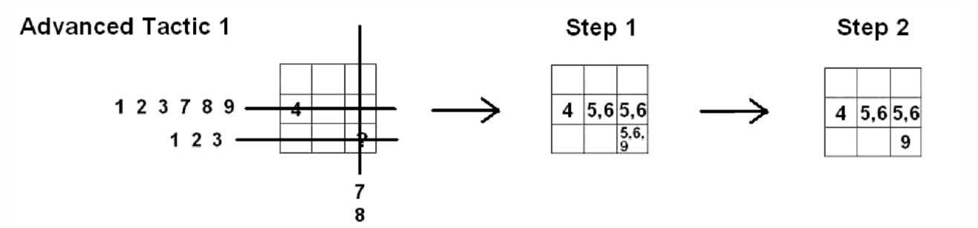

Step 1: Use exclusion tactic to eliminate all but n possibilities in a set by considering the occurrence of digits in other sets.

Step 2: Exclude these n values from any cell in the same set (where minimum value of $n = 2$).

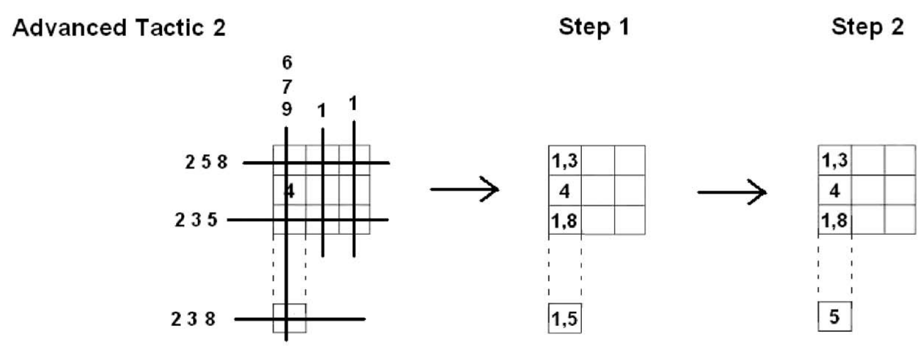

Step 1: Use exclusion tactics to yield two cells in the same set that are the only possibilities for a particular digit. This step, as shown, may also eliminate possibilities in a cell outside the box in the same set as the two cells.

Step 2: Exclude this digit from any other cell in the same sets.

Figure 4. An illustration of two advanced tactics. As in Figure 2, a line rules out the occurrence of the numbers at its end in the empty cells through which it strikes, because the numbers have occurred in the set already.

the figure shows, the digit 1 can occur only in the top or bottom cell in the left-hand column of the box, and so it can be eliminated from the cell in the same column below the box. It can also be eliminated from any other cells in this column outside the box that contain it as a possibility.

These advanced tactics yield valid conclusions about possibilities, and so their logical analysis calls for a modal predicate logic, which has a greater expressive power than standard predicate logic (see, e.g., Hughes & Cresswell, 1996). Psychological theories based on formal rules of inference do not as yet contain accounts of modal predicate logic, but the general principles of these theories stipulate that inferences depend on the translation of all premises into a formal language (see, e.g., Braine & O'Brien, 1998; Rips, 1994). However, the translation of the current state of the array into a single one-dimensional linguistic expression is likely to render proofs corresponding to advanced tactics rather cumbersome. It is an interesting intellectual challenge to formulate the appropriate rules of inference or axioms needed for such proofs. An array of digits appears to be a much more efficient representation—it acts as an external model, because it makes the constraints of row, column, and digit, easy to perceive (see Bauer & Johnson-Laird, 1993, on what makes a diagram helpful in reasoning).

Two factors should affect the difficulty of advanced tactics. The first factor is relational complexity, because it affects the ease of deducing the possible digits in the first step of advanced tactics. The second factor is more subtle. In an advanced tactic, individuals deduce that certain digits, which are possible for some cells only, cannot occur in any other cells in a set. A tactic should be easier to use when the number of possible digits is the same as the number of cells for which they are possible. For example, it is easier to detect that a pair of digits are the only possibilities for a pair of cells (as in the first advanced tactic in Figure 4), than to detect that one digit is among the varied possibilities in a pair of cells (as in the second advanced tactic). The aim of our final experiment was to test both this prediction and the prediction about the relational complexity of advanced tactics.

EXPERIMENT 3

In order to test the theory's predictions about advanced tactics, Experiment 3 presented problems in two ways. One group of participants tackled problems in which the cells contained lists of their possible digits (the ‘‘possibilities’’ group), and one group tackled the problems without such information (the ‘‘blank’’ group). The experiment examined two different instances of the first advanced tactic that differed in their relational complexity, and one instance of the second advanced tactic. In the blank group there should be a direct trend in difficulty over the three sorts of problems, whereas in the possibilities group the effect of relational complexity in the first step cannot occur because the results of the first step are already displayed by the possibilities in the array, and so the two instances of the first tactic should not differ reliably but be easier than the second tactic. In general, the problems in the possibilities group should be easier than in the blank group, because its participants first have to infer the possibilities.

Method and participants

The participants were 20 Chinese University of Hong Kong students (mean age = 21.5 years); none of the participants reported any prior experience with Sudoku puzzles. They were assigned at random to either the possibilities group or the blank group. The procedure was the same as in Experiment 2. The participants tackled a total of eight problems, two different instances of each of the three advanced tactics, and two filler problems calling for the use of an inclusion tactic with a relational complexity of four. Figure 5 presents three examples of the problems used in the experiment.

Results and discussion

Table 2 presents the percentages of correct responses and the mean latencies for all responses, correct and incorrect. The possibilities group solved more problems than the blank group (Mann-Whitney $U = 25.0$, $z = 1.93$, $p < .05$), and they responded faster ($U = 11.0$, $z = 2.95$, $p < .005$). As the theory predicts, the first advanced tactic was easier the second advanced tactic, as shown in both accuracy (Wilcoxon's test, $z = 1.96$, $p < .05$) and latency (Wilcoxon, $z = 2.32$, $p < .02$). Likewise, as predicted, this difference held for the possibilities group, both for accuracy (Wilcoxon test, $z = 2.3$, $p < .02$) and latency (Wilcoxon test, $z = 2.26$, $p < .01$). In the blank group the data showed a trend in the predicted direction, but it was not reliable for accuracy (Page's $L = 123.0$, $z = 0.67$, $p = .25$, *ns*) or for latency (Page's $L = 124.5$, $z = 1.01$, *ns*). The effect of relational complexity within the first advanced tactic was reliable for the blank group, as predicted, for both accuracy (Wilcoxon test, $z = 1.82$, $p < .05$) and for latency (Wilcoxon test, $z = 2.19$, $p < .02$).

An unintended consequence of the problems demonstrated the greater difficulty of the second advanced tactic over the first. It was possible for the participants to solve all the problems (apart from one investigating the first tactic with the higher relational complexity) by instead carrying out a sequence of three or more simple tactics. However, the participants' protocols showed that they had a preference to rely on the first advanced tactic rather than the sequence of simple tactics (16 participants did so, 1 relied on simple tactics, and 3 were ties, binomial test, $p < .0001$), whereas the participants' preferences split more evenly for the second advanced tactic (6 were biased in favour of the advanced tactic, 5 used simple tactics, and there were 9 ties, binomial test, $p = .5$, *ns*). In other words, this tactic was sufficiently difficult that the participants spontaneously discovered a different sequential strategy for solving the problems.

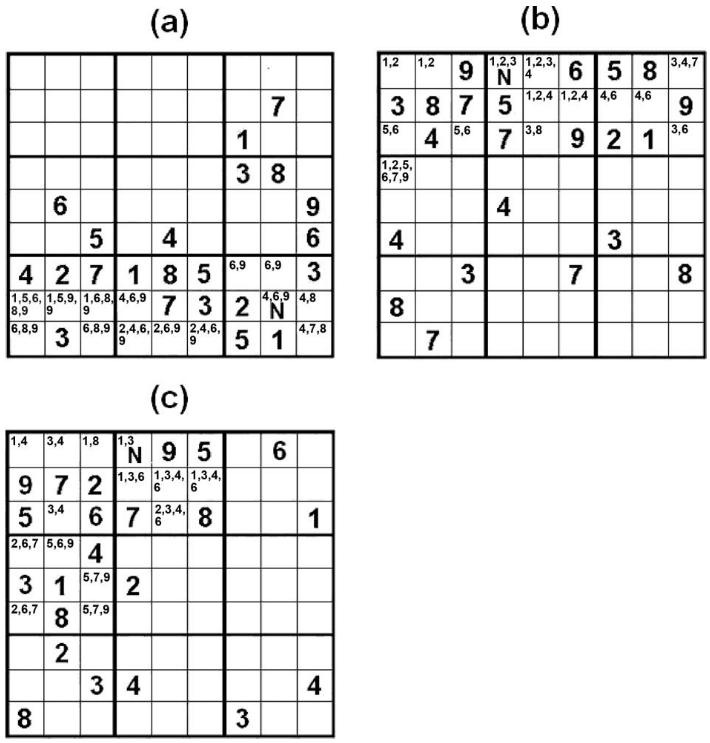

Figure 5. Three examples of problems in Experiment 3 shown for the "possibilities" condition, where the task is to infer the digit for the cell marked "N". (a) N = 4. The first step of the first advanced tactic in Figure 4 yields the result shown, which is that 6 and 9 must occur in the two hitherto empty cells of the seventh row; the second step eliminates them as possibilities for N, and so it must be 4. (b) N = 3. The first step of the first advanced tactic in Figure 4 yields the result shown, which is that 1 and 2 must occur in the two cells in the top left box; the second step eliminates them as possibilities for N, and so it must be 3. (c) N = 3. The first step of the second advanced tactic in Figure 4 yields the results shown for the top left box, which are that 1 and 4 must occur in its top left cell, and 1 and 8 must occur in its top right cell. The second step eliminates 1 from any cell in the same row as those two cells, including N, and so it must be 3. Other inferences can also be made from these arrays, although they are irrelevant to determining the value of N, e.g., in 3b the fifth cell on the third row can be inferred to contain 8 alone.

GENERAL DISCUSSION

Sudoku puzzles, as we have argued, have the striking property that their solution depends solely on deduction. In terms of dual-process theories, which we described in the Introduction, the intuitions of System 1 may affect performance, but what are decisive are the deliberations of System 2. At each point in an error-free solution of a puzzle, an individual makes a valid deduction from the current state of the puzzle and the general rule that each digit from 1 through 9 must occur once in each row, column, and 3-by-3 box in the array. That is all it takes—pure deductive reasoning from premises: no

TABLE 2
Percentages of correct responses and mean latencies of all responses, correct and incorrect, for the two groups in Experiment 3

		First advanced tactic of lower relational complexity	*First advanced tactic of higher relational complexity*	*Second advanced tactic*
Possibilities group	Percentages correct	85	75	50
	Overall latencies (s)	91.7	111.3	184.8
Blank group	Percentages correct	70	30	25
	Overall latencies (s)	174.3	258.9	227.2

arithmetic (the puzzles could just as well be framed in terms of alphabetical letters or any nine distinctive tokens), no inductive or probabilistic inferences (*pace* Oaksford & Chater, 2002), no fast and frugal heuristics (*pace* Hertwig et al., 1997), no pragmatic schemas (*pace* Cheng & Holyoak, 1985), no content-specific modules of evolutionary psychology (*pace* Cosmides et al., 2005), and no use of the atmosphere effect (*pace* Wetherick & Gilhooly, 1995). The puzzles are remote from the ecological and social structure of environments (*pace* Hertwig et al., 1997). Yet individuals all over the world, both Westerners and East Asians, are able to make the deductions necessary to solve Sudoku puzzles, even though they have had no training in logic. In sum, the puzzles establish that logically naïve individuals have the competence to make deductions about abstract matters, and that they enjoy exercising this ability, contrary to all the theories above that impugn it. Recent evidence also implies that the teaching of mathematics can be more effective when it too focuses on abstract matters rather than on concrete everyday examples (Kaminski, Sloutsky, & Heckler, 2008).

The difference in difficulty over the four grades of Sudoku—easy, mild, difficult, and fiendish—was elucidated by our computer program for solving the puzzles. The program measured difficulty in the number of definite digits deducible at each stage in a puzzle (see the Introduction). The use of the array as an external model allows individuals to develop various tactics, and it corroborates the theory that naïve reasoners rely on the meaning of the general rule, just as they appear to rely on the meaning of premises in straightforward deductions.

The difference in difficulty also depends on the strategy and tactics that are needed to solve a puzzle. Simple tactics take as their starting point definite digits in cells and they deliver as a result a new digit as the value of a cell. Most people who tackle a puzzle receive some instruction about simple tactics, which is often stated along with an introductory puzzle (see, e.g., Gould, 2005). However, with no instruction whatsoever the participants in Experiment 1 discovered most of the simple tactics for themselves. These tactics differ in difficulty. They exploit a certain number of distinct constraints, i.e., digits in rows, columns, or the same box as a target cell (see Figure 2). As Experiment 2 showed, the greater the number of these constraints in terms of sets rather than digits, i.e., the greater the *relational complexity* (Birney et al., 2006), then in general the harder the tactic is to use—it takes longer, and individuals are more likely to make a mistake. Only when an inclusion tactic depends on four occurrences of the same digit in an array, aligned in rows and columns that intersect a given box, does the trend in difficulty appear to reverse.

Relational complexity, as Experiment 3 showed, also affected the difficulty of using advanced tactics. The first step is a deduction about sets of possibilities in cells, and the second step is a deduction that eliminates the possibilities from other cells. What matters is the match between the number of possible digits and the number of cells available to put them in. Hence, it is easier to see that the same two digits are the only possibilities for two cells than to see that a single digit occurs in various lists of possibilities. Another advanced tactic that individuals may use in solving fiendish puzzles depends on postulating that one possible digit is the solution to a cell, and then following up this assumption until it either yields the solution or reaches an inconsistency calling for an alternative assumption to be made about the digit in the starting cell.

Two phenomena from Sudoku puzzles are most revealing about human reasoning, and call for a revision to standard theories of deduction. First, naive individuals are able to acquire simple deductive tactics in order to solve easy and mild Sudoku puzzles. Experiment 1 showed that naïve participants acquired explicit knowledge of such tactics in a quarter of an hour's experience in tackling a puzzle. Few empirical studies of deduction have reported the acquisition of explicit deductive tactics—explicit in the sense that participants can describe them to someone else. An analogy would be Aristotle enunciating rules for syllogisms (see Kneale & Kneale, 1962, p. 75). But Aristotle is not a typical participant in a psychological experiment, although one study of syllogisms has reported that individuals do discover some incomplete rules for them during the course of an experiment (Galotti, Baron, & Sabini, 1986).

Second, as individuals move on from easier puzzles to more difficult ones (see, e.g., Gould, 2005), they are bound to discover that their current strategy of stringing together simple exclusion and inclusion tactics fails. Some of them may give up the puzzles at this point, but others persevere and make a strategic shift: they begin to use a strategy that depends on a whole new set of tactics. They no longer aim to make inferences yielding only definite digits, but instead their strategy now depends on keeping a record of the possible digits in a cell. This shift in strategy is analogous to shifting from proofs in the first-order predicate calculus to proofs in the first-order *modal* predicate calculus.

Neither of these phenomena can be explained by any current theory of deductive reasoning. We can discount probabilistic theories and content-specific theories for the reasons that we have described earlier. Theories based on formal rules have no machinery for acquiring inferential tactics, i.e., new rules of inference, or for shifting from one inferential strategy to another, and tend to postulate a single fixed and deterministic strategy (Braine & O'Brien, 1998; Rips, 1994). The original model theory fares little better (e.g., Johnson-Laird, 1983). However, the model theory has recently been expanded in order to account for individuals' ability to learn how to solve problems that come in series, such as so-called ‘‘matchstick problems’’ (Lee & Johnson-Laird, 2004). These problems call for a certain number of pieces—the matchsticks—to be removed from an initial array in order to eliminate a certain number of squares. A typical problem, for instance, is to remove three pieces from a 2-by-3 array of squares in order to eliminate three squares from the array. To master these problems, individuals have to acquire a set of seven distinct tactics, and they do so by deducing the consequences of various moves, which depend on the number and location of the pieces that they remove. After they have deduced these consequences they can then deliberately apply a tactic whenever it is appropriate. Individuals tackle their first problem using a strategy of trial and error constrained by the initial array and the statement of the given problem. They soon learn that a better strategy is based on the ratio of the current number of pieces to be removed to the current number of squares to be eliminated. They then use this ratio to constrain their choice of tactics. In principle, the same acquisition mechanism applies to tactics in Sudoku. A proper account of human deductive reasoning should accordingly allow for the acquisition of inferential tactics, and for shifts from one sort of deductive strategy to another. Such shifts have also been observed in a study in which individuals reasoned from sets of three or four verbal premises (Van der Henst et al., 2002).

In conclusion, the solution to the puzzle of Sudoku yields an insight into human competence that is in stark contrast to many psychological theories: reasoners readily acquire the ability to make deductions about abstract contents, which are far removed from the exigencies of daily life and from the environment of our evolutionary ancestors. If they could reason only about this environment, then it is hard to see how they would ever have developed mathematics and logic, or science and technology. As Piaget recognised, our ability to make deductions about abstract matters remote from our mundane life is a fundamental human characteristic, and one that is essential to intellectual progress. Sudoku puzzles tap into the fun that we have in exercising this ability.

Manuscript received 19 November 2007
Revised manuscript received 12 May 2008
First published online 21 August 2008

REFERENCES

Ball, W. W. R., & Coxeter, H. S. M. (1987). *Mathematical recreations and essays*. New York: Dover.

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

Birney, D. P., Halford, G. S., & Andrews, G. (2006). Measuring the influence of complexity on relational reasoning: The development of the Latin Square Task. *Educational and Psychological Measurement*, *66*, 146–171.

Braine, M. D. S., & O'Brien, D. P. (Eds.). (1998). *Mental logic*. Mahwah, NJ: Lawrence Erlbaum Associates Inc.

Cheng, P. W., & Holyoak, K. J. (1985). Pragmatic reasoning schemas. *Cognitive Psychology*, *17*, 391–416.

Cosmides, L., Tooby, J., Fiddick, L., & Bryant, G. (2005). Detecting cheaters. *Trends in Cognitive Sciences*, *9*, 505–506.

Euler, L. (1849). Recherches sur une espèce de carrés magiques. *Commentationes Arithmeticae Collectae*, *II*, 301–361.

Evans, J. St. B. T. (2002). Logic and human reasoning: An assessment of the deduction paradigm. *Psychological Bulletin*, *128*, 978–996.

Evans, J. St. B. T. (2007). *Hypothetical thinking: Dual processes in reasoning and judgement*. Hove, UK: Psychology Press.

Galotti, K. M., Baron, J., & Sabini, J. P. (1986). Individual differences in syllogistic reasoning: Deduction rules or mental models? *Journal of Experimental Psychology: General*, *115*, 16–25.

Goodwin, G. P., & Johnson-Laird, P. N. (2005). Reasoning about relations. *Psychological Review*, *112*, 468–493.

Gould, W. (2005). *The Times Sudoku: The number-placing puzzle, Books 1, 2, and 3*. London: HarperCollins.

Halford, G. S., Wilson, W. H., & Phillips, S. (1998). Processing capacity defined by relational complexity: Implications for comparative, developmental, and cognitive psychology. *Behavioral and Brain Sciences*, *21*, 803–831.

Hertwig, R., Ortmann, A., & Gigerenzer, G. (1997). Deductive competence: A desert devoid of content and context. *Current Psychology of Cognition*, *16*, 102–107.

Hughes, G. E., & Cresswell, M. J. (1996). *A new introduction to modal logic*. London: Routledge.

Inhelder, B., & Piaget, J. (1958). *The growth of logical thinking from childhood to adolescence: An essay on the construction of formal operational structures*. New York: Basic Books.

Jeffrey, R. (1981). *Formal logic: Its scope and limits* (2nd ed.). New York: 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. (2006). *How we reason*. Oxford, UK: Oxford University Press.

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

Johnson-Laird, P. N., & Byrne, R. M. J. (2002). Conditionals: A theory of meaning, pragmatics, and inference. *Psychological Review*, *109*, 646–678.

Kaminski, J. A., Sloutsky, V. M., & Heckler, A. F. (2008). The advantage of abstract examples in learning math. *Science*, *320*, 454–455.

Kneale, W., & Kneale, M. (1962). *The development of logic*. Oxford, UK: Clarendon Press.

Lee, N. Y. L., & Johnson-Laird, P. N. (2004). Creative strategies in problem solving. In K. Forbus, D. Gentner, & T. Regier (Eds.), *Proceedings of the Twenty-Sixth Annual Conference of the Cognitive Science Society* (pp. 813–818). Mahwah, NJ: Lawrence Erlbaum Associates Inc.

Oaksford, M., & Chater, N. (2002). The probabilistic approach to human reasoning. *Trends in Cognitive Sciences*, *5*, 349–357.

Rips, L. J. (1994). *The psychology of proof*. Cambridge, MA: MIT Press.

Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, *119*, 3–22.

Stanovich, K. (2008). *The psychology of rational thought: What intelligence tests miss*. New Haven, CT: Yale University Press.

Van der Henst, J-B., Yang, Y., & Johnson-Laird, P. N. (2002). Strategies in sentential reasoning. *Cognitive Science*, *26*, 425–468.

Wetherick, N. E., & Gilhooly, K. J. (1995). "Atmosphere", matching, and logic in syllogistic reasoning. *Current Psychology: Developmental, Learning, Personality, Social*, *14*, 169–178.

Yato, T., & Seta, T. (2002). *Complexity and completeness of finding another solution and its application to puzzles*. Available at: http://www.phil.uu.nl/~oostrom/cki20/02-03/japansepuzzles/ASP.pdf