Title

The Limits of Supposing: Semantic Illusions and Conditional Probability

Permalink

https://escholarship.org/uc/item/6rs7s689

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 27(27)

ISSN

1069-7977

Authors

Elqayam, Shira
Evans, Jonathan St.B.T.
Handley, Simon J.

Publication Date

2005

Peer reviewed

The Limits of Supposing: Semantic Illusions and Conditional Probability

Shira Elqayam (selqayam@plymouth.ac.uk)
Simon J. Handley (shandley@plymouth.ac.uk)
Jonathan St.B.T. Evans (jevans@plymouth.ac.uk)
Centre for Thinking and Language, University of Plymouth
Plymouth PL4 8AA, UK

Abstract

Faced with extreme demands, hypothetical thinking runs the danger of total failure. Paradoxical propositions such as the LIAR ('I am lying') provide an opportunity to test it to its limits. Embedded in conditionals, they tended to occasion a breakdown of probabilistic inference (0% True-0% False pattern) demonstrating the vulnerability of hypothetical thinking when taxed by embedded suppositional processes. In contrast, items with the TRUTHTELLER ('I am telling the truth') were 'collapsed' to responses of conditional probability closely resembling estimates of control items.

Keywords: reasoning; conditionals; Liar paradox.

Introduction

Everyday assertions are not always simply true or false. They may refer to future events; they may refer to imaginary, mythological or fictional creatures or personae; they may be enigmatic or even paradoxical. In all these cases, we encounter indeterminate propositions – propositions to which we cannot assign a truth-value of either true or false. Whenever we are the recipients of an indeterminate assertion and have to make assumptions concerning its truth or falsity, we engage in hypothetical thinking. Consider, for instance, what happens if a politician and a journalist accuse each other of lying – you know one of them has to be lying, but you have to make an assumption in order to decide who does and who doesn't, follow up the possible consequences and decide if this assumption was justified.

The work we report here was designed to test some of the ways in which we reason when faced with indeterminate assertions and have to understand their hypothetical consequences. This sort of suppositional thinking of semantic concepts is mostly studied by metadeduction – the research paradigm studying reasoning with semantic concepts such as truth and falsity (e.g., Byrne & Handley, 1997; Byrne, Handley, & Johnson-Laird, 1995; Evans, 1990; Johnson-Laird & Byrne, 1990; Rips, 1989; Schroyens, Schaeken, & d'Ydewalle, 1999). Studies in metadeduction typically present participants with the island of knights and knaves, in which each of the inhabitants is either a knave, who tells nothing but lies, or a knight, who only tells the truth. Participants have to classify speakers as knights or knaves based on their utterances.

Consider the journalist from our example, writing in her column, 'What I am telling you right now is God's own truth'. Meta-deduction typically involves sentences of this sort. Some, like our example, are indeterminate because they can assume any truth value. You may assume that journalist was telling the truth; you may assume she was lying. In either case you will not run into contradiction. Sentences of this sort are called TRUTHTELLER-type sentences. But sentences that refer to their own truth status may also end up being paradoxical. This is the case with the infamous LIAR paradox: 'This sentence is false'. If you assume that the sentence is true, it turns out to be false, and vice versa. You end up with a contradiction each time.

Generally, sentences of the 'TRUTHTELLER' type ('I am telling the truth'; 'This sentence is true'; 'I am a knight') potentially evoke a semantic illusion that they are true rather than indeterminate, whereas sentences of the 'LIAR' type ('I am lying'; 'This sentence in false'; 'I am a knave') are resilient to this sort of semantic illusion, and are normally evaluated as indeterminate (Elqayam, in press). For instance, reasoners tend to evaluate 'I am a knave and snow is white' – a conjunction with a LIAR plus a true conjunct – as indeterminate (i.e., neither True nor False). In contrast, a matching conjunction with the TRUTHTELLER as the indeterminate constituent – 'I am a knight and snow is white' – is normally evaluated as true (i.e., determinate). In fact, such evaluation is often indistinguishable from evaluation of the matched item with a true constituent in place of the TRUTHTELLER (e.g., 'Sky is blue and snow is white'): the three-valued truth table produced by the indeterminate truth-value is thus 'collapsed' (Rescher, 1969) into a simpler, bivalent truth table containing just True and False as truth-values. This semantic illusion has thus been dubbed the 'collapse illusion' (Elqayam, in press).

The theory of hypothetical thinking (Evans, Over, & Handley, 2003) suggests that the LIAR and the TRUTHTELLER initially share a very similar process of representation (Elqayam, in press). The automatic, heuristic system 1 throws up the most relevant representation (*principle of relevance*) – this boils down to a single representation (*the principle of singularity*) of the surface form of the utterance. Hence, the TRUTHTELLER is represented as Knight(speaker), the LIAR as Knave(speaker). LIAR-type and TRUTHTELLER-type propositions critically diverge, though, when the evaluation process hits the *satisficing* phase mandated by the analytical system 2. For the TRUTHTELLER, the process ends then and there: the ensuing representation is satisficing and there is nothing to trigger further examination. For the LIAR, though, the inherent paradoxality of this representation violates satisficing, so reasoners are obliged to generate an

alternative model, i.e. Knight(speaker). This would in turn prove contradictory too, resulting in a vicious, tail-biting circle of supposition and contradiction, eventually leading up to the null model and an indeterminate evaluation – essentially a 'truth-value gap'.[1]

The LIAR is by no means unique in creating a truth-value gap in evaluation – the infamous 'defective' truth tables of the conditional (e.g., Evans, Newstead, & Byrne, 1993; Evans & Over, 2004) are another such case – that is, reasoners typically tend to evaluate *if p then q* as true in the p q case, false in the p not-q case, and irrelevant otherwise. For instance, 'If it rains I will take an umbrella' is typically evaluated as true in case it rains and I do take an umbrella, false in case it rains and I don't take an umbrella, and irrelevant if it does not rain at all.

The suppositional account of conditionals (Evans et al., 2004; Evans, Over, & Handley, 2005) invokes in this context the *Ramsey test* – Frank Ramsey's suggestion that when people evaluate the conditional *if p then q* they do so by adding *p* hypothetically to their stock of belief and evaluating *q* in this context (Ramsey, 1931). In other words: *if* triggers hypothetical thinking, a mental simulation that focuses on the antecedent (the 'p'). This focus on the antecedent leaves out of consideration cases in which the antecedent does not hold, thus creating a truth-value gap.

If the Ramsey test is viable as a psychological theory, then when participants are asked to evaluate the probability of a conditional, if p then q, we would expect them to respond with the conditional probability – the probability of q given p (abbreviated as q|p). This prediction has recently received abundant empirical support (e.g., Evans, Handley, & Over, 2003; Oberauer & Wilhelm, 2003).

Up until now conditionals have only been studied with determinate components; even the defective truth-tables were obtained for conditionals with perfectly innocuous, true-or-false materials: the truth-value gaps have been in the responses, not in the materials. But we have already seen that the suppositional process faces a unique challenge when LIAR-type constituents are involved. How, then, would reasoners respond if we place LIAR and TRUTHTELLER-type constituents in conditionals' antecedents? This is what the present paper endeavours to look into.[2]

As we saw, LIAR-type sentences produce, all by themselves, the intensive, 'hyper', vicious-circle hypothetical thinking described above. Moreover, *if* also produces hypothetical thinking. Combining *if* with LIAR-type antecedent produces, then, hypothetical thinking within hypothetical thinking, a simulation within simulation. This is bound to be a highly gruelling task, setting intolerable demands on working memory, and breaching tractability boundaries. A conditional with a LIAR-type antecedent tests hypothetical thinking to its limits.

How, then, could reasoners resolve this impossible tension? Our hypothesis is that they would tend to respond with a failure of hypothetical thinking or, in slightly more technical terms, the 'null model' (i.e., the 'vacuous' or 'empty' model): previous work demonstrates (Elqayam, in press) that this tends to be the most pervasive pattern. Reasoners typically evaluate conditionals with a LIAR-type antecedent (e.g., 'If I am a knave then snow is white') as neither true nor false (or, more precisely, judge them to be utterable by neither a knight nor a knave). The high proportion of such responses – about 70%-80% – clearly demonstrates the breakdown of hypothetical thinking faced with a paradoxical antecedent.

We still have to translate the null model to probabilistic response patterns. The best bet seems to be a response pattern of probabilistic inference breakdown. This would involve reasoners estimating the likelihood of the conditional being true as 0%, plus a high percent of '0%' ratings of the conditional being false – in effect, the null model. We call this the '0%-0%' hypothesis.

In contrast, for conditionals with TRUTHTELLER-type antecedents, we will expect a simple, straightforward conditional probability estimate: as TRUTHTELLER-type propositions are normally 'collapsed', that is, evaluated as if they were simply true, a conditional with a TRUTHTELLER-type antecedent should not be different from a conditional with any other antecedent. To illustrate, suppose we have the following conditional: *If I am a knight, then I live in Emerald City*; and suppose we know that there are 300 knights in Emerald City [pq], 100 knights in Sapphire City [p¬q], and 1000 knaves in each Emerald City [¬pq] and Sapphire City [¬p¬q]. The probability estimate we would expect in this case would simply be the conditional probability (.75 in this case), due to collapse. Whereas an equivalent conditional with a LIAR-type antecedent – *If I am a knave, then I live in Emerald City* – we would expect reasoners to estimate as 0%, regardless of the actual frequency distribution.

This study, then, endeavours to examine two hypotheses. The simpler one is that conditionals with a TRUTHTELLER-type antecedent would be collapsed, and hence estimates of their truth will reflect conditional probability, as if they are regular, everyday conditionals. The other hypothesis involves probability estimates of conditionals with a LIAR-type antecedent, and predicts a 0%-0% pattern in their evaluation, so that reasoners are expected to evaluate such conditional as 0% True but also 0% False, regardless of their actual probability. In other words, conditionals with TRUTHTELLER antecedent are expected to be collapsed, whereas conditionals with LIAR antecedents are expected to elicit probabilistic inference breakdown.

---

[1] This is perfectly in accord with Kripke's classic analysis (1975); cf. Elqayam (2005). Note that we are not proposing that the Liar constitutes a truth-value gap as a *solution* to the paradox – that approach only leads to the so-called 'strengthened Liar' (cf. Martin, 1984).

[2] Note that this is not a normative question: the proliferation of many-valued logic systems makes it effectively impossible to set normative standards in meta-deduction (Elqayam, 2003).

METHOD

Participants. 72 students of the University of Plymouth and 22 prospective students taking part in an open day participated on a paid volunteer basis and were tested in small groups. All participants were native speakers of English and none of them had had formal training in logic.

Materials and Procedure. Test materials were presented by software, beginning with a consent page, then a preamble that introduced the Island of Knights and Knaves and the format of conditional probability puzzles. The test phase consisted of control and test items presented in an individually randomised order. In all, each participant was presented with 16 items. After that, participants were asked for some personal data that included age, gender, and background in logic. The final page consisted of thanks and a detailed debriefing file.

Design. Figure 1 presents some matching sample items that illustrate item construction for the various test conditions. Design was mixed, with conditional probability and city size as repeated measure variables, and antecedent type as a between variable. Participants were randomly assigned either to a TRUTHTELLER condition, in which the test items consisted of conditionals with a TRUTHTELLER-type antecedent ('If I am a knight then I live in Emerald City'), or to a LIAR condition, in which the test items consisted of conditionals with a LIAR-type antecedent ('If I am a knave then I live in Emerald City'). Speaker and city names were randomly assigned from a list.

Another between test conditions was question format, which was either direct or indirect. Direct questions conformed to the format utilised in Evans et al. (2003), and were phrased as 'How likely is it that X is telling the truth / lying?' Indirect questions conformed to the format utilised in the previous collapse study (Elqayam, in press), and were phrased as 'How likely is it that X is a knight / knave?'

In all items, pq frequency was kept constant at 150, with the p¬q frequency set either at 50 (high conditional probability) or at 450 (low conditional probability). The ¬pq and ¬p¬q frequencies were set equal to each other, either at 200 or 400 for small city size, or 2000 / 4000 for large city size. Table 1 presents the design.

Table 1: Item frequencies

		pq	p¬q	¬pq	¬p¬q
High CP	Small city size	150	50	400	400
	Large city size	150	50	4000	4000
Low CP	Small city size	150	450	200	200
	Large city size	150	450	2000	2000

Each item was presented twice, once with a question about the probability of its being true ('What is the likelihood that X is a knight / is telling the truth?' and once with a question about the probability of its being false ('What is the likelihood that X is a knave / is lying?').

In addition to the test items, all participants were presented with matching control items, in which the knights / knaves were replaced by ordinary professions (e.g., 'If I am a teacher then I live in Emerald City'), with professions randomly assigned from a list of four: doctors, painters, teachers and musicians. Participants were instructed to assume that there were just two professions in each city. The control items served to replicate previous conditional probability findings, thus supporting the soundness of the paradigm utilised in the present study.

The mixed experimental design thus was 2 (TRUTHTELLER / LIAR -type antecedent of the test items) by 2 (direct / indirect question) by 2 (high / low conditional probability) by 2 (large / small city size), the latter two as repeated measures. For the control items, the antecedent condition was irrelevant, resulting in a 2x2x2 design. There were two types of dependent variables, True ratings and False ratings. Since the True and False questions were separated, this resulted in a set of 16 items for each participant.

RESULTS AND DISCUSSION

Since there were no significant differences between the direct and indirect test conditions, data from these conditions were pooled.

	High CP, Small city size, TRUTHTELLER Condition, Indirect Q		Low CP, Large city size, LIAR Condition, Direct Q	
	TRUE Control item	**TRUE Test item**	**TRUE Control item**	**TRUE Test item**
Cond.	Anne: If I am a teacher then I live in Emerald City.	Anne: If I am a knight then I live in Emerald City.	Ben: If I am a teacher then I live in Emerald City.	Ben: If I am a knave then I live in Emerald City.
p q **p ¬q** **¬p q** **¬p ¬q**	150 teachers in Emerald City 50 teachers in Opal City 400 doctors in Emerald City 400 doctors in Opal city	150 knights in Emerald City 50 knights in Opal City 400 knaves in Emerald City 400 knaves in Opal city	150 painters in Emerald City 450 painters in Sapphire City 2000 doctors in Emerald City 2000 doctors in Sapphire city	150 knaves in Emerald City 450 knaves in Sapphire City 2000 knights in Emerald City 2000 knights in Sapphire city
Q	How likely is it that Anne is a knight? ___%	How likely is it that Anne is a knight? ___%	How likely is it that Ben is telling the truth? ___%	How likely is it that Ben is telling the truth? ___%

p – antecedent (the 'if' part of the conditional); q – consequent (the 'then' part of the conditional); ¬ – 'not'

Figure 1: Matching Sample Items, True

First we analysed the control items, to ensure that they replicate previous conditional probability findings, and hence that using the knight-knave scenario in conjunction with estimates of the probability of conditionals is a viable paradigm. Table 2 shows the mean probability estimates for the control items. Two separate 2 (high / low conditional probability) x 2 (large / small city size) repeated-measure ANOVAs were conducted for the True and False ratings of the control items respectively. As predicted, the results of these analyses essentially replicated previous results reported for the conditional probability paradigm (Evans et al., 2003; Oberauer & Wilhelm, 2003), both analyses revealing a highly significant main effect of conditional probability. High conditional probability items elicited higher probability estimates than did low conditional probability items $(F(1,93)=67.9, MSE=612, p<.0005)$. This pattern establishes the soundness of the present paradigm.

Table 2: Mean Probability Estimates for True and False Ratings of Control Items (SD in parenthesis)

	True Ratings		False ratings	
Low CP, Small city size	55.4	(27.4)	43.8	(26.0)
Low CP, Large city size	55.4	(30.3)	49.4	(30.8)
High CP, Small city size	37.2	(25.7)	52.9	(27.2)
High CP, Large city size	31.5	(22.2)	58.1	(29.9)

CP – conditional probability

For the True ratings, there was no significant main effect of city size (i.e., no effect of conjunctive probability; $p > .05$). There was, however, a main effect of city size in the False ratings, in which large sample items tended to elicit higher False ratings than small sample items $(F(1,93) = 6.3, MSE = 442, p < .05)$. In both ANOVAs, The interactions between conditional probability and city size were non-significant $(p > .05.)$

Our second step was analysis of the test items, which we examined by two separate 2 (LIAR-type / TRUTHTELLER-type antecedent) 2 (high / low conditional probability) x 2 (large / small city size) mixed ANOVAs for the True ratings and the False ratings respectively. Figure 2 presents mean True and False ratings for the test items.

For the True rating, some of the predicted pattern emerged, with two main effects, of conditional probability and of antecedent type.

Mean True ratings of high conditional probability items were significantly higher than mean True ratings of low conditional probability items $(F(1,92) = 19.6, MSE = 703, p < .0005)$. In addition, the LIAR-type antecedent group produced True ratings of test items that were significantly lower than True ratings of test items produced by the TRUTHTELLER-type antecedent group $(F(1,92) = 4.6, MSE = 2183, p < .05)$. There were no significant city size effects and no significant interactions $(p>.05)$. The pattern revealed is entirely in line with the predictions as far as the TRUTHTELLER-type antecedent group is concerned, which was predicted to collapse, resulting in a pattern of

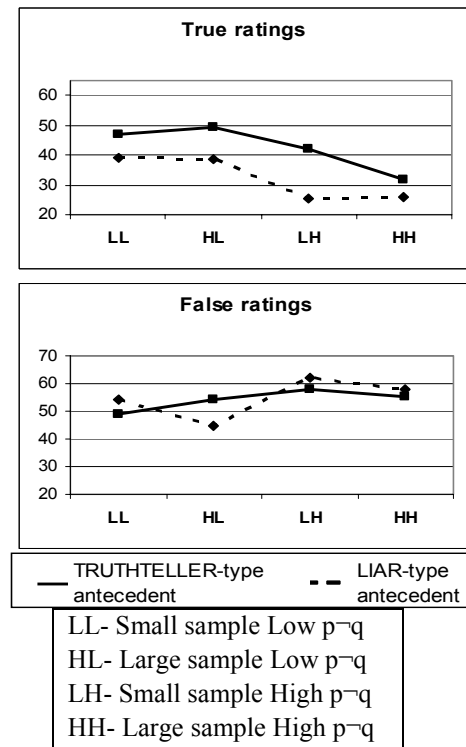conditional probability – this is just what the results demonstrate.

Figure 2: Mean Probability Ratings of Test items

The observed pattern for the LIAR-type antecedent group was somewhat surprising, though. The 0%-0% hypothesis (i.e., 0% True and 0% False) would predict a flattish line at or near the 0% mark for this group. Instead, we have response pattern that very closely echoes that of the TRUTHTELLER-type antecedent group's, only about 10% below it. How could this be explained?

Previous work (Elqayam, in press) has revealed a minority group of reasoners, who can work their way around the compelling tendency to produce the null model for LIAR-type constituents. Faced with conditionals such as 'If I am a knave then snow is black,' these reasoners typically evaluate them mainly as true (contrary to the predominate 'neither' evaluation of less able reasoners). If this group surfaced in the present context as well, we could reasonably surmise that they have been able to resist the null model, thus producing the equivalent of 'determinate' responses. To deal with the tractability problem they probably only kept one sort of hypothetical thinking, and that would have been the relatively easier one, the probabilistic hypothetical thinking triggered by 'if'. They thus ignored the paradoxality of the LIAR-type antecedent and responded as they would to any other conditional.

Hence, the lower line of conditional probability effect for the LIAR-antecedent group suggests that the mean True estimates reflect an averaging effect of two distinct response patterns: a conditional probability response pattern, and a

0%-0% response pattern. The 0% responses may have pulled down the means, resulting in lower probability estimates. We take up this analysis again later.

The picture that emerged for the False ratings was somewhat more ambiguous. There was a significant main effect of conditional probability, in which mean False ratings of high low conditional probability items were lower than mean False ratings of high conditional probability items $(F(1,92) = 6.1, MSE = 895, p < .05)$. However, there was no antecedent main effect such as was discerned in the True ratings $(p>.05)$; it may have been blurred by the added difficulty of False ratings (in effect, a sort of double negation effect; cf., e.g., Evans, Clibbens & Rood, 1995; Evans & Over, 2004) . Both True and False ratings analyses revealed no significant city size effects and no significant interactions $(p>.05)$.

To test the '0%-0%' hypothesis more directly, we computed two indices reflecting the number of 0% True and 0% False responses. First, we computed a '0% True index', consisting of the number of test items rated as 0% True. From this we reduced the number of control items rated as 0% True, to control for possible response bias towards 0%. As there were four test items with True ratings and four control items with True ratings, the 0% True index had a potential range of -4 to +4. The higher the score, the higher is the participant's tendency to rate test items as 0% True.

Similarly, we computed a '0% False index', which consisted of the number of test items rated as 0% False, again, minus the number of control items rated as 0% False, and again with a potential range of -4 to +4. The higher the score of the 0% False index, the higher is the participant's tendency to rate test items as 0% False.[3]

Participants who tend to respond with breakdown of probabilistic inference should have both a high 0% True index and a high 0% False index. Hence, we should expect both the 0% True index and the 0% False index to be significantly higher for the LIAR-type antecedent group than for the TRUTHTELLER-type antecedent group.

A repeated measure 2 (antecedent type) by 2 (True vs. False index) ANOVA supported the 0%-0% hypothesis, the results shown in Figure 3. As can be patently seen, both 0% indices for the LIAR antecedent group are higher, a significant main effect, $(F(1,92) = 17.7, MSE = 1.335, p < .0005)$. In addition to that, there were also a True-False main effect, the 0% False index being lower than the 0% True index $(F(1,92) = 16.1, MSE = .521, p < .0005)$; however, this was probably an artefact of the interaction effect $(F(1,92) = 12.0, MSE = .521, p < .001)$, in which 0% True index increased more dramatically than the False 0% index – probably due to the blurring effect associated with the False ratings and the double negation effect noted above. Although the gradient for the 0% False index was less steep, nevertheless the difference in 0% False index scores between the TRUTHTELLER antecedent group and the LIAR antecedent group was significant $(\underline{t}(1,54)=2.3, p<.05)$.

---

[3] We have also calculated 100% indices but reporting them and their results would be beyond the scope of present work.

Thus we can see that responses to the LIAR reflect the 0%-0% pattern of probabilistic inference breakdown, in which it is perceived as neither true nor false.
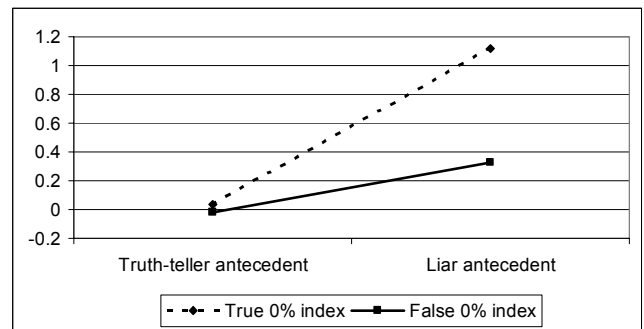

Figure 3: Mean 0% indices

GENERAL DISCUSSION

In this study we examined estimates of the probability of conditionals with indeterminate antecedents, such as the LIAR ('I am lying') and the TRUTHTELLER ('I am telling the truth'). We made two primary hypotheses: First, we predicted that conditionals with TRUTHTELLER-type antecedents would be 'collapsed' (Elqayam, in press), that is, regarded as if they were ordinary conditionals with nothing exceptional about them. Hence we predicted that their probability would be estimated as the conditional probability, q|p; in other words we predicted their estimation pattern to closely match that of control items with determinate antecedents and replicate the pattern previously established in the literature with ordinary conditionals (Evans et al., 2003; Oberauer & Wilhelm, 2003).

The second prediction we made was regarding conditionals with a LIAR-type antecedent, and was quite different. We hypothesised that, as the suppositional or hypothetical-thinking processes involved in judging the LIAR were immensely complex, and as conditionals involved hypothetical thinking by themselves, putting the two together would prove intractably intolerable for reasoners, and probabilistic inference breakdown would ensue. Thus, we predicted that probability estimates of conditionals with LIAR-type antecedent would be both 0% True and 0% False – what we called the 0%-0% pattern.

We also had a minor prediction involving the soundness the experimental paradigm. We predicted that probability estimates of control items involving ordinary conditionals with determinate antecedents such as 'If I am a painter…' would closely resemble probability estimates established in the literature. Hence we predicted the probability of control items to be influenced by conditional probability, q|p.

Of these three predictions, two were clearly supported and one partially so. First, our minor prediction of conditional probability estimate of control items was supported, lending credence to the viability of the experimental paradigm that conjoined meta-deductive components with probability

estimates of conditionals.

More importantly, the collapse prediction of conditionals with TRUTHTELLER-type antecedents was entirely borne out, as estimates of these conditionals were clearly affected by conditional probability, just as matching control items were affected by it. Thus the collapse affect predicted by hypothetical thinking theory has been supported.

One of the surprises of the present findings was how robust the conditional probability effect is. Not only did we replicate it in the control items, it has greatly affected the test items as well – including the LIAR-type antecedent group, where we did not expect to find it: conditionals with LIAR-type antecedents were affected by conditional probability, although their estimates were significantly lower. Clearly, at least some of the participants were responding to these items with conditional probability estimates. Previous work (Elqayam, in press) identified reasoners able to resist the pull of the null model in conditionals with LIAR-type antecedents and come up with determinate evaluations, and this pattern seems in accord.

More testing established the existence of conditional inference breakdown pattern as well, as participants did tend to respond to LIAR items with 0% estimates, both True and False – the predicted '0%-0%' pattern. This pattern was superimposed on the conditional probability pattern, resulting in significantly lowered probability estimates for the group presented with LIAR-type antecedents.

Hypothetical thinking at its best is only as good as our bounded rationality is good – we have seen some of these inherent limitations in this paper. In particular, we seem helpless to iterate suppositions, and embedding a supposition within supposition seems to exert an impossible demand on our capacity – our inferences just break down.

The TRUTHTELLER collapse may mean that the satisficing principle of bounded rationality is a device of 'self-defending rationality' – the built-in mechanisms that limit normative rationality actually protect us from its breakdown. Paradoxically, then, the very limitations on suppositional processes may actually support rationality.

REFERENCES

Edgington, D. (2003). What if? Questions about conditionals. *Mind & Language, 18*, 380-401.

Elqayam, S. (2003). Norm, error, and the structure of rationality: The case study of the knight-knave paradigm. *Semiotica*, *147,* 265-189.

Elqayam, S. (2005). The collapse illusion effect: a pragmatic-semantic illusion of truth and paradox. *Thinking and Reasoning*.

Evans, J. St. B. T., Clibbens, J., & Rood, B. (1995). Bias in conditional inference: implications for mental models and mental logic. *Quarterly Journal of Experimental Psychology, 48A,* 644-670.

Evans, J. St. B. T., Handley, S. H., & Over, D. E. (2003). Conditionals and conditional probability. *Journal of Experimental Psychology: Learning, Memory and Cognition, 29*, 321-355.

Evans, J. St. B. T., Newstead, S. E., & Byrne, R. M. J. (1993). *Human Reasoning: The Psychology of Deduction*. Hove & London: Erlbaum.

Evans, J. St. B. T. & Over, D. E. (2004). *If*. Oxford: Oxford University Press.

Evans, J. St. B. T., Over, D. E., & Handley, S. H. (2003). A theory of hypothetical thinking. In D. Hardman & L. Maachi (Eds.), *Thinking: Psychological perspectives on reasoning, judgement and decision making* (pp. 3-22). Chichester: Wiley.

Evans, J. St. B. T., Over, D. E., & Handley, S. J. (2005). Suppositions, extensionality, and conditionals: A critique of the mental model theory of Johnson-Laird & Byrne (2002). *Psychological Review*.

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

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

Kripke, S. (1975). Outline of a Theory of Truth. *Journal of Philosophy, 72*, 690-716.

Martin, R. L. (1984). (Ed.) *Recent essays on truth and the Liar paradox*. Oxford, UK: Oxford University Press.

Oberauer, K. & Wilhelm, O. (2003). The meaning(s) of conditionals: Conditional probabilities, mental models and personal utilities. *Journal of Experimental Psychology: Learning, Memory and Cognition, 29*, 680-693.

Ramsey, F. P. (1931). *The foundations of mathematics and other logical essays*. London: Routledge and Kegan Paul.

Schroyens, W., Schaeken, W. & d'Ydewalle, G. (1999). Error and bias in meta-propositional reasoning: A case of the mental model theory. *Thinking and Reasoning, 5*, 29-65.