The Reasoner

Volume 2, Number 1
January 2008

www.thereasoner.org

Contents

§1 Editorial

As a member of several editorial boards (e.g. The Reasoner) and program committees, I am often involved in the peer reviewing process of selecting the best papers from those submitted. Each submission is typically reviewed by 3 or 4 referees. In addition to the detailed reports, many journals and conferences expect the referees to score their assigned papers by a pair of values, one that reflects the overall quality of the paper and one that indicates the referee's own level of expertise or confidence. For the selection of the best papers, the editor-in-chief or the PC chair faces the problem of combining those scores to establish an overall ranking, a process that is typically done by hand. But what would be a reasonable procedure to combine the scores and to establish a ranking of peer-reviewed papers automatically?

By describing a set of so-called *process patterns*, Nierstrasz (2000: *Identify the champion*, Pattern Languages of Program Design, vol. 4, 539–556) gives an informal answer to this question. Examples of such patterns are "*Group papers according to their highest and lowest score*" or "*Take care to identify papers with both extreme high and low scores*". For the scores, Nierstrasz proposes four quality categories A="*Good paper*" to D="*Serious problems*" and three levels of expertise X="*I am an expert*" to Z="*I am not an expert*". Notice that Nierstrasz' pattern language has become something like the de facto standard for conferences in the field of Computer Science, and it is implemented in conference management tools such as CyberChair).

The success of Nierstrasz' patterns is perfectly comprehensible from the pragmatic point of view of an experienced PC chair, but from a more formal perspective, they give the impression of being constituted on an ad hoc basis and may therefore seem a bit rudimentary. So how would an expert in the area of formal reasoning and decision making under uncertainty respond to that question?

A partial answer can be found in the early literature on probability from the late 17th and early 18th centuries, as pointed out by Shafer (1993: *The early development of mathematical probability*, in *Companion Encyclopedia of the History and Philosophy of the Mathematical Sciences*, Routledge, London, 1293–1302) and Kohlas (2004: *Reliability of arguments*, in *Defining the Science of Stochastics*, Lemgo, Germany, 73–94). At that time, studying probability was often motivated by judicial applications, such as the reliability of witnesses in the courtroom, or more generally by the credibility of testimonies on past events or miracles. The first two combination rules for testimonies were published in an anonymous article (1699: *A calculation of the credibility of human testimony*, Philosophical Transactions of the Royal Society, 21, 359–365). One of them considers two independent witnesses with respective credibilities and then gives a formula for the combined credibility of both witnesses. The corresponding formula for n independent witnesses of equal credibility has been mentioned by Laplace (1820: *Théorie Analytique des Probabilités*, Courcier, Paris) and is closely related to the *Condorcet Jury Theorem* discussed in social choice theory. Boole (1854: *The Laws of Thought*, Walton and Maberley, London) mentioned a similar formula that includes a prior probability of the hypothesis in question.

In a recent article with Stephan Hartmann, we picked up these ancient ideas and turned them into a more general model of combining reports from partially reliable sources, see Haenni and Hartmann (2006: *Modeling partially reliable information sources: a general approach based on Dempster-Shafer theory*, International Journal of Information Fusion, 7(4), 361–379). The generality of the model allows it to be applied to situations of incompetent or even dishonest witnesses, who may deliver highly contradictory testimonies. At its core, the model presupposes a non-additive measure of belief, but Laplace's and Boole's formulae themselves are included as additive special cases. The model also includes various Bayesian approaches, which require a prior probability of the hypothesis in question to turn it into a corresponding posterior probability.

By interpreting the scores of a peer-reviewed paper as respective probabilities, another particular instance of the general model proposed in Haenni and Hartmann (2006: 361–379) arises. As we would expect the decision of an unbiased editor-in-chief or PC chair to be taken exclusively on the basis of the reports, it does not consider any form of prior knowledge (as would be required in a Bayesian setting). Furthermore, to be able to distinguish a paper reviewed by a group of experts from a paper reviewed by a group of non-experts, the procedure is supposed to output a pair of values (similar to its inputs), one that reflects the referees' combined quality judgement and one for the confidence level of the group of referees as a whole. This is exactly what is produced by our method. Additionally, it provides a quantitative measure of the conflict between the reports. Based on these values, the editor-in-chief or the PC chair can now decide to accept or reject a paper, or to send it to an additional referee in case of high conflicts or an insufficient confidence level. A prototype implementation of this scheme can be tested at

http://www.iam.unibe.ch/~run/referee,

where the referee reports and the results of the combination are visualized as points in a 2-simplex called *opinion triangle*, see Jøsang (2001: *A logic for uncertain probabilities*, International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems, 9(3), 279–311). It is interesting to observe that it almost perfectly reproduces Nierstrasz' classification patterns.

This month's issue of The Reasoner starts with an interview with Jürg Kohlas, who was the director of my PhD thesis in the early 1990s. He is the one who initially brought my interest and attention to the area of uncertain reasoning. He is also the one that pointed out the close connection from my own research to the above-mentioned early literature on probability.

Rolf Haenni
Computing, Bern University of Applied Sciences

§2
Features

Interview with Jürg Kohlas

Jürg Kohlas is full professor in the Department of Informatics (DIUF) at the University of Fribourg, Switzerland. He is the director of the Theoretical Computer Science (TCS) research group.

Rolf Haenni: In Kohlas (2004: *Reliability of arguments*, in *Defining the Science of Stochastics*, Lemgo, Germany, 73–94), you interpret some ancient views of probability (held by Laplace, Boole, J. Bernoulli, Lambert, and others) from a more modern perspective. What is the main lesson we should learn from them?

Jürg Kohlas: The main lesson is that it is worthwhile to read the original papers. Scientific progress is usually very selective. Often only parts of the original ideas are pursued. In the case of Jacob Bernoulli, for instance, the law of large numbers shaped modern statistics and reasoning, whereas interesting other parts of Bernoulli (1713: *Ars Conjectandi*, Thurnisiorum, Basel, Switzerland) went largely unnoticed. His very original ideas about pure and mixed arguments, which carried already the germ of some modern ideas of how to combine logic and probability, were only rediscovered recently in the context of the Dempster-Shafer theory of evidence.

RH: As you say, some of those ideas were picked up in the late 1960s and early 1970s by Arthur Dempster and Glenn Shafer, both statisticians. Until today, their theory has always been a controversial issue. Why this?

What went wrong?

JK: I think that although the work of Dempster and Shafer was mathematically very sound, its meaning was not quite clear.

RH: Your early work on reasoning and uncertainty was also influenced by Dempster and Shafer. When and why did you get interested in those topics and how did you encounter their work?

JK: I became aware of Shafer (1976: *The Mathematical Theory of Evidence*, Princeton University Press) in the 1980s. I was immediately intrigued by the originality and the rigorous mathematical flavor of this book. I wanted to understand what this theory meant exactly, in particular how it was related to probability theory. Some authors claim that it is different from probability, but I came to the conclusion that it can be interpreted as a theory of reliability of deduction. As I worked in reliability theory at that time, this idea suited me very much. It gives a clear semantics to the theory and it indicates also clearly how the theory can be applied.

RH: How exactly is your own *Theory of Hints* related to Dempster's and Shafer's work? Do you see your theory as an alternative or as a generalization of the Bayesian view of uncertainty?

JK: In Kohlas and Monney (1995: *A Mathematical Theory of Hints*, Springer, New York), the above viewpoint of the Dempster-Shafer theory is mathematically worked out as a theory of the probability of provability. I understand this work as a formal generalization of the Bayesian formalism; the latter appears as a special case. It is based however on a different philosophy than the modern view of Bayesian analysis as subjective probabilities.

RH: From your perspective of reasoning and inference, what is the connection between logic and probability?

JK: In my understanding of inference, logic is used to derive or prove hypotheses not only on the basis of knowledge and facts, but also on some more or less uncertain assumptions. Probability serves then to compute the reliabilites of the different possible arguments permitting to deduce a hypothesis. This measures how strongly a hypothesis is supported by the given knowledge and facts. In this way, logic and probability can be combined in a theory of "probability of provability".

RH: You have also a background in inferential statistics. Do you see a way to connect or unify Fisher's fiducial inference with Bayesian statistical inference?

JK: Dempster proposed his work as a means to generalize Bayesian inference for reasoning towards posterior distributions based on samples. His goal was to show that this is also possible without prior distributions. This corresponds to Fisher's fiducial probabilities, which can be obtained if a pivotal quantity exists. This requirement limits Fisher's approach to some extent, but not necessary Dempster's method. By defining statistical inference based on functional models, which describe how samples are generated, posterior probability statements in the sense of Dempster-Shafer belief functions (or hints) can be obtained. In such a framework, it is formally possible to reinterpret Fisher's fiduciual probabilities and also to reproduce Bayesian inference. In other words, both Fisher statistics and Bayesian inference appear formally as two special cases of an extended framework of statistical inference. This is the topic of Kohlas and Monney (2008: *Statistical Information: Assumption-Based Statistical Inference*, Sigma Series in Stochastics. Heldermann Verlag, Lemgo, Germany), which will soon be available.

RH: More recently, your principal area of interest shifted towards an algebraic theory of information. What is it all about? And how is it connected to your earlier work?

JK: This line of research emerged from the study of local computation schemes. Based on a fundamental paper by Lauritzen and Spiegelhalter (1988: *Local computations with probabilities on graphical structures and their application to expert systems*, Journal of the Royal Statistical Society, 50(2), 157–224), Shenoy and Shafer (1988: *Axioms for probability and belief-function propagation*, UAI'88, Proceedings of the 4th Conference on Uncertainty in Artificial Intelligence, 169–198) proposed an axiomatic system which is sufficient to enable local computation schemes. It appeared to me that their system essentially described very basic properties of information: (1) pieces of information refer to questions, (2) they come from different sources and must be combined or aggregated, and (3) they must be focused on questions of interest. In Kohlas (2003: *Information Algebras: Generic Stuctures for Inference*, Springer, London), these elements are captured in an algebraic framework called *information algebra*. Many instances of such information algebras are important in Computer Science, like relational databases, constraint systems, or various logics. They also cover many uncertainty formalisms including Bayesian networks, possibility theory, or belief functions. Moreover, uncertain information can be described in this general abstract framework by random variables with values in such algebras. Surprisingly, this simple scheme turns out to be the natural, most general framework for an abstract Dempster-Shafer theory and to offer the mathematical foundation of probabilistic argumentation systems.

RH: You are still teaching undergraduate and graduate students in Computer Science. How important is it for them to know something about reasoning and uncertainty?

JK: To know about reasoning and uncertainty is important in any academic discipline, since everybody has to act under uncertainty in her daily and professional life. In Computer Science, it should at least be taught in a course on Artificial Intelligence.

RH: What are the first 3 books you would recommend to a student who is interested in those topics?

JK: Well, in Computer Science, I would first recommend a book on AI, e.g. the excellent introductory textbook of Russell and Norvig (2004: *Artificial Intelligence: A Modern Approach*, Prentice Hall, 2nd edition). Second, I would propose a book which you, Rolf Haenni, should yet write, with a tentative title like *Probabilistic Argumentation Systems: Inference Based on Logic and Probability*. Finally, for those who are not afraid of some rather elementary mathematics and abstract thinking, I cannot resist to cite my own book Kohlas (2003), the only textbook that exists so far on the information algebraic perspective of reasoning and uncertainty.

RH: You are going to retire in the near future. When exactly will that be and what are your scientific plans for the time after your retirement?

JK: Well, officially I am supposed to retire sometimes in 2009. Then I hope to be able to continue my work on information algebras. The field is new, which means that the open problems are still not too difficult. But I am certainly going to miss my PhD students who contributed so much to my research, e.g. by detecting and correcting many of the errors I made.

The factivity failure of contextualist "knows"

In this paper we argue that *standard, indexical contextualism about "knows"* (ICK) is unable to account for the *factivity* of "knows".

If we suppose, for simplicity, that the basic parameters of a context c are just the world w_c and the speaker s_c of c, it is the principle that:

> An utterance of *S knows that* ϕ in a context c
> is true only if $[\![\phi]\!]^{c,w_c} = 1$

that a proper semantic treatment will have to make true before it can be said to have offered an account of the factivity of "knows".

Now, ICK is precisely intended to provide such a treatment by treating "knows" as an indexical in a broad sense, an expression whose content depends on some appropriate feature of the context. If, in the spirit of Lewis (1996: "Elusive Knowledge", *Australasian Journal of Philosophy*, 74, 549-567), we identify this feature with the *alternatives* that are somehow *relevant* in the context (say, those possibilities that the speaker is, or should be, attending to), we can model ICK by means of two functions:

1. for each subject x, a function $\mathcal{K}_x$ of *epistemic alternativeness*, mapping each world w into the set $\mathcal{K}_x(w)$ of worlds that are epistemic alternatives to w for x, i.e., compatible with x's information in w;

2. for each context c, a function $\mathcal{R}_c$ of *contextual relevance*, mapping each world w into the set $\mathcal{R}_c(w)$ of worlds that are relevant alternatives to w in c.

This allows for two different ways to capture ICK.

The most straightforward is to make the extension of "knows" at a context c and a world w, $[\![knows]\!]^{c,w}$, consist in this:

$$\lambda p.\lambda x.\ \forall w' \in \mathcal{K}_x(w) \cap \mathcal{R}_c(w) : p(w) = 1,$$

where x ranges over subjects, and p over sentential *contents*, functions from worlds to truth values. This entry makes "knows" an indexical in the broad sense, by making its content depend on the relevant alternatives associated with the context. Let us label the form of ICK that results from endorsing this analysis as *normal indexical contextualism about "knows"* (NICK). By imposing the following constraint on epistemic alternativeness and contextual relevance:

> *Reflexivity:*
> $\forall x, c, w : w \in \mathcal{K}_x(w) \cap \mathcal{R}_c(w),$

factivity follows immediately.

Despite this, contextualists should go for an alternative analysis, according to which $[\![knows]\!]^{c,w}$ consists in this:

$$\lambda P.\lambda x.\ \forall w' \in \mathcal{K}_x(w) \cap \mathcal{R}_c(w) : P(c_x)(w) = 1,$$

where (i) c_x is now the context associated with x instead of the initial context c, and (ii) P ranges over sentential *characters*, functions from contexts to sentential contents. The differences with NICK look like small differences, but they do make a difference, for now the truth of "S knows that ϕ" at an attributor's context depends on the truth of "ϕ" at another context, that associated with subject S. "Knows" is treated not only as an indexical, but also as a *context-shifter*, what Kaplan (1989: "Demonstratives". In J. Almog *et al.* (eds.), *Themes from Kaplan*, Oxford: OUP, 481-563) would have considered a *monster*. The form of ICK that results from endorsing this alternative analysis may thus be called *monster indexical contextualism about "knows"* (MICK).

One reason why contextualists should go for MICK rather than NICK is that MICK fares better than NICK with respect to the *subject-boundedness* of knowledge: the subject's context plays a crucial role in deciding whether or not she counts as knowing.

Indeed, suppose that f_1, f_2, and f_3 are all the Fs there are. Bob has observed them all and discovered that they all were also Gs. Al is informed that Bob has made these observations, but is very anxious about, therefore is attending to, the possibility that there might be Fs besides f_1–f_3. But I myself do not pay any attention to this possibility, the only one I am attending to being that Bob might not know of each of the Fs that they are Gs. Would it then be correct for me to utter:

> (1) Al knows that Bob knows that all Fs are Gs?

Intuitively, the answer is "no". However, the proponent of NICK will have to answer "yes". For Al *can* rule out the possibility that Bob cannot rule out the only possibility relevant in my context, that f_1–f_3 are not all Gs; so, by NICK, my possible utterance of (1) would be true! On the contrary, the proponent of MICK will answer "no". For the possibility that there is an F that is not-G is relevant in Al's context, and it is one that Bob *cannot* rule out. So, by MICK, my possible utterance of (1) would be false, in accordance with intuition. MICK is thus in a better position than NICK to take the subject-boundedness of "knows" into account.

Unfortunately, MICK is unable to account for the factivity of "knows". Indeed, a reasonable way to conceive of "knows" as a context-shifter is to augment contexts with an additional *salient subject* parameter, and to tell the following story: given a context c, the subject who is salient in c is, by default, identified with the speaker s_c of c; "knows" is then allowed to shift the salient subject parameter of a context *and only this parameter*. This story, however, will not preserve factivity, as second-order knowledge ascriptions reveal.

Suppose I make a true utterance of:

> (2) Keith knows that (2*) his wife knows that the bank will be open on Saturday morning.

Then, by MICK, the content of (2*) in the context associated with *Keith* must be true at the world of my utterance. But this is perfectly compatible with the content of (2*) in *my* context being false at the same world. (Just assume that a possible world in which the bank is closed on Saturday morning is both relevant in my context and an epistemic alternative for Keith's wife.) We would then have a case where in the same context in which I make a true utterance of (2), (2*) expresses a false proposition. So, MICK makes "knows" non-factive.

To sum it all up in four words: *neither NICK nor MICK!* Therefore, ICK cannot provide a proper semantic treatment for "knows".

Franck Lihoreau
Philosophy, Universidade Nova de Lisboa

Manuel Rebuschi
Philosophy, Nancy University

Reasoning: How We're Doing It

In the course of becoming rational animals, human beings have had eons to develop language and to hit upon a method of deductive reasoning with the sentences of their natural languages. What has actually evolved is very different from and strikingly more efficient than any conventional logic we learn in school. A seven year old child moves intuitively (and instantly) from

$$\text{'not all dogs are friendly'} \tag{1}$$

to

$$\text{'some dogs aren't friendly.'} \tag{2}$$

How does this untutored intuition work? The child's inference is obviously valid and students of college logic formally prove its validity by translating (1) and (2) in the standard quantifier/variable" notation as

$$\neg(\forall x \text{ if } x \text{ is a dog, than } x \text{ is friendly}), \tag{3}$$

$$\exists x \text{ such that } x \text{ is a dog and } \neg\ (x \text{ is friendly}), \tag{4}$$

after which, by applying laws of "quantifier interchange," they show that (4) follows from (3). The proof, including the translations, takes about 40 seconds. Of course, no one in real life—child or adult—normally thinks in sentences like (3) and (4). So the logic we use in formally justifying everyday inferences casts no light on how children (and adults) reasoning with variable-free English sentences, are getting from (1) to (2) in a fraction of a second.

A second example: Anyone who comes across a sentence like:

$$\text{All colts are horses but someone who is riding a colt is not riding a horse,} \tag{5}$$

instantly judges it to be inconsistent and false. The logician proves (5) inconsistent by a lengthy process of translating it into quantifier/variable notation, assigning values to the variables of the translated formulas, and deriving contradictions of form 'p and $\neg p$.' But such proofs have never moved anyone to exclaim, 'Aha! So that's why everyone rejects this sentence on sight.'

Why then do we instantly reject (5)? Is there a cognitively veridical logic that illuminates the process of how we mentally arrive at our everyday deductive judgments? One question we want answered is: How, using natural language in real-life reasoning, are we able to reckon with "the speed of thought," when standard logic, using the powerful technical apparatus of quantifiers and bound variables, typically requires so much reckoning time? My research here focused on the common logical words that figure crucially in our everyday reasoning words like 'not,' 'non-,' 'some,' 'every,' 'and,' 'if,' and 'then.' I eventually discovered that the logical particles (as I shall refer to them) are positively or negatively charged and that we reckon with them as we reckon with the '+' and '-' operators of elementary algebra. Some particles behave and are treated by us as "plus-words," others behave and are treated by us as "minus words." The following is a partial but representative list of logical particles that specifies their charged characters:

> **'some'** ('a'), **'is'** ('was' 'are', etc.), **'and'** are **plus-words**.
> **'every'** ('all' ... etc), **'not'** ('no', 'non-' ...), **'aren't,' 'if',** are **minus-words**.

We can account for the celerity and accuracy of a child's intuitive reasoning if we attribute to it an operative knowledge of the charged character of the logical particles and some command of elementary algebraic reckoning. At some point in its cognitive development, the child begins to reckon with 'no' as a minus-word, with 'and' as a plus-word, with 'some' as a plus-word, and with 'every' as a minus-word. On being told that "**Not**$^-$ **all**$^-$ **Dogs are**$^+$ **Friendly**," the child, reckoning with the $+/-$ charges, can transform '$-(-$**Dogs**$+$**Friendly**)' into '$+$**Dogs**$-$**Friendly**' [Some$^+$ Dogs aren't$^-$ Friendly]. Of course, the child is unaware of reasoning this way. At no point, for example, is the child conscious that it is reckoning with 'some' and 'all' as operators opposed to each other as '+' to '$-$', an oblivion that persists into adult life. The innate human ability to reason algebraically with the charged natural logical particles explains the deja vu feeling (noted by Plato) that children experience when first taught the elementary truths of algebra and geometry.

We unconsciously exploit the $+/-$ algorithm in all our deductive reasoning:

Syllogistic Reasoning

In syllogistic reasoning, we conjoin two or more premises by adding them and reckoning conclusions. For example, by conjoining 'All$^-$ natives are$^+$ citizens' and All$^-$ citizens are$^+$ voters', we quickly and easily derive the conclusion 'All$^-$ natives are$^+$ voters.'

	All$^-$ natives are$^+$ citizens	$-Natives + Citizens$
&$^+$	All$^-$ citizens are$^+$ voters	$-Citizens + Voters$
∴	All$^-$ natives are$^+$ voters	

Propositional Reasoning

Basic inference patterns in propositional logic such as modus ponens, modus tollens and the hypothetical syllogism are $+/-$ transparent. Since 'if' is a minus-word and 'then' is a plus-word, 'if p then q' transcribes as '$-p + q$:'

Modus Ponens

$-p + q$
p
$\therefore\ q$

Modus Tollens

$-p + q$
$-q$
$\therefore -p$

Hypothetical Syllogism

$-p + q$
$-q + r$
$\therefore -p + r$

An equivalence such as $p\&(\neg q) \equiv \neg(p \supset q)$ transcribes as $p + (-q) = -(-p + q)$.

Reasoning with Relational Terms

We remarked earlier that (5) *All colts are horses but some rider of a colt isn't a rider of a horse* is universally rejected as logically false. Since (5) is blatantly inconsistent, exposing a contradiction it entails ought to be quick and easy. But conventional logicians, oblivious to the charged character of the natural logical particles, routinely "translate" (5) into canonical quantifier/variable notation. The process of instantiating the variables and deriving contradictions is then neither easy nor quick. By contrast, anyone representing (5) in natural language confronts a contradiction immediately:

(i) All$^-$ colts are$^+$ horses;
(ii) Some$^+$ rider of a$^+$ colt isn't$^-$ a rider of a$^+$ horse;
∴ (iii) Some$^+$ rider of a$^+$ horse isn't$^-$ a rider of a$^+$ horse;

This formal proof of inconsistency is cognitively veridical. (For more detailed expositions of the +/- account of deductive reasoning, see Sommers, F. T. (1982: The Logic of Natural Language, Oxford: Clarenden Press) and Sommers, F.T. and Englebretsen, G. (2000: An Invitation to Formal Reasoning, Ashgate), and Sommers, F.T. (2008: "Ratiocination." *Ratio*).

Conclusion

To fulfill its traditional mission of exposing the "Laws of Thought" Logic must directly address the cognitive puzzle of how untutored human beings reason so well. The solution on offer is that we reason by exploiting the charged, +/-, character of the natural logical particles.

Fred Sommers
Philosophy, Brandeis

Bourne's Negation: No Equivocation

Bourne (2004: 'Future contingents, non-contradiction, and the law of excluded middle muddle', *Analysis* 64(2), 122–8) has argued for a 3-valued logic according to which the negations of indeterminate propositions are true, rather than indeterminate. But Schang (2007: 'Truth and truthmakers: A reply to Bourne's negation, *The Reasoner* 1(8), 5–6) claims to have found an equivocation in Bourne's argument, 'between two distinct senses of 'truth' ... the former concerns *truthmakers* ... whereas the latter concerns *truth-bearers*' (Schang: 5). I see no such equivocation.

Schang claims that there is 'a clear-cut difference between two distinct senses of truth-values, depending upon whether they are about states of affairs or sentences expressing them' (Schang: 5). But truth-*values* are neither *about* states of affairs or sentences. They are not *about* anything; rather, they are the values that truth-bearers take, depending on whether the truth-bearer in question has a truthmaker. A truthmaker, such as the concrete state of affairs *Gordon's being in trouble*, is not the kind of thing that can be true (the identity theory of truth aside). Rather, the proposition *that Gordon is in trouble* is made true by this state of affairs; it is the proposition's truthmaker.

So Schang's claim that '[Bourne's] matrix makes a confusion between two sorts of 'truth', i.e., *being* true and *telling* the truth' (Schang: 5) is unfounded. Bourne does at one point couch his argument in terms of saying, e.g., 'to say that it is not the case that p is clearly to say something true' (2004: 124) but this is merely a way of expressing the point that, if p is evaluated to $\frac{1}{2}$, then the proposition *that p is not the case*, i.e., $\neg p$, is true. Bourne's point has to do with the conditions under which negated propositions are true; it has nothing to do with 'telling the truth'.

Bourne does not couch his argument in terms of truthmakers but it could be put as follows. A proposition is true iff it (determinately) has a truthmaker. Negation has the property that: $\neg p$ is (determinately) true iff p is not (determinately) true, i.e., iff p does not (determinately) have a truthmaker. So if p evaluates to any value other than 1, then it is not the case that it (determinately) has a truthmaker (it may determinately lack one, or there may be no determinate fact either way), in which case $\neg p$ evaluates to 1. This line of argument may be resisted and it is not my intention here to defend it; my point is merely that it does not make the equivocation claimed by Schang.

Mark Jago
Philosophy, Nottingham
IEG, Oxford &
Philosophy of Information, Hertfordshire

Can we *really* falsify truth by dictat?

In a talk given on 25/5/96 at a BSPS conference in Oxford, on the Gödelian argument, J. R. Lucas (1996. *The Gödelian Argument: Turn Over the Page*) commented:

> ... in the case of First-order Peano Arithmetic there are Gödelian formulae (*many, in fact infinitely many, one for each system of coding*) which are not assigned truth-values by the rules of the system, and which could therefore be assigned either TRUE or FALSE, each such assignment yielding a logically possible, consistent system. These systems are random vaunts, all satisfying the core description of Peano Arithmetic.

Can we *really* falsify Arithmetical truth by such a dictat?

In other words, if $[(\forall x)R(x)]$ is the PA-unprovable Gödelian formula, which Gödel (1931. *On formally undecidable propositions of Principia Mathematica and related systems I*. In M. Davis. 1965. *The Undecidable*) interprets as Tarskian-true intuitively under the standard interpretation, can we *really* add its negation, $[\neg(\forall x)R(x)]$, as an axiom to PA, and *still* obtain a consistent system with a putative non-standard model of Arithmetic?

Here, $[(\forall x)R(x)]$ is the formula that Gödel (1931. p25) defines, and refers to, by its Gödel-number, 17*Genr*.

Further, $[(\forall x)R(x)]$ is Tarskian-true under the standard interpretation if, and only if, the arithmetical relation, $R(x)$, holds for any given natural number.

Prima facie, Lucas's assumption appears to be based on a counter-intuitive interpretation, and extension, of

Gödel's (1931. p27) assertion that, if an arithmetic such as PA is ω-consistent, then the system PA+$[\neg(\forall x)R(x)]$, say PA*, is consistent, but not ω-consistent.

Gödel defines PA as ω-consistent if, and only if, there is no PA-formula such as $[R(x)]$ for which:

(i) $[\neg(\forall x)R(x)]$ is PA-provable,

and:

(ii) $[R(n)]$ is PA-provable for any given numeral $[n]$ of PA.

Classically, any first order theory with equality is consistent if, and only if, it has a model. Gödel's postulation of the consistency of PA*, therefore, implies that it has a model.

Further, an implicit belief of classical theory is that any consistent first order mathematical theory can be interpreted *suitably* within a set theory such as Zermelo-Fraenkel (ZF), under which every model of ZF is a model of the theory.

Now, the *suitable* interpretation of the primitive symbols of PA in ZF which *transforms* the axioms of PA into theorems of ZF, whilst preserving its rules of inference, is the one (e.g., 1964. Elliott Mendelson. *Introduction to Mathematical Logic*. Van Norstrand. p192) that restricts the range of the interpreted variables to Cantor's first limit ordinal, ω, so that the PA-formula $[(\forall x)R(x)]$, for instance, *transforms* as $[(\forall x)((x \in \omega) \rightarrow R(x))]$.

Clearly, every model of ZF is, then, a model of the *transformed* axioms of PA.

However, does such an interpretation also assure us of a ZF model for the, similarly *transformed*, axioms of PA*?

Now, for PA* to have a model in ZF, in the above sense, we would need $[\neg(\forall x)((x \in \omega) \rightarrow R(x))]$ to be a theorem of ZF.

This, however, is not possible if ZF is consistent.

The reason: Since $[(\forall x)R(x)]$ is Tarskian-true under the standard model of PA, it follows, from Gödel's Completeness Theorem, that the ZF-formula $[(\forall x)((x \in \omega) \rightarrow R(x))]$ is already a theorem of ZF, as it is true in every model of ZF.

Gödel's Completeness Theorem: In any first-order predicate calculus, the theorems are precisely the logically valid well-formed formulas (*i.e., those that are true in every model of the calculus*).

The above argument holds for every interpretation of PA*, since (*see following section*) the Induction Axiom of PA would hold if, and only if—as in ZF—we can introduce an element, in the domain of the interpretation, which restricts the range of the interpreted variables to natural numbers.

Induction Axiom of PA: For any formula $F(x)$ of PA:

$F(0) \rightarrow ((\forall x)(F(x) \rightarrow F(x')) \rightarrow (\forall x)F(x))$

So, was Gödel's postulation a case of a falsifiable conjecture, or are there alternative arguments for concluding that PA+$[\neg(\forall x)R(x)]$ has a non-standard model, and is, therefore consistent?

The following argument suggests that the question does not admit a simplistic answer.

Does First-Order PA **Really** Admit a Non-Standard Model?

Let $G(x)$ denote the PA-formula:

$[x = 0 \vee \neg(\forall y)\neg(x = y')]$

This translates, under every interpretation of PA, as:

Either x is 0, or x is a 'successor'.

Now, in every interpretation of PA:

(a) $G(0)$ is true;

(b) If $G(x)$ is true, then $G(x')$ is true.

By Gödel's completeness theorem:

(c) PA proves $[G(0)]$;

(d) PA proves $[G(x) \rightarrow G(x')]$.

By Generalisation:

(e) PA proves $[(\forall x)(G(x) \rightarrow G(x'))]$;

Generalisation in PA: $(\forall x)A$ follows from A.

By Induction:

(f) $[(\forall x)G(x)]$ is provable in PA.

Hence, except 0, every element in the domain of any interpretation of PA is a 'successor'.

Further, x can only be a 'successor' of a unique element in any interpretation of PA.

Now, since Cantor's first limit ordinal, ω, is not the 'successor' of any ordinal in the sense required by the PA axioms, and there are no infinitely descending sequences of ordinals, every set-theoretical interpretation of PA must, therefore, be restricted to the domain that consists only of the ordinal 0, and the ordinals that are the 'successors' of 0.

Although we *can* define a model of Arithmetic with an infinite descending sequence of elements (e.g., Boolos, Burgess and Jeffrey 2003: *Computability and Logic*, 4th ed., CUP, Section 25.1, p303), any such model is isomorphic to the "*true arithmetic*" (ibid. p150. Ex. 12. 9) of the integers (*negative plus positive*), and *not* to any model of first-order PA (ibid. Corollary 25.3, p306).

Moreover, since we cannot add a *non-successor* constant, say c, to PA such that $c \neq 0, c \neq 1, c \neq 2, \ldots$, we *cannot* apply the Compactness Theorem and the Löwenheim-Skolem Theorem (ibid. p306) to conclude that first-order PA has a non-standard model!

Hence PA admits no non-standard models!

Bhupinder Singh Anand
Mumbai

§3
News

Scalable Uncertainty Management, 10–12 October 2007

During the second week of October, a group of researchers from two different communities got together at the University of Maryland College Park to hold the First International Conference on Scalable Uncertainty Management. For some time now, there has been an apparent need for researchers in Artificial Intelligence and those in Databases to be aware of the work that is carried out in each other's communities, as well as the problems that are faced when formulating their approaches. It is not surprising then to observe that there is no unified community working on managing huge amounts of uncertainty when processing huge amounts of data. The goal of this first SUM conference was to get the ball rolling in this direction, and in my opinion it was a success.

Researchers from all over the world gathered to present and discuss their work with others, covering topics as varied as schema matching, data integration, database repairs, consistent query answering, probabilistic logic programming, semantic web, probabilistic graph identification, temporal and probabilistic databases, and top-k retrieval, among others. Because of their strong orientation towards real-world applications, work carried out in these topics is very likely to run into issues concerning both Databases and AI. As an illustrative example, data integration is a topic that exhibits these properties. Researchers in AI may not be familiar with the vast amount of work carried out in the Databases community, such as the development of architectures and techniques for Datawarehousing, as well as the all-important hands-on experience that many individuals have picked up as a result of having been involved in the development of working systems. On the other hand, researchers in the Databases community can surely gain from familiarizing themselves with research done in the AI community for developing algorithms and systems for reasoning in the Semantic Web, since these tools can certainly inform the process of integrating heterogeneous databases successfully. It is with this spirit that this conference came to be.

All sessions contained presentations of great quality, and in many opportunities I witnessed how people gathered during the breaks to enthusiastically ask each other questions about their work. I believe this attitude is a clear sign that the objective of the conference was met with success, which will surely have a positive effect on the work being done in these shared areas. The next edition of SUM, to be held next year in October at the University of Napoli, will undoubtedly shed more light on how this newfound collaboration between communities is coming along.

Finally, I would like to invite the interested readers to browse the conference proceedings, which can be found at `http://www.springerlink.com/content/978-3-540-75407-7/`.

Gerardo Ignacio Simari
University of Maryland College Park

Calls for Papers

Competing Fusion Methods: Real-world Performances, special issue of Information Fusion, deadline 15 January 2008.

Hybrid Logic: Special Issue of the Journal of Logic, Language and Information, deadline 1 March 2008.

Machine Learning in Space: Special Issue of the Machine Learning Journal, deadline 31 March 2008.

Conditionals and Ranking Functions: Special issue of Erkenntnis, franz.huber@uni-konstanz.de, deadline 31 May 2008.

§4
Introducing ...

In this section we introduce a selection of key terms, texts and authors connected with reasoning. Entries will be collected in a volume *Key Terms in Logic*, to be published by Continuum. If you would like to contribute, please click here for more information. If you have feedback concerning any of the items printed here, please email thereasoner@kent.ac.uk with your comments.

Paradox

A paradox is a piece of reasoning that leads from apparently true premises, via apparently acceptable steps of inference, to a conclusion that is contradictory or in some other way unacceptable. Typically, the reasoning is utterly simple, so it is very alarming that we should be led astray in this way, and we have to confront the real possibility that some of our deepest beliefs or most fundamental principles of inference are wrong. This is the reason why the study of paradoxes is so important. A satisfactory solution will not only expose the basic error afflicting a given paradox, but will also account for how we were taken in by it. Here is W.V. Quine's lovely formulation of *The Barber of Alcala* paradox:

> Logicians tell of a village barber who shaves all those villagers—and only those—who do not shave themselves. The question of the barber's own toilet holds a certain fascination

> for the logical mind. For it has been agreed that the barber shaves any villager, x, if and only if x does not shave himself; hence when we let x be the barber, we conclude that he shaves himself if and only if he does not.

The assumptions here are: There is a village where shaving is *de rigeur* for adult males. Some shave themselves; all those who don't shave themselves get shaved by the sole barber in the village, himself an adult male needing to be shaved. Yet the assumption that he shaves himself leads to the conclusion that he does not (for his job is to shave only those who do not shave themselves); and the assumption that he does not shave himself leads to the conclusion that he does shave himself (because it's his job to shave precisely those who do not shave themselves).

Some philosophers hold that there is an easy solution to this paradox, namely to reject the assumption that there can be a village containing one adult male barber who shaves all and only those adult male villagers who do not shave themselves. Well, perhaps, but *Russell's paradox* has a structure similar to *The Barber*, yet no such easy solution seems available for it. The class of horses contains only horses as its members; in particular, it does not contain classes and so does not contain the class of horses. So the class of horses does not contain itself as a member; it is a non-self-membered class. By contrast, the class of all things that are not horses does contain itself as a member. Now consider the Russell Class R that contains all and only the non-self-membered classes. Does it contain itself as a member? The situation (compare the last sentence of Quine's formulation of *The Barber*) is that R contains any class x as a member if and only if x does not contain itself; hence when we let x be the Russell Class R, we conclude that R is a member of R if and only if it isn't! Can we now suggest, in parallel with what was suggested as a solution to *The Barber*, that R does not exist? Well, the class of horses exists, as do may other non-self-membered classes. And R just collects up all these non-self-membered classes. So how could it *not* exist?

Perhaps the most famous paradox, the discovery of which is attributed to the ancient Greek Eubulides, is called *The Liar* and concerns a person who says 'This statement is false'. You can quickly see that if the statement is true, then it is false and if false, is true. Let us make the assumption, then, that it is neither true nor false, i.e., not true and *not false*. But then, since the statement claims itself to be false, it must be false—contrary to our assumption. One way out of this impasse is the Dialetheist proposal that the Liar statement, and others like it, are at once *both* true and false. A somewhat less outrageous solution (though it needs a lot of careful defending) is that the Liar *sentence* fails to make a *statement* and so does not get into the true/false game. A near-relation of the Liar is the *Curry-Löb Paradox*, one version of which takes as its starting point the statement 'If this statement is true, then pigs can fly'. Or consider this closely related version: 'Either this statement is false or pigs can fly'. That statement cannot be false, for if it were then the first disjunct would be true, hence the whole statement would be true. So it must be true, therefore the first disjunct is false, so the whole statement can be true only if the second judgement is true. Hence pigs can fly!

Of the ancient paradoxes attributed to Zeno of Elea, the most compelling concerns a race between swift *Achilles* and a tortoise who starts half way up the racetrack. The race begins, and Achilles quickly reaches the point where the tortoise was—but, of course, the tortoise has then moved a little bit ahead to a new point. When Achilles reaches that new point, the tortoise has moved ahead, if only a small distance. And so on. The argument seems to show that Achilles can never catch the tortoise, for whenever he reaches where the tortoise was, the tortoise is ahead. Yet common sense, elementary mathematics, or watching a re-creation of the event discloses that Achilles *does* catch and overtake the tortoise.

The *Sorites Paradox*, in its original form, concerns a large number of grains of sand piled together, indisputably a heap. Surreptitiously remove one grain and you still have a heap, in fact the removal of one grain makes no perceptible difference. Remove another grain. Indistinguishable difference; still a heap . . . At no point will the removal of one grain transform a heap into a non-heap. But that seems to show that, continuing to remove one grain at a time, you will have to call a heap three grains of sand. . . two. . . one. . . nought!

Laurence Goldstein
Philosophy, Kent

§5
Letters

Dear Reasoners,

It has been pointed out to me that on page 7 of *The Reasoner* 1(8), December 2007, para (c) of my article 'A constructive definition of the intuitive truth of the Axioms and Rules of Inference of Peano Arithmetic' is inaccurately worded, and should read:

(c) If a Turing Machine T computes '$F(0)$' as TRUE, and T also computes '$(F(x) \rightarrow F(x'))$' as always TRUE, then T computes $F(x)$ as always TRUE.

Sincerely,

Bhupinder Singh Anand
Mumbai

§6
Events

January

ISAIM: Tenth International Symposium on Artificial Intelligence and Mathematics, Fort Lauderdale, Florida, 2-4 January.

3rd IMS and ISBA meeting: The third joint international meeting of the IMS (Institute of Mathematical Statistics) and ISBA (International Society for Bayesian Analysis), Bormio, Italy, 9–11 January.

Perspectives on Truth: University of Nottingham, 11–12 January.

Graduate Conference: 1st Cambridge Graduate Conference on the Philosophy of Logic and Mathematics, St. John's College, Cambridge, 19–20 January.

Bayesian Biostatistics: Houston, Texas, 30 January – 1 February.

February

Logic Meeting: University of California, Los Angeles, 1–3 February.

FoIKS: Foundations of Information and Knowledge Systems, Pisa, Italy, 11–15 February.

March

Relativism and Rational Reflection: 10th Annual Pitt–CMU Graduate Student Philosophy Conference, University of Pittsburgh, 1 March.

Artificial General Intelligence: The First Conference on Artificial General Intelligence, Memphis, Tennessee, 1–3 March.

Science and Pseudoscience: University of Birmingham, UK, 15 March.

Constraint-Sac: Track on Constraint Solving and Programming, at the 23rd Annual ACM Symposium on Applied Computing, Fortaleza, Brazil 16–20 March.

Causation: 1500-2000: King's Manor, University of York, 25–27 March.

UnClog: International Workshop on Interval / Probabilistic Uncertainty and Non-Classical Logics, Ishikawa, Japan, 25–28 March.

April

AISB: Artificial Intelligence and Simulation of Behaviour, Aberdeen, 1–4 April.

Subjective Bayesian Methods: Department of Probability and Statistics, University of Sheffield, 2 April.

RelMiCS10-AKA5: 10th International Conference on Relational Methods in Computer Science & 5th International Conference on Applications of Kleene Algebra, Frauenwörth, Germany, 7–11 April.

Reduction and the Special Sciences: Tilburg Center for Logic and Philosophy of Science, 10–12 April.

FLOPS: Ninth International Symposium on Functional and Logic Programming, Ise, Japan, 14–16 April.

Workshop: XVIII Inter-University Workshop on Philosophy and Cognitive Science, Madrid, luis.fernandez@filos.ucm.es, 22–24 April.

Practical Rationality: Intentionality, Normativity and Reflexivity, University of Navarra, 23–25 April.

SDM: 8th Siam International Conference on Data Mining, Hyatt Regency Hotel, Atlanta, Georgia, USA, 24–26 April.

May

SBIES: Seminar on Bayesian Inference in Econometrics and Statistics, University of Chicago Graduate School of Business Gleacher Center, 2–3 May.

SIG16: 3rd Biennial Meeting of the EARLI-Special Interest Group 16—Metacognition, Ioannina, Greece, 8–10 May.

CLE, EBL & SLALM: 30th Anniversary of the Centre for Logic, Epistemology and the History of Science (CLE), UNICAMP, 15th Brazilian Logic Conference, and 14th Latin-American Symposium on Mathematical Logic, Paraty, Brazil, 11–17 May.

ArgMAS: Fifth International Workshop on Argumentation in Multi-Agent Systems, Estoril, Portugal, 12–13 May.

UR: Special Track on Uncertain Reasoning, 21st International Florida Artificial Intelligence Research Society Conference (FLAIRS-21), Coconut Grove, Florida, 15–17 May.

AI Planning and Scheduling: A Special Track at the 21st International FLAIRS Conference (FLAIRS 2008), Coconut Grove, Florida, 15–17 May.

RSKT: Rough Sets and Knowledge Technology, Chengdu, 17–19 May.

ISMIS: The Seventeenth International Symposium on Methodologies for Intelligent Systems, York University, Toronto, Canada, 20–23 May.

COMMA: Second International Conference on Computational Models of Argument Toulouse, France, 28–30 May.

AI: 21st Canadian Conference on Artificial Intelligence, Windsor, Ontario, 28–30 May.

Expressions of Analogy: Faculty of Social and Human Sciences, New University of Lisbon, 29–31 May.

June

WCCI: IEEE World Congress on Computational Intelligence, Hong Kong, 1–6 June.

CSHPS: Canadian Society for History and Philosophy of Science, University of British Columbia, Vancouver, 3–5 June.

CiE: Computability in Europe 2008: Logic and Theory of Algorithms, University of Athens, Athens, 15–20 June.

DM: SIAM Conference on Discrete Mathematics, University of Vermont, Burlington, Vermont, 16–19 June.

Logica: Hejnice, Czech Republic, 16–20 June.

IEA-AIE: 21st International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems, Wroclaw, Poland, 18–20 June.

HOPOS: Seventh Congress of the International Society for the History of Philosophy of Science, Vancouver, Canada, 18–21 June.

EPISTEME: Law and Evidence, Dartmouth College, 20–21 June.

IPMU: Information Processing and Management of Uncertainty in Knowledge-Based Systems, Malaga, Spain, 22–27 June.

Med: 16th Mediterranean Conference on Control and Automation, Ajaccio, Corsica, 25–27 June.

July

LOFT: 8th Conference on Logic and the Foundations of Game and Decision Theory, 3–5 July.

ICML: International Conference on Machine Learning, Helsinki, 5–9 July.

CAV: 20th International Conference on Computer Aided Verification, Princeton, 7–14 July.

Bayesian Modelling: 6th Bayesian Modelling Applications Workshop, Helsinki, 9 July.

UAI: Uncertainty in Artificial Intelligence, Helsinki, 9–12 July.

COLT: Conference on Learning Theory, Helsinki, 9–12 July.

IKE: International Conference on Information and Knowledge Engineering, Las Vegas, 14–17 July.

NorMAS: 3rd International Workshop on Normative Multiagent Systems, Luxembourg, 15–16 July.

DEON: 9th International Conference on Deontic Logic in Computer Science, Luxembourg, 15–18 July.

NCPW: 11th Neural Computation and Psychology Workshop, Oxford, 16–18 July.

ISBA: 9th World Meeting, International Society for Bayesian Analysis, Hamilton Island, Australia, 21–25 July.

Model Selection: Current Trends and Challenges in Model Selection and Related Areas, University of Vienna, 24–26 July.

First Formal Epistemology Festival: Conditionals and Ranking Functions, Konstanz, 28–30 July.

August

Conference: Language, Communication and Cognition, University of Brighton, 4–7 August, Brighton, UK.

ESSLLI: European Summer School in Logic, Language and Information, Freie und Hansestadt Hamburg, Germany, 5–15 August.

IJCAR: The 4th International Joint Conference on Automated Reasoning, 10–15 August.

ICT: The Sixth International Conference on Thinking, San Servolo, Venice, 21–23 August.

Compstat: International Conference on Computational Statistics, Porto, Portugal, 24–29 August.

September

IVA: The Eighth International Conference on Intelligent Virtual Agents, Tokyo, 1–3 September.

COMSOC: 2nd International Workshop on Computational Social Choice, Liverpool, 3–5 September.

10th Asian Logic Conference: Kobe University, Kobe, Japan, 1–6 September.

Soft Methods for Probability and Statistics: 4th International Conference, Toulouse, France, 8–10 September.

ICAPS: International Conference on Automated Planning and Scheduling, Sydney, 14–18 September.

§7 Jobs

Cambridge Machine Learning: Two Postdocs, deadline 15 January 2008.

Imperial Statistics: Chair and two lectureships, deadline 31 January 2008.

SAMSI Statistics: 7 Postdoc positions, The Statistical and Applied Mathematical Sciences Institute, NC, deadline 31 January 2008.

§8 Courses and Studentships

Courses

Second Indian Winter School on Logic: IIT Kanpur, 14–26 January 2008.

MA in Reasoning

An interdisciplinary programme at the University of Kent, Canterbury, UK. Core modules on logical, causal, probabilistic, scientific and mathematical reasoning and further modules from Philosophy, Psychology, Computing, Statistics and Law.

MLSS: 10th Machine Learning Summer School, Kioloa Coastal Campus, Australian National University, 3–14 March 2008.

Logic School: State University of Campinas, Brazil, 7–9 May.

Probabilistic Causality: Central European University, Budapest, 21 July–1 August.

ESSLLI: European Summer School in Logic, Language and Information, Hamburg, 4–15 August 2008.

Studentships

Cognitive Science: PhD Studentships, Human-Level Intelligence Laboratory, Rensselaer Department of Cognitive Science, deadline January 2008.

Computational Neuroscience: 4-year PhD studentships, Gatsby Computational Neuroscience Unit, University College London, deadline 6 January 2008.

Arché Postgraduate Studentships: The Arché Research Centre at the University of St Andrews is offering up to six three-year PhD studentships for uptake from September 2008, deadline 1 February 2008.

Leeds Philosophy: 2-3 postgraduate studentships in philosophy and history and philosophy of science, deadline 1 March 2008.

Acknowledgements

The Reasoner is a development of the progicnet academic network supported by the Leverhulme Trust.