
THE REASONER

VOLUME 10, NUMBER 4

APRIL 2016

thereasoner.org

ISSN 1757-0522

CONTENTS

Editorial	26
Features	27
News	30
What's Hot in ...	31
Events	34
Courses and Programmes	34
Jobs and Studentships	35

EDITORIAL

Constructing a causal Bayesian network (CBN) of the situation at issue is often a good technique for clarifying informal arguments. In particular, such networks can clarify how much confidence we should have in the conclusion based on the evidence. Here we sketch very briefly how this is done and what its relative advantages are to alternatives. We follow this by interviewing one of the leading proponents of this technique, computer scientist [Norman Fenton](#), who has notably applied it to legal cases. This technique is a valuable tool for any keen reasoner, and we hope to see it applied more broadly.



Kevin Korb

Consider the following simple argument:

26
27
30
31
34
34
35

We believe that Smith murdered his wife. A large proportion of murdered wives turn out to have been murdered by their husbands. Indeed, Smith's wife had previously reported to police that he had assaulted her, and many murderers of their wives have such a police record. Furthermore, Smith would have escaped from the scene in his own blue car, and a witness has testified that the car the murderer escaped in was blue.

Unlike many informal arguments, the expression is already simple and clear: the conclusion is stated upfront, the arguments are clearly differentiated, and there is no irrelevant verbiage. Like most informal arguments, however, they are probabilistic enthymemes: they support the conclusion probabilistically rather than deductively, and rely on unstated premises. So, it's hard to give a precise evaluation until we make both probabilities and premises more explicit, and combine them appropriately.



Erik P. Nyberg

We can use this simple CBN to assess the argument:

Wife reported assault → Smith murdered wife → Car blue
→ Witness says car blue

The arrows indicate a causal influence of one variable on the probability distribution of the next variable. In this case, these

are simple Boolean variables, and if one variable is true then this raises the probability that the next is true, e.g., if Smith did assault his wife then this caused him to be more likely to murder his wife. (It could be that spousal assault and murder are actually correlated by common causes, but this wouldn't alter the probabilistic relevance of assault to murder, so we can ignore the possibility here.)

First, we can research crime statistics to find that 38% of murdered women were murdered by their intimate partners, and so get our probability prior to any other evidence.

Second, we can establish that 30% of women murdered by their intimate partners had previously reported to police being assaulted by those partners. Admittedly, as O. J. Simpson's lawyer argued, the vast majority of husbands who assault their wives do not go on to murder them. However, his lawyer was wrong to claim that Simpson's assault record was therefore irrelevant! We just need to add some additional probabilities, which a CBN forces us to find, and combine them appropriately, which a CBN does for us automatically. Suppose that in the general population only 3% of women have made such reports to police, and this factor doesn't alter their chance of being murdered by someone else. Then it turns out that the assault information raises the probability of Smith being the murderer from 38% to 86%.

Third, suppose we accept that if Smith did murder his wife, then the probability of him using his own blue car is 75-95%. Since this is imprecise, we can set it at 85% (say) and vary it later to see how much that affects the probability of the conclusion (in a form of sensitivity analysis).

Fourth, we can test our witness to see how accurate they are in identifying the colour of the car in similar circumstances. When a blue car drives past, they successfully identify it as blue 80% of the time. Should we conclude that the probability that the car was blue is 80%? This would be an infamous example, due to Tversky and Kahneman, of the Inverse Fallacy. In fact, we also need to know how successfully the witness can identify non-blue cars as non-blue (say, 90%) and the base rate of blue cars in the population (say, 15%). Then it turns out that the witness testimony alone would raise the probability that Smith was the murderer from 38% to 69%. If this is combined with the assault information, then the updated probability that Smith is the murderer rises to 96%.

Even this toy example illustrates that building a CBN forces one to think about how the main factors are causally related and to research all the necessary probabilities. Assuming the CBN is correct for the variables considered, and is built in one of many good BN software tools, it acts as a useful calculator: it combines these probabilities appropriately to calculate the probability of our conclusion. Thus, it helps prevent much of the vagueness and fallacious reasoning that are widespread, even in important legal arguments.

Versus alternatives

Although there are genuine difficulties in using this technique, we believe that much of the resistance to it is based on imaginary difficulties, while the (*italicized*) rival techniques below have difficulties of their own. In our toy example, *the prose version of the argument* doesn't quantify the probabilities involved, doesn't specify the missing premises, doesn't indicate how the various factors are related to each other, and it's far from clear how to compute an appropriate probability for the

conclusion. The fact that the probabilities and premises aren't specified doesn't really make the argument non-probabilistic, it just makes it vague. Prose is often the final form of presenting an argument, but it is far from ideal for the prior analysis of an argument.

Resorting to techniques from *formal logic*, diagrammatic or otherwise, requires even more effort than CBN analysis, while typically losing information. It is really appropriate only for the most rigorous possible examination of essentially deductive arguments.

A more recent approach with some promising empirical backing is *the use of argument maps*. These are typically unparameterised non-causal tree structures in which the conclusion is the trunk and all branches represent lines of argument leading to it. (See Tim van Gelder's '[Critical Thinking on the Web](#)'.) Arguably, these are equivalent to a restricted class of Bayesian network without explicit parameters. Thus, they have many of the advantages of BNs but they don't provide much guidance in computing probabilities, so they can be vague and subject to the kinds of fallacious reasoning that are avoided with actual BNs. Also, as they are typically not causal, they can actually encourage misunderstanding of the scenario.

(For a more detailed version of this editorial, see our '[Arguments Using Causal Bayesian Networks](#)'.)

KEVIN KORB

ERIK P. NYBERG

Monash University

FEATURES

Interview with Norman Fenton

Erik P. Nyberg & Kevin Korb: As a computer scientist, how did you get involved in legal argument analysis?

Norman Fenton: It was through using Bayesian networks (BNs) in the 1990's in a very highly related area, namely the assessment of safety critical systems. In that area, we had the problem of determining whether or not a system was 'safe enough' to release, and the decision had to be made based on a combination of evidence—some of it quantitative, like test results, but much of it more subjective, like the quality of the design process. The situation is highly analogous to a typical legal case where we have to determine whether a defendant is guilty; some of the evidence (e.g., forensic match evidence) is quantitative and some of it is subjective (e.g., an alibi). We found that BNs are the ideal formalism for combining evidence in the safety critical case and so started to apply them to the law. At first, I used them to examine—and resolve—fallacies of legal reasoning (such as the [Prosecutor's Fallacy](#)), but I then got involved in a number of live cases where the lawyers needed help in understanding statistical and probabilistic evidence.



EPN & KK: Tell us about your [Bayes and the Law](#) project. How did that get started? Has it had the impact you were hoping for? What's in its future?

FN: Through my work as an expert witness/consultant on a number legal cases after 2005 I got to know a lot of the people who had been either using Bayes in the context of legal arguments or who were interested in the issues and controversies surrounding its use. After the [R v T case in 2011](#) when the judge on appeal ruled that Bayes should not be used to assess evidence (with the exception of DNA—an irrational exception in my view) a lot of people felt that there were many misunderstandings in the judgment. Whereas there had been previous initiatives to promote the use of Bayes in the law, I felt that lawyers themselves had not been sufficiently involved. Hence, with a very small grant from Queen Mary I set up the Bayes and the Law consortium primarily to improve communication between Bayesians, forensic scientists and lawyers in order to properly understand the role of Bayes in the Law. The project has been quite successful in terms of raising awareness and we have hosted a number of meetings—some of which involved several practicing lawyers (including the defense barrister in the R v T case). The project has fed into a number of other initiatives, the most important of which is the prestigious 6-month Programme on Probability and Statistics in Forensic Science at the Isaac Newton Institute for Mathematical Sciences, University of Cambridge, which is running from July to December 2016 (see [here](#)). Part of this also includes a number of events specifically targeted at lawyers. I am hopeful that by the end of the programme there will be a better understanding of where and how Bayes should be used in the law and the lawyers will be much more supportive of it.

EPN & KK: What are the most common objections when a Bayesian net is presented in the context of argumentation? Do these accurately reflect the real difficulties you have in constructing one?

NF: The most common objections are that prior probabilities have to be explicitly included, something which has little to do with the difficulty of constructing a model.

EPN & KK: Do these objections tend to be motivated by some kind of anti-Bayesian philosophy (e.g., frequentist statistics)? Are these misgivings usually allayed when you are able to provide empirically-based prior distributions?

NF: There is definitely a resistance to the Bayesian idea of using prior probabilities and obviously this is especially strong when subjective judgment is used. The use of empirically-based prior distributions does not always help, because there is often criticism that the empirical data is not sufficiently representative/rigorous or is based on inadequate sample sizes. What the objectors often fail to understand is that this is a limitation of all statistical sampling and that (in contrast to the frequentist approach) the Bayesian approach enables us to easily incorporate our uncertainty about the data. So, for example, while data for DNA profiles is widely recognised as being sufficiently “statistically rigorous” for use in courts—in contrast, say, to data on shoe prints—people are unaware that

in both cases (DNA and shoe prints) there are limitations in the sample databases. And in both cases the uncertainty about these limitations can be incorporated using Bayes.

EPN & KK: How do you go about putting numbers (unconditional and conditional probabilities) into a Bayesian net? How accurate do these need to be?

NF: This is indeed a major challenge for those variables (i.e., nodes) for which there is no data except expert judgment, which is why I always present the results with a range of priors for the most critical/controversial nodes.

EPN & KK: What problems have you encountered in choosing variables or graph structure?

NF: We have tried to adopt a legal idiom-based approach (see, e.g., Fenton, N. E., D. Lagnado and M. Neil 2013: ‘A General Structure for Legal Arguments Using Bayesian Networks’ *Cognitive Science* 37, 61-102) to minimise the ‘choice’ difficulty, but there remains a problem in incorporating the notions of ‘opportunity’ and ‘motive’. The idioms have these nodes as parents of offence-level (or activity level) hypotheses like “Defendant committed the crime”, which is correct from a causal perspective. However, ideally we would prefer such hypotheses to have no parents so that we can assign unconditional priors. In fact, ideally, the unconditional prior for “Defendant committed the crime” would be one contextualised to the relevant population (and every piece of evidence—even things like the sex of the person who committed the crime—would be incorporated and would update that prior).

EPN & KK: Do your Bayesian nets include all the relevant evidence? Do they calculate when the probability of guilt is greater than 90% (or any other threshold)?

NF: It depends on the case. I have done quite a lot of work on the impact of forensic evidence in particular cases. Although I would like to include *all* evidence, I have been forced to restrict it in those cases. So, e.g., the key unknown variable might be “defendant is source of DNA trace found at scene” (the model would *not* include offense-level variables like “Defendant committed crime”) and the model incorporates evidence about the DNA matching process, the possibility of errors (including contamination of samples), the quality of samples, etc. The model will calculate the probability of any unknown variable after observing the evidence, so if the ‘guilty’ node is included (which, as I said, is not always the case) then, yes, it will calculate whether or not a particular threshold has been met.

EPN & KK: How do Bayesian nets help with understanding legal argument? In which cases have they made the most difference?

NF: I have written a review of this in a paper about to be published: Fenton, N. E., M. Neil and D. Berger (forthcoming: ‘Bayes and the Law’). In summary, there have been no publicised uses of BNs in court, and almost all the published articles talking about the use of BNs in real cases provide examples of how BNs could have avoided problems, improved the arguments, etc. However, based on my own experience, there may be many other unpublicised cases where BNs have

been used ‘in the background’. From my own experience I have used BNs to help lawyers understand the impact of evidence in a range of criminal cases including murder, rape, assault, and theft, and civil cases of medical negligence. In each case, the lawyers have subsequently used a completely informal presentation of what the BN is saying. The most important one is an ongoing rape case where my analysis of the errors in the presentation of the DNA evidence will hopefully lead to the conviction being overturned.

EPN & KK: Likelihoods have been infamously abused in legal argument, as in the Sally Clark case. How can Bayesian nets help there?

NF: Using a BN easily avoids the Prosecutors Fallacy and the kind of errors made in the Sally Clark case, but more importantly a BN can avoid errors that even some ‘Bayesians’ make when using the likelihood ratio (LR). In particular, the practice has been to simplify the argument to such an extent that manual calculation of the LR is possible. Essentially this means everything gets reduced to a 2-node model (*Hypothesis* → *Evidence*). But the reality is that even in the simplest cases there is more complexity than this (e.g., for the simplest DNA evidence you need a minimum 5-node model if you wish to incorporate even the crudest possibility of errors). Now building and running the 5-node model (and getting the LR for the evidence) in a BN tool is absolutely trivial. But if you tried to do the calculations by hand it is more or less impossible—which is why people who are unaware of BNs and BN tools simply ignore it and get it all wrong as a result. This is discussed extensively in these two papers:

Fenton, N. E., Neil, M., and Hsu, A. (2014). ‘Calculating and understanding the value of any type of match evidence when there are potential testing errors’. *Artificial Intelligence and Law* 22, 1-28.

Fenton, N. E., D. Berger, D. Lagnado, M. Neil and A. Hsu, (2014). ‘When ‘neutral’ evidence still has probative value (with implications from the Barry George Case)’. *Science and Justice* 54(4), 274-287.

EPN & KK: Can Bayesian nets help you (or anyone) avoid other common probabilistic fallacies?

NF: Yes, we have used BNs to show how to avoid a number of fallacies including: Defendant Fallacy, Confirmation Bias Fallacy, Base Rate Neglect, treating dependent evidence as independent, Coincidences Fallacy, various evidence utility fallacies, Cross Admissibility Fallacy, ‘Crimewatch UK’ Fallacy. Some of these were explained in the paper:

Fenton, N. E. and Neil, M. (2011), ‘Avoiding Legal Fallacies in Practice Using Bayesian Networks. *Australian Journal of Legal Philosophy* 36, 114-151.

EPN & KK: How would you compare argument analysis using argument diagrams (e.g., the unparameterised trees supported by Tim van Gelder’s Rationale software) to using Bayesian nets? What are the relative advantages and disadvantages?

NF: I see these methods as complementing the BN approach, especially as BNs are not well suited to modelling genuinely

contradictory ‘narratives’. For example, when the prosecution and defence both propose complex—and mutually exclusive—narratives to explain a crime then this becomes a problem to model in a BN because most of the evidence will only be meaningful for one of the narratives but with a BN model you have to still consider its impact on the other. Argument diagrams, while unable to quantify the probability of the unknown hypotheses, do enable you to easily model such alternative narratives.

EPN & KK: How can Bayesian net technology be improved to better support argument analysis?

NF: Following on from the previous question, we are currently looking at a new approach to BN argumentation which enables different narratives to be modelled as separate BN models while using Bayesian model comparison to determine which one better supports the evidence. There are also many ways in which GUIs could be configured to interface with legal BN models, making it easy for lay people to see and input only the things relevant for them.

EPN & KK: What are the major challenges to getting BNs accepted as a standard method for evaluating the impact of evidence in legal arguments?

NF: First, we have to make sure that statisticians (and even some Bayesians) are actually properly aware of BNs and the state-of-the-art tools that support them, because otherwise lawyers will keep hearing from statisticians that it is not possible/feasible to do proper Bayesian reasoning and that the only thing we can do is the simplest LR calculations by hand.

Second, we need to focus on the use of BNs where they are most effective, namely in analysing evidence pre-trial (to determine what is important and what is not) and even to help prosecutors to determine whether a case should proceed to trial. I would not expect BNs to be used routinely ‘in the courtroom’ in the foreseeable future.

Third, we need a strategy that convinces lawyers that nobody in a case (be it lawyers, forensic scientists or even juries, if it does come to the courtroom) needs to understand the underlying Bayesian inference calculations in a BN model. All they need to know are the assumptions about the model structure and the prior probabilities—so we need the ‘calculator analogy’ whereby people accept the result of, say, a long division in a calculator without having to understand the underlying algorithm and circuit level calculations that produce the result (only the inputs need to be discussed and agreed).

Fourth, we also need more standard BN idioms and templates that capture commonly occurring legal arguments (such as for DNA evidence, alibis, etc).

EPN & KK: Thank you!

Against the Brogaard-Salerno Stricture

‘It is widely agreed that contraposition, strengthening the antecedent and hypothetical syllogism fail for subjunctive conditionals’, write Brogaard and Salerno (2008: Counterfactuals and context, *Analysis* 68(1), 39–46). In that article they argue that the putative counterexamples to these principles are actually no threat, on the grounds that they involve a certain kind of

illicit contextual shift.

Here I argue that this particular kind of contextual shift, if it is properly so called, is not generally illicit, and that therefore the counterexamples shouldn't be blocked with the kind of blanket restriction Brogaard and Salerno advocate. The idea that the reasoning patterns in question can be vindicated given restrictions still seems promising; the purpose of this note is to show that the simple restriction proposed by Brogaard and Salerno isn't the right way of going.

Brogaard and Salerno conduct their discussion within the framework of the standard Lewisian account of counterfactuals, which says that

a subjunctive of the form 'if A had been the case, B would have been the case' is true at a world w iff B is true at all the A-worlds closest (or most relevantly similar) to w .

(This is the formulation used by Brogaard and Salerno. It is adapted from Lewis 1973: *Counterfactuals*, Oxford, Blackwell.) They introduce the term 'background facts' to designate 'the respects in which A-worlds are relevantly similar to w '. Thus every counterfactual in a particular context, on the standard theory, is attached to a set of background facts. Now, the central claim of their article is that "the set of contextually determined background facts must remain fixed when evaluating an argument involving subjunctives for validity". One set of background facts per argument. Let us call this the *Brogaard-Salerno Stricture*. Brogaard and Salerno say that to flout this stricture is to make an illicit contextual shift, and that since the putative counterexamples to contraposition etc. flout the stricture, they should not be accepted. (While Brogaard and Salerno use Lewis's account, it is important to note that their Stricture, and my argument against it, can be carried over to other accounts which differ from Lewis's in detail but still involve background facts or something like them.)

For an argument to comply with the Brogaard-Salerno Stricture, all counterfactuals occurring within it have to be alike in background facts. What I wish to point out is that this condition is unsatisfied by many valid arguments, including the following:

(P1) If Mary hadn't had breakfast, she would have lunched sooner.

(P2) If John had worn black shoes, he would have worn black socks.

(C) Therefore, if Mary hadn't had breakfast, she would have lunched sooner, and if John had worn black shoes, he would have worn black socks.

The conclusion follows from the premises by conjunction introduction. For the first premise, one of the background facts might be that Mary has a normal appetite. Another might be that she does not like to go hungry. These are plainly irrelevant to the second premise. Conversely, John's sense of style has nothing to do with the first. So the *salient* background facts are different for each premise. More acutely: with (P1) we are certainly *not* including Mary's having had breakfast—let's assume she did, i.e., that (P1)'s antecedent is false—as a background fact, since the conditional is about what she would have done had she not had breakfast. Likewise, with (P2), we are certainly not including the fact—let's assume it is one—that John did not wear black shoes. (I owe this way of making the point to an anonymous referee.) The point is, we cannot stip-

ulate that these premises are attached to the same set of background facts without doing obvious violence to their meaning. These two premises, if they are to be understood the way they are meant to be understood, cannot figure in the same argument without flouting the Brogaard-Salerno Stricture. But the above argument is obviously valid. Therefore the stricture is not appropriate.

That is my argument against the Brogaard-Salerno Stricture. That the Stricture is too strong is a negative result, but there is no reason to think we have reached a dead end here. Brogaard and Salerno's basic idea, that the inference patterns at issue can be vindicated once proper restrictions are observed, has not been seriously threatened. What I have shown is just that their particular approach to the restricting is too simple.

One alternative approach which suggests itself is to place restrictions regarding background facts on particular inference rules, rather than all deductive reasoning occurring within a given argument. We might do well to start with contraposition, strengthening the antecedent and hypothetical syllogism. Other rules may be fair game too. In this connection, consider this passage:

But suppose we are wrong about this. Suppose shifting context mid-inference is no fallacy at all. Then a rather surprising consequence follows. Modus ponens—which many possible world accountants love and cherish—fails too. (Brogaard and Salerno 2008: p. 44.)

On the present suggestion, the evidence for the claim of the last sentence might motivate the view that modus ponens needs to be restricted too—but still, not all deductive reasoning within a given argument. Conjunction introduction, for example, is *prima facie* OK without such a restriction.

This rule-by-rule approach may also be mistaken (or insufficient by itself, or adequate in principle but inelegant compared to some other approach). Also, requiring sameness of background facts, even at the level of particular rules, may in some cases be too simple; for instance, perhaps bringing in a counterfactual whose set of background facts is a proper superset of another's—adding but not subtracting background facts, so to speak—is sometimes allowable. These are no more than suggestions, but their availability indicates that there is a good opportunity for further work here.

TRISTAN HAZE
University of Sydney

NEWS

Science vs Common Sense? 25–27 February

This conference was dedicated to common sense philosophy and its relation to science. The [conference](#) was part of the Beyond Scientism Project, funded by the Templeton Foundation and conducted by a research group at the Free University of Amsterdam (VU). The conference was organized by Rik Peels, Jeroen de Ridder, Irma Verlaan and René van Woudenberg. The key note lectures were favourable to common sense and maintained that science is not really antithetical to common sense beliefs.

On the first day, Noah Lemos laid the groundwork for a discussion of the epistemic status of common sense beliefs. He

discussed the possible positions on common sense beliefs, using the familiar quantificational devices:

(EC) All common sense beliefs are true

(SC) No common sense beliefs are true

(MC) Some common sense beliefs are true.

Lemos deems the extreme claim (EC) to be false, because many of our previous common sense beliefs have been overturned by science. Take the previously widely held views that the earth was flat, that the sun revolves around the earth. Even some of the views that are widely held in common sense ontology nowadays can be deemed false pointing to the rise of quantum mechanics and general relativity theory. So, to say that all common sense beliefs are true is false. The skeptic claim (SC) can be deemed implausible at least. Science relies on common sense beliefs like the veracity of sense perception and the like. Accepting SC would lead to outright skepticism and is antithetical to science at large. Thus only the modest claim (MC) remains. Even science has to admit that at least some common sense beliefs are true. Thus, MC stands as the only plausible position to take in the debate.

In the second keynote lecture, René van Woudenberg argued that our belief in free will at least has some positive epistemic status: it is commonsensical.

The first argument concludes that belief in free will is *practically rational*. Practical deliberation is predicated on the premise that we are able to decide between different courses of action, i.e. that there is a real choice to be made. Therefore, using a definition of practical rationality used by Alston, belief in free will is practically rational because it is indispensable for a practical deliberation.

The second argument argues that belief in free will is *properly basic*: i.e., (i) not held in virtue of other beliefs, (ii) held even if one abides by one's epistemic duties and (iii) held if the belief lacks defeaters. And because we may presume properly basic beliefs to be true, we are entitled to believe in libertarian free will. Most of the supposed defeaters of free will come from neuroscience, but were (without much regard) dismissed out of hand. I am not even sure that (i) and (ii) were argued for, but there did not seem to be a lot of opposition from the audience on these points.

The third argument claims that empirical evidence can show that we were able to do otherwise. It does so by pointing to variety of circumstances and instances where a person, at an earlier time, did in fact do otherwise. Van Woudenberg thus rejects the proposition that belief in libertarian free will violates the standard of naturalistic compatibility proposed by Vargas.

The second day kicked off with a lecture by Russ Shafer-Landau, who discussed two epistemic evolutionary debunking arguments against moral realism and why they fail. The first argument concerns an empirical premise: evolutionary processes have exerted doxastic pressures on some if not all of our moral beliefs and makes all of our moral beliefs suspect. Shafer-Landau deems this premise implausible because we either have a way to discern which moral beliefs are thus pressured or we do not. If we do, then there is a way for us to have moral knowledge. If not, then evolutionary influences are untestable.

The second argument concerns an epistemological premise: we are not justified to believe that we have reliable moral knowledge, because it is not plausible that evolutionary pressures have left us reliable moral faculties. This unreliability can

be defended by (1) appealing to the distortion that evolutionary pressures exert on our moral belief-forming mechanisms or (2) that there are no viable reasons to believe that those pressures hold reliable outcomes for moral beliefs. Shafer-Landau argues against (1) by appealing to a standard of truth, because if those pressures produce unreliable results, then there has to be a way to assess those problematic pressures. He comes out against (2) by saying that what supposedly goes for moral beliefs (i.e., that evolutionary pressures are not adaptive) goes for arithmetic knowledge as well. Shafer-Landau's suggestion is that such a position is absurd.

Katia Vavova posed a dilemma for the Darwinian debunkers. If true moral beliefs and fitness enhancing beliefs come apart and evolution has exerted pressures on our moral beliefs, then some of our moral beliefs may come out false. Now, the debunker can choose to accept whether our moral assumptions are (a) legitimate or (b) illegitimate. If (a), then we cannot say why the evolutionary pressures affect our moral beliefs badly. If (b), then we cannot suppose that there is a gap between true moral beliefs and our accepted moral beliefs, because the gap cannot be established. Hence, both (a) and (b) support the conclusion that there is no good reason to suppose that our moral beliefs are mistaken.

The picture emerging from the keynotes is that common sense has *prima facie* justification. In order to be overturned, viable defeaters have to be produced. Both theoretical and moral skeptics face a similar challenge: to show that our faculties are unreliable without appealing to an independent standard of truth. If they appeal to such a standard, then the inability to gain knowledge can be circumvented. If they do not, then we have no reason to believe that we are mistaken. This makes me wonder, is the burden on the skeptics' side? Isn't arguing for skepticism, moral or otherwise, a fool's errand?

ANDRIES DE JONG
Utrecht University

Calls for Papers

LOGICAL PLURALISM AND TRANSLATION: special issue of *Topoi*, deadline 30 April.

EXPERIMENTAL PHILOSOPHY: special issue of *Teorema*, deadline 30 April.

LOGIC AS TECHNOLOGY: special issue of *Philosophy and Technology*, deadline 1 May.

MEANING AND COMPUTER GAMES: Special issue of *Journal of the Philosophy of Games*, deadline 15 May.

STATISTICAL SIGNIFICANCE AND THE LOGIC OF HYPOTHESIS TESTING: special issue of *Entropy*, deadline 30 May.

A HUNDRED YEARS OF DONALD DAVIDSON. HIS INFLUENCE ON CONTEMPORARY PHILOSOPHY: Special issue of *Argumenta*, deadline 30 June.

THE BACKGROUND OF CONSTITUTIVE RULES: Special issue of *Argumenta*, deadline 10 November.

WHAT'S HOT IN . . .

Uncertain Reasoning

The Oscar winning documentary *Citizenfour* brought the concept of *metadata* to the attention of general audiences. As one scene of the film explains, we leave, mostly unwillingly, many

digital traces of our daily activities. Most Londoners, for instance, use an Oyster card to travel across the city. When they top-up their Oyster online or opt in for the convenient auto top-up, they effectively allow whoever has access to the data to track their routine. (And the recent introduction of contactless payment on the London transport system clearly made this even simpler.) This can then be linked to what people buy, what they read on the internet, what they post on social networks, and indeed, to what other people do. That's metadata.



It goes without saying that metadata is syntax with no semantics. There are many reasons as to why people do what they do, and there are many people traveling independently on the same journey. Quite obviously then, the dots representing their digital traces can be joined in a number of distinct ways, and specific but wrong pictures can be drawn. That's why the Orwellian idea that someone possesses a wealth of metadata about us is indeed frightening. But knowing that governments may kill based on that, is rather hard to accept.

The opening of [this recent piece](#) by C. Grothoff and J.M. Porup on Arstechnica UK is chilling:

In 2014, the former director of both the CIA and NSA proclaimed that “we kill people based on metadata.” Now, a new examination of previously published Snowden documents suggests that many of those people may have been innocent.

The article refers to the US National Security Agency's SKYNET programme which monitors massively Pakistan's mobile phone networks to obtain metadata. The goal is to quantify the likelihood of any particular individual being a terrorist. Data scientist and human right activist Patrick Ball dubs the method used by NSA as “ridiculously optimistic” and “completely bullshit.” The reported result is appalling:

...thousands of innocent people in Pakistan may have been mislabelled as terrorists by that “scientifically unsound” algorithm, possibly resulting in their untimely demise.

As the piece then explains, the methods used by the NSA are very similar to those used by Big Data business applications and spam filters. With a twist: instead of selling products, the output of the machine learning algorithm is a death-sentence for those who are labelled “terrorists” by it. (Needless to say the details are politically quite involved, so I refer interested readers to Grothoff and Porup's rich list of links to find out more.) Whilst an irrelevant suggestion to buy a certain book or an email labelled wrongly as spam can be at most annoying, giving the wrong label to a target of the SKYNET programme may have dreadful consequences. And yet, all those mistakes boil down to nothing more sophisticated than the base-rate fallacy.

In a nutshell, this very well-known problem in the calculus of probability shows that in testing for a property which is not frequently observed in a population, even very accurate tests may


lead to a great proportion of false positives, i.e., individuals who are wrongly attributed the property tested for. This fallacy is so well-known that it features in textbooks, with the typical example being the disproportionate number of false positives which arise from a 99% accurate HIV test run on randomly selected samples. An example from the SKYNET programme mentioned in the Grothoff and Porup article is the Al-Jazeera journalist and longtime bureau chief in Islamabad, who scores very high on the NSA terrorist ranking because of his frequent journeys in areas known for terrorist activities.

It is quite unbelievable that such a macroscopically flawed piece of reasoning is being used in SKYNET, thereby threatening the lives of thousands. For the fact that terrorists are a tiny minority of the (Pakistani) population doesn't require proof. And even an otherwise remarkably low rate of false positives can potentially lead to thousands and thousands of false-positive executions. Indeed many are reluctant to believe that no one at NSA is able to spot this gigantic mistake, see for instance the discussion on [Andrew Gelman's blog](#). So it is quite likely that the latest installment of the [Snowden documents](#) is just showing one very incomplete fragment of the story. Be this as it may, it's certainly a story which shouldn't have existed in the first place.

(Many thanks to Teddy Groves for pointing this out to me.)

HYKEL HOSNI

Philosophy, University of Milan



smbc-comics.com

Evidence-Based Medicine

Study 329 has gone down in history as one of the most infamous clinical trials in medicine. The study was a double-blinded randomized controlled trial testing paroxetine and imipramine against placebo in adolescents diagnosed with major depression. The conclusion of the study was that '[p]aroxetine is generally well tolerated and effective for major depression in adolescents'. Soon after, on the basis of this study, paroxetine, a selective serotonin reuptake inhibitor, was widely prescribed by doctors for off-label use in children. (Some of the figures are given in this Medicines and Healthcare Products Regulatory Agency (MHRA) [report](#).)



The study has become infamous because the MHRA arrived at the opposite conclusion when they later took a look at the study. The MHRA concluded that the study had failed to demonstrate the effectiveness of paroxetine for treating depression, and in fact demonstrated an increased risk in suicidal ideation and self harm in teenagers. Given this, it was recommended that doctors stop prescribing the drug to adolescents.

Although the drug looked to perform well in terms of a number of outcomes, these were not among the eight outcomes that the study had initially intended to measure. In fact, the drug performed poorly in terms of the initially specified eight outcomes. This is an instance of outcome switching, which is problematic because there is a possibility that a correlation between the drug and an outcome is simply the result of chance. Given this, choosing which outcomes to measure after the fact makes it possible to pretend that a fluke is a significant result.

But switching outcomes need not be a bad thing, so long as the switch is pointed out—at least, that is the recommendation in the CONSORT guidelines for reporting randomized controlled trials. The problem with Study 329 was that it looked designed to mislead in this respect. In part as a response to Study 329, there began an initiative to restore invisible and abandoned trials (RIAT). In particular, there has been a reanalysis of Study 329. More information is available at [Restoring Study 329](#).

Recently, some results on the prevalence of outcome switching have been published. And a group at the Centre for Evidence Based Medicine at the University of Oxford has a strategy to remedy this state of affairs. The COMPare team have begun systematically checking trial results published in the top five medical journals for evidence of undeclared outcome switching. The group's exact methods are given [here](#). The methods involve comparing the outcomes of the published trial results with a trial registry or the trial protocol. In cases where they find a discrepancy, they send a letter to the journal in question pointing out the discrepancy, recommending that the instance of outcome switching be made clear.

The group have found outcome switching in the vast majority of the trials recently published in the top five medical journals. (The results so far are listed [here](#).) They have received some interesting and varied responses from the journals that they have contacted, and they are listing these responses on their [blog](#). The responses have varied from an acknowledgement and correction of the instances of outcome switching to an unwillingness to publish the group's letters. In an interview to [Retraction Watch](#), the project members have said:



Until we began writing to journals, we only knew that outcome switching was highly prevalent, despite

most journals promising to adhere to high reporting standards. Now, from the responses we've had, we're learning why it continues to be so prevalent, we are identifying the recurring misunderstandings and systemic shortcomings. Essentially we've solicited qualitative data on the reasons why outcome switching occurs in journals, and it could only have been done by writing these letters.

The group are now writing up a paper with an analysis of these results. But until then, readers can follow the action over at the COMPare [blog](#).

MICHAEL WILDE
Philosophy, Kent

EVENTS

APRIL

FE: Workshop in Mathematical Philosophy: Formal Epistemology, Munich Center for Mathematical Philosophy, 7–9 April.

RoD: The Roots of Deduction, University of Groningen, 7–9 April.

CI: Causal Inference Meeting, London, 13–15 April.

HL: Conference on Hegel's Logic, University of Valencia, Spain, 20–22 April.

PoKR&R: Conference on Principles of Knowledge Representation and Reasoning, Cape Town, 25–29 April.

SMI&P: Scientific Models: Imagination and Practice, University of Exeter, 26 April.

KDW: Knowledge in a Digital World, University of Lund, Sweden, 27–29 April.

ICMMP: Conference on Intelligent Computing, Mechanical and Production Processes, Pattaya, Thailand, 28–29 April.

RR&RE: Reasons, Rationality, and Rationalising Explanation, University of Warwick, 29 April.

MAY

ADR: Aspects of Defeasible Reasoning, Konstanz University, Germany, 4 May.

PROCESSES: Bringing Analytic and Continental Traditions Together, University of Kent, Canterbury, 12 May.

MS: Models and Simulations, Barcelona, 18–20 May.

PSP: Probabilities in Science and Philosophy, The Hebrew University of Jerusalem, 19–20 May.

E&U: Workshop on Explanation and Understanding, Aarhus University, Denmark, 19–20 May.

NPV: Non-physicalist Views of Consciousness, University of Cambridge, 24–26 May.

RP&P: Rationality, Probability, and Pragmatics, Berlin, 25–27 May.

FoD: Faces of Disagreement, Montreal, 26–28 May.

TE&E: Truth, Existence & Explanation, University of Chieti-Pescara, Chieti, Italy, 26–28 May.

JUNE

T&PR: Workshop on Theoretical and Practical Reasoning, University in Leipzig, Germany, 2–4 June.

MCMP5: Five Years MCMP: Quo Vadis, Mathematical Philosophy?, Ludwig Maximilian University of Munich, 2–4 June.

IDIS: Infinite Idealizations in Science, Ludwig Maximilian University of Munich, 8–9 June.

GEM: Ground, Essence and Modality, Helsinki, 8–10 June.

PoI: Workshop on the Philosophy of Information: The Role Of Data In Biomedical Sciences, University of Ferrara, Italy, 13–14 June.

CE: Chance Encounter, University in Groningen, Netherlands, 23–24 June.

MI: Mechanistic Integration and Unification in Cognitive Science, Warsaw, Poland, 23–26 June.

BD&DL: Big Data and Deep Learning in High Performance Computing, Porto, Portugal, 30 June.

COURSES AND PROGRAMMES

Programmes

APHIL: MA/PhD in Analytic Philosophy, University of Barcelona.

MASTER PROGRAMME: MA in Pure and Applied Logic, University of Barcelona.

DOCTORAL PROGRAMME IN PHILOSOPHY: Language, Mind and Practice, Department of Philosophy, University of Zurich, Switzerland.

HPSM: MA in the History and Philosophy of Science and Medicine, Durham University.

MASTER PROGRAMME: in Statistics, University College Dublin.

LoPHISC: Master in Logic, Philosophy of Science & Epistemology, Pantheon-Sorbonne University (Paris 1) and Paris-Sorbonne University (Paris 4).

MASTER PROGRAMME: in Artificial Intelligence, Radboud University Nijmegen, the Netherlands.

MASTER PROGRAMME: Philosophy and Economics, Institute of Philosophy, University of Bayreuth.

MA IN COGNITIVE SCIENCE: School of Politics, International Studies and Philosophy, Queen's University Belfast.

MA IN LOGIC AND THE PHILOSOPHY OF MATHEMATICS: Department of Philosophy, University of Bristol.

MA PROGRAMMES: in Philosophy of Science, University of Leeds.

MA IN LOGIC AND PHILOSOPHY OF SCIENCE: Faculty of Philosophy, Philosophy of Science and Study of Religion, LMU Munich.

MA IN LOGIC AND THEORY OF SCIENCE: Department of Logic of the Eotvos Lorand University, Budapest, Hungary.

MA IN METAPHYSICS, LANGUAGE, AND MIND: Department of Philosophy, University of Liverpool.

MA IN MIND, BRAIN AND LEARNING: Westminster Institute of Education, Oxford Brookes University.

MA IN PHILOSOPHY: by research, Tilburg University.

MA IN PHILOSOPHY, SCIENCE AND SOCIETY: TiLPS, Tilburg University.

MA IN PHILOSOPHY OF BIOLOGICAL AND COGNITIVE SCIENCES: Department of Philosophy, University of Bristol.

MA IN RHETORIC: School of Journalism, Media and Communication, University of Central Lancashire.

MA PROGRAMMES: in Philosophy of Language and Linguistics, and Philosophy of Mind and Psychology, University of Birmingham.

MRRES IN METHODS AND PRACTICES OF PHILOSOPHICAL RESEARCH: Northern Institute of Philosophy, University of Aberdeen.

MSC IN APPLIED STATISTICS: Department of Economics, Mathematics and Statistics, Birkbeck, University of London.

MSC IN APPLIED STATISTICS AND DATAMINING: School of Mathematics and Statistics, University of St Andrews.

MSC IN ARTIFICIAL INTELLIGENCE: Faculty of Engineering, University of Leeds.

MA IN REASONING

A programme at the University of Kent, Canterbury, UK. Gain the philosophical background required for a PhD in this area.

Optional modules available from Psychology, Computing, Statistics, Social Policy, Law, Biosciences and History.

MSC IN COGNITIVE & DECISION SCIENCES: Psychology, University College London.

MSC IN COGNITIVE SYSTEMS: Language, Learning, and Reasoning, University of Potsdam.

MSC IN COGNITIVE SCIENCE: University of Osnabrück, Germany.

MSC IN COGNITIVE PSYCHOLOGY/NEUROPSYCHOLOGY: School of Psychology, University of Kent.

MSC IN LOGIC: Institute for Logic, Language and Computation, University of Amsterdam.

MSC IN MIND, LANGUAGE & EMBODIED COGNITION: School of Philosophy, Psychology and Language Sciences, University of Edinburgh.

MSC IN PHILOSOPHY OF SCIENCE, TECHNOLOGY AND SOCIETY: University of Twente, The Netherlands.

MRES IN COGNITIVE SCIENCE AND HUMANITIES: LANGUAGE, COMMUNICATION AND ORGANIZATION: Institute for Logic, Cognition, Language, and Information, University of the Basque Country (Donostia San Sebastián).

OPEN MIND: International School of Advanced Studies in Cognitive Sciences, University of Bucharest.

PHD POSITION: in Bayesian Statistics, Trinity College Dublin, Ireland, deadline 15 April.

PHD POSITION: in Philosophy of Science, University of Groningen, Netherlands, deadline 28 April.

PHD POSITION: in Inference, Testimony and Memory, University of Aberdeen, deadline 29 April.

PHD POSITION: in epistemology of computer simulation, Clermont University, deadline 15 May.

PHD POSITION: in philosophy of mathematics, Clermont University, deadline 15 May.

JOBS AND STUDENTSHIPS

Jobs

RESEARCH ASSISTANT: in Machine Learning, University of Cambridge, deadline 6 April.

SENIOR LECTURER: in Machine Learning, University of Sheffield, 8 April.

RESEARCH ASSOCIATE: in Medical Statistics, Kings College London, deadline 10 April.

LECTURER: in Theoretical Reasoning, University of Kent, deadline 11 April.

POST-DOC: in Topic Modeling, University of Skövde, deadline 20 April.

ASSISTANT PROFESSOR: in Philosophy of Science, Tilburg University, Netherlands, deadline 22 April.

ASSOCIATE PROFESSORSHIP: in Statistics, University of Bath, deadline 22 April.

POSTDOCTORAL FELLOWSHIP: in Philosophy, University of Milan, deadline 26 April.

Studentships

PHD POSITION: in applied mathematics, Plymouth University, deadline 6 May.

PHD POSITION: in Scientific Metaphysics, University of Calgary/University of Geneva/University of Minnesota, deadline 15 April.