
THE REASONER

VOLUME 7, NUMBER 1
JANUARY 2013

www.thereasoner.org

ISSN 1757-0522

CONTENTS

Editorial

Features

News

What's Hot in ...

Introducing ...

Events

Courses and Programmes

Jobs and Studentships

EDITORIAL

One ideal of mathematical proof views proof as derivation from axioms. Each step in a proof should follow from axioms and other assumptions by means of a sound rule

of inference. Some proofs do look like that; even without fully formalizing, one can still see how the reasoning steps accord with standard natural deduction rules. And with a little training in goal analysis, one can learn how to construct such proofs so that they fall out almost automatically; filling in the steps can become an almost mechanical procedure, something a computer could do.

But one does not have to look very far in mathematics to find proofs that do not easily fit this mould. These are proofs whose completion requires the introduction, into the flow of reasoning, of a ‘clever idea’, plucked apparently from nowhere. As an undergraduate maths student, these were the proofs that led me to despair of my own abilities as a creative mathematician. I could do goal analysis, I could do deduction from axioms, I could get so far with proofs. But how on earth could I expect to be able to come up with the proof ideas that appeared as if from nowhere in the textbooks? This appeared to require a level of mathematical insight / creativity that I simply didn’t have. So I ditched the maths and became a philosopher!

In my PhD research, however, I revisited this puzzle about creative mathematical activity, by means of a case study of the development of a mathematical proof. I attended a research seminar at Toronto’s Fields institute, where I had the privilege to observe while a mathematician (George A. Elliott) talked through his progress towards proving a classification theorem for C^* -algebras. One thing this unique peek ‘behind the scenes’ showed me was how proof ideas and definitions that appeared fully formed and as if from nowhere in the final published proof can have their genesis in comprehensible subgoals and understandable proof strategies. Put in the context of a proof strategy, apparently arbitrary introductions could appear to be the obvious way to proceed, even if their obviousness disappeared in the cleaned up published presentation of a proof.



Is all mathematics like that, behind the scenes? I don’t know, but at the 2012 Cambridge Conference on [Foundations of Mathematics](#) I was pleased to hear distinguished mathematician Sir W.T. Gowers suggest that it could be. In particular, Gowers raised the question of whether there could be a ‘foundation’ for mathematical discovery, in the form of a smallish number of rules about how to proceed in discovering a mathematical proof. In support of a positive answer to this question he presented an example of a proof containing an ‘as if from nowhere’ proof idea, and then explained how, by breaking the proof attempt into small enough steps, the proof idea could be seen as just another ‘automatic’ move. In theory, Gowers suggested, if he is right about mathematical discovery we should be able to teach computers to come up with such apparently ‘creative’ mathematical proofs.

This might seem plausible for some forms of mathematical discovery, but surely humans have access to sources of mathematical insight that are unavailable to computers? Visualization, for example: sometimes a picture or diagram is all that is needed for us to see why a mathematical result holds. Surely it would be crazy to think that we could teach computers that kind of route to mathematical discovery? Or so one might think. But coincidentally, before hearing Gowers’ conjecture about the automation of mathematical discovery at the Foundations conference, I had caught up with [Mateja Jamnik](#) at

the [Cambridge Computer Laboratory](#), whose research involves precisely this problem: teaching computers to reason visually.

We started with a discussion of formal versus informal conceptions of mathematical proof. . .

MARY LENG

Philosophy, University of York

FEATURES

Interview with Mateja Jamnik

Mateja Jamnik: I think that a proof is a social construct in mathematics. I think the history of mathematics shows us that where mathematicians—famous mathematicians—came up with a solution or a proof of a problem or a theorem, and presented it to other famous mathematicians, as long as they convinced them, they trusted that that was the correct proof. And it was only when the logicians came along that they formalised the logic of that domain and were able to either prove or disprove the proof, to actually formalise it. And that's what we mean when we talk about the formal proof, that you can verify that it is correct, that it follows from a set of axioms and so on. Whereas mathematicians, I don't think, are very interested in that. They are really more interested in the insight in the proof, and trying to understand why the theorem holds, and that's why I say it's a social construct because as long as they convince enough people that it's correct, nobody's going to go and check it out to see whether it is or not. I mean, as long as fellow mathematicians believe that they understand and that they trust the process of the proof, then they're fine. And history of mathematics is full of examples like that, where there are proofs that were thought to be proofs for fifty years and even by very famous people, and they were disproved and it was shown that they were not proofs at all. Whereas, from a formalist point of view, formal proof has a very technical meaning, which is that it follows from a set of axioms.



Mary Leng: There's certainly a line of thought in philosophy that says that real proof is formal derivation, and what mathematicians are doing is approximating or sketching that. But fundamentally, what justifies mathematical results is derivations, and in informal maths you just get some indications that there is a formal derivation behind that. And some of the recent interest in pictures and diagrams—certainly James R. Brown's thinking about pictures and picture proof and diagrams and so on—is to try to push against that picture, because as a Platonist he wants to suggest that we have a direct insight into mathematical truths and a lot of this is not via formal derivation. You can just see that it's true.

MJ: That's absolutely correct. You know, logic is only, what? A hundred and something years old. So for thousands of years, mathematicians have done proofs with different methods. And that's, I suppose, where I come in. I'm really interested in those

types of methods that people—and mathematicians in particular—have been using for thousands of years.

ML: So what interests me in your work is that you're saying that these methods, though visual, can be formalized.

MJ: Yes, that's exactly what I'm trying to do.

ML: Whereas some of the thought in thinking about diagrammatic reasoning in philosophy is to say that there's this element of our cognition of mathematics that isn't formalizable. I suppose that's something that Penrose is pushing at as well, in his claim that there are diagrammatic proofs that cannot be automated, because he wants to push the idea that there's something special about us, that we're not purely computational.

MJ: That's right. It is. But what I'm trying to show is that all these so-called 'informal' methods, they can be formalized. And so I'm looking into the use of diagrams. In fact I just came from a conference on diagrams where everybody was coming from areas of mathematics, philosophy, cognitive science, cognitive psychology, computer science, artificial intelligence, and basically we have a common interest, which is to study diagrams, and the applications and theoretical foundations of the use of diagrams. And there are lots of people who come up with representations—diagrammatic representations—that they formalize and they use reasoning on. So it's totally possible. And also my work for my PhD, on DIAMOND [Jamnik's interactive theorem prover, which made use of diagrams to construct proofs of theorems], was in the domain of natural number arithmetic, where the proofs were not at all like the normal logical proofs. In fact, I would say that one of the hypotheses that came out of that work was that people use something like what we call 'schematic proof' to find a solution to a problem. So basically, you look at a few examples—concrete examples—of your problem and you solve them, and then you spot the pattern and you generalize that pattern, and you try to make an argument about how that pattern is a justification for the general statement for all cases. So what I did with natural number arithmetic was that I would represent these theorems and statements in mathematics using diagrams, and then use just manipulations of concrete cases of diagrams—for some natural number like 5, 6, or whatever—and then spot the pattern and generalize this into a program which, basically upon input, will produce a solution for that particular case. We call this—this program, this general pattern, this procedure—we call this 'schematic proof'. My hypothesis is that this is one possible model of how people do proofs, and there's plenty of evidence of that from history.

ML: Yes, because we're very good at seeing how special cases can be general.

MJ: Yes, and then what mathematicians can do is refine their proof procedure. So if they come across some counterexample, they will refine either the proof or the definitions that the theorem is based on. And there's this process of refinement until, again, most people are convinced that that is a proof—until something else turns up. But what I did in my work was that, yes, I did this generalization—I managed to get the program to spot the general pattern automatically. The users would do concrete cases so that they have an idea of what the proof should look like for specific cases, and then the program would come up with a general pattern, which is a proof for the general theorem.

ML: Which means, given a *particular* number as input it can produce a proof?

MJ: A proof for that number, that's right. But the logicians, the formalists, would

still not be satisfied because, they would ask, how do you know that your generalization is correct? So I used logic, basically, to do meta-level reasoning, which is not reasoning about the theorem any more, but reasoning about the proof that came up. And I used standard techniques, standard theorem provers, to build up a theory about the proofs and then prove that that procedure is correct. So at the end of the process you get something which I think resembles how mathematicians would come up with a proof, but you also get something which satisfies the formalists—the logicians—that your proof is formally correct. Now the interesting thing that allows me to do this step from concrete cases to the general case is something called the constructive omega-rule. There's the omega-rule, which is basically an alternative to induction, but obviously you can't capture it on a machine because you'd have to prove an infinite number of cases. So there's a constructive version, which says that if you can find an effective procedure, then you can conclude the general statement. And usually an effective procedure would be something like a recursive function or something like that. So as long as my program—procedure for the proof—is effective, then it's a justification that I can conclude the general statement. But we still do the verification in the formal system then. But that's nothing to do with the theorem itself, it's to do with the proof. The actual coming up with the proof is done with concrete diagrams.

ML: So when we have the insight that this special case stands for the general, presumably we're not doing some metalogical justification of that?

MJ: Oh no. What I'm saying is that people omit that last step altogether, right? So they do the first two steps. They do the concrete cases and then they spot the pattern, and that's their proof. And even with the machine you could stop at that point. But I do a third stage, which is to satisfy the logician that it is formally correct. But I think that the first two stages model more closely what people do when they spot patterns.

ML: So people are reasoning in accord with the constructive omega rule by seeing the pattern as something that they could replicate.

MJ: Yes I think so yes. That's right. Or I would put it differently. I would say that the omega-rule—the constructive omega-rule—nicely captures why people reason in this way. Or describes—you can use that vocabulary to model or to talk about how people do things. So I'm really much more interested in how people do things, and how to capture that. So basically to give computer systems this ability. To capitalize on this human ability.

ML: From the perspective of AI—bearing in mind that I don't really know the field at all—has the tendency been to try to model the way humans actually reason, or is there an approach within that which says no, that we want to do computer reasoning, a computer form of artificial intelligence, and that's not about modelling human intelligence?

MJ: I think AI is interested in both really. There are people who are interested in making machines better, faster, using machine type methods to mimic the type of reasoning that people do, and there's a lot of people who are interested in modelling what people do on computers in order to make computers better, faster, more intelligent, in fact. To have these eureka steps—to enable machines to come up with these eureka steps that people obviously can do but with machines it's only every so often that there is a breakthrough.

ML: So we can use our attempts to model human reasoning to make better ma-

chines. . .

MJ: And I think we're basically modelling what people do—at least I am interested in modelling what people do—to understand better what people do.

ML: That was what was going to be my question. So there's obviously the practical issue of getting better machines, but there's also the hope of getting some insight into us.

MJ: And that's the starting point really. Once you understand better what people do you can also model that, mimic that, emulate that on machines. That's absolutely fundamental. And that's where I come from—I come from an AI perspective. I did my PhD on AI.

ML: This brings me to the issue about Penrose, because Penrose wants to say that we're fundamentally different from machines. You mention in your book (*Mathematical Reasoning with Diagrams*, 2001) Penrose's scepticism about the possibility of modelling diagrammatic reasoning in computers, and I suppose behind all that is the thought that we want to find things that we can do that computers can't. So if it turns out that we can model this reasoning in automated settings, then that speaks against this idea that we're so different.

MJ: Yes absolutely. I think that we don't understand reasoning enough to be able to make claims like this. So what spurred me on to say something about Penrose in my book is because he presented this example about a cube—about something which is innately human reasoning that machines wouldn't be able to do. He presented this example of a proof that the sum of the first n hexagonal numbers is n -cubed. He presented this visual proof which basically says, "I'll give you an example for a cube, that is of size three, and if you split it up in this way you can see that those are the first three hexagonal numbers, and if you continue doing that, then you understand that the general theorem holds." And that's precisely what, for example, the theorem about the odd natural numbers, in 2-D, is. It's an analogue of that in 2-D. And I thought, there's something very procedural about this. You know, and we see, I suppose he's appealing to the fact, the visual effect, that you sort of 'splatter' the 3-D onto 2-D to see that those three shells of the cube form a hexagonal number. But you can think of it in the same way about the odd natural numbers. You have that the sum of the first n odd natural numbers is n -squared. You can take a square of size three and then you split it into 'ells', and you see that each ell is the subsequent odd natural number. And that's exactly an analogue of what Penrose presented. And obviously I showed we can capture that, you know? DIAMOND could easily do that. OK it doesn't have a 3-D interface, so technically you can't, but in principle, there's nothing that would stop it really. So I suspect that Penrose is asking, really, "Could the computer come up with an idea like that?" and we don't know that yet.

ML: So just to clarify, DIAMOND is a proof checker rather than a theorem prover?

MJ: It's an interactive theorem prover, so it means that the user constructs the sample cases, so it means the user has the insight, really, into what the proof should look like. So the computer's not coming up with an insight.

ML: In your book you mention the hope that you could actually develop a computer program that could do the insight steps. How have things moved in that regard?

MJ: I haven't moved much in that direction yet, because I've been looking at a

different direction with reasoning with diagrams, but that would definitely be the next step, to put some sort of search procedure on top of all these visual methods and geometric manipulations of diagrams and check whether anything interesting comes up. Now of course Penrose would probably say, the computer doesn't have the insight. But where does that come from in a mathematician? It comes from experience, it comes from. . . well we don't know. That's why I'm interested in modelling this type of reasoning.

ML: But in other contexts—I'm just thinking about what little I know of theorem provers in non-diagrammatic contexts—there are programs that do non-diagrammatic reasoning where they can apparently be creative and come up with insight in that area. So they can search a problem space and pick out potentially interesting proofs and so on. So if, in the non-diagrammatic case you've got the insight, and in the diagrammatic case you've got the ability to deal with diagrams it looks like it shouldn't be impossible to put the two together.

MJ: Yes. In fact this is my current research agenda. I'm trying to marry the two. I'm trying to bring diagrammatic reasoning and combine it with symbolic reasoning. Because obviously you can't do everything with diagrams, and clearly we can't do everything—we don't have the insight—in the solutions that machine oriented methods can come up with, necessarily. So that's exactly what I'm looking into. I'm trying to combine the two, to come up with a system where the proofs could be done with a mixture of steps—whichever is more appropriate at that point. At the moment we're still looking at the interactive case, where the user comes up with the proof and the machinery then checks it—but you can see, if you put some automation onto it, you know some search techniques, you would hope that you could come up with something interesting. So yes, that's the focus of my current research.

ML: I'll look forward to hearing the results!

Belief Revision, Uniqueness, and the Equal Weight View

Thomas Kelly has argued that the Equal Weight View of peer disagreement is committed both to the idea that the subjects should revise their opinions by *splitting the difference*, and to the Uniqueness Thesis, which claims that for any hypothesis H and body of evidence e , e rationalizes only one doxastic attitude towards H . I contend that Kelly's interpretation of the Equal Weight View is doomed to failure and I argue that the Equal Weight View and the Uniqueness Thesis can be independent of each other.

A popular view in contemporary epistemology (see for instance Christensen, D. (2007) 'Epistemology of Disagreement: The Good News', *The Philosophical Review* 116(2): 187–217; Elga, A. (2007) 'Reflection and Disagreement', *Noûs* 41(3), pp. 478–502) maintains that disagreement with an epistemic peer—someone who has the same evidence and general intellectual skills you have—counts as further evidence on what the original shared body of evidence supports. The *Equal Weight View* (henceforth EWV) regulates the epistemic weight of peer disagreement thus:

In cases of peer disagreement, one should give equal weight to the opinion of a peer and to one's own opinion. (Kelly 2010: 'Peer Disagree-

ment and Higher Order Evidence’, in *Disagreement*, (eds.) R. Feldman and T. Warfield, Oxford: Oxford University Press, 111–74. See p. 187.)

Thomas Kelly has argued that (EWV) is necessarily committed to the *Uniqueness Thesis* (henceforth UT):

(UT) For any hypothesis H and a total body of evidence *e*, there is some one doxastic attitude that is uniquely rational to adopt towards H.

Let me quote Kelly’s argument (2010, pp. 118–9):

(KELLY’S CASE)

At time *t*₀, my total evidence with respect to some hypothesis H consists of E. My credence for H stands at 0.7. Given evidence E, this credence is perfectly reasonable. Moreover, if I was slightly less confident that H is true, I would also be perfectly reasonable.

At time *t*₀, your total evidence with respect to H is also E. Your credence for H is slightly lower than 0.7. Given evidence E, this credence is perfectly reasonable. Moreover, you recognize that, if your credence was slightly higher (say, 0.7), you would still be perfectly reasonable.

At time *t*₁, we meet and compare notes. How, if at all, should we revise our opinions?

Kelly’s answer to the question is in three steps (pp. 119 and ff.). First step: he claims that from (EWV) it should follow that I am rationally required to decrease my credence while you are rationally required to increase your credence. And yet (second step) that seems awkward, for we both acknowledge the rationality of our respective doxastic attitudes. So (third step), a proponent of (EWV), in order to motivate our doxastic change, has to accept (UT), which says that a body of evidence rationalizes only one doxastic attitude.

Kelly takes (EWV) as saying that we should revise our beliefs in order to adopt a new credence that is the average of our initial credences. To put it roughly, we should split the difference. If we spell out (EWV) as directly stating something about how we should update probability functions, it follows that to give equal weight means that both credences have the same probability of being right. For, as Adam Elga puts it (2007, fn. 21), disagreement with an epistemic peer constitutes evidence for the fact that we are equally likely to get it right. Linear updating offers us a clear algorithm for implementing this idea:

$$\frac{x_1 \cdot n + x_2 \cdot m}{n + m},$$

where *x*₁ and *x*₂ are the subjects’ opinions; *n* and *m* are their probabilities of being right. Here the value of *n* and *m* is the same, since the subjects are equally likely to be right.

In probability calculus nothing prevents one from saying that the probability of being right is zero (that is, *n* = *m* = 0). Moreover, to assign zero as probability of being right to both opinions is consistent with Kelly’s interpretation of (EWV). And yet, in ordinary real number arithmetic we cannot divide by zero, for the division by zero is undefined.

Therefore, the averaging function $\frac{x_1+x_2}{2}$ does not follow from the above formula over the closed interval $[0, 1]$. Hence, the explanation of (EWV) in terms of giving equal probability of being right to both credences is not valid for every value one can assign to n and m . If $n = m = 0$, the explanation is meaningless.

Does this license us to conclude that we'd better give up (EWV)? Not quite so.

In my view, the foregoing analysis shows that we should not conceive of (EWV) as encapsulating a principle of split-the-difference revision. However, we can avail ourselves of an alternative interpretation of (EWV), recently advanced by Christensen (see Christensen, D. (2011) 'Disagreement, Question-Begging and Epistemic Self-Criticism', *Philosophers' Imprint*, 11(6), see p. 4), to the effect that (EWV) states how to take correct account of one bit of evidence, i.e., the peer's opinion. As I see it, the discovery of disagreement changes the shared body of evidence bearing on the hypothesis H: the new body of evidence should include both peers' opinions. This reading doesn't commit (EWV) to the split-the-difference revision endorsed by Kelly and Elga, for it doesn't interpret (EWV) as directly encapsulating that algorithm for belief revision.

There is a last wrinkle to notice in Kelly's argument. In the diagnosis of the case, Kelly maintains that belief revision is required in order to satisfy (UT). Yet there is a plausible reading of (UT) that is compatible with Kelly's case and with the idea that the subjects should not revise their doxastic attitudes in order to satisfy it. In the formulation I gave above, (UT) says that the uniquely rational response to a body of evidence is a single credence that can be represented with a single probability function. However, as Christensen (2007, fn. 9) and others pointed out, there is a more permissive reading of (UT) which instead of taking a single probability function takes sets of these functions as representing the rational attitude to have. Thus, instead of saying that a probability function must have a single value between 0 and 1, one can pick out a subinterval of the interval $[0, 1]$ and say that a doxastic attitude is rational if it is represented by a function that takes values from that particular subinterval. Once this reading is adopted, it is possible to claim that little differences in credences don't involve a violation of (UT), for if the values are so close as to be in the same subinterval that counts as rational (as it is plausible to think in Kelly's case), then (UT) isn't violated. This interpretation of (UT) entitles us to conclude that (EWV) and (UT) can be independent of each other.

Many thanks to Valeria Vignudelli and two anonymous referees for helpful comments.

MICHELE PALMIRA

COGITO, University of Modena & Reggio Emilia

More Intuition Mongering

In 'Intuition Mongering' (*The Reasoner*, 6(6)), I argue that, for appeals to intuition to be strong arguments, the relevant philosophers (the "experts") must agree on the intuition in question. If the "experts" disagree, then an appeal to intuition is weak. This Principle of Agreement on Intuition (PAI) is a necessary, not a sufficient, condition for strong appeals to intuition. Another necessary condition is the reliability of intellectual intuition. If intellectual intuition doesn't track truth, then appeals to intuition are weak.

Appeals to intuition look like this (cf. Huemer, 2007, ‘Compassionate Phenomenal Conservatism,’ *Philosophy and Phenomenological Research*, 74, 30–55):

(I1) It seems to *S* that *p*.

(I2) (Therefore) *p*.

The Mary thought-experiment (Jackson, 1982, ‘Epiphenomenal Qualia,’ *Philosophical Quarterly*, 32, 127–136) is an example:

(M1) Before her release, Mary has complete physical information about color vision.

(M2) (*It seems that*) Mary learns something new upon her release.

(M3) (Therefore) What Mary learns upon her release must be non-physical.

For Jackson’s argument to be strong, the “experts” must agree on (M2). Even if there’s agreement about (M2), one could challenge the belief-forming-process by which the “experts” come to believe (M2). Even if it *seems* to the “experts” that (M2), why should that count as a strong reason to believe (M2)? I will argue that it doesn’t because intellectual intuition is an unreliable belief-forming-process.

In ‘Intuition Mongering,’ I show that appeals to intuition (whether the intuition of one or many) presuppose the following:

(A1) If, in response to case *C*, it seems to *S* that *p*, then *p*.

(A1) is unwarranted, since the inference from (I1) to (I2) is unreliable. For example:

- It *seems* to me that the correct answer to the multiple-choice question is (a), so the correct answer is (a). [But (a) is incorrect. I confused (a) with (b) because they’re similarly worded.]
- It *seems* to most students in the class that the correct answer to the multiple-choice question is (a), so the correct answer is (a). [But (a) is incorrect. Most students confused (a) with (b) because they’re similarly worded.]

If this is correct, then intellectual seemings (Casullo, 2002, ‘A Priori Knowledge,’ in P. Moser (ed.), *The Oxford Handbook of Epistemology* (pp. 95–143), NY: OUP) by *themselves* don’t provide a strong basis for inference about what is the case.

In the case of philosophical thought-experiments, there’s another reason to be suspicious about appeals to intuition. To see why, consider the Mary thought-experiment. To Jackson (1982, 130), “*It seems just obvious* that [Mary] will learn something about the world,” so he infers that Mary will learn something. As I show in ‘Intuition Mongering,’ however, other philosophers have different intuitions about this case. This state of affairs is quite common in philosophy. One philosopher intuitively judges that *p*, whereas another intuitively judges that not-*p*. For example:

- Jackson (1982) and Dennett (1991, *Consciousness Explained*, Boston: Little Brown) have conflicting intuitions about Mary.

- Chalmers (1996, *The Conscious Mind: In Search of a Fundamental Theory*, NY: OUP) and Dennett (1995, 'The Unimagined Preposterousness of Zombies,' *Journal of Consciousness Studies*, 2, 322–326) have conflicting intuitions about zombies.
- Bird (2008, 'Scientific Progress as Accumulation of Knowledge,' *Studies in the History and Philosophy of Science*, 39, 279–281) and Rowbottom (2008, 'N-Rays and the Semantic View of Scientific Progress,' *Studies in the History and Philosophy of Science*, 39, 277–278) have conflicting intuitions about progress.

But both philosophers cannot be right. Since the same belief-forming-process, namely, intellectual intuition, gives incompatible verdicts about the same case, it's unreliable. More explicitly:

1. In response to case *C*, philosopher *A* intuitively judges that *p*, whereas philosopher *B* intuitively judges that not-*p*.
2. It's not the case that *p* and not-*p*.
3. (Therefore) Either *A* is wrong or *B* is wrong.
4. But both *A* and *B* came to judge that *p* and not-*p*, respectively, by relying on the same belief-forming-process, namely, intellectual intuition.
5. If the same belief-forming-process yields incompatible verdicts about the same case, then it's unreliable.
6. (Therefore) Intellectual intuition is unreliable.

In support of (5), consider the following: suppose I use a Litmus Test to form beliefs about whether a solution is acidic or basic. When testing the same solution, however, my blue Litmus paper sometimes turns red (thereby indicating an acidic solution) and sometimes stays blue (thereby indicating a base solution). In this case, I wouldn't put much trust in my Litmus paper as a basis for forming beliefs about the pH of the solution.

Applied to the case of zombies, the aforementioned reasoning looks like this:

1. In response to zombie thought-experiments, philosopher David Chalmers intuitively judges that zombies are conceivable, whereas philosopher Daniel Dennett intuitively judges that zombies are inconceivable.
2. It's not the case that zombies are conceivable and inconceivable.
3. (Therefore) Either Chalmers is wrong or Dennett is wrong.
4. But both Chalmers and Dennett came to judge that zombies are conceivable and that zombies are inconceivable, respectively, by relying on the same belief-forming-process, namely, intellectual intuition.
5. If the same belief-forming-process yields incompatible verdicts about the same case, then it's unreliable.

6. (Therefore) Intellectual intuition is unreliable.

Even if some philosophers share Chalmers' intuition (e.g., Webster, 2006, 'Human Zombies are Metaphysically Possible', *Synthese*, 151, 297–310), whereas others share Dennett's intuition (e.g., Marcus, 2004, 'Why Zombies are Inconceivable', *Australasian Journal of Philosophy*, 82, 477–490), the fact that the same belief-forming-process yields incompatible judgments about the same cases is enough to cast doubt on the reliability of that belief-forming-process.

If my argument is sound, then appeals to intuition are weak arguments because intellectual intuition is an unreliable belief-forming-process (since it yields incompatible verdicts in response to the same cases, and since the inference from (I1) to (I2) is unreliable). So, although *in principle* the (PAI) could be met, *in practice*, the track record of appeals to intuition in philosophy provides strong reasons to believe that intellectual intuition is unreliable. Since the reliability of intellectual intuition is a necessary condition for strong appeals to intuition, it follows that appeals to intuition are weak arguments.

MOTI MIZRAHI

Philosophy, St. John's University

NEWS

Alternatives, Belief and Action, 15–16 November

From the 15th to the 16th of November 2012, the University of Valencia (Spain) hosted the [International Symposium on Alternatives, Belief, and Action](#), organized by Carlos Moya and Sergi Rosell, under the auspices of Phronesis Philosophy Group. This is the second Valencia Symposium on these themes—after the one organized in 2008 under the heading of International Workshop on Belief, Responsibility, and Action—and is part of a research project with the same title, run by researchers of the University of Valencia but also by others working at several universities in the UK and Belgium, and is funded by the Spanish Ministry of Economy and Competitiveness (ref. FFI2009-09686).

There were eight invited speakers. On Thursday morning, Neil Levy (University of Oxford) opened the meeting by talking about consciousness and moral responsibility, and, after setting out and defending a particular account of consciousness (the global workspace account), he argued that agents who are in disordered states of consciousness are not morally responsible for their actions, and that moral responsibility of ordinary agents who fail to be conscious of facts that causally impact on their behaviour is also reduced. In the next talk, Michael McKenna (University of Arizona) defended a compatibilist reasons-responsiveness account of moral responsibility, which, against Fischer and Ravizza, is agent-based rather than mechanism-based. Particularly, he put forward the difficulties that theorizing about freedom in terms of mechanisms of action lead to, and underlined the advantages of returning to an agent-based view.

On Thursday afternoon, Helen Steward (University of Leeds) defended, against other ways (e.g., PAP) of formulating the Control Principle, the advantages of formulating this principle in terms of the very ordinary 'can't/couldn't help it' locution. She

examined alternative ways of attempting to capture the principle and argued that certain sorts of examples reveal that they do not work. She also presented the distinctive virtues of her favoured account. Later on, Carlos Moya (University of Valencia) started his talk by presenting briefly his recent defence of the Principle of Alternative Possibilities against Frankfurt cases, on the basis of an assumption that, he argued, underlines our intuitive judgments about some ordinary examples of morally blameworthy behaviour; and, afterwards, he went on to consider several possible objections to this way of defending PAP and argued that they do not succeed.

On Friday morning, George Sher (Rice University) critically assessed a central claim made by McKenna in his new book, *Conversation and Responsibility*, namely that being responsible and holding responsible do not bear metaphysically priority over one another, but are instead interdependent elements of a single larger family of practices. He drew particular attention to some of the difficulties that McKenna's position appears to raise and opposed to it the idea that the agent's responsibility partially depends on how she understands the ways others would or should respond to her action. Afterwards, Constantine Sandis (Oxford Brookes University) examined the relation between the various factors and forces which underlie human behaviour, on the one hand, and agents' verbal reports on their reasons for acting as they did, on the other. He argued that over a century of psychological literature on purported confabulations rests on a dangerous conflation of the reasons for which people act with the factors that motivate them to act on those reasons (or otherwise explain why they acted as they did).

In the afternoon, Andrei Buckareff (Marist College) held that the tenability of Humean compatibilism—which combines both a Humean conception of laws of nature that takes laws to be non-governing with a strong dual ability condition for free will that requires that free agents possess the ability to decide differently when they make a free decision—depends in part upon what theory of time is correct. And he argued that Humean compatibilism is untenable in a deterministic universe if eternalism is true. Finally, the last speaker, Maria Alvarez (King's College London), opposed the traditional intentionalist conception of human agency, and argued for an alternative account in terms of our dual active powers, which can already be seen at work in Aristotle. No movement, no change can be an exercise of human agency unless the agent can refrain from performing it. She then went on to consider and rejoin putative counterexamples to her proposal.

The commentators were: Santiago Amaya (Humboldt Universität, Berlin), Claudia Compte (UV), Josep Corbí (UV), Tobies Grimaltos (UV), Edgar Maragat (UV), Carlos Patarroyo (Universidad del Rosario, Bogotá), Josep L. Prades (Universitat de Girona), Sergi Rosell (UV), Pablo Rychter (UV), and Carmen Santander (UV). Michael McKenna kindly agreed to give, on Wednesday evening, a Special Pre-Symposium Session on his new book *Conversation and Responsibility* (OUP, 2012). The organizers are grateful to all the participants for their excellent contributions.

SERGI ROSELL

Philosophy, University of Valencia

Intentions: Philosophical and Empirical Issues, 29–30 November

On 29 and 30 November, the [Italian Institute for Cognitive Sciences and Technologies](#) (ISTC-CNR) hosted in Rome the 1st *Topoi* Conference, focused on “Intentions: Philosophical and Empirical Issues”. This event inaugurated a series of biannual conferences, organized by the journal *Topoi* and sponsored by Springer. This first edition was also supported by the European Network for Social Intelligence ([Sintelnet](#)), and co-chaired by Fabio Paglieri (Rome) and Markus Schlosser (Groningen).

The focus of the conference was on intentions, with a strong emphasis on integrating both philosophical analysis and empirical findings. Theorizing about human action has a long history in philosophy, and the nature of intention and intentional action has received a lot of attention in recent analytic philosophy. At the same time, intentional action has become an empirically studied phenomenon in psychology, cognitive neuroscience, artificial intelligence, and robotics. As a result, the study of intentions is nowadays a thriving enterprise, where both conceptual and empirical issues are discussed in a dialogue across disciplines. This variety was well reflected in the conference programme, comprising 5 invited presentations and 7 contributed papers (out of 29 submissions).

On the first day, the meeting was opened by Marcel Brass (Ghent), discussing the neuroanatomy of intentional action: he first distinguished different components of intentional action and related them to different parts of the medial frontal cortex; then he showed that disbelief in free will affects intentional motor preparation.

This was followed by Kevin Tobia (Oxford) presenting recent work in experimental philosophy on the “side-effect effect”: the fact that people tend to judge a good side-effect as unintentional, while a bad side-effect is judged as intentional. Based on empirical studies, Tobia argued that the side-effect asymmetry may not be as robust or as troubling as initially thought.

Ariel Furstenberg (Jerusalem) argued for the existence of non-executed unconscious proximal intentions, i.e., unconscious proximal intentions to act that do not turn into doing. He first presented a philosophical perspective to help conceptually account for this phenomenon, and then described a specific EEG signal as the neural signature of a non-executed proximal intention.

In the next talk, Zoe Drayson (Stirling) reviewed recent findings on how a certain pattern of neural activity may be indicative of intentional mental action, thus allowing to attribute conscious awareness to patients in vegetative states. However, she questioned the ability of this technique to discriminate between mental events and mental actions, thus undermining its value.

The first day was concluded by Corrado Sinigaglia (Milan): given that motor representations recruited during action observation are known to facilitate action understanding, he discussed how exactly this process works. He posited that these representations underpin a motor experience of action effects as related to goals, thus facilitating goal ascription and action understanding.

The second day was opened by Élisabeth Pacherie (Paris), addressing the relation between intentions and actions. She argued that both the traditional causal view and some of the objections against it rest on an over-simplified conception of agency, whereas considering the hierarchical nature of intentions and control processes gives a new role

to consciousness in action production.

Lilian O'Brien (Cork) first showed that Bratman's thought experiment on mutually exclusive planning does not undermine the claim that an intention to A is necessary if I am to A intentionally. Then she discussed how we can rationally intend mutually exclusive ends, suggesting that the subjective authority of intentions includes also conditions on when the agent is criticizable.

In the next paper, Till Vierkant (Edinburgh) argued that resisting temptation by "sheer willpower" is not importantly different from using external props that simply prevent one from undermining previous resolutions, since the same kind of agentic control over one's own mental states is involved in both forms of self-control. This in turn paves the way to externalist views of the will.

The morning session was concluded by Gregory Strom (Sydney), who first refuted Davidson's classic account of weakness of the will, or *akrasia*, and then argued that akratic agents balk from doing what they know they should because they merely have existential knowledge that there is some decisive practical reason to act in a certain way, rather than also knowing what that reason is.

In the afternoon, the conference continued with Cristiano Castelfranchi (Roma), presenting an account of intentions based on goals. He first analyzed the various steps of goal processing and intention creation, regulated by specific sets of beliefs; then he described intentions as two-layered goal structures, and finally compared intentions "in agenda" and intentions under execution.

Marco Mazzone (Catania) focused on the claim that ordinary intentions are much more constrained than language meanings (communicative intentions), thus lacking the generative nature of the latter. He argued that this view underestimates the complexity of non-communicative intentions, and presupposes structural identity between communicative intentions and linguistic utterances.

The conference came to an end with Bruno Verbeek (Leiden), who discussed why failing to do as intended is not only a causal failure, but also a normative one. He defended the view that intending to A creates additional reasons to do A, due to the fact that the agent formulating the intention has the authority (epistemic or coordinative) to impose a behavioral obligation on all subsequent selves.

Further information on the conference, including a book of abstract, are available online [here](#). Extended and revised versions of all contributions will appear in a future issue of *Topoi*, forthcoming in late 2013 or early 2014.

FABIO PAGLIERI

Cognitive Science, ISTC-CNR Roma

Models and Mechanisms, 6–7 December

The aim of this [Tilburg workshop](#) (previously held in Edinburgh 2011) was to bring together two significant areas of research, on mechanistic explanation and the roles of scientific models, to better understand research practices in cognitive science.

The first keynote speaker, Stuart Glennan (Butler) began by pointing out that it is one thing to say that a certain phenomenon is produced by the operations of mechanism

and another to say that we have offered a mechanistic explanation. Using case-studies, Glennan suggested that mechanistic models bear at least some resemblance, in some respects to the phenomenon, its parts and internal relations, though questions were subsequently raised about the resemblance relation.

Several talks then focused on the short-comings of mechanistic explanation. Catherine Stinson suggested how to square the idea of explanatory mechanism schemas with ontic explanation. Eric Hochstein argued that phenomenological models are essential for some explanatory goals, and Raoul Gervais explored the use of covering law-type generalisations in explanations. Maria Serban suggested that mathematical models contribute to explanation independently of their role in identifying mechanisms. Scott Thomas outlined the problems in identifying mechanisms in generative linguistics.

On the theme of cross-level constraints in research within the cognitive science, Carlos Zednik and Frank Jäkel suggested an alternative framework for thinking about levels of analysis ('what' 'why' and 'how' models/mechanisms), and the role that Bayesian task analysis plays in this. Stephan Güttinger outlined how models guide the development and interpretation of experiments in biology. Mieke Boon argued that models should not be seen as representations, but as epistemic tools for investigating and controlling phenomena.

In relation to robotic models in particular, Serge Thill suggested that if human cognitive is best seen as embodied and embedded, then modeling practices should reflect this. Marcin Miłkowski suggested that gradual de-idealisation of robotic models, and the use of multiple robotic platforms, highlight what is explanatory about a (robotic) mechanism.

Andreas Hüttemann (Cologne), the second keynote speaker, looked at different concepts of emergence, and analysed in what sense mechanisms may be said to give rise to emergent phenomena. On theme of cross-level interactions, Felipe Romero argued that treating mutual manipulability relations as non-ideal causal interventions dissolves problems about top-down causation in mechanisms.

Iris van Rooij (Radboud/Donders Nijmegen) gave the final keynote talk. She convincingly argued that treating Bayesian models as instrumental 'as-if' models does not allow these models to escape from basic mechanistic constraints, such as tractability. Exploring a theme of this talk more closely, Johan Kwisthout (with van Rooij) argued that modelers should be more careful when claiming that approximation methods ensure that their models are tractable. Mark Blokpoel used a communication game to suggest how the addition of computational constraints into a Bayesian model can make it more mechanistically plausible.

All in all, the workshop nicely achieved its aim. Under the snowy sky of Tilburg, we had two days of fruitful discussion between researchers from different areas—including philosophy, psychology, computer science, and machine learning—curious to learn from each other.

The workshop was generously supported by the NWO Internationalisation grant on

ELIZABETH IRVINE

CIN, Tuebingen

MATTEO COLOMBO

TiLPS, Tilburg University

Belief Change in Social Context, 14–15 December

Belief Change in Social Context workshop took place December 14–15th in Amsterdam, and was organized by Sonja Smets, Zoé Christoff, and Nina Gierasimczuk, at the Institute for Logic, Language and Computation, University of Amsterdam. The Workshop was the opening event of Sonja Smets' ERC project *The Logical Structure of Correlated Information Change* (LogiCIC).

One theme of the conference was the necessity of combining probabilistic and logical reasoning. Hannes Leitgeb (Munich) presented a theory connecting the two realms, based on a thesis he attributes to Hume: that full belief can be equated with a degree of belief that is sufficiently high and stable relative to other beliefs. The resulting view depicts probability theory, belief revision, and epistemic logic as mutually supporting each other. In his talk, Christian List (LSE) showed how plausibility orderings could help systematically assign prior beliefs. Alexandru Baltag (Amsterdam) discussed how different combinations of logic and probabilistic reasoning performed in information cascades, and indicated that a version of belief revision could provide a qualitative alternative to the existing Bayesian model of such situations. Vincent Hendricks and Rasmus Rendsvig (Copenhagen) presented their concept of social proof, and its use in the analysis of bystander effect and other phenomena. Erik Olsson (Lund) presenting studies on how communication between scientists affects group performance, finding that dense communication links lead to faster convergence but poorer end results.

A second theme of the conference was connecting computation and interaction. Johan van Bentham (Amsterdam) provided a historical perspective on the presence of interaction into theoretical computer science, and set forth the challenge of finding a computational paradigm, which could serve as a general model of agency. Nina Gierasimczuk (Amsterdam) talked about recasting formal learning as an interaction between agents, in the simplest case viewing our data as coming from another agent—a “teacher”—instead of Nature. She presented results on the computational complexity of learning and teaching within the conclusive update paradigm. Patrick Girard (Auckland) presented a model of social influence taking epistemic agents as automata, with transitions triggered by the state of the network in which each agent was embedded.

A third theme of the conference was mistaken assumptions about others. In the game theory arena, Andrés Perea (Maastricht) showed how plausibility orderings could be used to find an alternative assumption about a player when the assumption of rationality is contradicted. Amanda Friedenberg (Tempe) discussed how bargaining proceeds if players are uncertain about each other's strategy, showing that even mere uncertainty about surprise moves can cause impasses. Ziv Hellman (Tel Aviv) presented a model of agreement based on belief rather than knowledge, which allows for agents to mistakenly

believe they had reached an agreement. Hans van Ditmarsch (Nancy) gave a dynamic epistemic game that allowed players to lie and recognize the occurrence of lies. Hans Rott (Regensburg) presented an open problem about disagreements where the agents differ both in their set of facts and in how they use their vocabulary.

A distinguishing feature of the conference was the diversity of approaches being brought to a single topic of inquiry. Various points of intersection were pointed out during the course of the conference, suggesting future collaborative work.

ROB CARRINGTON

ILLC, University of Amsterdam

Calls for Papers

HYPERINTENSIONALITY: special issue of *Synthese*, deadline 1 March.

WEIGHTED LOGICS FOR AI: special issue of *Journal of Approximate Reasoning*, deadline 15 January.

THE QUESTION OF BIO-MACHINE HYBRIDS: special issue of *Philosophy and Technology*, deadline 28 February.

INFINITE REGRESS: special issue of *Synthese*, deadline 1 July.

WHAT'S HOT IN . . .

Uncertain Reasoning

It often happens that a classic book becomes the standard reference for ideas which are in fact articulated elsewhere. In this respect, F. Knight (1921: *Risk, Uncertainty and Profit*, Houghton Mifflin) appears not to be an exceptional classic. (The volume is [freely available here](#).)

The received view, which features in virtually all introductory sections of research papers which purport to challenge the probabilistic representation of uncertainty, is that Knight (1921) introduces a fundamental distinction between those uncertainties which are probabilistically quantifiable and those which are not. The idea is roughly as follows. Probability is a suitable representation of an agent's rational degrees of beliefs only in the specific sort of situations in which a (unique) distribution of such degrees of belief is either given with the problem, or it is immediately derivable from its features. An obvious case in point is how strongly should I believe in 18 being the outcome of the next spin of the casino's roulette. It goes without saying that many situations of interest in "the real world" are not casino-like. A case in point here concerns how strongly should I believe that Greece will exit the Eurozone by May 2013. The received view credits Knight (1921) for introducing the distinction between *risk* (the roulette case) and *uncertainty* (GREXIT) with the understanding that the latter is not probabilistically quantifiable. Indeed, the expression *Knightian uncertainty* is often used to mark the distinction between "probabilistic" and "non-probabilistic" uncertainty.

Nearly three decades ago, F. Leroy and L. Singell (1987: "Knight on Risk and Uncertainty," *Journal of Political Economy*, 95(2): 394–406) posed a serious challenge

to the received view. Yet it appears to have gone pretty much unnoticed.

Leroy and Singell's main point is that Knightian uncertainty (as outlined by the received view) is very unlikely to be compatible with Knight's own view. In their words:

The thesis of this paper is that Knight did not intend the risk-uncertainty distinction to refer to whether or not agents are able to form subjective probabilities. On the contrary, Knight explicitly stated that in his view agents can be assumed to have subjective probabilities even in cases of uncertainty. Rather, Knight designated by risk situations in which insurance markets do exist and by uncertainty situations in which they do not. (396)

I refer the interested reader to the original paper for the justification of the above claim. One point to which I'd like to draw attention though is the following. Some opponents to the Bayesian approach in decision theory argue that "Knightian uncertainty" exposes the normative inadequacy of (subjective) probability (see my December 2011 column for some recent references). Yet a closer look at *Risk, Uncertainty and Profit* reveals (i) how Knight's argument is entirely framed in terms of the (potential) failures of the insurance market and (ii) how those failures largely depend on the stakeholders' practical inability to reach a consensus on whether "the event insured against has occurred or to evaluate the magnitude of the loss" (Leroy and Singell 1987, p. 400).

Interestingly enough, the interpersonal agreement between "gamblers" and "book-makers" on the conditions under which an event of interest has occurred or not, constitutes the key defining feature of de Finetti's notion of "event" in his construction of the Dutch Book argument. The extent to which this intriguing parallel between Knight's and de Finetti's conceptualisation of uncertainty can be pushed, emerges as a very fascinating question.

HYKEL HOSNI

Scuola Normale Superiore, Pisa
CPNSS, LSE

INTRODUCING . . .

Epistemic Utility Theory

If Yasho believes that Sonya is an accountant and an activist more strongly than he believes she is an accountant, we judge him irrational. Similarly, we judge him irrational if, having had a low degree of belief that Sonya is an activist conditional upon her being an accountant, he then learns only that she is an accountant and comes to have a high degree of belief that she is an activist. What justifies these judgments? More generally, how can we justify the general norms that govern our degrees of belief?

Traditionally, philosophers have given pragmatic arguments for these norms. For instance, the Dutch Book arguments for Probabilism and Conditionalization purport to show that an agent who violates either norm will be led by her degrees of belief to make choices that will result in a sure loss for her.

However, there is something unsatisfactory about these arguments. Epistemic states, such as degrees of belief, have two roles: their *pragmatic role* is to guide action, while their *alethic role* is to represent the world accurately. The Dutch Book argument for Probabilism shows that an agent who violates that norm must thereby have degrees of belief that play the pragmatic role poorly. But it says nothing about their ability to play the alethic role. Epistemic utility theory seeks to rectify this situation by providing arguments for epistemic norms that show that epistemic states that violate them play the alethic role poorly.

The idea behind epistemic utility theory is this: An epistemic state can be treated as an epistemic act. The rationality of such an epistemic act can then be assessed using the same techniques we use to assess the rationality of non-epistemic acts, namely, the techniques of utility theory. In traditional utility theory, we appeal to an agent's utility function U , which takes an action a together with a possible world w , and returns a measure $U(a, w)$ of the value of the outcome of act a at world w . And we state norms that govern which acts are rationally permissible in terms of her utility function and sometimes also in terms of her epistemic state. In epistemic utility theory, the acts are epistemic states, and the utility function is taken to measure how well an epistemic state plays the alethic role at a given world.

Consider, for instance, Joyce's argument for Probabilism (1998: 'A Nonpragmatic Vindication of Probabilism', *Philosophy of Science*, 65(4):575–603). He represents an agent's epistemic state by her *credence function*, which takes each proposition that the agent entertains and returns a measure of her degree of belief in that proposition. Thus, for Joyce, an epistemic utility function EU takes a credence function c and a world w and returns a measure $EU(c, w)$ of the epistemic utility of having c at w . He begins by enumerating necessary conditions on a measure of epistemic utility—let's call a function *Joycean* if it satisfies those conditions. He then identifies a norm of standard utility theory:

Dominance It is irrational to choose action a if there is another action b such that $U(a, w) < U(b, w)$ for any world w .

He then proves a theorem:

Theorem 1 Suppose EU is Joycean and c is a credence function. Then

$$c \text{ satisfies Dominance relative to } EU \\ \iff \\ c \text{ satisfies Probabilism.}$$

That is, if c violates Probabilism, then there is c' that has greater epistemic utility than c however the world turns out; but if c satisfies Probabilism, this does not happen.

Graham Oddie's argument for Conditionalization has the same structure (1997: 'Conditionalization, Cogency, and Cognitive Value', *British Journal for the Philosophy of Science*, 48:533–41.). Again, he represents an agent's epistemic state by her credence function. But this time he also represents an agent's *updating rule*: he represents it by a function \mathbf{R} that takes her current credence function c , the partition \mathcal{E} of

propositions from which she knows her evidence will come, and a particular proposition E in \mathcal{E} that she obtains as evidence, and returns c' , the credence function that the rule recommends in the light of this evidence. He enumerates necessary conditions on a measure of epistemic utility for credence functions—we call a function *proper* if it satisfies Oddie's conditions. And, for each proper epistemic utility function EU for credence functions, he defines an epistemic utility function EU^* for updating policies as follows: $EU^*(\mathbf{R}, w) = EU(c', w)$ where c' is the credence function that \mathbf{R} recommends at world w . He then identifies a norm of standard utility theory:

Maximize Expected Utility If an agent has a probabilistic credence function c , it is irrational for her to choose an action a if there is another action b such that

$$\text{Exp}_U(a|c) = \sum_w c(w)U(a, w) < \sum_w c(w)U(b, w) = \text{Exp}_U(b|c)$$

He then proves a theorem:

Theorem 2 *Suppose EU is proper and \mathbf{R} is an updating rule. Then*

For all c , \mathbf{R} satisfies Maximize Expected Utility relative to c and EU

\iff

\mathbf{R} is the updating rule of Conditionalization.

Thus, Conditionalization is the unique updating rule that maximizes expected epistemic utility given any initial credence function.

Thus, a standard argument for an epistemic norm in epistemic utility theory has the following structure:

- (1) Choose a formal representation of an agent's epistemic states.
- (2) Enumerate necessary conditions on epistemic utility functions.
- (3) Identify a norm of standard utility theory.
- (4) Prove the following theorem: For any epistemic utility function EU that satisfies the conditions in (2), epistemic states represented as in (1) satisfy the epistemic norm in question iff they satisfy the norm of standard utility theory from (3) relative to EU .

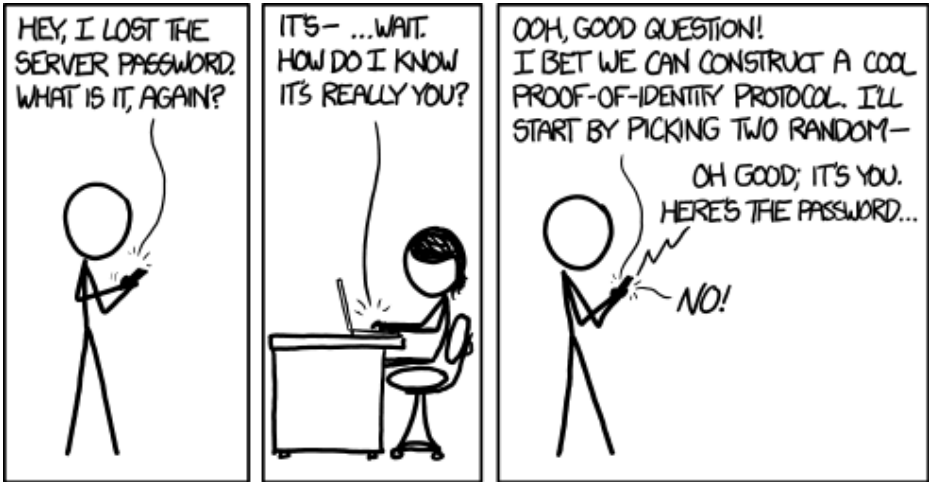
Understanding the argument strategy in this way allows us to see how we might extend the project begun by Oddie and Joyce:

- We might consider different ways of representing an agent's epistemic state: e.g., a set of full beliefs; a set of credence functions.
- We might explore different conditions on epistemic utility functions.

- We might explore the consequences of different norms of standard utility theory: e.g., norms that appeal to known objective chances or norms that try to accommodate attitudes to risk.

Many of these extensions will be explored over the coming four years in my ERC research project *Epistemic Utility Theory: Foundations and Applications*, which will run at the University of Bristol from January 2013 until December 2016.

RICHARD PETTIGREW
Philosophy, University of Bristol



xkcd.com

EVENTS

JANUARY

SODA: ACM-SIAM Symposium on Discrete Algorithms, New Orleans, Louisiana USA, 6–8 January.

LFCS: Symposium on Logical Foundations of Computer Science, San Diego, California, USA, 6–8 January.

TARK: 14th Conference on Theoretical Aspects of Rationality and Knowledge, Chennai, India, 7–9 January.

FOMCAF: Foundation of Mathematics for Computer-Aided Formalization, Padova, Italy, 9–11 January.

ICLA: 5th Indian Conference on Logic and its Applications, Chennai, India, 10–12 January.

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

[SAPoSC](#): South African Philosophy of Science Colloquium, Durban, South Africa, 13 January.

[A&N](#): Aims and Norms: Reasoning, University of Southampton, 18 January.

[CGCotPoM&L](#): 6th Annual Cambridge Graduate Conference on the Philosophy of Mathematics and Logic, Cambridge University, 19–20 January.

[ULTRA-COMBINATORICS](#): Pisa, Italy, 24–25 January.

FEBRUARY

[ICIIN](#): 2nd International Conference on Intelligent Information Networks, Maldives, 2–3 February.

SPIM: Workshop on Semantic Personalized Information Management, Rome, Italy, 4 February.

LAFLANG: 2nd International Workshop on Learning, Agents and Formal Languages, Barcelona, Spain, 15–18 February.

ICAART: 5th International Conference on Agents and Artificial Intelligence, Barcelona, Spain, 15–18 February.

CSEE: 2nd International Conference on Advances in Computer Science and Electronics Engineering, New Delhi, India, 23–24 February.

SAPHIR: Systematic Analytic Philosophy and Interdisciplinary Research, Ruhr-Universität Bochum, 25–27 February.

MARCH

THEORETICAL AGENCY: Auburn, Alabama, 1–2 March.

PTS: 2nd Conference on Proof-Theoretic Semantics, Tübingen, Germany, 8–10 March.

METAPHYSICAL VIRTUES: Western Michigan University, Kalamazoo, Michigan, 15–17 March.

SIMRIDE: 1st workshop on Uncertainty Quantification and Data Assimilation in Numerical Simulation of Physical Systems for Risk-Informed Decision Making, Durham, 18–21 March.

INFORMATION: 5th Workshop on Philosophy of Information, University of Hertfordshire, UK, 27–28 March.

UNIOLOG: 4th World Congress and School on Universal Logic, Rio de Janeiro, Brazil, 29 March–7 April.

APRIL

SBP: International Conference on Social Computing, Behavioral-Cultural Modeling, & Prediction, UCDC Center, Washington DC, USA, 2–5 April.

LATA: 7th International Conference on Language and Automata Theory and Applications, Bilbao, Spain, 2–5 April.

AISB: 6th AISB Symposium on Computing and Philosophy: The Scandal of Computation—What is Computation?, University of Exeter, 2–5 April.

THE ANALYSIS OF THEORETICAL TERMS: Munich, Germany, 3–5 April.

UNIOLOG: 4th World Congress on Universal Logic, Rio de Janeiro, Brazil, 3–7 April.

IMLA: 6th Workshop on Intuitionistic Modal Logic and Applications, Rio de Janeiro, 3–7 April.

ICANNGA: 11th International Conference on Adaptive and Natural Computing Algorithms, Switzerland, 4–6 April.

PERCEPTION, MODELS, AND LEARNING: 15th Annual Pitt-CMU Graduate Conference, Carnegie Mellon University, 5–6 April.

ADS: Agent-directed Simulation Symposium, Bahia Resort, San Diego, CA, USA, 7–10 April.

INFORMATION: SPACE, TIME, AND IDENTITY: Milton Keynes, 8–10 April.

PHDs IN LOGIC: Munich, 8–10 April.

MODELS & DECISIONS: 6th Munich-Sydney-Tilburg Conference, Munich, 10–12 April.
IDENTITY AND PARADOX: Lille, France, 11–12 April.
PAKDD: 17th Pacific-Asia Conference on Knowledge Discovery and Data Mining, Gold Coast, Australia, 14–17 April.
IEEE-SSCI: Symposium Series on Computational Intelligence, Singapore, 15–19 April.
GCTP: Graduate Conference in Theoretical Philosophy, Groningen, Netherlands, 18–20 April.
R&R: Reasons and Reasoning, Georgetown University, 20 April.
IMPLICIT BIAS: University of Sheffield, 20–21 April.
SOoSI: The Social Organization of Scientific Inquiry, Center for Philosophy of Science, University of Pittsburgh, 20–21 April.
GIRL@LUND: 2nd Conference on Games, Interactive Rationality, and Learning, Lund, 23–26 April.
EXPLANATORY POWER: Understanding Through Modeling. Epistemology, Semantics, and Metaphysics of ‘Inadequate’, Ruhr-Universität Bochum, 25–26 April.
NU/NDGC: 4th Annual Northwestern / Notre Dame Graduate Epistemology Conference, University of Notre Dame, South Bend, IN, 26–27 April.
AISTATS: 16th International Conference on Artificial Intelligence and Statistics, Scottsdale, AZ, USA, 29 April–1 May.

MAY

SDM: 13th SIAM International Conference on Data Mining, Austin, Texas, USA, 2–4 May.
O&M: Ontology and Methodology, Virginia Tech, 4–5 May.
CTFoM: Category-Theoretic Foundations of Mathematics, Irvine, California, 4–5 May.
AAMAS: 12th International Conference on Autonomous Agents and Multiagent Systems, Saint Paul, Minnesota, USA, 6–10 May.
ADMI: 9th International Workshop on Agents and Data Mining Interaction, Saint Paul, USA, 6–10 May.
PHILANG: 3rd International Conference on Philosophy of Language and Linguistics, University of Lodz, Poland, 9–11 May.
UK-CIM: Causal Inference in Health and Social Sciences, University of Manchester, 15 May.
MCS: 11th International Conference on Multiple Classifier Systems, Nanjing University, China, 15–17 May.
MATHEMATISING SCIENCE: University of East Anglia, Norwich, 16–17 May.
SLACRR: St. Louis Annual Conference on Reasons and Rationality, St Louis, MO, 19–21 May.
NIDISC: 16th International Workshop on Nature Inspired Distributed Computing, Boston, Massachusetts USA, 20–24 May.
TAMC: 10th Conference on Theory and Applications of Models of Computation, Hong Kong, China, 20–22 May.
UNCERTAIN REASONING: St. Pete Beach, Florida, USA, 22–24 May.
EI&I: Evolution, Intentionality and Information, University of Bristol, 29–31 May.

AIME: Artificial Intelligence in Medicine, Murcia, Spain, 29 May–1 June.
LoQI: Logic, Questions and Inquiry, Paris, France, 30 May–1 June.
GRADUATE EPISTEMOLOGY CONFERENCE: University of Edinburgh, 31 May–1 June.

JUNE

BSPS: British Society for the Philosophy of Science Annual Conference, University of Exeter, 4–5 June.
BAYSM: Bayesian Young Statistician Meeting, Milan, Italy, 5–6 June.
BISP: 8th workshop on Bayesian Inference in Stochastic Processes, Milan, Italy, 6–8 June.
CADE: 24th International Conference on Automated Deduction, Lake Placid, USA,, 9–14 June.
ICAIL: 14th International Conference on Artificial Intelligence & Law, Rome, Italy, 10–14 June.
INEM: Conference of the International Network for Economic Method, Erasmus University Rotterdam, The Netherlands, 13–15 June.
SocPHILPSYCH: 39th meeting of the Society for Philosophy and Psychology, Brown University, Providence, RI, 13–15 June.
TRoREC: The Reach of Radical Embodied or Enactive Cognition, University of Antwerp, 17–19 June.
LOGICA: Hejnice, Czech Republic, 17–21 June.
TAP: 7th International Conference on Tests and Proofs, Budapest, Hungary, 18–19 June.
GP@50: The Gettier Problem at 50, University of Edinburgh, 20–21 June.
ICFIE: 2nd International Conference on Fuzzy Information and Engineering, Kanyakumari, India, 22–23 June.
ISF: 33rd International Symposium on Forecasting, Seoul, Korea, 23–26 June.
HDIA: High-Dimensional Inference with Applications, University of Kent, Canterbury, 24–25 June.
CSR: 8th International Computer Science Symposium in Russia, Ekaterinburg, Russia, 25–29 June.
BW8: 8th Barcelona Workshop on Issues in the Theory of Reference, Barcelona, 26–28 June.
APPLIED PHILOSOPHY: Society for Applied Philosophy Annual Conference, University of Zurich, 28–30 June.
AIME: Artificial Intelligence in Medicine, Murcia, Spain, 29 May–1 June.

COURSES AND PROGRAMMES

Courses

BFAS: Spring School on Belief Functions Theory and Applications, Carthage, Tunisia, 20–24 May.
NORDIC SPRING SCHOOL IN LOGIC: Nordfjordeid, Norway, 27–31 May.

ACAI SUMMER SCHOOL 2013: Computational Models of Argument, King's College London, UK, 1–5 July.

ESSLLI: 25th European Summer School in Logic, Language and Information, Heinrich Heine University in Düsseldorf, Germany, 5–16 August.

ETHICSCHOOL: Virtual Summerschool on Ethics of Emerging Technologies, 9–13 September.

Programmes

APHIL: MA/PhD in Analytic Philosophy, 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.

MASTER PROGRAMME: Philosophy of Science, Technology and Society, Enschede, the Netherlands.

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 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 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.

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

MRES 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 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 MATHEMATICAL LOGIC AND THE THEORY OF COMPUTATION: Mathematics, University of Manchester.

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 Sebastian).

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

PHD SCHOOL: in Statistics, Padua University.

JOBS AND STUDENTSHIPS

Jobs

ASSISTANT PROFESSOR: in Logic or Analysis, Department of Mathematics, University of Connecticut, until filled.

POST-DOC POSITION: in Artificial Intelligence, Institute for Artificial Intelligence, University of Georgia, until filled.

POST-DOC POSITION: on Data Analysis for Knowledge Discovery and Decision Making, Department of Electrical, Computer, and Systems Engineering at Rensselaer Polytechnic Institute (RPI), Troy, NY, until filled.

ASSOCIATE PROFESSOR OR PROFESSOR: in Logic and the Philosophy of Science, University of Calgary, until filled.

POST-DOC POSITION: in Probabilistic Reasoning, Vienna University of Technology, Austria, until filled.

POST-DOC POSITION: in Cognitive Psychology and/or Computational Modelling at the Center of Experimental Psychology and Cognitive Science, Justus Liebig University Giessen, until filled.

ASSISTANT PROFESSOR: in Cognitive Psychology, Center of Experimental Psychology and Cognitive Science, Justus Liebig University Giessen, until filled.

POST-DOC POSITION: in Graphical Models / Structural Learning, Uncertainty Reasoning Laboratory, Queens College / City University of New York, until filled.

POST-DOC POSITION: in Artificial Intelligence / Biomedical Informatics, Stevens Institute of Technology, until filled.

LECTURER: in Philosophy of Language, Philosophy, UCL, deadline 4 January.

LECTURER: in Logic and Philosophy of Language, University of Edinburgh, deadline 4 January.

LECTURER: in Philosophy of Mind and Cognition, University of Edinburgh, deadline 4 January.

PROFESSOR: in Statistics, University of New South Wales, Australia, deadline 13 January.

LECTURER: in Statistics, University of Oxford, deadline 14 January.

POST-DOC POSITIONS: Israeli Center of Research Excellence in Algorithms, deadline 15 January.

LECTURER: in Philosophy of Mathematics, University of Oxford, deadline 18 January.

POST-DOC POSITION: in Machine Learning at the Intelligent Systems Laboratory, University of Bristol, deadline 20 January.

LECTURER: in Statistics, University of Kent, deadline 20 January.

POST-DOC POSITIONS: in Bayesian Inference, Department of Statistics, University of Oxford, deadline 8 February.

LECTURER: in Probability or Statistics, School of Mathematics, University of Bristol, deadline 11 February.

POST-DOC POSITIONS: in Philosophy of Social Science, TINT Centre of Excellence in the Philosophy of the Social Sciences, Helsinki, deadline 15 February.

POST-DOC POSITION: in Metaphysics of Science, Institut d'Histoire et de Philosophie des Sciences et des Techniques, Paris, deadline 15 February.

POST-DOC POSITION: in Theoretical Philosophy working on "Infinite Regress" project, University of Groningen, The Netherlands, deadline 8 April.

Studentships

PHD POSITION: on project "Non-Classical Foundations of Mathematics," Department of Mathematics and Statistics, University of Canterbury, New Zealand, until filled.

PHD POSITION: on the project "Models of Paradox," Philosophy, University of Otago, until filled.

PHD POSITION: on Data Analysis for Knowledge Discovery and Decision Making, Department of Electrical, Computer, and Systems Engineering at Rensselaer Polytechnic Institute (RPI), Troy, NY, until filled.

PHD POSITIONS: in the Statistics & Probability group, Durham University, until filled.

Two PhD Positions: for research project on “Managing Severe Uncertainty,” Department of Philosophy, Logic and Scientific Method at the London School of Economics and Political Science, deadline 11 January.

PhD Position: at the Institute for Logic, Language and Computation (ILLC), University of Amsterdam, deadline 14 January.

PhD Position: in Belief Functions Theory, Department of Computing and Communications Technologies, Oxford Brookes University, deadline 15 January.