The Reasoner

Volume 6, Number 8
August 2012

www.thereasoner.org
ISSN 1757-0522

Contents

Editorial

This month we cover the nature of scientific understanding; utility monsters and miserable clowns; computational logic; models and simulations; decisions, games and logic; the hardest logic problem ever; and betting on the Palace of Minos. *The Reasoner* con-

tinues to cover an eclectic mix of topics related to reasoning, inference and method broadly construed.

Please consider contributing to this mix. Features aim to put forward an original argument or perspective. Submitted features are refereed, with a decision often made within two weeks thanks to the tireless efforts of the editorial board. We primarily aim to publish features that are comprehensible and of interest to a broad audience. More specialised pieces are also welcome but should have an introductory paragraph that can be generally understood.

Please also consider submitting items of news and 'what's hot' columns. The aim of these pieces is to help readers stay aware of what is going on in the reasoning community. These pieces are not refereed and a decision can normally be made very quickly.

With a thousand regular subscribers and well over a thousand downloads of each issue on average, this gazette reaches a substantial audience—an audience from a variety of academic disciplines and other walks of life, and from more than a hundred different countries. But please encourage others to take a look and to contribute too.

Jon Williamson
Philosophy, Kent

Features

Two Gaps in the Contextual Theory of Scientific Understanding

In a previous paper in this journal (volume 6, issue 4) I have criticised the contextual theory of scientific understanding as developed by Henk De Regt & Dennis Dieks (2005: "A Contextual Approach to Scientific Understanding", *Synthese* 144, 137–170) and further developed by Henk De Regt (2009: "The Epistemic Value of Understanding", *Philosophy of Science* 76, 585–597) at a meta-level: I have argued that in some respects it is not contextual and that its scope is more limited than the authors suggest. Here I discuss the theory at the object level, i.e., I go into what they say about scientific understanding. I will point at two gaps in the theory. My main point is that there is important unfinished business with respect to this theory. My critique is at the same time a call for action: if De Regt & Dieks or someone else fills the gaps, important progress can be made in the philosophical analysis of scientific understanding and scientific explanation.

De Regt & Dieks first develop an argument for the idea that achieving understanding is one of the epistemic aims of science (pp. 139–143) and then investigate what scientific understanding is. For convenience, I summarise the theory into the following crucial claim:

(QD) Understanding as an epistemic aim of science consists in the capacity (of the scientists in the field at a particular time) to make qualitative derivations with a theory.

This claim summarises the central tenets of the theory (CUP and CIT, see pp. 150–151). De Regt & Dieks do not deny that exact calculations are important. But these are not

the only thing that matter in science: understanding as defined in (QD) is an *additional* epistemic aim of science (p. 153).

In order to provide a convincing argument for (QD), De Regt & Dieks have to answer the question "What makes qualitative derivation interesting so that it can count as an additional aim of science?" The problem can be clarified by means of their example: the explanation of Boyle's law by the kinetic theory of gases. We have an empirically established quantitative version of Boyle's law, and this triggers the need for understanding. In order to clarify what understanding consists in in this case, they use the qualitative analysis which can be found in the introductory sections of Ludwig Boltzmann's *Lectures on Gas Theory* to make their case. They conclude as follows:

> Together these conclusions lead to a qualitative expression of Boyle's ideal gas law. It is important to note that the above reasoning does not involve any calculations. It is based on general characteristics of the theoretical description of the gas. Its purpose is to give us *understanding* of the phenomena, *before* we embark in detailed calculations. (pp. 152–153; italics in original)

The crucial question is: what is so interesting about being able to derive a "qualitative expression" of this law from the theory? Why do we want that thing? A positivist would answer: these qualitative derivations are only heuristically useful, they lead the way to exact predictions. But that road is not open to De Regt & Dieks: they want understanding to be an additional *aim* of science, heuristic value is not enough. In order to establish the surplus value of qualitative derivations, there are two options. The first is to adopt the following position:

> Qualitative derivations are important for understanding. Quantitative accuracy is important in practical contexts (manipulation or forecasting) but not for understanding.

The second option is the following position:

> Both qualitative derivations and quantitatively exact calculations are important for understanding.

De Regt & Dieks choose the first position. They do not argue against this second position, but that is only a minor problem: this alternative would preserve the core idea that qualitative derivations are crucial for understanding. However, there is also a third position which completely contradicts their views:

> Qualitative derivations are useless for understanding, only quantitatively exact calculations provide understanding.

The first important gap in the theory is that they do not argue against this view (the fact that Boltzmann wrote a qualitative analysis is not an argument against this view; what we need is vindication of why and how this analysis provides understanding).

The second gap can be introduced by means of a quotation from the seminal paper of the complex-system mechanist tradition:

> Descriptions of mechanisms render the end stage intelligible by showing how it is produced by bottom out entities and activities. (Machamer, Darden & Craver 2000: "Thinking about Mechanisms", *Philosophy of Science* 67, 21)

In the example of De Regt & Dieks, the behaviour of gases is explained by referring to the entities of which the gas is constituted (molecules) and the activities of these entities. So it is a good explanation according to a complex-system mechanist. The question now is: could it have had another format? For the complex-system mechanist the answer is negative: the fact that the explanation uses lower-level entities and their activities is essential. This answer has certain advantages, e.g., that it allows one to deal with explanatory asymmetries: we can understand Boyle's law by means of the kinetic theory of gases, but not the principles of the latter theory by means of Boyle's law because the part-whole relation is asymmetric and it has to be used in the bottom-up direction in explanations.

In order to solve this ambiguity in the theory and in its relations to complex-system mechanism, there are two options. Both options are worth exploring because they would lead to progress in the philosophical analysis if scientific understanding. First, one can regard entities and activities as essential, which entails that the theory fits into the complex-system mechanist tradition. The main contribution of De Regt & Dieks then lies in (QD): while other complex-system mechanists focus on the content of the explanans (what do mechanisms looks like?) they focus on the relation between explanans and explanandum (how do mechanisms explain?).

The other option is to argue that entities and activities are not essential: theoretical models can look completely different. In commenting on the explanation of Boyle's law they write:

> In this case, the general picture of moving gas *particles* allows us to make qualitative predictions of *macroscopic* properties of gases in particular situations. (p. 153, emphasis mine)

Either one claims that the part-whole relation to which the italicised terms refer is essential for providing understanding (first option) or not (second option). Philosophers who choose the second option have to argue that complex-system mechanists are mistaken. Furthermore, they should either propose and defend an alternative (a different restriction that saves the asymmetry of understanding) or argue that there are no restrictions on the nature of the theoretical model that is used as explanans.

My conclusion is that the contextual theory of understanding has two important gaps: it has no argument for its crucial claim about the value of qualitative derivations and it is not clear whether and how it differs from complex-system mechanism.

Erik Weber
Centre for Logic and Philosophy of Science
Ghent University

What's the Use of Utility Monsters?

In a famous thought experiment, Robert Nozick (1974: *Anarchy, State, and Utopia*, Oxford: Basil Blackwell, 41) challenges total utilitarianism because of the presumably possible existence of beings called 'utility monsters': "Utilitarian theory is embarrased by the possibility of utility monsters who get enormously greater gains in utility from any sacrifice of others than these others lose. For, unacceptably, the theory seems to require that we all be sacrificed in the monster's maw, in order to increase total utility." Here, I argue that Nozick's thought experiment is unable to provide a reason for rejecting total utilitarianism because it does not meet a requirement that all thought experiments have to meet in order to challenge an ethical theory.

One feature of an acceptable ethical theory is that it is action-guiding: It guides us to perform actions that rational persons would perform. It seems that there exists an 'Action-Guidingness Principle for Ethical Theories' (AGPET):

> **AGPET**: If an ethical theory is acceptable, then it is action-guiding.

It is easy to reconstruct an argument against total utilitarianism that incorporates AGPET and would probably have been accepted by Nozick:

> **MONSTER** 1. If the existence of utility monsters is possible and individual or collective self-sacrifice cannot be reasonably demanded from a person or a group of persons, then total utilitarianism is not action-guiding. (Monster I)
> 2. The existence of utility monsters is possible. (Monster II)
> 3. Individual or collective self-sacrifice cannot be reasonably demanded from a person or a group of persons. (Assumption)
> Therefore 4. Total utilitarianism is not action-guiding. (From Monster I, Monster II, Assumption)
> 5. If an ethical theory is acceptable, then it is action-guiding. (AGPET)
> Therefore 6. Total utilitarianism is unacceptable. (From 4., AGPET)

Monster I and Monster II are provided by Nozick. The third premise, called 'Assumption', is needed because a staunch defender of total utilitarianism could argue that his theory tells us that we ought to commit self-sacrifice and is, therefore, action-guiding. However, the defender's objection should not be accepted because rejecting Assumption seems to imply the rejection of the principle that 'ought' implies 'can' in at least one specific reading that conceives of 'can' as 'is psychologically able to'. But many, perhaps most, people are psychologically unable to sacrifice themselves which is why they do not have this obligation. The total utilitarian has to look for another premise to attack if he wants to challenge MONSTER.

It is crucial to know exactly which kind of possibility is involved in Monster I and Monster II. (They obviously have to use the same kind of possibility to avoid an equivocation.) The candidates are logical and nomological possibility. As I believe, several reasons speak against nomological possibility although it cannot be ruled out with certainty. First, a utility monster has not yet been observed although it would have had

enough time to evolve during the last billions of years. Although we cannot be sure that no utility monsters will evolve in the future, there might be some deeper reasons why they have not evolved other than a mere non-occurrence of the sufficient genetic mutations. Second, being the subject of utility and disutility requires the satisfaction of some necessary conditions, for example, being able to process a certain amount of biological information. However, the information processing power of biological organisms is limited, and the proponents of MONSTER have to show that the information processing power needed for a utility monster to exist can be realised within the boundaries of biological organisms. Third, if nomological possibility within MONSTER is defended by claiming that a utility monster could also be a non-biological being, then it still has to be shown that non-biological systems can have an information processing power that is high enough. Furthermore, it seems to be the case that a non-biological system can only be the subject of utility in the sense discussed here if it is a conscious system. But we do not know whether a conscious non-biological system can exist at all.

Now that we have reason to believe that MONSTER is committed to logical possiblity, Monster I can be identified as its weak spot because it implausibly presupposes that a merely logically possible state of the world is sufficient to render an entire ethical theory unacceptable. Imagine, for example, another counterargument to total utilitarianism that involves the logically possible existence of a 'miserable clown'. This unfortunate being's utility is always and severely decreased whenever another being's utility is increased, and the decrease in the miserable clown's utility is vastly greater than the increase in other people's utility. I guess that few people believe that total utilitarianism is unacceptable and that we should not increase other people's utility just for the mere logical possibility of the existence of a miserable clown.

The reason why few people would accept the case of the miserable clown as an unsound objection to total utilitarianism is the same reason why we should not accept MONSTER as an objection: Both thought experiments present us with logically possible states of the world that can be doubted to be nomologically possible and will never be relevant for any of our actions because our actions are not situated in a world that is solely governed by concepts and logical rules but by the laws of nature. These thought experiments violate what I call the 'Action-Relevance Principle for Ethical Thought Experiments' (ARPETE):

> **ARPETE**: If a thought experiment successfully challenges the action-guidingness of an ethical theory, then it is action-relevant.

Action-relevance is conceived of as a necessary requirement for a thought experiment to pose a threat to the action-guidingness of an ethical theory. Mere logical possibility of a scenario presented in a thought experiment is not sufficient for the thought experiment to successfully challenge an ethical theory by declaring it to fail to be action-guiding. If the ethical theory appears to be action-guiding, then its action-guidingness cannot be doubted just on the basis of a scenario that will never occur because it is ruled out for its nomological impossibility and, therefore, will never be relevant to any of our actions. In another world, though, where the laws of nature are different, ethical theories that are

acceptable and action-guiding in our own world might indeed be questionable.

Tobias Hainz
Philosophy
Heinrich-Heine-University Düsseldorf

News

Italian Convention on Computational Logic, 6–7 June

The 9th Italian Convention on Computational Logic (CILC 2012) has been held in Rome, Italy, organized by Francesca A. Lisi (Università degli Studi di Bari "Aldo Moro") and Fabio Patrizi ("Sapienza" Università di Roma). CILC 2012 encompasses the 27th annual meeting of the Italian Association for Logic Programming (GULP) affiliated to ALP. Since 1986, the GULP meeting has continually widened its scope by using the general term *Computational Logic* for integrating the various research fields, which use direct or indirect, practical or theoretical ways or just addresses the ideas or techniques of logic as a tool for representation and calculation. Today it has become the most important occasion for meeting and exchanging ideas and experiences between users, researchers and developers, who work in the field of Computational Logic in Italy.

The programme of CILC 2012 featured one invited talk and twenty-two regular talks. The invited speaker, Umberto Straccia, reviewed the state of the art in representing and reasoning with fuzzy knowledge in Semantic Web languages. The regular talks spanned from Description Logics (DLs) to Answer Set Programming (ASP) and Constraint Logic Programming (CLP).

Baldoni *et al.* presented an operational semantics for the commitment protocol language 2CL. Cantone *et al.* discoursed on the elimination of quantifiers through descriptors in FOL. Cattafi *et al.* illustrated a home health care application of CLP. Colucci and Donini faced non-standard reasoning problems in DLs with Prolog. Costantini and Formisano formalised preferences and priorities in ASP. Costantini and Tsintza introduced temporal meta-axioms in logical agents. Dal Palù *et al.* reported on promising computing resources for parallelizing SAT solvers. De Angelis *et al.* presented a CLP-based method for model checking of imperative programs. Dowek and Gabbay proposed a nominal semantics for FOL. Febbraro *et al.* described novel features of the ASPIDE system. Ferilli *et al.* discoursed on similarity assessment of first-order Horn clauses. Flaminio and Hosni discussed the epistemic structure of de Finetti's betting problem. Gavanelli *et al.* illustrated an ASP approach for the valves positioning optimization in a water distribution system. Gentilini *et al.* talked about rank-based simulation on acyclic graphs. Giordano *et al.* provided a general semantic framework for nonmonotonic reasoning. Giordano *et al.* proposed a model checking approach for the verification of action theories in ASP. Kerhet and Franconi discoursed on checking domain independence. Maratea *et al.* presented a multi-engine solver for ASP. Minervini *et al.* faced the problem of learning Bayesian classifiers in DLs. Riguzzi *et al.* introduced probabilistic ontologies in Datalog+/−. Riguzzi illustrated an optimized inference for probabilistic logic programs. Senni and Fioravanti proposed the use of CLP

for the generation of test data structures.

A special issue of the "Journal of Logic and Computation" with selected contributions from CILC 2012 is in preparation.

Francesca A. Lisi
Dipartimento di Informatica
Università degli Studi di Bari "A. Moro"

Models and Simulations 5, 14–16 June

The 5th conference on Models and Simulations (MS5) took place at the University of Helsinki, 14–16 June 2012. As with the previous conferences in Paris, Tilburg, Charlottesville and Toronto, MS5 brought together an interdisciplinary audience sharing an interest in epistemological, methodological and policy issues concerning modelling and simulation. Although the major part of the participants consisted of philosophers of science and philosophically minded modellers, also represented were the perspectives of cognitive science and sociology of science.

All the Models and Simulations conferences have managed to offer a great mix of presentations ranging from specific issues of modelling in various branches of science to the more general questions on the nature of modelling and the various uses of models. The Helsinki conference was no different in this respect. There were 45 presentations, and the themes of idealization, representation, model explanation, robustness, and the specific features of statistical and mathematical models were discussed and developed based on the foundations of systematic study and various interesting scientific case studies. Of the relatively new themes in the discussion on modelling and simulation, the use of fictions in modelling and the nature and role of model-based evidence especially spurred a lot of interest.

Climate science, as well as modelling in biology, ecology and physics, was discussed at MS5. A new thematic flavour to the conference series was added by the special focus on models and simulations within and across social sciences—a theme close to the heart of the local hosting institution, The Finnish Centre of Excellence in the Philosophy of the Social Sciences. The nature of economic models has been a hot topic for discussion for several decades now, and with the current financial crises at hand, there seems to be no end in sight for this endeavor. The fairly new trend of doing computer simulations in social sciences has not been left unnoticed among philosophers of science, either. In addition to several papers on modelling in economics, there were also presentations on agent-based modelling in sociology and cognitive science.

The three keynote presenters took up big issues in philosophy of modelling although none of the three presenters is primarily working in the field of philosophy of science. Mary Morgan, an economic historian from London School of Economics and Political Science, presented features that, according to her, make models appropriate "working objects" in gaining scientific knowledge. Models should be small-scale enough to be manageable, they should contain resources for investigation, and finally, they should be justifiable as typical, representative or exemplary for some class of phenomena.

Rosaria Conte, a social and cognitive scientist (Institute of Cognitive Sciences and

Technologies, CNR, Rome), with extensive work experience on agent-based modelling, targeted the role of mathematics in social sciences. Is mathematics needed in all modelling of social and economic phenomena or are there cases where non-mathematical models might bring more insight than mathematical ones? Conte's view is that mathematic modelling can indeed be irrelevant or even harmful in understanding social phenomena if it becomes "autarchic," that is, when it is meant to replace any other type of models, explanations or evidence.

Tim Benton, a population ecologist from the University of Leeds specialized in food security issues, closed the conference with a serious note. He presented a wide array of alarming evidence concerning the climate change, population growth and food security issues. What stood out clearly was the need to get appropriate model-based evidence for making societally difficult decisions in the next decades. Benton defended the importance of modelling socio-ecological systems in all their complexity, and strongly contrasted this view with the traditional praise of the fruitfulness of simple models in understanding phenomena. Another point that Benton stressed was that philosophers of science are seriously needed in the process of increasing the policy-relevance of scientific knowledge gained from modelling and simulation. The use of model-based evidence in shaping and implementing policies in a reliable way needs to be based on the reflection of what kind of knowledge can be gained by modelling and how model-based knowledge could be assessed. The conference participants seemed to widely agree with Benton that philosophers of science should take upon themselves this important task of consulting modellers and their audiences in shaping future policies to cope with the severe environmental and other challenges facing humanity.

Päivi Seppälä
Tarja Knuuttila
Philosophy, University of Helsinki

Decisions, Games & Logic, 28–30 June

The sixth Workshop in Decisions, Games & Logic was held at the Munich Center for Mathematical Philosophy (Ludwig-Maximilians-Universität) from June 28th to June 30th. It was organised by Thomas Meier (Munich), Martin Rechenauer (Munich), Olivier Roy (Munich, Chair) and Conrad Heilmann (Rotterdam). Three invited speakers delivered instructive tutorials at the beginning of each conference day covering the three main topics. Three contributed papers followed by a short comment were also presented each day.

Richard Pettigrew (Bristol) opened the workshop by giving the first tutorial. He talked about the application of decision theory in epistemology where the focus is not on how bad credences lead to bad actions, but rather on how they lead to bad epistemic states.

Whereas Pettigrew applied decision theory in epistemology, Amanda Friedenberg (Arizona) went the other way and applied epistemology in game theory. In this second tutorial, she talked about enriching classical game theory with epistemic elements like the beliefs of the agents about the rationality or the strategies of the other agents.

The third tutorial was given by Sonja Smets (Amsterdam), who introduced the audience to dynamic epistemic logic and models of belief upgrades. She introduced an updating rule called Action-Priority Rule. It is a belief-revision rule in the spirit of AGM revision which doesn't require the performance of arithmetic operations on degrees of belief.

A central moment of the workshop was the round table on 'Beliefs and degrees of beliefs' with Richard Bradley (LSE), Branden Fitelson (Rutgers), Amanda Friedenberg (Arizona), Hannes Leitgeb (MCMP / LMU) and Sonja Smets (Amsterdam). In a first round, each of the participants made a brief presentation of his point of view on the topic. It was followed by a lively discussion between the participants and with the audience.

Unusual in formal philosophy, the workshop included a poster session where the participants briefly presented their work in front of the audience and could then discuss in front of their poster while enjoying a cold buffet. This style of discussion allows for more personal and detailed comments and the participants found it very constructive.

Most of the talks and tutorials were recorded and will shortly be available on the Center's iTunes channel.

Karine Fradet
Philosophy, Université de Montréal
Seraphin A. Frimmer
MCMP, Ludwig-Maximilians-Universität München

Calls for Papers

The Aim of Belief: special issue of *Teorema*, deadline 15 September.
Science vs. Society? Social epistemology meets the philosophy of the humanities: special issue of *Foundations of Science*, deadline 31 October.
Grammatical Inference: special issue of *Machine Learning*, deadline 1 December.

What's Hot in ...

Logic and Rational Interaction

Among the the first cases of rational interaction studied by logicians were fictive stories they created themselves: Logic puzzles such as the following

> A hiker arrives at a junction with two roads, only one of them leading to his destination. There are two natives around, one of them a knight *always* telling the truth, the other one being a knave *always* lying. However, it is not clear who of the two natives is who. Can the hiker find out in one question which road to follow?

One of the most well-known generalisations of this puzzle is the *Hardest Logic Problem Ever* (HLPE), originally published by Smullyan in 1983:

> There are three gods standing at the road, one being a truth teller, one a liar and the third one giving random answers. The gods do understand English, but answer the question with *da* and *ja* meaning *yes* and *no*, but not necessarily in this order.

Over the last three years HLPE has received increased attention by logicians, attempting to find suitable formalisations. For instance Rabern and Rabern have pointed out the necessity to differentiate between two different ways of spelling out the random god's behaviour, one depending on the answer to the question, the other one not. A second important factor is whether the other gods can predict random's behaviour or not. In the latter case the possibilites *da* and *ja* are not sufficient. When answering a question would require determining random's behaviour the gods need to have 'remaining silent' as a third option.

The central question about HLPE is: How many questions are necessary to find out which road to take? A bunch of recent results show that the solution depends upon the exact formulation of the problem. For example Wheeler and Barahona published a paper last year titled "Why the [HLPE] Cannot Be Solved in Less than Three Questions". On the other hand, Uzquiano (2009) has a two question solution for the case of the other gods not being able to predict the random god's behaviour. Further two-question solutions can be obtained when allowing for self-referential questions.

In his recent PhD thesis "Playing with Truth," Stefan Wintein gives a formalization of the different versions of HLPE, presenting all the solutions mentioned above. More generally, he shows that allowing for self-referential statments can increase the computational power of a query language.

LORIWEB is always happy to publish information on topics relevant to the area of Logic and Rational Interaction—including announcements about new publications and recent or upcoming events. Please submit such news items to Rasmus Rendsvig, our web manager or to the loriweb address.

Dominik Klein
TiLPS, Tilburg University

Uncertain Reasoning

In the most immediate sense of the word, uncertainty is future directed: if we must make a decision now, we often feel that such a decision should be based on a sensible weighing of the uncertainty related to the future consequences of our present behaviour. Financial transactions are a paramount example of this and the question as to whether sensible quantitative predictions, say about asset pricing, can be made at all is heatedly debated.

More generally, it appears as if a combined temporal-epistemic dimension is common to all genuine cases of uncertainty. I'm saying "genuine" here because nothing prevents people from gambling on known facts, although such a quantification of uncertainty can hardly be related to rational behaviour. Indeed two conditions appear to be necessary in order for the gambling interpretation to make sense as a model of uncertainty quantification. First, at the time of betting (say t_0), neither the bookmaker nor the

gambler should know the truth value of the event E of interest. Second, at time t_0 they must agree on what conditions will have to be met at time $t > t_0$ in order for the event to have (not) occurred. Both conditions can be seen as defining properties of events as in de Finetti's theory of subjective probability—see, e.g., de Finetti (1972: *Probability, Induction and Statistics*, Wiley).

One fundamental aspect of this interpretation is that events admitting of a genuine quantification of uncertainty must be decided at some time $t > t_0$. Suppose on 1st July 2012 a financial broker bet on Greece exiting the Euro by May 2013 (an event usually referred to as GREXIT). Then, on 1st June 2013 GREXIT will be decided by some truth of the matter concerning Greece, and the broker will cash their win if Greece returned to the Drachma and pay their losses otherwise. Note that the generality of the model depends on the fact that we can somehow relax the reference to the underlying truth of the matter giving the arrow of time an epistemic interpretation based on the conditions which are sufficient for gamblers and bookmakers to agree on whether event E occurred or not. Thus, as recently reported by the BBC, Stephen Hawking lost 100 dollars in his 2008 bet against the event that the LHC experiment would have decided positively the existence of any Higgs boson-like particle. Hawking acknowledged losing his bet because the existence of the boson-like particle has now been inferred with a probability which falls within the confidence interval accepted by the community.

This raises the question as to whether we can completely detach our quantification of uncertainty from any underlying truth of the matter. Put otherwise, does it make sense to quantify our uncertainty about propositional entities which are not events in the sense that we know at time t_0 that E will never be determined as either true or false? Archeology seems to be one research field in which uncertainty comes naturally in this form. A particularly interesting example is illustrated in R. Castleden (1990: *The Knossos Labyrinth: A new view of the 'Palace of Minos' at Knossos*, Routledge).

Sir Arthur Evans's excavation of the Minoan complex at Knossos, a few miles south of Crete's capital Heraklion, began in March 1900 and soon became a highly controversial scientific enterprise. Owing to the excitement caused by Schliemann's identification of ancient Troy during the late 1870s, Evans's main motivation was to set the legendary Minoan kingdom of Minos on a firm historical footing. Evans's goal was to discover the "Palace of Minos," so he started by assuming that the ruins he was excavating at Knossos were in fact those of the "Palace" from where Minos peacefully ruled the Aegean some four millennia ago. This assumption didn't constrain Evans too much, for there is no agreement as to whether Minos actually existed beyond Greek mythology. The structural, mostly concrete-based, reconstructions at the Knossos site, the rather arbitrary arrangement of largely guessed frescoes and the dramatic labelling of the site's areas—e.g., The Throne Room, The Queen's Bathroom, up to the very unlikely Piano Nobile—were all concocted by Evans to create a very definite suggestion that the complex at Knossos could not be other than Minos's palace. Castleden's book, however, illustrates how our best interpretation of the excavation evidence, mostly the pictures taken by Evans *before* his reconstruction, is largely inconsistent with Evans's own interpretation and fits with an alternative, perhaps less exciting, interpretation of the Labyrinth as a religious complex.

A fascinating aspect of Castleden's line of argument is that consistency is the only

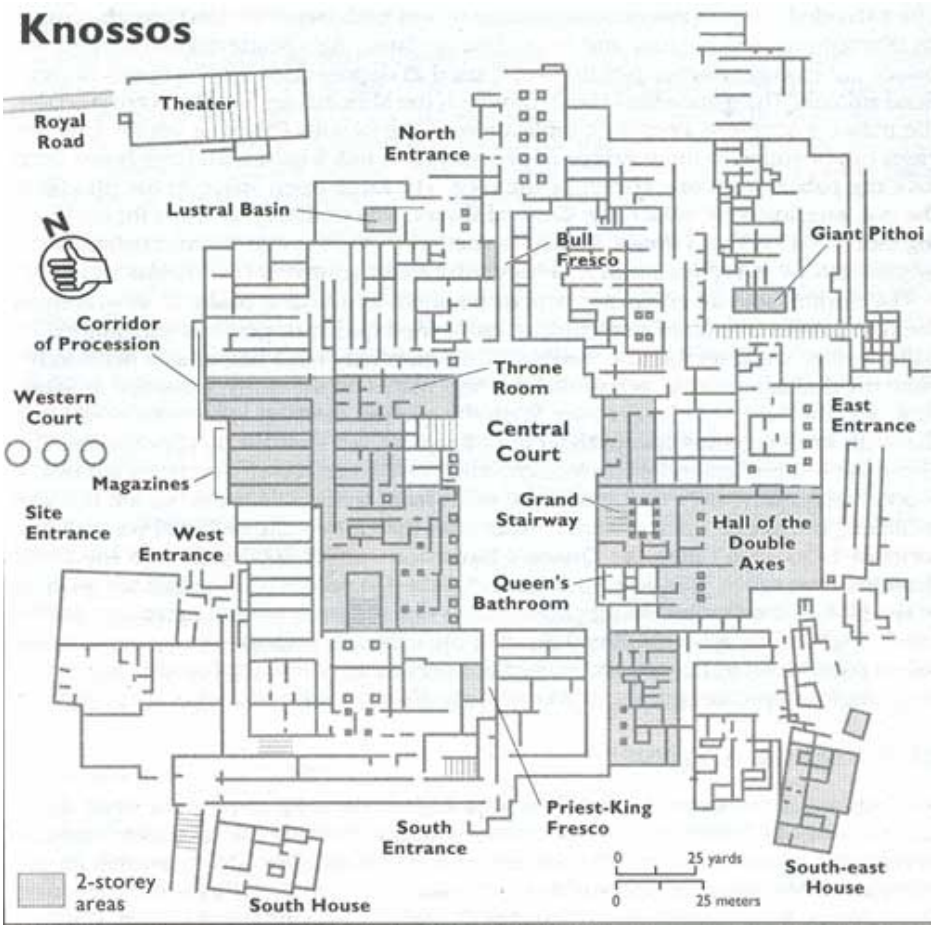

uncontroversial principle when reasoning must take into account extremely heterogeneous sources of information and unreliable sources of evidence. In the specific case of Knossos things are even worse, as there is often no agreement on what counts as evidence in the first place. The many mythological accounts of the lives of Minos, his wife Pasiphae and her monstrous son—the Minotaur—diverge a good deal as they were preserved orally for about one millennium before the classical Greek and then Roman writers fixed them in the forms we now know. The Linear B written records which were found at Knossos and in the excavation of other "palatial" complexes in Crete, have been deciphered only in 1952, but do not provide particularly useful information. Tablets with Linear A inscription still remain undeciphered today.

In the light of this, it is certainly not surprising that Castleden's arguments are all purely qualitative. Evans, on the other hand, invested all his personal wealth in the excavation of his Palace (a French commentator of the time referred to Sir Arthur as the "builder of ruins"). So the standard Bayesian interpretation would lead us to think that

Evans's had a very high degree of belief in the Knossos site being Minos's palace. This is a fact that Castleden stresses abundantly in his book: Evans really believed that the Labyrinth was the Palace of Minos and to make sure he didn't lose his bet, he provided all the required evidence.

Hykel Hosni
Scuola Normale Superiore, Pisa
CPNSS, LSE

Events

August

NAFIPS: 31th North American Fuzzy Information Processing Society Annual Conference, Berkeley, 6–8 August.
CLAM: Logic and Computability Session, Latin American Congress of Mathematicians, Argentina, 6–10 August.
PMUV: Philosophy and Mathematics of Uncertainty and Vagueness, Brazil, 6–15 August.
KDD: 18th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Beijing, China, 12–16 August.
StaRAI: 2nd Statistical Relational AI Workshop, Cataline Island, USA, 13 August.
ITP: 3rd Conference on Interactive Theorem Proving, Princeton, NJ, 13–16 August.
Logic and Cognition: Logic and Cognition Workshop, Opole, Poland, 13–17 August.
Historical Counterfactuals: Workshop, Bristol, 14 August.
UAI: Conference on Uncertainty in Artificial Intelligence, Catalina Island, USA, 15–17 August.
BMAw: 9th Bayesian Modeling Applications Workshop, Catalina Island, 18 August.
UiNI: Uncertainty in Natural Intelligence Workshop, Catalina Island, 18 August.
SLS: 8th Scandinavian Logic Symposium, Roskilde University, Denmark, 20–21 August.
ALFAn: Latin American Analytic Philosophy Conference, Buenos Aires, Argentina, 21–24 August.
CTF: Concept Types and Frames in Language, Cognition, and Science, Düsseldorf, 22–24 August.
AiML: Advances in Modal Logic, Copenhagen, 22–25 August.
EASLLC: International Conference and the Second East-Asian School on Logic, Language and Computation, Chongqing, China, 25–31 August.
FLINS: 10th International FLINS Conference on Uncertainty Modeling in Knowledge Engineering and Decision Making, Istanbul, Turkey, 26–29 August.
CLIMA: 13th International Workshop on Computational Logic in Multi-Agent Systems, Montpellier, France, 27–28 August.
STeDy: Spatio-Temporal Dynamics Workshop, Montpellier, France, 27–28 August.
WL4AI: Weighted Logics for AI Workshop, Montpellier, France, 27–28 August.

AIGM: Algorithmic Issues for Inference in Graphical Models, Montpellier, France, 27–28 August.
ARCOE: 4th International Workshop on Acquisition, Representation and Reasoning with Contextualized Knowledge, Montpellier, France, 27–28 August.
ECAI: 20th European Conference on Artificial Intelligence, Montpellier, France, 27–31 August.
COMPSTAT: 20th International Conference on Computational Statistics, Cyprus, 27–31 August.
Collective Intentionality: University of Manchester, 28–31 August.
Bayes Lectures: University of Edinburgh, 29–30 August.
CNL: Workshop on Controlled Natural Language, Zurich, 29–31 August.
FoR&D: Conference on Frontiers of Rationality and Decision, University of Groningen, 29–31 August.

September

CSL: 21st EACSL Annual Conference on Computer Science Logic, Fontainebleau, France, 3–6 September.
WoLLIC: Workshop on Logic, Language, Information and Computation, Argentina, 3–6 September.
ABS: Applied Bayesian Statistics School, Italy, 3–7 September.
ICLP: 28th International Conference on Logic Programming, Budapest, 4–8 September.
iKNOW12: 12th International Conference on Knowledge Management and Knowledge Technologies, Graz, Austria, 5–7 September.

ECiTS
Evidence and Causality in the Sciences, University of Kent, 5–7 September

GAMES: Games for Design and Verification, Napoli, Italy, 7–12 September.
Intuitions, Experiments and Philosophy: University of Nottingham, 8–9 September.
Logic and Relativity: 1st International Conference on Logic and Relativity, Budapest, 8–12 September.
WEO-DIA: 1st Workshop on Well-founded Everyday Ontologies–Design, Implementations & Applications, Wroclaw, Poland, 9 September.
COMMA 2012: 4th International Conference on Computational Models of Argument, Vienna, Austria, 10–12 September.
ITA13: 5th International Conference on Internet & Applications, Glyndwr University, Wrexham, North Wales, UK, 10–13 September.
LATD: Logic, Algebra and Truth Degrees, Japan, 10–14 September.
WPMSIIP: 5th Workshop on Principles and Methods of Statistical Inference with Interval Probability, Munich, Germany, 10–15 September.
Datalog 2.0: 2nd Workshop on the Resurgence of Datalog in Academia and Industry, Vienna, Austria, 11–14 September.
L&R: Workshop on Lattices and Relations, ILLC, University of Amsterdam, 12–14 September.

WUPES: Workshop on Uncertainty Processing, Czech Rep., 12–15 September.
ENFA: 5th Meeting of the Portuguese Society for Analytic Philosophy, University of Minho, Braga, 13–15 September.
SOPhiA: Salzburg Conference for Young Analytic Philosophy, University of Salzburg, Austria, 13–15 September.
Colloquium Logicum: Paderborn, Germany, 13–15 September.
SIFA: 10th National Conference of the Italian Society for Analytic Philosophy, Alghero, 13–15 September.
SUM: 6th International Conference on Scalable Uncertainty Management, Marburg, Germany, 17–19 September.
ILP: 22nd International Conference on Inductive Logic Programming, Dubrovnik, 17–19 September.
GAP8: 8th Conference of the Society for Analytic Philosophy, Konstanz, Germany, 17–20 September.
Logical Form: University of Cambridge, 18–19 September.
SemDial: 16th Workshop on the Semantics and Pragmatics of Dialogue, Université Paris-Diderot, 19–21 September.
PGM: 6th European Workshop on Probabilistic Graphical Models, Granada, Spain, 19–21 September.
Formal Methods in Argument Reconstruction: Konstanz, Germany, 20–21 September.
CaLintSS: Causation and Laws in the Special Science—Metaphysical Foundations, Konstanz, 21–22 September.
Philosophical Issues in Belief Revision, Conditional Logic and Possible World Semantics: Konstanz, Germany, 21–22 September.
ENPOSS: 1st European Network for the Philosophy of the Social Sciences Conference, University of Copenhagen, 21–23 September.
MLSP: IEEE Workshop on Machine Learning for Signal Processing, special session on Causal Discovery, Spain, 23–26 September.
Structure and Uncertainty: Workshop on Modelling, Inference and Computation in Complex Stochastic Systems, Bristol, 24–27 September.
ECML-PKDD: European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases, Bristol, UK, 24–28 September.
JELIA: 12th European Conference on Logics in Artificial Intelligence, Toulouse, 26–28 September.
Consciousness and Volition: 1st International Krakow Conference in Cognitive Science, Krakow, Poland, 27–29 September.
LNK: 5th Conference on Non-Classical Logic. Theory and Applications, Poland, 27–29 September.
MEW6: 6th annual Midwest Epistemology Workshop, Indiana University, Bloomington, 28–29 September.
LSFA: 7th Workshop on Logical and Semantic Frameworks with Applications, Rio de Janeiro, 29–30 September.

October

Departing from Sainsbury: University of Barcelona, 1–2 October.
SMPS: 6th International Conference on Soft Methods in Probability and Statistics, Konstanz, 4–6 October.
FPMW: 4th French PhilMath Workshop, Collège de France, Paris, 4–6 October.
Philosophy of Scientific Experimentation: University of Colorado, Boulder, 5–6 October.
TiC2: Turing in Context II: Historical and Contemporary Research in Logic, Computing Machinery and AI, Brussels, 10–12 October.
Formal Ethics: Munich, 11–13 October.
CoMiC: Graduate Conference in Philosophy of Mind and Cognitive Science, Edinburgh, 12–13 October.
The Roles Of Experience in A Priori Knowledge: University of Cologne, Germany, 13–14 October.
PhiloSTEM: Midwest Workshop in Philosophy of Science, Technology, Engineering, and Mathematics, Indiana University-Purdue University, Fort Wayne, IN, 19–20 October.
Numbers & Truth: The Philosophy and Mathematics of Arithmetic and Truth, University of Gothenburg, Sweden, 19–21 October.
ATAI: Advanced Topics in Artificial Intelligence, Bali, Indonesia, 22–23 October.
ECREA: 4th European Communication Conference, Istanbul, Turkey, 24–27 October.
IDA: 11th International Symposium on Intelligent Data Analysis, Helsinki, Finland, 25–27 October.
ISELL: International Symposium Of Epistemology, Logic And Language, Lisbon, Portugal, 29–30 October.

November

MAgg: AAAI Fall Symposium on Machine Aggregation of Human Judgment, Arlington, VA, USA, 2–4 November.
ACML: 4th Asian Conference on Machine Learning, Singapore, 4–6 November.
BotB: Bayes on the Beach, Queensland, Australia, 6–8 November.
Cultures of Mathematics and Logic: Guangzhou, China, 9–12 November.
URSW: Uncertainty Reasoning for the Semantic Web, Boston, USA, 11–12 November.
Arché/CSMN: Graduate Conference, University of Oslo, Norway, 17–18 November.
SILFS: Italian Society of Logic and Philosophy of Science Conference, University of Milan-Bicocca, 20–21 November.
Modal Logic in the Middle Ages: University of St Andrews, 22–23 November.
CogSc: ILCLI International Workshop on Cognitive Science, Donostia, San Sebastian, 28–30 November.
René Descartes Lectures: Tilburg Center for Logic and Philosophy of Science, 28–30 November.
LEMMing: Graduate Conference, Cologne, Germany, 29 November–1 December.

Courses and Programmes

Courses

Einstein's Philosophy of Science: Summer School, University of Tübingen, 30 July–3 August.
ESSLLI: 24th European Summer School in Logic, Language and Information, Opole, Poland, 6–17 August.
FMIP: Munich / Groningen Summer School: Formal Methods in Philosophy, Groningen, 23–28 August.
TISS: Tübingen International Summer School—How do we make decisions?, 24–27 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.

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

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.
Associate Professor: in Cognitive Modelling, Institute for Logic, Language and Computation, University of Amsterdam, deadline 1 August.
Assistantship: in Theoretical Philosophy, University of Zurich, deadline 10 August.
Post-doc Positions: on project "Contextual and formal-logical approach to scientific problem solving processes," Centre for Logic and Philosophy of Science at Ghent University, deadline 12 August.
Post-doc Position: on the project "Explaining Language: Philosophical Perspectives on Computational Linguistics," TiLPS, Tilburg University, deadline 15 August.
Post-doc Position: in Machine Learning for Autonomous Systems, Department of Engineering, University of Cambridge, deadline 20 August.
Post-doc Position: in Computational Linguistics, Kings College London, deadline 27 August.
Lecturer: in Philosophy, AOS: Logic, Formal Linguistics & Philosophy of Language, Kings College London, deadline 27 August.
Research Programmer: on project "Logics for Autonomous Systems," Department of Computer Science, University of Oxford, deadline 7 September.

Studentships

PhD Positions: in the Statistics & Probability group, Durham University, until filled.
PhD Position: in Logic and Theoretical Philosophy at the Institute for Logic, Language and Computation at the University of Amsterdam, until filled.
Assistantship: in Theoretical Philosophy, University of Zurich, deadline 10 August.
Marie Curie Early Stage Researcher: in Mathematical Logic, Mathematics & Physical Sciences, University of Leeds, deadline 24 August.
Four PhD Positions: in "Foundations of the Life Sciences and their Ethical Consequences," European School of Molecular Medicine, University of Milan, deadline 3 September.
PhD Position: on the project "Knowledge Representation and Inference Based on Type-2 Fuzzy Sets and Systems," School of Computer Science, University of Nottingham, deadline 30 December.

FOR THE LAST HUNDRED YEARS, SWISS PATENT CLERKS HAVE BEEN UNDER SOME WEIRD PRESSURES.

University of Kent CENTRE FOR REASONING
www.kent.ac.uk/reasoning