The Reasoner

Volume 6, Number 3
March 2012

www.thereasoner.org
ISSN 1757-0522

Contents

Editorial

Features

News

What's Hot in ...

Events

Courses and Programmes

Jobs and Studentships

Editorial

While the importance of experiments in the natural sciences is not particularly controversial, their particular roles and uses for supporting theoretical claims remain a fertile ground for discussion and investigation by scientists and philosophers alike. Philosophical investigations about the nature of experiments and their role in scientific reasoning

concern largely epistemological issues about the growth of knowledge, its reliability and validity in relation to the specific circumstances and material settings in which experiments are carried out, and the status of problematic dichotomies such as theory vis-à-vis empirical evidence, observable *versus* unobservable entities or operational versus theoretical descriptions.

Sabina Leonelli

Another way of thinking about experiments has recently become prominent in several parts of philosophy—namely, the idea of using experiments to foster and support philosophy itself. This idea parallels older traditions of providing empirical grounding for philosophical claims by means of thought experiments and historical case studies (which have been particularly prominent and effective within the philosophy of science, thanks to advocates of 'experimental realism' such as Ian Hacking, but have also found forceful advocates among naturalists, operationalists and pragmatists, not to speak of the influence of the late Wittgenstein in encouraging philosophers to take the material world seriously). This approach challenges Marx's extensively quoted critique of philosophy as merely interpreting, rather than changing, the world: philosophers are here encouraged to assist in the production of 'facts' or 'data', which can in turn help to determine the validity of specific philosophical arguments. One way to do this is through so-called 'experimental philosophy'.

Opposed to what had been criticized as 'armchair philosophy', experimental philosophy gathers and evaluates data by conducting surveys and experiments on human subjects, which aim to explore the extent to which intuitions about philosophical problems such as determinism and free will are shared by large groups of individuals (and possibly influenced by social circumstances). Another way to introduce experiments into philosophy is by means of existing social science research. By merging social scientific methods of inquiry into philosophy, philosophers can take advantage of centuries of developments in the skills and methods needed to carry out research on human subjects and their activities—including qualitative methods such as interviews, discourse analysis and participant observation as used in sociology and anthropology. This trend is being particularly encouraged by the Society for the Philosophy of Science in Practice, which acknowledges the importance of studying actual scientific practices, past and present, through collaborations with practicing scientists; and is exemplified by the research carried out at the ESRC Centre for Genomics in Society in Exeter, where we are based, where collaborations with both social and natural scientists are routinely used to inform philosophical analysis. A third way to conduct experiments of use to philosophy is through another kind of interdisciplinary collaboration: that with the cognitive sciences.

Ann-Sophie Barwich

In the interview below, Giovanna Colombetti explains how she has successfully worked with a cognitive scientist to carry out an experiment to inform her philosophical view of the relationship between language and emotional feelings. Her work, as

well as the trends mentioned above, exemplifies the intellectual wealth to be gained by philosophers through interactions with researchers in the natural and social sciences. Such interactions do not weaken philosophy as a discipline and a worthwhile intellectual activity: on the contrary, they confirm the important role to be played by philosophy alongside other forms of knowledge in contemporary society, at the same time as they enable us philosophers to question and revise our own practices to keep up with our ever-evolving world.

Sabina Leonelli
Ann-Sophie Barwich
University of Exeter

Features

Interview with Giovanna Colombetti

Dr Colombetti is Senior Lecturer in the Department of Sociology and Philosophy of the University of Exeter, and the holder of a Starting Grant from the European Research Council on 'Emoting the Embodied Mind'. This grant is for her current project, which consists in proposing a reconceptualization of the emotions and other affective phenomena from an 'embodied-mind' perspective. Her work is highly interdisciplinary and empirically minded, and draws extensively on current theoretical and experimental work in the so-called 'affective sciences' (primarily the psychology and neuroscience of emotion). She is currently working on a book for MIT Press provisionally titled *The Feeling Body: Emoting the Embodied Mind*. Jo Donaghy and Sabina Leonelli have approached her with eight questions spanning several aspects of her work.

Giovanna Colombetti

Jo Donaghy & Sabina Leonelli: Given the diverse readership of *The Reasoner* could you start by providing us with a brief introduction to the 'embodied approach' in cognitive science?

Giovanna Colombetti: There is no single characterization of this approach. To keep things simple, I usually say that, according to the embodied approach, cognition and mind more generally do not depend on the brain only, but on the complex interrelation of neural, bodily and environmental processes. I also add that strictly speaking, from this perspective cognition is not just 'embodied', but also 'embedded' or 'situated' in the world. The embodied-embedded approach is best seen as a departure from the view that cognition is a 'higher' faculty realized by some kind of central representational processor located between perception and action. According to this view, the body is merely *extrinsic* to cognition: it provides inputs to the central processor via the sensory system, and carries out motor instructions. By contrast, according to the embodied-embedded approach, the body *constitutes* cognitive activity. For example, perception is not passive encoding of sensory information, but active exploration of the world, and as such it is

not separate from action. The extent to which the embodied-embedded approach still needs to posit some kind of 'internal representations' to account for cognitive activity is one major philosophical issue in this area. The embodied-embedded approach is also importantly related to phenomenology. Husserl, Merleau-Ponty, Heidegger, Sartre and others all emphasized, although in different ways, the place of the body in shaping and structuring our experience—from self-awareness to our dealings with the world, including other people.

JD & SL: How have interactions between philosophical and experimental approaches in cognitive science been important in the development of the embodied approach?

Jo Donaghy

GC: They have been crucial. All philosophical texts on embodied cognition make ample reference to experimental research. The development of the field of situated robotics, for example, has played a very important role in supporting the philosophical view that symbolic structures are not necessary for cognition. A properly situated robot is a robot that behaves intelligently in the world (e.g., picks up cans of coke in a messy office) without possessing any complete internal representation of its environment; rather its sensors are directly connected to its limbs, and built in such a way that different perception-action subroutines interact with one another to enable the robot to accomplish the task it was designed for. Another discipline highly relevant for the embodied approach has been developmental psychology, in particular the dynamical systems approach to development. According to the latter, development does not consist in behavioural changes brought about at specific stages by pre-determined instructions. Rather, behavioural changes result from complex nonlinear interactions among neural, bodily and environmental processes as they unfold over time. Again, the idea is that behaviour is not the output of internal instructions, but rather 'emerges' as the embodied organism interacts with the world. Experimental evidence supporting this view indicates for example that the development of the capacity to reach for an object that has been hidden in a certain location *A* and then moved, in full sight, to another location *B*, depends on the interaction of a variety of factors spanning brain, body and world (against the idea that it is a skill regulated by internal genetic or neural structures that mature at pre-determined stages). Philosophers of embodied cognition rely all the time on this kind of evidence to support the view that knowledge is 'performed' rather than 'represented'. Of course they do not *only* rely on experimental evidence to develop their arguments. Part of the enterprise is to critically assess ongoing experimental work and related theoretical claims made by scientists, and to highlight the crucial conceptual issues that are raised by a certain strand of empirical work.

JD & SL: What is the status of research on emotions in contemporary philosophy, particularly in relation to cognition?

GC: The philosophy of emotion, I would say, is still very much a field apart from the philosophy of cognition. The latter is now largely empirically minded, often guided by the latest experimental research, and engaged in interdisciplinary exchanges with cognitive scientists. The philosophy of emotion on the other hand is still very much

conducted 'from the armchair' (with notable exceptions of course). There is no such thing (yet) as 'the philosophy of affective science', and philosophical accounts of emotion rarely engage with the empirical literature. There is also not much in particular in terms of philosophical critiques of affective science—of its assumptions and theoretical constructs, which affective scientists themselves recognize as highly controversial and problematic. Finally, the philosophy of emotion, at least in the Anglo-American world, has hardly come in touch with philosophical phenomenology, with the result that there has been little attention to, for example, the experiential dimension of emotion and affectivity, their lived bodily nature, the extent to which they may 'incorporate' aspects of the world etc.

JD & SL: What particular questions are raised by taking an embodied approach to emotion?

GC: As I see it, emotion theory is still largely characterized by a dichotomy between mind and body. Whereas most affective scientists and philosophers of emotion today acknowledge that bodily changes (physiological, autonomic, musculoskeletal) are an important part of emotion, still they tend to distinguish quite neatly between these changes, and the purely 'brainy' cognitions or *appraisals* that allegedly cause them. In this sense then, emotions remain partially 'disembodied'. From an embodied-mind perspective, one is led to ask: what are these appraisals that are supposed to excite the body in emotion? How are they realized in the brain? Can we draw a clear separation between the brain-bases of appraisal, and non-neural bodily processes? Another question raised by taking an embodied approach to emotion regards the extent to which emotion *extends* not just beyond the skull, but the skin too. The view that the mind is extended (about the technologies that we rely on everyday to perform intelligently, for example), and not just embodied, is an important aspect of the embodied approach, and quite a lot has already been written about the possibility of extending cognition. What about emotion however? Is it possible to extend emotion beyond the skin? How? And what about specifically emotional *feelings*—can they be extended? What is the role of language in scaffolding our emotional life? And then there are all sorts of phenomenological questions one can ask from an embodied-mind perspective, which lead us beyond the long-debated issue of whether or not an emotion necessarily involves bodily feelings. Once the body is taken for granted, so to speak, then one can ask finer-grained questions such as: how is the body felt in emotion experience? Does the body always need to be explicitly attended to be experienced? Are there bodily feelings that are not feelings *of* the body? What is the relationship between bodily feelings and experience of the world?

JD & SL: What kind of experiments do you use as part of your research, and how do they help your philosophical project?

GC: I refer to all sorts of empirical evidence in my research, and I use it in different ways depending on the context. For example, I use experimental research in developmental psychology on the variability of emotional expressions in infants to argue against the idea that facial expressions are outputs of genetically instructions or 'affect programs' (a view held by several affective scientists), and to defend the alternative view that facial expressions are dynamically assembled coordinative motor structures. I then argue that similar considerations can be applied to other bodily and behavioural aspects of emotional episodes, and thus that emotional episodes altogether should be

seen as dynamical patterns. This is a case in which I use empirical research to defend one conceptualization of emotions rather than another. In other cases I present certain strands of experimental research as manifesting a philosophical view that I consider misleading, and I try to propose alternative approaches. For example, I think that the current way in which affective neuroscience looks at how different emotions depend on different brain areas is inadequate for the study of subjective emotion experience (i.e., emotional consciousness or feelings). The current affective-neuroscientific method relies on a predominantly behaviourist methodology that presents subjects with emotional stimuli (e.g., scary pictures, sad music, etc.) and looks at which brain and/or bodily areas are significantly activated in response. This methodology does not inquire into the nature of emotion experience (e.g., how is the body experienced in fear? Are there different experiences of fear? How are differences in intensity experienced?). For the latter, more qualitative methods are needed, together with specific ways to correlate experience and neuro-physiological activity.

JD & SL: Could you tell us a bit about your current collaboration with Kim Wright, a senior lecturer in clinical psychology at the Mood Disorders Centre (MDC) at the University of Exeter. How did the collaboration come about? What is Kim's main area of research and why is she interested in embodied cognition from a philosophical perspective?

GC: The collaboration came about as I was working on a paper on the relation between language and feelings. In it I distinguished among various ways in which putting one's feelings into words can affect them—for example I argued that in some cases verbalizing can help the subject become explicitly aware of her otherwise tacit experiences, whereas in other cases language can induce novel and original feelings, and so on. The reason why I decided to address this question is related to the embodied approach, although perhaps this is not immediately apparent ... some arguments for the view that the mind is 'extended' refer to language, which in this context is characterized as an external tool for enhancing our cognitive capacities. I thus wondered whether a similar point could be made about the relationship between language and feelings. In any case, this paper was written entirely from the proverbial philosophical armchair, and while I was working on it I realized that the question of whether language affects feelings is also an empirical one, that would be interesting to address experimentally. I mentioned this issue over lunch to Kim, whom I had met at one of those introductory courses on 'academic practice' for new lecturers, and who had already kindly invited me to present my work at the Mood Disorders Centre. Kim thought the question of the relationship between language and emotion experience was an interesting one, and relevant for her own work as well. Kim is a clinical psychologist and her main research area is bipolar disorder. One of the things she is interested in is how subjects affected by this disorder can learn to regulate their mood. Becoming increasingly aware of one's mood may be a first step, and labelling one's feelings may be one way of enhancing this self-awareness. I was lucky that Kim decided to give the issue more thought, and eventually came up with an experimental design that we discussed with other members of the MDC and refined accordingly, and eventually was financed as part of my ERC project. We are still analyzing the data so I cannot say more about the study ... but in terms of the collaboration with Kim, I can say it's been very stimulating—I learnt a lot about different

clinical approaches to emotion, and she has learnt about the embodied-extended mind view. I don't know that I have convinced her that the latter is 'the right approach'—but the interesting thing for me is that, in spite of coming from such different backgrounds, we were able to find a common research question and investigate it together.

JD & SL: Philosophical and experimental approaches in cognitive science constitute distinct bodies of disciplinary expertise, assumptions and research agendas. What challenges does this pose for people working in cognitive science who want to use both philosophical and experimental approaches?

GC: One main challenge is to remain competent in one discipline whilst opening up to a different one. Philosophy and experimental cognitive/affective science are different disciplines and as such ask different questions and use different methods. As a philosopher, I read many experimental papers, and it is difficult (impossible, in effect) to keep up with the burgeoning literature in affective science, including its more refined methodologies, the understanding of which often requires specific technical expertise. I do my best to identify relevant aspects of current empirical work, and whether and how it is relevant for my own philosophical views, but this takes a lot of time of course, which is taken away from reading 'purely' philosophical works. I don't think the kind of philosophy I am interested in can survive without engaging with the empirical sciences, but one challenge is to strike the right balance between remaining viewed as 'a philosopher' and acquiring sufficient expertise in another discipline to be able to discuss and criticize it competently.

JD & SL: How do you see the relationship between philosophy of mind and philosophy of science, especially in view of movements towards empirically-informed analysis in both fields?

GC: Very briefly, to the extent that philosophy of mind is philosophy of neuroscience and/or psychology, it is also philosophy of science. In this sense it needs to involve an evaluation and critique of empirical work (and not just an appropriation of empirical results). It is interesting that the philosophical view of the mind one favours typically influences the type of science one thinks is most relevant for studying the mind! In the embodied approach in particular, philosophy of mind comes in very close contact with philosophy of biology and even overlaps with it. This is because the embodied approach posits a deep continuity between life and mind, and consequently in it philosophical issues about the nature of the mind are often connected with philosophical issues about the nature of living systems—for example: what is the relationship between cognition and the organization of living systems? Are there principles of biological organization that can be used to describe and explain mental phenomena? Can non-biological systems have a mind or even be conscious?

Repeated St Petersburg Two-Envelope Trials and Expected Value

Consider the repeated St Petersburg two-envelope game. An even number of envelopes are filled with money, one-by-one, according to the following procedure: for each envelope flip a coin until it comes up tails, and if it comes up tails on the n^{th} flip, put $\$2^n$ into that envelope. The expected value of each envelope is infinite. Then let *a trial* consist of two envelopes, initially envelopes 1 and 2. Envelope 1 is opened. Should you switch

to envelope 2? The standard reasoning is that you should, because the finite value you have, having opened envelope 1, is less than the infinite expected value of envelope 2. In this paper we argue that it is not in one's interest to switch envelopes, even after opening one envelope. Put more cautiously, we argue that there is no advantage in switching envelopes over repeated trials. Imagine that finitely many envelopes (e.g., envelopes numbered 1, 2, 3, ..., $2,097,152 = 2^{21}$) are filled according to the St Petersburg procedure described above. An agent then selects a strategy. The agent will open envelope 1, and then have the option of switching to envelope 2; next the agent opens 3, and has the option of switching to 4, etc. In general, *trial n* consists of envelopes $2n - 1$ and $2n$; $2n - 1$ is opened, and the agent then has the option of switching to $2n$. One possible strategy is to *always-switch*; another possible strategy is to *always-stay*. The former strategy always opts to trade the finitely valued, opened, odd-numbered envelope for the even envelope; the latter strategy always keeps the odd envelope. Let us imagine that $1,048,576 = 2^{20}$ trials are run (this number is half of the number of envelopes filled, $2,097,152$, which we are using as an example). The always-stay strategy results in the agent's having all of the odd envelopes ($1, 3, 5, \ldots, 2,097,151$); the always-switch strategy results in the agent's having all of the even envelopes ($2, 4, 6, \ldots, 2,097,152$). But there can be no reason why the evens should be preferable to the odds, that is, switching is not, in general, beneficial. Each strategy outperforms the other strategy almost half of the time over the $1,048,576$ trials, with some small percentage resulting in ties. Ending up with all of the even envelopes (the result of always-switch), in general, cannot be preferable to ending up with all of the odds (the result of always-stay). The strategy of opening an envelope, finding that it contains some finite amount of money, and yet refusing to switch to an envelope with an infinite expected value, performs just as well as the strategy of switching to an infinite expected value over repeated trials. Computer simulations reinforce this conclusion. We found that the always-switch strategy outperformed the always-stay strategy about half of the time, and vice versa.

We now show that it is possible to always opt for a better expected value, and yet do worse over any finite number of trials. We believe that this result is interesting and surprising. Let us compare two strategies. One strategy is, as above, always-stay. The other strategy we call the *always-switch-limited-gain* strategy. In this strategy the agent always switches to the even envelope with the following condition: if the even envelope contains more than $\$2^{\text{ODD_ENV}+1}$, then \$0 is given, where ODD_ENV is the odd envelope's value. For example, if the odd envelope contains \$8, then if the even envelope contains more than $2^{8+1} = 2^9 = \$512$, then \$0 is given. Note two facts. 1) This always-switch-limited-gain strategy is worse than the always-switch strategy, in that sometimes \$0 is given, instead of the even envelope's value. 2) Switching is recommended based on expected value, in that the expected value of switching is ODD_ENV + 1, as compared with a certain value of ODD_ENV. That is, the expected gain from switching is \$1. From 1), it follows that the always-switch-limited-gain strategy is worse than the always-stay strategy, as

> always-stay = always-switch > always-switch-limited-gain, and so always-stay > always-switch-limited-gain

That is, in the first part of this paper it was shown that always-stay is the same

as the always-switch strategy ('always-stay = always-switch'). Then, as always-switch-limited-gain sometimes returns \$0 instead of the amount in the even envelope, always-switch is better than always-switch-limited-gain ('always-switch > always-switch-limited-gain'). *Therefore always-stay is better than always-switch-limited-gain* ('always-stay > always-switch-limited-gain'). Put in reverse, always-switch-limited-gain is worse than always-stay, even though always-switch-limited-gain *is recommended* over always-stay based on a consideration of expected value. Recall that always-switch-limited-gain has an expected gain of \$1 in each trial relative to always-stay. By opting for an expected gain of \$1 in each trial, an agent most likely fares worse, over any number of trials. Computer simulations supported the conclusion that always-stay outperforms always-switch-limited-gain.

We can also argue to this conclusion using two agents and repeated trials. Let the first agent begin with the odd envelopes, and use the always-switch-limited-gain strategy to switch to the even envelopes (or \$0 if the condition is met). Let the second agent begin with the even envelopes, and use the always-switch-limited-gain strategy to switch to the odd envelopes (or \$0 if the condition is met). Though both agents have an expected gain of \$1 on every trial (as compared with always-stay), it is clear that both agents fare worse by employing the always-switch-limited-gain strategy over repeated trials, because sometimes \$0 is given instead of one envelope's value.

We have not been able to arrive at general conditions under which maximizing expected value fails. One factor playing a role in the examples above is that the sequence of gains and losses over repeated trials is not independent and identically distributed (iid), and so neither the weak nor the strong law of large numbers holds. The examples above do show that maximizing expected value is not always all that it is cracked up to be. Repeatedly opting for an infinite expected value (always-switch) fares no better than staying with a finite value (always-stay). Also, repeatedly opting for a larger, finite expected value (always-switch-limited-gain) fares worse than staying with a finite value (always-stay). Given that it is not, in general, advantageous to switch to a higher expected value, two envelope paradoxes lose some of their force.

JEREMY GWIAZDA
Graduate Center of the City University of New York

An argument for not equating confirmation and explanatory power

Quantitative explicata of confirmation and of explanatory power are often defined as functions of probability values involving an event e and a hypothesis h. These two kinds of formal constructs also tend to display common features and salient analogies (see Schupbach J. and Sprenger J., 2011, "The logic of explanatory power", *Philosophy of Science*, 78, pp. 105–27, and Crupi V. and Tentori K., "A second look at the logic of explanatory power", *Philosophy of Science*, forthcoming, for some relevant remarks and further references). In view of this, investigating the connection between a probabilistic measure of confirmation $C(h, e)$ and of explanatory power $E(e, h)$ appears appropriate, if not pressing, as a source of theoretical clarification. An instructive possibility to explore is the statement of outright identity between the two notions. As a

concrete illustration, consider that I.J. Good's formal representation of the explanatory power of candidate explanans h with regards to explanandum e (in "Weight of evidence, corroboration, explanatory power, information and the utility of experiments", *Journal of the Royal Statistical Society*, Series B, 1960, 22, pp. 319–31) is identical to the "one true measure" of the degree of confirmation from e to h as once advocated by Peter Milne ("Log $[P(h/eb)/P(h/b)]$ is the one true measure of confirmation", *Philosophy of Science*, 1996, 63, pp. 21–6)—of course with the caveat that the hypothesis h at issue be in some explanatory relation with evidence e at all. Always keeping this latter proviso in mind, the general form of a "reductionist" claim to identity would be as follows:

Reduction (R): For any e, h and any P, $C(h, e) = E(e, h)$.

(I'll be assuming throughout that statements are contingent and the probability function P is regular.) Statement R may seem overly strong to begin with, but notice that it neatly conveys so-called *inference to the best explanation* (IBE). After all, for advocates of the IBE view, "observations support the hypothesis *precisely because* it would explain them" (Lipton P., "Inference to the best explanation", in W.H. Netwon-Smith, ed., *A Companion to the Philosophy of Science*, Blackwell, 2000, p. 185, emphasis added). And R is of concern even beyond that, if only because it would arguably trivialize the division of labour between two otherwise distinct branches of formal epistemology and philosophy of science. For short, R is not to be dismissed too quickly, i.e., unless a relevant argument is provided to undermine it. Such an argument is put forward in what follows.

A compelling principle of a model of explanatory power seems to be that the better hypothesis h succeeds in explaining the occurrence of a state of affairs e, the worse it fails in explaining the occurrence of its complementary, $\neg e$. Formally, such an inverse ordinal correlation between explanatory success and explanatory failure with regards to a pair of complementary statements e and $\neg e$ is spelled out as follows:

Symmetry (S): For any e, h_1, h_2 and any P, $E(e, h_1) > / = / < E(e, h_2)$ iff $E(\neg e, h_1) < / = / > E(\neg e, h_2)$.

On the other hand, consider the following condition concerning confirmation:

Final probability incrementality (F): For any h, e_1, e_2 and any P, $C(h, e_1) > / = / < C(h, e_2)$ iff $P(h|e_1) > / = / < P(h|e_2)$.

Condition F states that, for any given hypothesis h, confirmation is an increasing function of the posterior probability conditional on the evidence at issue—a virtually unchallenged assumption in contemporary probabilistic analyses of confirmation.

Notably, the following can be proven (see below):

Theorem: $\{S, F\}$ is consistent, but $\{R, S, F\}$ is not.

Relying on both S and F as sound, the theorem above discredits the reductionist claim to identity R. Apparently, probabilistic confirmation and explanatory power cannot be identified, for the two notions are constrained by genuinely distinct principles on

a quite basic level. They are irreducible, or "independent", much in the sense of Peano's disciple Padoa (on so-called Padoa's method in axiomatics, see Suppes P., *Introduction to Logic*, Van Nostrand, 1957, pp. 169 ff.). For all its tempting simplicity, thus, the reductionist thesis R turns out to be a naïve view of the connection between confirmation and explanatory power. This is not to say, of course, that there cannot be other meaningful and systematic relationships. This does mean, however, that one natural candidate formal rendition of IBE is flawed.

Proof of the Theorem.

For the first clause of the theorem, posit $C(h,e) = P(h|e) - P(h)$ and $E(e,h) = P(e|h) - P(e)$. This demonstrably makes S and F jointly true, so S, F is consistent. [Note: This is only one choice of measures yielding the relevant result, namely, that S and F be jointly satisfied. Another nice example is as follows: $C(h,e) = P(h|e)/P(h|\neg e)$ and $E(e,h) = P(e|h)/P(e|\neg h)$.]

For the second clause of the theorem, let P be such that $P(e|h) \neq P(e)$ and x is probabilistically independent from e, h, and their conjunction. We thus have $P(h \wedge e)/P(e) = P(h \wedge e)P(x)/P(e)P(x) = P(h \wedge e \wedge x)/P(e \wedge x)$, so that:

$$(*)\ P(h|e) = P(h|e \wedge x)$$

We also have $[P(e|h) - P(e)]P(x) \neq P(e|h) - P(e)$, which holds iff $P(e|h)P(x) - P(e)P(x) \neq P(e|h) - P(e)$ iff $-P(e|h) - P(e)P(x) \neq -P(e) - P(e|h)P(x)$ iff $1 - P(e|h) - P(e)P(x) + P(e)P(x)P(e|h) \neq 1 - P(e) - P(e|h)P(x) + P(e)P(x)P(e|h)$ iff $[1 - P(e|h)][1 - P(e)P(x)] \neq [1 - P(e)][1 - P(e|h)P(x)]$ iff $[1 - P(e|h)]/[1 - P(e)] \neq [1 - P(e|h)P(x)]/[1 - P(e)P(x)]$ iff $[1 - P(e|h)]/[1 - P(e)] \neq [1 - P(e \wedge x|h)]/[1 - P(e \wedge x)]$ iff $P(\neg e|h)/P(\neg e) \neq P(\neg(e \wedge x)|h)/P(\neg(e \wedge x))$ iff $P(\neg e|h)P(h)/P(\neg e) \neq P(\neg(e \wedge x)|h)P(h)/P(\neg(e \wedge x))$, that is (by Bayes's theorem):

$$(**)\ P(h|\neg e) \neq P(h|\neg(e \wedge x))$$

Now, by F and (*), $C(h,e) = C(h, e \wedge x)$. Thus, by R, $E(e,h) = E(e \wedge x, h)$. By S, the latter implies $E(\neg e, h) = E(\neg(e \wedge x), h)$, so that, again by R, $C(h, \neg e) = C(h, \neg(e \wedge x))$. Yet F and (**) imply the opposite, i.e., $C(h, \neg e) \neq C(h, \neg(e \wedge x))$. So R, S, F is inconsistent. *QED*.

Vincenzo Crupi
Philosophy, Turin
Munich Center for Mathematical Philosophy,
Ludwig Maximilian University

News

New book: *The Limits of the Self: Immunology and Biological Identity*

What counts as an individual in the living world? What does it mean for a living thing to remain the same through time while constantly changing? Immunology, one of the most

dynamic fields of today's biology, considers these questions its province, and answers them through its crucial concepts of "self" and "nonself."

Though immunology has been dominated since the 1940s by the self-nonself theory, this book argues that this theory is inadequate, because immune responses to self constituents and immune tolerance of foreign entities are the rule, not the exception. An alternative theory, the continuity theory, is advanced instead. This theory offers a new way to answer the question of what triggers an immune response. It also echoes the recent realization that all organisms, and not higher vertebrates only, have an immune system.

This book's main thesis is that the self-nonself theory should be abandoned, but that immunology still proves to be decisive for delineating the boundaries of the organism. Articulating an evolutionary and an immunological perspective, it offers an original conception of the organism. Tolerance of the fetus by the mother and of countless bacteria on the body's surfaces proves that every organism is heterogeneous, that is, made of entities of different origins. In other words, every organism appears as a *chimera*, a mixed living thing the cohesiveness of which is ensured by the constant action of its immune system. *The Limits of the Self* thus offers a contribution to the definition of biological individuality and to the understanding of the immune system.

Thomas Pradeau
Philosophy, Paris-Sorbonne

Ceteris paribus Laws and Reasoning, 4–5 November

On November 4–5, Lund Department of Philosophy hosted a workshop on "*Ceteris paribus* Laws and Reasoning". The main aim of this meeting was to bring together researchers from Philosophy of Science, Logic and Philosophy of Language, mainly focusing on the analysis of the nature and the logical structure of *ceteris paribus* laws and clauses. Ceteris paribus assertions are non-strict generalisations: they are laws and general statements allowing for exceptions, or which only hold given certain specific conditions. Their role is central in many sciences (e.g., Economics, Medicine and Psychology) as well as in everyday reasoning, but their nature and their conditions of validity is far from clear and often non univocal.

The modal logic of ceteris paribus necessity and possibility was the object of the countribution of Patrick Girard (Auckland) "Being flexible in ceteris paribus reasoning". This paper presents an expansion and generalisation of the logical framework already available in the previous "Everything else being equal: a modal logic for ceteris paribus preference" (with J. Van Benthem and O. Roy).

An extensive logico-linguistic analysis of the different senses of ceteris paribus was presented by Sheldon Smith (UCLA) in his talk "What is a ceteris paribus law?".

Annika Wallin's (Lund) "External validity and ceteris paribus laws" opened the scope to cognitive sciences and investigated the role of ceteris paribus assumptions and background conditions in experimental settings in the area of cognitive psychology.

Frank Zenker (Lund) argued for ceteris paribus laws being best understood as reflecting some epistemic deficit of scientists, rather than constituting a distinct type of

law or generalization. On this view, ceteris paribus laws may be viewed as transitory formulations *en route* to better theories and a role of logic would be to address the conditions under which a ceteris paribus law should be given up in favor of another.

Davide Grossi (Liverpool) proposed a view of deontic logic blending the "semantic" view with the "syntactic" view of norms. Its technical underpinnings lie in recent contributions to the modal logic of preference. He discussed deontic logic as resulting from both an ideality ordering on states and a syntactic priority ordering on propositions. He mainly addressed the issue of norms as resulting from ascriptions of "constituted" properties (e.g., good, bad, right, etc.) emphasizing technical similarities with established logical representation of ceteris paribus statements.

A structural analysis of enthymematic reasoning and rules for enthymematic validity was offered by Hans Rott (Regensburg).

In "Natural Laws and Physical Necessity", Lars-Göran Johansson (Uppsala) argues that natural laws are implicit definitions rather than universal axiomatic statements and that their necessity is therefore a kind of conceptual necessity: necessity relative to the concepts and predicates employed. The process of theory formation and concept formation goes hand in hand.

Gregory Wheeler's (Lisbon) talk "On the possibility of measuring coherence" discussed the scope of Erik Olsson's impossibility result undermining the link between probabilistic measures of coherence and higher likelihood of truth. He argued that Olsson's specific witness model, like bayesian witness models in general, does not give charitable ceteris paribus conditions for every possible theory of probabilistic coherence.

In his talk on "Logic and ceteris paribus reasoning" Carlo Proietti (Lund) offered an overview of some possible applications of ceteris paribus modal logic to problems in temporal logic, game theory and metalogic.

Some Presentations remain available at the workshop website.

Carlo Proietti
University of Lund

Causality and Mechanisms in Philosophy of Science, 24–25 November

A small, two-day workshop was held at the University of Bern. Its key topics were the conceptual foundations, the empirical confirmation and the explanatory role of causal and mechanistic claims in natural science.

The programme allotted equal time to presentations and discussions: 45 minutes each. This restricted the number of talks to two per morning and afternoon, but it allowed for deep and thoughtful discussions of each topic. The participants felt unanimously that the format was congenial and perfectly suited for small informal workshops.

The workshop started on Thursday morning with a session on conceptual foundations. Jens Harbecke (Witten) gave a talk in which he proposed a criterion (based on the notions of minimal sufficiency and minimal necessity) by which we can distinguish between levels of organisation in mechanistic explanations. Lorenzo Casini (Kent) offered

a comparison between the mechanistic and the inferentialist accounts of causality.

The Thursday afternoon session was about the empirical confirmation of causal and mechanistic hypotheses. Marcel Weber (Geneva) presented his paper on Meselson's and Stahl's "most beautiful experiment in biology", which he reconstructed as an inference to the best mechanistic explanation. Raphael Scholl (Bern) gave a talk about Semmelweis's discovery of the cause of childbed fever, arguing for a methodological reconstruction in terms of Mill's methods and mechanisms.

Explanation and prediction were the focus of the presentations on Friday morning. Jon Williamson (Kent) asked "How can causal explanations explain?", contrasting in particular difference-making causal claims and mechanistically grounded causal claims. Alex Broadbent (Johannesburg) then developed an account of prediction in epidemiology based on the idea that a prediction is robust if we have good grounds for believing it "could not easily be wrong".

The final session on Friday afternoon returned to the topics of confirmation and conceptual foundations. Michael Baumgartner (Konstanz) presented a simple example in which Mill's method of difference fails in a deterministic system. Tim Räz (Lausanne) gave an account of the relationship between mathematics and causality which tries to account for both the metaphysical dispensability and the theoretical indispensability of mathematics in empirical science.

The workshop was generously funded by the University of Bern Intermediate Staff Association (Mittelbauvereinigung der Universität Bern).

Raphael Scholl
Philosophy, Bern

Probability and Statistics in Physics, 16 December

The one-day international workshop "Probability and Statistics in Physics" took place at Complutense University of Madrid on the 16th of December 2011. The workshop was organised by the *Methods of Causal Inference and Scientific Representation* (MCISR) research group. It brought together researchers working on different issues related to the role, significance and methodological application of probabilities and statistics in physical theories, most notably in statistical mechanics or quantum physics.

The workshop had four sessions, each of which included a presentation of some recent work and extra time allocated for discussion. Participants included Janneke van Lith (Utrecht), Orly Shenker (Jerusalem), Mauricio Suárez (Madrid) and Meir Hemmo (Haifa).

The workshop opened with Janneke van Lith presenting her own approach to the Second Law of thermodynamics. In "How to Derive the Second Law from Internal Dynamics alone" van Lith offered an alternative solution to the well known problem of how to account for the Second Law of thermodynamics in Gibbsian statistical mechanics. (The fact that the Gibbsian fine-grained entropy function remains constant under Hamiltonian dynamical evolution is at the origin of the problem.) The usual solutions—offered by coarse-graining or interventionist approaches—reject, in some way or another, either the fact that the system evolves according to Hamiltonian dynamics, or

the very idea of fine-grained entropy. In contrast, van Lith suggested that the Gibbsian fine-grained approach can still make sense as long as equilibrium is not taken to require stationary probability distributions. Hamiltonian dynamics then applies all the same to derive the Second Law of thermodynamics.

In the second session the talk "Objective Probability in Classical Physics" by Orly Shenker discussed certain thermodynamical phenomena displaying objective probabilities. The expansion of a gas in a container is one such phenomena. The fact that gases are likely to expand in containers constitutes an objective probabilistic regularity, which does not depend on the amount information about the gas and its initial conditions we may have (or lack). This is not however counterintuitive and can be reconciled with the standard view that probabilities are associated to a certain ignorance about initial conditions. This is done by referring to the Second Law of thermodynamics, which underlies such objective regularities. Thus, insofar as statistical mechanics explains the Second Law of thermodynamics in terms of probabilities associated to the lack of information, it will also constitute an explanation of such objective probabilities along similar lines.

Mauricio Suárez gave in "Propensities and Pragmatism" an account, from a pragmatist perspective, of the relation between probabilities and propensities, and how should the latter be understood. A central assumption in most propensity theories of probability is what Suárez calls the "identity thesis", i.e., that propensities are indeed probabilities. In his opinion, however, this identification must not be a requirement on propensities—i.e., the notion of propensity does not need to presuppose it at all. As a result, the "identity thesis" must be rejected. Suárez vindicates as an alternative a pragmatist account of propensities, based on Peirce's views on probabilistic dispositions. By rejecting the "identity thesis" moreover, Suárez's proposal is able to avoid common objections against propensity accounts of probabilities. In particular, arguments based on Humphreys' Paradox will not stand any more, for they require the "identity thesis" to go through.

In the final session, Meir Hemmo presented the talk "Probability in Bohmian Mechanics", where he discussed the significance of probabilities in Bohmian mechanics. Despite being a deterministic theory, Bohmian mechanics is able to recover the probabilistic predictions of standard quantum mechanics by assuming that there is a certain degree of ignorance about the actual positions of the particles in a quantum system. This ignorance is then expressed in terms of a probability distribution, set over the positions of the particles. In this view, the main difference between probabilities in Bohmian mechanics and standard quantum mechanics is just the difference between epistemic and objective probabilities. In Hemmo's view however, there is a further (fundamental) difference between both these theories. Namely, in contrast to standard quantum mechanics, in Bohmian mechanics there are certain "zero probability" events—i.e., events that will never take place—which must not be considered to be impossible. In other words, these theories entail different event structures, although their probability assignments coincide. Hemmo's conclusion is then that Bohmian mechanics and standard quantum mechanics are different theories, already at the level of events.

In sum, despite the limited size of the event and small number of participants, we can say the "Probability and Statistics in Physics" international workshop was a complete success, both for the high quality of the works presented, as well as from the point of

view of discussion and scientific interaction that it promoted.

Iñaki San Pedro
Fernanda Samaniego
Department of Logic and Philosophy of Science
Universidad Complutense de Madrid

Perspectivalism Workshop, 19–20 January

The topic of perspective and perspectivalism recently gained a lot of attention in analytic philosophy. Accordingly, the international workshop "Perspectivalism", held at the University of Ghent, invited a range of renowned philosophers from Europe and America to a very pulsating debate.

After the introductory speech, "Carving The World As We Please", given by Jan Willem Wieland, there was an exciting discussion that focused on the different contributions to the workshop. Anna-Sofia Maurin presented her criticism of Boghossian's argument against global cognitive relativism. James McAllister argued for a "Perspectivalism About Structures in the World". He defended the view that we should consider different structures to be real, where 'real' means 'being in the world'.

Some contributions focused on Giere's *Scientific Perspectivism* from 2006. Michela Massimi analysed this approach. Ioannis Votsis argued for a "perspectival realism". Particularly the question whether perspectivalism implies or is merely compatible with a sort of antirealism was a major point of interest in the debate. One important talk targeted the relationship between perspectivalism and realism. While the standard view would connect perspectivism and perspectivalism to antirealism, Anjan Chakravartty presented a version of a realist perspectivism. He focused on synchronic perspectives of taxonomic systems in science rather than on diacronic perspectivism of historical paradigms. In particular, he argued that "Taxonomic Monism", the thesis that a unique structure of entities and processes exists in the world, is undermined by modern sciences. According to him, the species concepts in biology represent a clear counterexample to Taxonomic Monism. In response to this challenge, reductionism has been formulated, which tries to reduce all science to chemistry or physics. Chakravartty claims that this strategy fails to save Taxonomic Monism. Instead, he argues for "Taxonomic Pluralism", which is the thesis that there is more than one structure of entities in the world. His main arguments for that point are the argument from detection and the argument from inconsistent models, as he calls them. His idea is that non-perspectival facts are those facts that are dispositions of a theory's target system to behave differently in different contexts of observation and detection. Chakravartty therefore argues for a realism regarding that system's properties, which are represented by different but inconsistent models. The properties within this system share "sociability", which supposedly allows for realism despite inconsistency. Gurpreet Rattan showed how disagreement can be understood in terms of perspectival metarepresentation. Metarepresentation is the representation of representation. These latter representations can be arguments. Furthermore, Jesse Mulder, Filip Kolen, Eric Weber and Leen De Vreese contributed to the debate.

In sum, the workshop offered an overview over the range of topics related to perspectivalism in the current debate. The focus clearly lay on the philosophy of science, but metaphysics and epistemological issues were also addressed. Throughout the workshop, it became evident that the perspectival take on philosophy is a promising emerging approach in different areas of discourse.

Jorn Wustenberg
University of Konstanz

Calls for Papers

Input & Output Analysis for Simulation: special issue of the *Journal of Simulation*, deadline 1 March.
Non-classical Modal and Predicate Logics: special issue of Logic Journal of IGPL, deadline 31 March.
Game Theoretic Models of Communication: special issue of *Erkenntnis*, deadline 31 March.
Disagreements: special issue of *Erkenntnis*, deadline 1 April.
Probability, Logic and Learning: special issue of *Theory and Practice of Logic Programming*, deadline 2 April.
Formal and Intentional Semantics: special issue of *The Monist*, deadline 30 April.
The Mind-Body Problem in Cognitive Neuroscience: special issue of *Philosophia Scientiæ*, deadline 1 May.
Inforgs and the Infosphere: Themes from Luciano Floridi's Philosophy of Artificial Intelligence: special issue of *The Journal of Experimental & Theoretical Artificial Intelligence*, deadline 1 July.
Mind and Paradox: special issue of *Journal of Experimental & Theoretical Artifical Intelligence*, deadline 1 July.
The Aim of Belief: special issue of *Teorema*, deadline 15 September.

What's Hot in ...

...Mathematical Philosophy

A topic of major and growing interest for the last decade in mathematical philosophy is formalized truth theories, studied either via an axiomatic approach or a model-theoretic one. On the axiomatic approach, the basic idea is to examine systems of axioms for a primitive truth predicate, usually combined with Peano arithmetic as the underlying theory of the syntactical truth bearers. Given that the simplest truth theory—just the disquotational T-sentences—is inconsistent in this setting (because of the liar paradox), the usual approach is to modify the axioms to achieve a consistent and intuitively plausible truth theory. There is now a range of fairly well-understood truth theories with distinctively modified axioms and rules (the classic in this field being Friedman & Sheard (1987, "An Axiomatic Approach to Self-Reference Truth")); fixed-point and revision-theoretic models; recent work on notions of "dependence" (Leitgeb 2005, "What Truth

Depends On"); and theories that relate truth to other notions, such as provability, validity and necessity. On the model-theoretic approach, one aims (usually assuming the natural number structure as the given base structure) to specify a truth set S for the language which contains the truth predicate. The two classic examples of this approach are Kripke's fixed-point models, and the revision models of Gupta and Belnap.

Throughout, a significant impetus has been the work of Volker Halbach, whose *Axiomatic Theories of Truth* (OUP, 2010) has become a classic, bringing together the major results from Tarski, through to Kripke, Feferman, Gupta, Friedman, Sheard, McGee, Cantini and more recent developments. Leon Horsten's new monograph *The Tarskian Truth: Deflationism and Axiomatic Truth* (MIT, 2011) is an important addition to this literature. The recurrent central philosophical theme of such work is deflationism about truth. In 2011, conferences in Amsterdam and Barcelona were devoted to this topic. The next few months of 2012 will see a number of conferences and workshops focusing on this topic. In partcular, a three-day conference will be held in Vienna, March 9–11, organized by the Vienna Forum for Analytic Philosophy, featuring talks by Paul Horwich, Leon Horsten and myself, and submitted talks. This is followed by a three-day conference at MCMP, Munich, March 14–16, featuring Volker Halbach, Andrea Cantini, Stewart Shapiro and many others.

Jeff Ketland
Munich Center for Mathematical Philosophy

. . . Uncertain Reasoning

Decision theorists have long used top-down classifications of decision problems. The canonical taxonomy, put forward by Luce and Raiffa (1957: *Games and Decisions*, Wiley and Sons), distinguishes between certainty, risk and uncertainty according to how the outcomes of a decision depend on the agent's choice. If the outcome is a fixed deterministic function of the choice, then the agent faces a problem under certainty. In this case "rational choice" is reduced to the mathematical theory of (linear) programming. Otherwise the outcome depends stochastically on the agent's choices. If the stochastic dependence is fixed (i.e the probability distribution on the state space is given in the description of the decision problem), the agent faces a decision under *risk*. If such a distribution is not given *a priori*, the decision problem is one under *uncertainty*. This distinction goes hand in hand with the "objective" and "subjective" interpretations of probability, respectively.

Whilst popular, the Luce and Raiffa taxonomy can give rise to inconsistencies, as the authors themselves notice (Luce and Raiffa 1957, p.300). Partly owing to this, there has been a recent revival of the distinction between "risk" and "uncertainty" put forward by Knight (1921: *Risk, Uncertainty, and Profit*, Hart, Schaffner & Marx; Houghton Mifflin Co.). According to Knight, a decision problem involves *risk* if the state of ignorance of the agent facing it can be quantified probabilistically, and involves *uncertainty* otherwise. In an attempt to avoid confusion, the expression *Knightian uncertainty* is often used to designate this latter kind of "non probabilistically quantifiable" uncertainty.

This top-down approach appears to have little or no practical relevance, as two rather

heterogeneous experimental studies, namely the BBC Big Risk Test and the RiskLiteracy.org project seem to indicate.

The Big Risk Test rather ambitously "aims to be the biggest study of risk ever undertaken". It has been designed by an interdisciplinary team of Cambridge academics led by David Spiegelhalter (Winton Professor of the Public Understanding of Risk) and Mike Aitken (Experimental psychologist at Cambridge Behavioural and Clinical Neuroscience Institute).

The experiment consists of six tasks.

1. Rate the risk. The "game" requires the subject to order—strictly—rather dishomogeneous items. In the first round, they include gardening, alcohol and air-travel, whilst in the second vaccines, terrorism and climate changes. No further information is provided, thus making some choices rather arbitrary. (What does it mean, for instance, to say that alcohol is risky? It clearly depends on how one (ab)uses it!)

2. How risky is it?. Similar task, but now the labels to which subjects must attach a degree of riskiness are described quantitatively. Example: *Of a million people going hang-gliding once, how many would you expect to die?* A reference point for all the eight questions is provided, e.g., "16 out of a million people riding 100 miles by motorbike are expected to die". It is implied, but not stated, that the deaths are causally related to the activities.

3. Doctor's advice. The subject is presented with a Doctor's hypothetical description of the risk, the degree to which the Doctor's advice was clear, and the extent to which we would be happy to follow it.

4. The risk spinner. The subject is presented ten times with a choice between two wheels holding identical expectations. The task aims at assessing the subject's attitude towards risk.

5. The numbers quiz. Three levels: easy, medium, difficult. The difficult include applying Bayes theorem.

6. Risk personality. Subjects must classify themselves in relation to their attitute towards risk.

A rather detailed feedback is provided at the end of the Big Risk Test, but it doesn't appear to be very reliable. For example I adopted the rather natural rule "always choose the first wheel" in the "Risk spinner" test, and my choice has been misinterpreted as risk aversion!

The Risk Literacy project takes a much more direct approach than the BBC-Cambridge test. As the website points out it was "[i]nitially developed and validated at the Max Planck Institute for Human Development, the Berlin Numeracy Test provides a fast and psychometrically sound instrument for the assessment of statistical numeracy and risk literacy. The Berlin Numeracy Test was created to help increase public awareness and to improve research conducted with commonly used samples from diverse

cultures and backgrounds". The Berlin Numeracy Test is in fact centered on testing the subjects' mathematical, and in fact statistical, abilities. The test follows a sequential structure.

Q1 Out of 1,000 people in a small town 500 are members of a choir. Out of these 500 members in a choir 100 are men. Out of the 500 inhabitants that are not in a choir 300 are men. *What is the probability that a randomly drawn man is a member of the choir?*

Question 2 is chosen according to the subject's performance. If the answer to Q1 was correct, the question is

Q2a Imagine we are throwing a loaded die (6 sides). The probability that the die shows a 6 is twice as high as the probability of each of the other numbers. Now imagine you would throw this die 70 times. *On average, out of these 70 throws how many times would the die show the number 6?*

Otherwise the question is

Q2b Imagine we are throwing a five-sided die 50 times. *On average, out of these 50 throws how many times would this five-sided die show an odd number (1, 3 or 5)?*

Unlike the BBC Test, the online version of the Berlin Numeracy Test provides, as yet, no feedback. However, the rationale behind the experiment and the statistical analysis of previous paper-based tests are detailed by Cokely, E.T., Galesic, M., Schulz, E., Ghazal, S., & Garcia-Retamero, R. (2012: Measuring risk literacy: The Berlin Numeracy Test, *Judgment and Decision Making*, 7, 25–47).

Hykel Hosni
Scuola Normale Superiore, Pisa
LSE Choice Group, London

Events

March

Laws and Chances Workshop: Cologne, 5 March.
FoIKS: 7th International Symposium on Foundations of Information and Knowledge Systems, Kiel, Germany, 5–9 March.
LATA: 6th International Conference on Language and Automata Theory and Applications, La Coruña, Spain, 5–9 March.
Dispositions, Causes, Modality Workshop: Cologne, 7–9 March.
Graduate Conference in Philosophy of Science: Erasmus University Rotterdam, 8–9 March.
Nothing but the Truth: Vienna Forum for Analytic Philosophy, University of Vienna, 9–11 March.

ICMLC: 4th International Conference on Machine Learning and Computing, Hong Kong, 10–12 March.
LPAR: 18th International Conference on Logic for Programming, Artificial Intelligence and Reasoning, Merida, Venezuela, 11–15 March.
Axiomatic vs Semantic Truth: Munich, 14–16 March.
Knowledge and Disagreement: Universidade Nova de Lisboa, Portugal, 15–16 March.
&HPS4: Integrated History and Philosophy of Science, Department of Philosophy and History of Science, University of Athens, 15–18 March.
The Scope and Limits of Mathematical Knowledge: University of Bristol, 17–18 March.
Reflection Principles in Set Theory: Oxford, 19–20 March.
Empirical Philosophy of Science. Qualitative Methods: Sandbjerg, Denmark, 21–23 March.
Workshop on Philosophical and Formal Theories of Truth: Amsterdam, 23–25 March.
Pragmatism, Law, and Language: University of Idaho, 23–25 March.
LABCII: Logical Approaches to Barriers in Complexity II, Newton Institute, Cambridge, UK, 26–30 March.
CIFEr: Computational Intelligence for Financial Engineering and Economics, New York, 29–30 March.
Philosophy of Risk: Center for Philosophy of Science, University of Pittsburgh, 30–31 March.
PiPoS: Pragmatism in Philosophy of Science, University of San Francisco, 30–31 March.
DICE: 3rd Workshop on Developments in Implicit Complexity, Tallinn, Estonia, 31 March–1 April.

April

YSM: Young Statisticians' Meeting, Cambridge, 2–3 April.
DARC: Dynamics Of Argumentation, Rules, and Conditionals workshop, Luxembourg, 2–3 April.
BCTCS: British Colloquium for Theoretical Computer Science, Manchester, UK, 2–5 April.
SBP: International Conference on Social Computing, Behavioral-Cultural Modeling, & Prediction, University of Maryland, 3–5 April.
Mind, Method and Morality: Pittsburgh, 6–7 April.
CNCS: International Conference on Computer Networks and Communication Systems, Malaysia, 7–8 April.
EMCSR: European Meetings on Cybernetics and Systems Research, Vienna, 10–13 April.
Time for Causality: Workshop on Causal Inference and Dynamic Decisions in Longitudinal Studies, Bristol, 10–13 April.
evoSTOC: Evolutionary Algorithms in Stochastic and Dynamic Environments, Malaga, Spain, 11-13 April.
PhDs in Logic IV: Ghent, 12–13 April.

Objects, Kinds and Mechanisms in Biology: One Day Workshop, University of Leeds, 13 April.
Northwestern/Notre Dame Graduate Epistemology Conference: Northwestern University, Evanston, IL, 13–14 April.
philoSTEM: 3rd Midwest Workshop in Philosophy of Science, Technology, Engineering, and Mathematics, Indiana, USA, 13–14 April.
BMC2012: Workshop on Turing's Legacy in Mathematics and Computer Science, University of Kent, 16–19 April.
Confronting Intractability in Statistical Inference: University of Bristol, 16–19 April.
Collective Intelligence: MIT, Cambridge, MA, 18–20 April.
Being Free, Doing Free: Freedom Between Theoretical and Practical Philosophy, University of Freiburg, Germany, 19–21 April.
GIRL: 1st Conference on Games, Interactive Rationality and Learning, Lund, 19–21 April.
Psychology, Emotion, and the Human Sciences: University of Windsor, Windsor, Ontario Canada, 20–21 April.
MAICS: 23rd Midwest Artificial Intelligence and Cognitive Science Conference, Ohio, 21–22 April.
AISTATS: 15th International Conference on Artificial Intelligence and Statistics, La Palma, Canary Islands, 21–23 April.
Agents and Causes: Interdisciplinary Aspects in Mind, Language and Culture, Bielefeld, 21–23 April.
ENTIDENTIC: Entity and Identity in Bioethics, Paris, France, 23–24 April.
The Progress of Science: Tilburg Center for Logic and Philosophy of Science, 25–27 April.
SDM: 12th SIAM International Conference on Data Mining, Anaheim, California, USA, 26–28 April.

May

SOPHA: Société de philosophie analytique, Paris, 4–6 May.
ICDDM: International Conference on Database and Data Mining, Chengdu, China, 5–6 May.
ICFCA: 10th International Conference on Formal Concept Analysis, Leuven, Belgium, 6–10 May.
Belief Functions: Compiégne, France, 9–11 May.
Naturalism and Normativity in the Social Sciences: University of Hradec Králové, Czech Republic, 10–12 May.
Philosophy and Computation: Lund University, Sweden, 12–13 May.
CASI: 32nd Conference on Applied Statistics Ireland, 16–18 May.
Games, Game Theory and Game Semantics: 8th International Symposium of Cognition, Logic and Communication, Riga, Latvia, 18–20 May.
LMP: 12th Annual Philosophy of Logic, Mathematics, and Physics Conference, University of Western Ontario, 20–21 May.

SLACRR: St. Louis Annual Conference on Reasons and Rationality, 20–22 May.
IPDPS: 26th IEEE International Parallel and Distributed Processing Symposium, Shanghai, China, 21–25 May.
JdS: 44th Journées de Statistique, Brussels, 21–25 May.
PhML: Philosophy, Mathematics, Linguistics: Aspects of Interaction, St. Petersburg, Russia, 22–25 May.
UR: Uncertain Reasoning, Special Track at FLAIRS-25, Marco Island, Florida, USA, 23–25 May.
SSHAP: Mind, Language and Cognition, McMaster University, Canada, 24–26 May.
PhilMiLCog: 10th Annual Graduate Conference on the Philosophy of Mind, Language and Cognitive Science, University of Western Ontario, 24–26 May.
The Aims of Inquiry and Cognition: Edinburgh Epistemology Research Group, University of Edinburgh, 25–26 May.
CSAE: IEEE International Conference on Computer Science and Automation Engineering, Zhangjiajie, China, 25–27 May.
ICKD: 2012 International Conference on Knowledge Discovery, Indonesia, 26–27 May.
AI2012: Canadian Conference on Artificial Intelligence, 28–30 May.
RTA: 23rd International Conference on Rewriting Techniques and Applications, Japan, 28 May–2 June.
FEW: 9th Annual Formal Epistemology Workshop, Munich, 29 May–1 June.
ICCC12: Third International Conference on Computational Creativity, Dublin, 30 May–1 June.
StochMod: 4th meeting of the EURO Working Group on Stochastic Modeling, Ecole Centrale Paris, 30 May–1 June.
Human Complexity: The University of North Carolina, Charlotte, 30 May–1 June.
Cambridge Pragmatism: a Research Workshop, Cambridge, UK, 31 May–1 June.
Rudolf Carnap Lectures: Ruhr-Universität Bochum, 31 May–2 June.

June

Incommensurability 50: Taipei, Taiwan, 1–3 June.
ICFIE: International Conference on Fuzzy Information and Engineering, Hong Kong, 2 June.
Trends in Logic XI: Advances in Philosophical Logic, Ruhr University Bochum, 3–5 June.
LAMAS: 5th Workshop on Logical Aspects of Multi-Agent Systems, Valencia, 4–5 June.
WCSB: 9th International Workshop on Computational Systems Biology, Ulm, Germany, 4–6 June.
FEW: Formal Epistemology Week, Konstanz, 4–6 June.
AAMAS: 11th International Conference on Autonomous Agents and Multiagent Systems, Valencia, Spain, 4–8 June.
CILC: 9th Italian Convention on Computational Logic, Sapienza University of Rome, 6–7 June.

MFPS: 28th Conference on the Mathematical Foundations of Programming Semantics, University of Bath, 6–9 June.
Minds, Bodies, and Problems: Bilkent University, Ankara, 7–8 June.
Edinburgh Epistemology Graduate Conference: University of Edinburgh, 8–9 June.
Foundations of Logical Consequence: University St Andrews, 8–10 June.
NMR: 14th International Workshop on Non-Monotonic Reasoning, Rome, Italy, 8–10 June.
RATS: Recent Advances in Time Series Analysis Workshop, Cyprus, 9–12 June.
NORDSTAT: 24th Nordic Conference in Mathematical Statistics, Northern Sweden, 10–14 June.
MS5: Conference on Models and Simulations, Helsinki, 14–16 June.
CSam: Classification Society Annual Meeting, Carnegie Mellon University, Pittsburgh, PA, 14–16 June.
Basic Knowledge: Conference on the A Priori, Aberdeen, 16–17 June.
SAT: International Conference on Theory and Applications of Satisfiability Testing, Trento, Italy, 17–20 June.
LOFT: 10th Conference on Logic and the Foundations of Game and Decision Theory, Sevilla, Spain, 18–20 June.
DM: Discrete Mathematics, Dalhousie University, Halifax, Nova Scotia, Canada, 18–21 June.
LOGICA: Hejnice, northern Bohemia, 18–22 June.
CiE: Computability in Europe, University of Cambridge, Cambridge, 18–23 June.
SISSM: Scientific Meeting of the Italian Statistical Society, Rome, Italy, 20–22 June.
Philosophical Insights: Senate House, University of London, 21–23 June.
MBR12: Model-Based Reasoning in Science and Technology, Sestri Levante, Italy, 21–23 June.
SPP: Annual Meeting of the Society for Philosophy and Psychology, University of Colorado at Boulder, 21–24 June.
HOPOS: Halifax, Nova Scotia, Canada, 21–24 June.
CCA: 9th International Conference on Computability and Complexity in Analysis, Cambridge, UK, 24–27 June.
COLT: 25th Annual Conference on Learning Theory, Edinburgh, 25–27 June.
MPC: 11th International Conference on Mathematics of Program Construction, Madrid, Spain, 25–27 June.
Artificial Intelligence and Soft Computing: Naples, Italy, 25–27 June.
VaNiM: Values and Norms in Modeling, Eindhoven, The Netherlands, 25–27 June.
LICS: 27th ACM/IEEE Symposium on Logic In Computer Science, Dubrovnik, Croatia, 25–28 June.
Square of Opposition: American University of Beirut, 26–29 June.
ICML: 29th International Conference on Machine Learning, University of Edinburgh, 26 June–1 July.
IJCAR: 6th International Joint Conference on Automated Reasoning, Manchester, UK, 26 June–1 July.
DGL12: Sixth Workshop in Decisions, Games & Logic, LMU Munich, 28–30 June.

EEN: European Epistemology Network Meeting, Universities of Bologna and Modena, Italy, 28–30 June.
Evolution and Function of Consciousness: Summer School in Cognitive Science 2012, Montreal, Canada, 30 June–9 July.

July

Uncertainty in Computer Models: Sheffield, UK, 2–4 July.
AISB/IACAP: Birmingham, UK, 2–6 July.
LASR: 31st Leeds Annual Statistical Research Workshop, University of Leeds, 3–5 July.
Bounded Rationality: Summer Institute on Bounded Rationality, Berlin, Germany, 3–10 July.
IIBM: 5th International Workshop on Intelligent Informatics in Biology and Medicine, Palermo, Italy, 4–6 July.
History and Philosophy of Programming: Ghent University, 5–6 July.
BSPS: Annual Conference of the British Society for the Philosophy of Science, University of Stirling, 5–6 July.
CAV: 24th International Conference on Computer Aided Verification, Berkeley, 7–13 July.
ISSCSS: International Summer School in Cognitive Sciences and Semantics, Latvia, 8–18 July.
ASC: 21st Australian Statistical Conference, Adelaide, 9–12 July.
IPMU: 14th International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems, Catania, Italy, 9–13 July.
ICALP: 39th International Colloquium on Automata, Languages and Programming, University of Warwick, 9–13 July.
Foundations of Mathematics: University of Cambridge, 10–12 July.
TViTC: Theoretical Virtues in Theory-Choice, University of Konstanz, 12–14 July.
DEON: 11th International Conference on Deontic Logic in Computer Science, University of Bergen, Norway, 16–18 July.
WorldComp: The 2012 World Congress in Computer Science, Computer Engineering, and Applied Computing, Nevada, USA, 16–19 July.
DMIN: 8th International Conference on Data Mining, Nevada, USA, 16–19 July.
SIPTAss: Society for Imprecise Probability: Theories and Applications Summer School, Pescara, Italy, 16–20 July.
ISA: IADIS International Conference Intelligent Systems and Agents, Lisbon, Portugal, 21–23 July.
Paradox and Logical Revision: LMU, Munich, 23–25 July.
WoMO: 6th International Workshop on Modular Ontologies, Graz, Austria, 24 July.
FOIS: 7th International Conference on Formal Ontologies in Information Systems, Graz, Austria, 24–27 July.

NAFIPS: 31th North American Fuzzy Information Processing Society Annual Conference, Berkeley, 6–8 August.
ESSLLI: 24th European Summer School in Logic, Language and Information, Poland, 6–17 August.
KDD: 18th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Beijing, China, 12–16 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.
UAI: Conference on Uncertainty in Artificial Intelligence, Catalina Island, USA, 15–17 August.
SLS: 8th Scandinavian Logic Symposium, Roskilde University, Denmark, 20–21 August.
ALFAn: Latin American Analytic Philosophy Conference, Buenos Aires, Argentina, 21–24 August.
AIML: Advances in Modal Logic, Copenhagen, 22–25 August.
FMIP: Munich / Groningen Summer School: Formal Methods in Philosophy, Groningen, 23–28 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, 26–29 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.
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.

Courses and Programmes

Courses

LI: Logic and Interactions, Winter School and Workshops, CIRM, Luminy, Marseille, France, 30 January–2 March.
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.

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.
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 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

An interdisciplinary programme at the
University of Kent, Canterbury, UK.
Core modules provided by Philosophy and further modules 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

Post-doc Position: in Probabilistic Reasoning, Vienna University of Technology, Austria, until filled.
Post-doc positions: in all areas of speech and language processing at the Human Language Technology Center of Excellence at Johns Hopkins University, until filled.
Post-doc position: on the project "Explanatory Reasoning: Normative and Empirical Considerations," Tilburg Center for Logic and Philosophy of Science, until filled.
Post-doc position: in History and Philosophy of Science, University of Sydney, deadline 1 March.
Post-doc Position: in Artificial Intelligence, Monash University, deadline 1 March.
Post-doc position: in History and Philosophy of Science, Technology and Medicine, Hebrew University of Jerusalem, deadline 5 March.
Professor: in Philosophy of the Social Sciences and Social Philosophy, Leibniz Universität Hannover, Germany, deadline 8 March.
Post-doc Position: in Philosophy, Department of Philosophy, Logic and Scientific Method, LSE, deadline 8 March.

Post-doc Positions: in Philosophical Logic, University of Oslo, deadline 13 March.
Lecturer: in Statistics, University of Lancaster, deadline 16 March.
Post-doc Position: in Philosophy AOS: Epistemology, Logic, and the Theory of Reasoning, Ruhr University Bochum, deadline 20 March.
Post-doc Position: in Philosophy, AOS: Logic or Philosophy of Science, University of Calgary, deadline 22 March.
Post-doc Position: in the Philosophy and History of Science and Medicine, University of Saskatchewan, deadline 1 April.

Studentships

Three Doctoral Training Grants: School of Computing, Faculty of Engineering, University of Leeds, until filled.
PhD position: in Bayesian Decision Theory, School of Computer Science and Statistics, Trinity College Dublin, until filled.
Two PhD positions: in the project "Designing and Understanding Forensic Bayesian Networks with Arguments and Scenarios", Utrecht University / University of Groningen, to be filled asap.
PhD positions: in the Statistics & Probability group, Durham University, until filled.
PhD positions: in Statistical Methodology and its Application, University College London, until filled.
PhD position: in Logic and Theoretical Philosophy at the Institute for Logic, Language and Computation at the University of Amsterdam, until filled.
PhD Position: in Cognitive Science at the Philosophy Department, Lund University, deadline 1 March.
PhD Position: on project "Causal modelling in public health," School of Medicine, University of Manchester, deadline 6 March.
PhD Position: in advanced Bayesian computational methods, School of Mathematics, Statistics and Actuarial Science, University of Kent, deadline 9 March.
PhD Positions: in Mathematical Philosophy, Munich Center for Mathematical Philosophy, deadline 12 March.
PhD Position: in Philosophy AOS: Epistemology, Logic, and the Theory of Reasoning, Ruhr University Bochum, deadline 20 March.
PhD Position: in Statistics, Department of Mathematics, University of Oslo, deadline 1 April.