Second Reasoning Club Conference

Scuola Normale Superiore, Pisa

17-19 June 2013
http://reasoningclubpisa.weebly.com/

Organisers:
Hykel Hosni (Scuola Normale Superiore, Pisa and CPNSS, London School of Economics)
Julien Murzi (University of Kent and Munich Centre for Mathematical Philosophy)

With the support of the *Centro di Filosofia*

Monday 17 June

09:45 - 10:00	Welcome
10:00 - 11:00	George Wilmers **Collective Probabilistic Reasoning: A Survey of Some Results and an Open Problem**
11:00 - 11:15	Break
11:15 - 12:00	Juergen Landes **Objective Bayesianism and the Maximum Entropy Principle**
12:00 - 12:45	Jacopo Amidei **The Turing-Good analysis of weight of evidence**
12:45 - 14:00	Lunch
14:00 - 14:45	Rossella Marrano **From degrees of truth to degrees of consequence**
14:45 - 15:30	Graziana Ciola **Relations and Relational Inferences in Medieval and Early Modern Logic**
15:30 - 16:00	Break
16:00 - 16:45	Orri Stefánsson **The Relationship Between Centring and Supervenience**
16:45 - 17:30	Alex Marcoci **In defense of fair lotteries**
19:00 -	Reception

Tuesday 18 June

09:30 - 10:30	Julia Tanney **Achilles and the Tortoise, Rule-Following, and the Intellectualist Legend**
10:30 - 11:15	Michael Wilde **Timothy Williamson on Evidence and Probability**
11:15 - 11:45	Break
11:45 - 12:30	Jun Lai **The statistical explanation of learning complex recursive structure**
12:30 - 14:00	Lunch
14:00 - 14:45	Liz Howarth **The Theory of Spectrum Exchangeability**
14:45 - 15:30	Malte Kliess **Extending Pure Inductive Logic to Second Order Language**
15:30 - 16:00	Break
16:00 - 16:45	Lorenzo Casini **Interventions, Simulation and Causal Inference**
16:45 - 17:30	Janine Reinert **Ontological Omniscience: The Case of Lewisian Modal Realism**

Wednesday 19 June

09:30 - 10:30	Jan Sprenger **Theory Confirmation: From Popper to Bayes (and back again)**
10:30 - 11:15	Teddy Groves **An application of Carnapian inductive logic to philosophy of statistics**
11:15 - 11:45	Break
11:45 - 12:30	Tahel Ronel **Rational Principles based on Symmetry in Polyadic Inductive Logic**
12:30 - 14:00	Lunch
14:00 - 14:45	Filippo Ferrari **The Value of Deflationary Truth**
14:45 - 15:30	Andreas Fjellstad **Tonking around with implicit definitions**

Abstracts

Collective Probabilistic Reasoning: A Survey of Some Results and an Open Problem

George Wilmers

(Joint work with Martin Adamcik)

Within the framework of discrete probabilistic uncertain reasoning an agent's *credence base* $\mathbf{K}$ may be defined as a set of *constraints* on the possible (probabilistic) *belief functions* on some fixed finite propositional language $\mathcal{L}$ which are consistent with that agent's beliefs based on her experience of the world. $\mathbf{K}$ thus uniquely determines a set of probability functions $\mathbf{V_K}$ any one of which the agent *may* consistently choose to be her personal belief function[1]. If the agent is now required to choose a single probability function as her belief function, which one should she choose in the absence of any other information? An *inference process* is a function $\mathcal{I}$ which, given any credence base $\mathbf{K}$, chooses a single probability function $\mathcal{I}(\mathbf{K})$ from the set $\mathbf{V_K}$. There is a considerable literature justifying the maximum entropy inference process, **ME**, as being optimal in this or similar contexts (see Paris(1994), Jaynes(1979), Paris and Vencovská (1989), Paris and Vencovská (1990), Williamson (2010)). The principal methods of justification of **ME** are of three very different kinds: (i) deriving **ME** from an epistemic possible worlds interpretation in which, in the absence of any other information, each possible world is considered a priori equally likely, (ii) characterising **ME** by using arguments about minimising expected betting loss or minimising expected gain in information, and finally (iii) characterising **ME** axiomatically as the *unique* inference process satisfying a simple set of formal principles which in turn are justified by intuitive notions of symmetry, consistency, or

[1] In the literature a credence base in this context is usually referred to, perhaps a little misleadingly, as a *knowledge base*. A technical aside: it is assumed throughout that $\mathbf{V_K}$ is a non-empty closed convex set of probability functions. Non-emptiness here corresponds to assuming that the agent's probabilistic beliefs are consistent.

semantic invariance.

More recently the axiomatic approach of (iii) has been extended to an analysis of a *social* inference process in the context of n agents whose respective credence bases $\mathbf{K}_1, \ldots, \mathbf{K}_n$, while individually consistent, may be collectively inconsistent (Wilmers (2010), Wilmers (2010), Adamčík and Wilmers (2012;2013)). Such a social inference process, (abbreviated to SIP), is a function, say $\mathcal{S}$, which operates on the multiset of credence bases $\mathbf{K}_1, \ldots, \mathbf{K}_n$ to yield a single "social" belief function $\mathcal{S}(\mathbf{K}_1, \ldots, \mathbf{K}_n)$ which best represents the beliefs of all the agents. Here a consideration of the general characteristics of the process by which $\mathcal{S}$ extracts the social belief function $\mathcal{S}(\mathbf{K}_1, \ldots, \mathbf{K}_n)$ raises some subtle questions. *A priori* we might consider three possible approaches to how an SIP $\mathcal{S}$ might work. We will assume here that **ME** is the accepted inference process for a single agent, but our categorisation below of three types of SIP's holds with obvious modifications even if this assumption is not made:

1. Using **ME** each agent i selects from $\mathbf{V_K}$ a personal belief function $\mathbf{w}^{(i)} = \mathbf{ME}(\mathbf{K}_i)$ which represents her personal credence base $\mathbf{K}_i$, thus ignoring the credence bases of the other agents. These representative probability functions $\mathbf{w}^{(1)} \ldots \mathbf{w}^{(n)}$ are then combined using a *pooling operator* (e.g. arithmetic mean or normalised geometric mean) to yield the "social" belief function $\mathcal{S}(\mathbf{K}_1, \ldots, \mathbf{K}_n)$.

2. The agents' credence bases $\mathbf{K}_1, \ldots, \mathbf{K}_n$ are *merged* into a single "intersubjective" credence base $\mathbf{K}$ by some process such that the probability functions $\mathbf{v} \in \mathbf{V_K}$ are exactly those which *globally* minimise some particular function $F(\mathbf{v}, \mathbf{w}^{(1)} \ldots \mathbf{w}^{(n)})$ subject only to the conditions that for each i $\mathbf{w}^{(i)}$ is consistent with $\mathbf{K}_i$. The intuitive idea here is that F should represent some notion of combined distance, or sum of distances, between $\mathbf{v}$ and the $\mathbf{w}^{(1)} \ldots \mathbf{w}^{(n)}$. If more than one probability function is consistent with the resulting merged $\mathbf{K}$ then $\mathcal{S}(\mathbf{K}_1, \ldots, \mathbf{K}_n)$ is simply defined to be $\mathbf{ME}(\mathbf{K})$.

3. $\mathcal{S}(\mathbf{K}_1, \ldots, \mathbf{K}_n)$ is determined by some type of procedure *other* than (1) or (2) above.

It can readily be seen that the axiomatic framework of SIP's provides at once a synthesis and generalisation of both the axiomatic framework of inference processes as in Paris (1994), and also of the axiomatic framework of discrete probabilistic pooling operators from decision theory[2], both of which are included as marginal cases.

Let us call an **SIP** *obdurate* if it is of type (1) and *deliberative* if it is of type (2). There are good reasons for believing that an *obdurate* SIP is inherently unsatisfactory, since it appears gratuitously to "throw away" information. Furthermore, except possibly in pathological cases, obdurate SIP's always violate elementary consistency desiderata including the following two natural principles:

The Consistency Principle

If $\mathbf{K}_1, \ldots, \mathbf{K}_n$ are jointly Nation, consistent then

$$\mathcal{S}(\mathbf{K}_1, \ldots, \mathbf{K}_n) \in \bigcap_{i=1}^{n} \mathbf{V}_{\mathbf{K}_i}$$

.

The Ignorance Principle

If agent $n+1$ has empty credence base (i.e. she has "not a clue"), then

$$\mathcal{S}(\mathbf{K}_1, \ldots, \mathbf{K}_n, \mathbf{K}_{n+1}) = \mathcal{S}(\mathbf{K}_1, \ldots, \mathbf{K}_n)$$

Nonetheless, despite this disastrous failure of obdurate SIP's, we can learn something from them. It turns out that if **LogOp**, the normalised geometric mean pooling operator, is used to define an obdurate **SIP** as in (1) above, then the resulting **SIP**, called the Obdurate Entropy Process, (**OEP**), satisfies generalisations of some powerful and highly desirable principles associated with **ME**, which appear otherwise very difficult for an SIP to satisfy.

[2]See French (1985) for a survey of classical results on pooling operators,

One such principle inherited from Paris and Vencovská's work (1990) on **ME** is:

The Irrelevant Information Principle

Let $\mathcal{L}_1$ and $\mathcal{L}_2$ be two disjoint propositional languages such that $\mathcal{L} = \mathcal{L}_1 \cup \mathcal{L}_2$. Suppose that for each $i = 1 \dots n$ agent i has credence base $\mathbf{K}_i \cup \mathbf{T}_i$ where $\mathbf{K}_i$ is formulated using only propositional variables from $\mathcal{L}_1$ while $\mathbf{T}_i$ uses only propositional variables from $\mathcal{L}_2$. Then if θ is any sentence of $\mathcal{L}_1$

$$\mathcal{S}(\mathbf{K}_1 \cup \mathbf{T}_1, \dots, \mathbf{K}_n \cup \mathbf{T}_n)(\theta) = \mathcal{S}(\mathbf{K}_1, \dots, \mathbf{K}_n)(\theta)$$

This result, which is due to Martin Adamčík (2013), suggests that there exists a natural association between **ME** and the pooling operator **LogOp** in the SIP context, which is perhaps independent of the assumption of obdurateness. The suggestion is reinforced by the following. In [**?**] the author formulated a deliberative SIP called the Social Entropy Process (**SEP**), defined by taking F in (2) to be the sum of Kullback-Leibler divergences from the $\mathbf{w}^{(i)}$ to $\mathbf{v}$. Now **SEP** has very strong consistency properties, satisfying both the Consistency and Ignorance Principles above as well as many other desirable principles (Wilmers (2010,2011), Adamčík and Wilmers (2013)). It turns out that **LogOp** also plays an essential role in the structural properties of **SEP** in the sense that as a consequence of the definition of **K** in (2) we find that $\mathbf{v} = \mathbf{LogOp}(\mathbf{w}^{(1)} \dots \mathbf{w}^{(n)})$ whenever $F(\mathbf{v}, \mathbf{w}^{(1)} \dots \mathbf{w}^{(n)})$ is minimised. However, despite this, **SEP** does *not* satisfy the Irrelevant Information Principle.

So a natural question arises: can we improve on **SEP**? In particular is there there *any* SIP which satisfies the the Consistency Principle and the Ignorance Principle and which also satisfies the Irrelevant Information Principle? More ambitiously, noting that both **SEP** and **OEP** marginalise to the single agent inference process **ME** and to the pooling operator **LogOp**, is there some SIP which also marginalises to **ME** and to **LogOp** and which satisfies the union of those desirable principles which are known to be satisfied either by **SEP** or by **OEP**? It seems likely that if such an SIP exists then such a list of principles could be sufficient characterise it. One reason why this last

question seems interesting is that there do exist epistemic possible worlds interpretations which justify both the inference process **ME** and the pooling operator **LogOp**, but it does not seem immediately obvious how to construct a similar epistemic model for any SIP. The interplay between the study of epistemic interpretations and the use of axiomatic methods seems potentially as fruitful in the study of SIP's as it has been in the study of inference processes.

References

M. Adamčík and G. M. Wilmers, (2012) *The Irrelevant Information Pinciple for Collective Probabilistic Reasoning*, to appear. Preprint at http://www.academia.edu/1942166/The -Irrelevant-Information-Principle-for-Collective-Probabilistic-Reasoning

M. Adamčík and G. M. Wilmers, (2013) *Probabilistic Merging Operators*, Logique et Analyse, to appear. Preprint available at http://personalpages.manchester.ac.uk/staff/martin.adamcik/mergingpaper.pdf

M. Adamčík, (2013) PhD Thesis, Manchester, to appear.

S. French, (1985) *Group Consensus Probability Distributions: A Critical Survey*, in J. M. Bernardo, M. H. De Groot, D. V. Lindley, and A. F. M. Smith (Eds.), Bayesian Statistics, Elsevier, North Holland, pp. 183-201.

E.T. Jaynes, (1979) *Where do we Stand on Maximum Entropy?* in "The Maximum Entropy Formalism", R. D. Levine and M. Tribus (eds.), M.I.T. Press, Cambridge, MA.

J. B. Paris, (1994) *The uncertain reasoner's companion*, Cambridge University Press, Cambridge

J. B. Paris and A. Vencovská, (1989) *On the Applicability of Maximum Entropy to Inexact Reasoning*, International Journal of Approximate Reasoning, 3

J. B. Paris and A. Vencovská, (1990) *A Note on the Inevitability of Maximum Entropy*, International Journal of Approximate Reasoning, 4, pp. 183-224

Williamson, J. (2010), *In Defense of Objective Bayesianism*, Oxford University Press

G. M. Wilmers (2010), *The Social Entropy Process: Axiomatising the Aggregation of Probabilistic Beliefs* , Probability, Uncertainty and Rationality,

pp 87-104, edited by Hykel Hosni and Franco Montagna, 10 CRM series, Scuola Normale Superiore, Pisa

G. M. Wilmers (2011), *Generalising the Maximum Entropy Inference Process to the Aggregation of Probabilistic Beliefs*, available from http://manchester.academia.edu/GeorgeWilmers/Papers

Objective Bayesianism and the Maximum Entropy Principle

Jürgen Landes

(Joint work with Jon Williamson)

Objective Bayesian epistemology as developed by Williamson (2010) invokes three norms that lead together to the maximum entropy principle: the strengths of our beliefs should be probabilities, they should be calibrated to our evidence of physical probabilities, and they should otherwise equivocate sufficiently between the basic propositions that we can express. However, these three norms are usually justified in rather different ways. In this paper we show that the three norms can all be subsumed under a single justification in terms of minimising worst-case expected loss.

The Turing-Good analysis of weight of evidence

Jacopo Amidei

In this talk I will present the concept of weight of evidence defined by Alan Turing and subsequently investigated by Irving John Good. After a quick historical overview of how the concept of weight of evidence was developed during the World War II, I will define the central notions (odds, Bayes factor and weight of evidence) and I will point out some of their properties. I will then illustrate the relation between the concept of weight of evidence and the measure of information defined by Shannon in 1948. Then I will discuss a result by Good which shows how the weight of evidence is a more refined measure of the amount of information. I will conclude by putting forward two proposals that use weights of evidence to quantify second-order uncertainty, and relate it to imprecise probabilities.

From degrees of truth to degrees of consequence

Rossella Marrano

(Joint work with Hykel Hosni)

The intuitive plausibility of the relation "less wrong than" motivates the investigation into how degrees of truth can be interpreted and modelled. The standard formalisation based on infinite-valued Łukasiewicz logic turns out to be inadequate, both intuitively and formally. First, degrees of truth suffer from an interpretational problem: what does it mean for a sentence to be 0.704638346 true? Second, the argument known as Suszko's reduction shows a formal incompatibility between logical "many-valuedness" and the Tarskian notion of consequence. I argue that suitably defined *graded consequence relations* provide an adequate model of the intuitive notion of "less wrong than" without running into the above difficulties. By shifting the focus from degrees of truth to degrees of consequence we can trigger a positive feedback which sharpens our intuition on the initial problem. In turn, this raises the question as to how to interpret the resulting degrees of consequence. I will conclude by suggesting that objective probabilities may find an interesting application here.

Relations and Relational Inferences in Medieval and Early Modern Logic

Graziana Silvia Ciola

(Joint work with Massimo Mugnai)

According to the mainstream historiographical interpretations, until the mathematization of logic relations could not be treated in traditional Aristotelian logic. This is roughly true, but only partially so. In Medieval and Early Modern ontological and logical traditions, on the one hand, the metaphysical problem of relations is thematised and widely discussed; on the other hand, there are some attempts to deal with relational inferences. Yet such attempts were often not successful or at least quite problematic themselves. I shall sustain that these two sides of the problem are quite intertwined and condition each other: medieval and early modern logical treatments of relations appear to be crushed between ontological presuppositions/desiderata and the predominance of syllogistic logic. Therefore, after presenting the general problem and introducing briefly the ontological side, I shall examine – on the side of relational inferences – the example of oblique syllogisms.

In defence of fair lotteries

Alex Marcoci

Taurek [1977] raised the following problem: given two unequal groups of equally worthy individuals facing mortal peril and given that we desire to save at least one group, who should we save? Taurek then argued that we should decide in such a way that we offered all individuals (irrespective of their numbers and groups) an equal opportunity of being saved and claimed that the only way to do that would be by distributing life between the two groups through a fair lottery. In a very influential passage, Scanlon [1998] argues that if we were to employ a fair lottery to distribute life among two groups (as Taurek recommends) we would not be giving equal moral consideration to all the individuals concerned. Scanlon's argument goes by comparing the following two cases. 1 vs. 2 We have two single-individual groups. In this case, if we assume that 1 and 2 are perfectly interchangeable then it is obvious that we should distribute life among them through a fair lottery. 1,3 vs. 2 Now we add person 3. If we follow Taurek then we should again distribute life between these two groups through a fair lottery. However, Scanlon remarks, given that this is the same method we employed in 1 vs. 2, following Taurek means acting, in fact, as if person 3 did not exist. And this means treating 3 as if he/her were not as important as 1 and 2. Scanlon puts the argument thus: "The presence of the additional person, however, makes no difference to what the agent is required to do or to how is required to go about deciding what to do" (Scanlon [1998], p. 232) In this paper I show that Scanlon is not right: distributing life through a fair lottery does take into account the presence of the third person. The argument goes as follows:

1. I present an allocation procedure which is sensitive to the relative sizes of the two groups: for any difference in size the steps we take

to determine how to distribute life between the two groups will be different.

2. I prove that this allocation procedure is formally equivalent to the fair lottery no matter of the sizes of the two groups. In effect, this means that the fair lottery is equivalent to infinitely many lotteries (one for every difference in the relative size of the two groups)

3. Because the allocation procedure I present recommends different steps in deciding who to save for any difference in the sizes of the two groups I will argue that such a procedure gives equal consideration to all individuals concerned: they all make a like difference to the decision. However, because the allocation is always equivalent to the fair lottery, I argue that Scanlon is not right in claiming that the fair lottery does not take into account the additional person.

In other words, my argument is based on the observation that there is a difference between what the agent is required to do and how she is required to go about deciding what to do: while according to the allocation procedure outlined in this paper the agent is deciding in a different manner for each difference in the relative sizes of the two groups, she will always be required to do the same thing (i.e. distribute through a fair lottery).

The Relationship Between Centring and Supervenience

H. Orri Stefansson

According to *Strong Centring*, we can infer the *truth* of the counterfactual AB from the conjunction $A \wedge B$. According to *Weak Centring*, we can infer the *falsity* of AB from the conjunction $A \wedge \neg B$. Both principles are implied by David Lewis' (and Robert Stalnaker's) classical possible worlds semantics for counterfactuals (and the first implies the latter given 'counterfactual consistency'). In this paper I prove that together these two centring conditions imply that *facts supervene on counterfactuals*. (I also prove a stronger result: weak and strong centring together imply that any counterfactual is *equivalent* to a factual sentence.) These centring conditions are thus incompatible with a strong version of Lewis' *Humean Supervenience* thesis (HS), according to which certain non-modal facts are more fundamental than all other facts (in particular modal facts) and do not supervene on any less fundamental facts. Some will perhaps take the aforementioned result as establishing that the supervenience relation is not strong enough to sustain metaphysically significant distinctions in 'fundamentality'. But I suggest that if we combine Carl Hoefer's recent Humean account of objective chance with a probabilistic semantics for counterfactuals (as e.g. formulated by Hannes Leitgeb), then we might be able to satisfy the strong version of HS (but at the expense of centring).

Achilles and the Tortoise, Rule-Following, and the Intellectualist Legend

Julia Tanney

This talk will examine the ways in which Lewis Carroll's seminal paper anticipates discussions in Wittgenstein and Ryle on the importance of recognizing a type-distinction between rules and performances; a distinction that requires the fundamental (normative, explanatory, and justificatory) priority of the performances over the rules that 'govern' them.

As with the relation between knowledge of meaning-rules and language use, or between priori theoretical reasoning and the manifestation of intelligence, the vicious regress Carroll spells out thwarts the construal of 'propositional' logical knowledge as fundamentally prior to logical prowess (e.g, deductive reasoning).

In the course of this talk I hope to cast light on common misunderstandings in the literature about Ryle's discussions of dispositions, their relation to subjunctive conditionals or hypotheticals, as well as about the distinction he introduced between 'knowing how' and 'knowing that'.

Timothy Williamson on Evidence and Probability

Michael Wilde

In Knowledge and Its Limits (2000) Williamson gives an account of the nature of evidence in terms of knowledge and an account of probability similar to the logical theory of probability due to Keynes. In this talk I consider and respond to a number of objections to both Williamson's account of evidence and his account of evidential probability.

The statistical explanation of learning complex recursive structure

Jun Lai

Hierarchical center-embedded structures pose a large difficulty for language learners due to their complexity. A recent artificial grammar learning study (Lai & Poletiek, 2011) demonstrated a starting-small (SS) effect, i.e. staged-input and sufficient exposure to 0-level-of-embedding exemplars were the critical conditions in learning A^nB^n structures. The current study aims to test: 1) a more sophisticated type of SS (a gradually rather than discretely growing input); 2) the frequency distribution of the input. The results indicate that SS optimally works under other conditional cues, such as a skewed frequency distribution with simple stimuli being more numerous than complex ones.

The Theory of Spectrum Exchangeability

Liz Howarth

(Joint work with Jeff Paris)

Spectrum Exchangeability (Sx) is a symmetry principle of Inductive Logic, and arguably the most natural (but not the only) extension of atom exchangeability to polyadic languages. It has been shown that all probability functions which satisfy Sx are comprised of a mixture of two essential types of probability functions; heterogeneous and homogeneous functions. We determine the theory of Spectrum Exchangeable functions, which for a fixed language L is the set of sentences of L which must be assigned probability 1 by every function satisfying Sx, by examining separately the theory of heterogeneous and homogeneous functions. An unexpected consequence of adopting Sx as a principle of 'rational' reasoning emerges.

References

Landes, J., Paris, J.B. & Vencovská, Representation Theorems for Probability Functions satisfying Spectrum Exchangeability in Inductive Logic, *International Journal of Approximate Reasoning*, 2009, 51(1):35-55.

Nix, C.J., *Probabilistic Induction in the Predicate Calculus*, Doctorial Thesis, University of Manchester, UK, 2005.

Nix C.J. & Paris, J.B., A Note on Binary Inductive Logic, *Journal of Philosophical Logic*, 2007, 36(6):735-771.

Extending Pure Inductive Logic to Second Order Languages

Malte Kleiss

(Joint work with Jeff Paris)

Let w be a probability function on a purely unary language L and suppose w satisfies Unary Language Invariance. Then there exists an extension w' of w to the language L_∞, containing countably many predicates. One can then extend w' to second-order L sentences, treating second-order quantifiers in a similar way as their first-order counterparts.

The reason one might be interested in studying such extensions could be to ask whether a definable property is already represented by one of the predicates in L_∞.

We will present George's Principle which states that probability functions should give such statements a probability of 1, and then give an account on how George's Principle relates to principles like Regularity and how one might go about obtaining a Representation Theorem for such functions.

Interventions, Simulation and Causal Inference

Lorenzo Casini

According to Woodward, X causes Y iff there is an 'ideal' intervention on X that would change Y (as defined in Woodward 2003, p. 84) whilst leaving the relation invariant (INT). Whilst I do agree that INT sheds light on one important aspect of the meaning of causal claims, I argue that it fails to be an all-encompassing criterion. I buttress my claim with reference to simulations of the asset pricing mechanism. To the extent that such simulations provide genuine causal explanations, not interpretable by means of INT, they invite us to rethink not only the role of interventions for establishing causal claims, but also the meaning of causal claims more generally.

Woodward motivates INT as follows. INT makes precise the connection between 'causality' and the concept with which 'causality' is most intuitively and fruitfully associated, viz. 'intervention' (Woodward 2003, §2.1). INT applies to all contexts in which causal talk is appropriate or, at any rate, always better than other criteria. If ideal interventions are not practically possible, INT works as a "regulative ideal" (ibid, 114), such that the other criteria by which we establish causality (e.g., RCTs, observational studies) may always be (re-)interpreted as attempts to 'approximate' ideal interventions (ibid, 34-35). And if ideal interventions are not well-defined (e.g., they violate physical laws), INT works as a conceptual criterion, such that causal inference is granted by the interventions being logically/conceptually possible. Here, by providing a principled basis to answer counterfactual questions, INT gives us "a purchase on what we mean or are trying to establish when we claim that X causes Y" (ibid, 130). However, Woodward's arguments do not always apply.

Consider the following example. The asset pricing mechanism is known to produce financial time series with statistical features (e.g., fat-tailed distributions of returns, volatility clustering and persistence), known as 'stylised facts', that cannot be explained by traditional economics. Recently, agent-

based models (ABMs) built on non-classical assumptions (prices are determined by aggregating the demands of heterogeneous, non-maximally-rational agents) have provided support to the hypothesis that 'stylised facts are caused by the endogenous and destabilising mechanism generated by the agents' heterogeneity and interactions'. For instance, the ABMs in (Lux and Marchesi, 1999) and (LeBaron et al, 1999) interpret the market as, respectively, a many-particle system undergoing phase transition and a population undergoing natural selection. They identify two potential causes of the stylised facts, respectively, the agents' switching process and learning speed. Also, they show that the effect obtains robustly across variations of modelling assumptions, provided the agents are heterogeneous and influence each other. Admittedly, such causal explanations are coarse-grained – e.g., they don't pin down any specific event as the cause. Nonetheless, they convincingly identify causally relevant properties of the mechanism generating the stylised facts. Importantly, the inference to the causal claim does not depend – whether explicitly or implicitly – on INT.

On the one hand, since the ABMs depict highly idealised scenarios, they can hardly be interpreted as suggesting the existence of properties of real markets corresponding to the variables and parameters responsible for the stylised facts (e.g.: different 'kinds' of traders 'switching' from one class to another; changes of trading strategies as evolution of 'genotypes', and 'learning' as random mutation and crossover). So, it makes little sense to interpret the robustness of the results across variations in variable- and parameter-values in terms of invariance under ideal interventions. The variations that warrant the causal claim are not meant to be, or to approximate, ideal interventions in Woodward's sense. Thus, INT does not work as a regulative ideal.

On the other hand, to the extent that the ABMs warrant the claim, they do so in the absence of clear intuitions on the outcomes of ideal interventions. What would happen if one were to affect the switching process or the learning speed is either unclear or only clear with respect to the model and not the real system. But this is largely irrelevant. We infer to the causal claim because the ABMs reproduce the phenomena, and because the market is analogous in relevant respects (heterogeneity and interactions) to other mechanisms (phase transition and evolution) with which we are familiar, so that the ABMs also explain the phenomena. Thus, INT doesn't work as a

conceptual criterion either – it is silent on what we mean or try to establish when we say that X causes Y.

Now, Woodward allows that there may be cases where INT gives no good purchase on the meaning of causality. However, he contends that the less a causal claim is interpretable according to INT, the less it is meaningful (Woodward 2008, §12) – unless there is some alternative test criterion conflicting with INT and intuitively superior to INT (ibid, §14; cf. Woodward 2003, p. 91). But this is controversial, as the ABMs case suggests. For one thing, one may admit that interventions are the most intuitive test criterion for causality and yet deny that the meaning of causal claims always depends on the existence of clear intuitions on the outcomes of ideal interventions. For another thing, the justification for the causal inference need not necessarily depend on the existence of a superior test criterion. Having the right contextual reasons seems enough. (This is in line with a cluster-concept view of causality that I elaborate elsewhere.)

Causal inference is often warranted by interventions. But sometimes it isn't, nor need it be.

References

LeBaron, B. D., Arthur, W. B., and Palmer, R. G. (1999). Time Series Properties of an Artificial Stock Market. Journal of Economic Dynamics and Control, 23:1487–1516.

Lux, T. and Marchesi, M. (1999). Scaling and Criticality in a Stochastic Multi-agent Model of a Financial Market. Nature, 397:498–500.

Woodward, J. (2003). Making Things Happen. New York: Oxford University Press.

Woodward, J. (2008). Causation and Manipulability. E. N. Zalta (ed.): The Stanford Encyclopedia of Philosophy (Winter 2008 Edition).

Ontological Omniscience: The Case of Lewisian Modal Realism

Janine Reinert

A simple argument against Lewisian Modal Realism (MR) as portrayed in On The Plurality of Worlds ([Lew86]) arises from an interaction between its treat- ment of doxastic content and of modality. It is easily shown that it is either impossible to doubt the theory on ontological grounds, or equivalently, that if it is possible to maintain doubt about MR's existential postulate, it has to be false. The argument hinges on the fact that MR's main ontological hypothesis, if true, is necessarily true. The thesis that there exists a plurality of concrete and (causally as well as spa- tiotemprorally) isolated worlds, henceforth: the existential claim, EC, is what Divers has called an 'extraordinary claim' of MR [Div02, 48]. 'Extraordinary' means the following: While (pluralities of) other possibilia can be said to exist at (that is: as parts of worlds), each world is only – improper – part of itself. A plurality of at least two worlds cannot exist at any single world. EC thus calls for a different interpretation, which instead of parsing existence as existence-at understands it as existence simpliciter (cf. [Div99]). By reductio ad absurdum it is easily shown that if EC is true at some world, it has to be true at all worlds, and therefore is necessary in MR (cf. [Div99] on the redundancy interpretation in cases of 'advanced modalizing'). The necessity of EC now poses a problem for Lewis's treatment of doxastic content: While he is absolutely aware that many philosophers have found EC incredible, this disbelief turns out to be impossible on the Hintikka-style analysis Lewis offers in [Lew86, 27-9]. Doxastic content is given in terms of an agent's doxastically accessible worlds at a given time, i.e. "World W is one of those iff he believes nothing, either explicitly or implicitly, to rule out the hypothesis that W is the world where he lives." [Lew86, 27]. If Lewis's concrete worlds are the only worlds there are, then this class has to be empty if the agent's belief system includes belief in the falsity of EC. But this amounts to a

trivialization of belief due to doubt about MR, and this is inadequate with respect to actually (or merely possibly) held beliefs. We can therefore conclude that either, if MR is true, every doxastic agent with a non-trivial belief system in fact has implicit belief in EC, or that if someone (non-trivially) believes in the falsity of EC, MR has to be false.

After a detailed presentation of this argument, I will discuss the objection that it is a mere 'artefact' of the presupposed analysis of doxastic content and, as a typical case of logical omniscience, is not peculiar to MR. I will then apply Lewis's strategies to make seemingly impossible belief possible to the presented case and explain why none of these can explain away the problem at hand. I will conclude with a discussion of why also the introduction of impossible worlds could not block the presented argument, and argue that as a case of ontological omniscience, the argument from doubt discredits MR as an ontological founda- tion of doxastic semantics altogether.

References

[Div99] John Divers. A genuine realist theory of advanced modalizing. Mind, 108(430):217–239, 1999.

[Div02] John Divers. Possible Worlds. Routledge, 2002.

[Lew86] David Lewis. On the Plurality of Worlds. Blackwell Publishers, Oxford, 1986.

[Rei13] Janine Reinert. Ontological Omniscience in Lewisian Modal Realism. Analysis, (forthcoming), 2013.

Theory Confirmation: From Popper to Bayes (and back again)

Jan Sprenger

Karl R. Popper was, famously, a fervent opponent of inductive inference. Nevertheless he recognized the need to form preferences over competing theories and to measure how well a theory has stood up to test in practice. Such a measure of "degree of corroboration" does not only fill a gap in critical rationalism; it also helps to properly interpret the result of statistical tests.

In this contribution, I show how much Popper's own approach to quantifying corroboration relies on principles of Bayesian confirmation theory. Ironically, this turns out to be a major weakness of Popper's proposal. Instead, I suggest to measure corroboration in a way that is relative to a *partition* of alternative hypotheses. This proposal, inspired by some ideas of I.J. Good, does not only provide a satisfactory explication of degree of corroboration/confirmation: it also elucidates the relationship between Popper's degree of corroboration and Bayesian measures of evidential support. Thereby it unifies critical rationalist and Bayesian perspectives on theory appraisal in science.

An application of Carnapian inductive logic to philosophy of statistics

Teddy Groves

In my talk I claim that an argument in philosophy of statistics found in Gelman and Shalizi (2012) can be improved using Carnapian inductive logic. Gelman and Shalizi argue against the 'conventional philosophy' of Bayesian statistics, which stipulates that statistical models should be chosen only on the basis of how well they represent knowledge, and only draw conclusions on the basis of knowledge that is represented probabilistically. According to Gelman and Shalizi, these stipulations rule out practically important ways of choosing and evaluating such models. They claim that a hypothetico-deductive philosophy does better justice to the practice of Bayesian statistics. I argue that an approach informed by Carnapian inductive logic should play the role of hypothetico- deductivism in Gelman and Shalizi's arguments.

A Carnapian approach to choosing statistical models

Carnap sought to explicate—that is, usefully replace—everyday inductive discourse with statements of induc- tive logic. An inductive logic consists of a formal language together with a 'confirmation function' associating pairs of its sentences with positive real numbers. The focus of Carnap's research was to find formal conditions that inductive logics must satisfy in order to be satisfactory explications, and to identify which confirmation functions various combinations of such conditions pick out.

For example, Johnson's sufficientness postulate requires that the value assigned to pairs of sentences (h, e), where h describes the instantiation of some property and e describes a sample of observations, must depend only on the frequency of the property in the sample and the sample's size. Carnap argued that that this was a plausible condition of explication and showed that the confirmation functions satisfying it can be characterised by a single continuous parameter.

De Finetti's theorem identifies a correspondence between a class of inductive

logics that interested Car- nap and an important class of statistical models: any probability distribution over independent identically distributed random variables can be represented by a Carnapian inductive logic satisfying the condition constant exchangeability, and vice versa.

This correspondence suggests the following Carnapian approach to selecting statistical models: choose statistical models whose corresponding inductive logics are good explications. This is the approach that I believe should play the role of hypothetico-deductivism in Gelman and Shalizi's argument.

Advantages of the Carnapian approach

Agreement with Gelman and Shalizi's arguments In line with Gelman and Shalizi's criticism of the first stipulation of the 'conventional philosophy', Carnap argued that inductive logics need not represent beliefs in order to be good explications. Analogously to their second argument, Carnap also favoured the use of methods for evaluating inductive logics that are not 1themselves adequate . I argue that these inductive logical arguments are convincing and allow the Carnapian approach to accommodate Gelman and Shalizi's criticisms. It is harder for hypothetico-deductivism to accommodate Gelman and Shalizi's criticisms, as the latter depend on pragmatic reasoning about how statistical models are used. This kind of reasoning is awkward to construe hypothetico-deductively, in contrast to the Carnapian approach, which is inherently pragmatic.

Technical fruitfulness

The Carnapian approach is technically fruitful, allowing statisticians to identify the prior distributions corre- sponding to natural qualitative assertions about the processes underlying their data. I rehearse an argument to this effect from Zabell (2011), present the correspondence between the Dirichlet family of probability distributions and the Johnson's sufficientness postulate as a case study and point to further work in the inductive logic literature that could be relevant to applied Bayesian statistics.

Permissiveness

A final advantage of the Carnapian approach over hypothetico-deductivism is that it is in harmony with Gelman and Shalizi's reluctance to articulate formally precise procedures for evaluating statistical models. The Carnapian

approach offers a principled justification for this attitude: it says that such procedures should depend on which logics are good explications of everyday discourse and so fundamentally concern which form of expression it is preferable to adopt in a particular situation. In Carnap (1950), Carnap argues, in strikingly similar terms to Gelman and Shalizi, that philosophers should refrain from trying to give definitive answers to questions about which forms of expression to adopt. On the other hand I claim, following arguments in Mayo (2013), that Gelman and Shalizi's under- formulation of model-evaluation procedures is difficult to square with hypothetico-deductivism.

Addressing qualms about inductive logic In the last section of my talk I address two reasons why statisticians might be cautious of the Carnapian interpretation of Gelman and Shalizi's argument, arguing that neither is well-founded.

The problem of induction It might be thought that the Carnapian approach suffers particularly from the problem of induction due to its use of 'inductive' logic. An anti-inductive philosophy of statistics, according to which there is no need to reason inductively, might seem preferable. I argue that the Carnapian approach has no greater problem of induction than its alternative. On the one hand, while statements of Carnapian inductive logic are just as difficult to justify as the everyday inductive statements that they explicate and Hume problematised, no additional difficulty arises from the formalisation. On the other hand, anti-inductivism must address the 'practical problem of induction', as set out in Salmon (1981), by providing rational grounds to be confident in their predictions without employing any kind of inductive reasoning. This task seems at least as difficult as overcoming the original problem.

One true inductive logic

Secondly, statisticians might be wary of the Carnapian approach on the grounds that it is associated with a discredited attempt to find one formal method that explicates all of scientific discourse.

To allay such concerns I present quotes from Carnap's work which demonstrate that he did not have this ambition, and indeed considered it impossible.

References

Carnap, R. (1950). Empiricism, Semantics and Ontology. Revue Internationale de Philosophie, 4:20–40.

Gelman, A. and Shalizi, C. R. (2012). Philosophy and the practice of bayesian statistics. British Journal of Mathematical and Statistical Psychology, 66(1):8–38.

Mayo, D. G. (2013). The error-statistical philosophy and the practice of bayesian statistics. British Journal of Mathematical and Statistical Psychology, 66(1):57–64.

Salmon, W. C. (1981). Rational prediction. British Journal for the Philosophy of Science, pages 115–125.

Zabell, S. (2011). Carnap and the logic of inductive inference. Inductive Logic, 10:265.

Rational Principles based on Symmetry in Polyadic Inductive Logic

Tahel Ronel

Symmetry considerations in polyadic inductive logic have to date produced two key players - the Principle of Spectrum Exchangeability (Sx) and the Permutation Invariance Principle (PIP). While we know that Sx implies PIP (and not the converse), the complex relationship in the opposite direction still leaves much to be explored. In this talk I will describe current methods we have for producing probability func- tions that satisfy these principles and demonstrate how these meth- ods help us elucidate connections between polyadic rational principles, thus helping us to fill in some of the more general polyadic picture. In particular, I will show that fixing any number n, there are language invariant families of probability functions with PIP such that only the members of the family that involve relation symbols up to arity n satisfy Sx.

The value of deflationary truth

Filippo Quinto Ferrari

The focus of this paper is a discussion of one aspect of a familiar challenge raised against Deflationism about truth. The general tenor of such challenge is that deflating truth in the way deflationists wish to do is in contrast with the axiological and normative role that truth plays both in enquiry and in the assertoric practice. The specific challenge I am primarily interested in this paper concerns the value that truth possesses within our enquiry. The aim is, accordingly, not to defend Deflationsim against the general challenge from the normativity of truth, but rather to show that under a certain way of understanding the kind of value that truth exhibits in enquiry, deflationists can account for the value of truth. Accordingly, the main thesis I will defend in this paper is that if we can provide a coherent story of the value of truth as final but extrinsic value, Deflationists have no special problem in accounting for what I call here the 'axiological dimension' of the normativity of truth. The plan is first to provide some qualifications of the strand of Deflationism I will focus on. Secondly, to distinguish between different ways in which truth normatively shapes enquiry with the aim of clarifying the various aspects involved in the challenge from the normativity of truth. I will then discuss the axiological challenge, namely whether and how Deflationism can account for the value of truth, and reply to a recent objection by Michael Lynch. I will conclude by outlining a possible strategy our deflationist can endorse to accommodate the challenge.

Tonking around with implicit definitions

Andreas Fjellstad

Ever since the invention of Prior's infamous connective tonk, a connective supposedly governed by two inference rules in a natural deduction system, one allowing us to derive 'A-tonk-B' from 'A' and one allowing a derivation of 'B' from 'A-tonk-B', the literature on logical inferentialism is typically centred on finding a suitable notion of harmony in order to exclude tonk under the central assumption the connective is meaningless or in any case not a logical operator we could accept. Going against the trend, this paper explores how to make sense of tonk and a related connective as a way of showing that the inference rules for tonk and the related connective are not intrinsically bad implicit definitions. The paper presents a use-theoretic interpretation of the connectives to show how we can make sense of them, it elaborates on proof-theoretic aspects in a cut-free sequent calculus for both connectives to show that they satisfy Belnap's two criteria for implicit definitions, and finally develops a non-transitive two-valued model-theoretic semantics based on the use-theoretic interpretation.