Open Peer Commentary

Classical logic, conditionals and "nonmonotonic" reasoning

doi:10.1017/S0140525X09000296

Nicholas Allott and Hiroyuki Uchida
Department of Linguistics, University College London, London WC1N 1PF, United Kingdom.
n.allott@ucl.ac.uk **uclyhuc@ucl.ac.uk**
http://www.phon.ucl.ac.uk/home/nick/

Abstract: Reasoning with conditionals is often thought to be non-monotonic, but there is no incompatibility with classical logic, and no need to formalise inference itself as probabilistic. When the addition of a new premise leads to abandonment of a previously compelling conclusion reached by *modus ponens*, for example, this is generally because it is hard to think of a model in which the conditional and the new premise are true.

We doubt two claims made by Oaksford & Chater (O&C), one linguistic and one about inference, as they relate to conditionals (*Bayesian Rationality*, Oaksford & Chater 2007, henceforth *BR* as in target article). We do not think that the case has been made that sentences of the form "If A then B" have a semantics diverging from the material implication of propositional logics. We focus here on the related claim that human inference is probabilistic.

Classical Propositional Logic (CPL) is often claimed to be inadequate for explaining our spontaneous propositional inferences. The claim is based on the observation that, whereas CPL is monotonic, human inferences seem to be non-monotonic. For example, it is argued that the validity of some inferences, such as the one in (1a), may be cancelled by addition of another premise, such as the premise R in (1b). The claim is that now that John has broken his left leg, John will no longer run, even if the weather is fine. This is claimed to show the inadequacy of CPL as a tool to explain our propositional inferences, because in CPL, addition of another premise does not influence the provability of the sequent, as is shown in the inference from top to bottom in (1c).

(1a) (Premise 1) $P \rightarrow Q$: If the weather is fine, John runs for a mile.

<u>(Premise 2) P: The weather is fine.</u>
(Conclusion) 3. Q: John runs for a mile.

The valid inference: Whenever Premise 1 and Premise 2 are true, 3 is also true.

(1b) R: John breaks his left leg.
(1c) $P, P \rightarrow Q \vdash Q$
$P, P \rightarrow Q, R \vdash Q$ Weakening

However, it is not clear that the addition of the premise R in (1b) to the set of premises in (1a) actually cancels out the validity of the initial sequent, $P, P \rightarrow Q \vdash Q$, in our propositional inferences. The point becomes clearer if we add the new proposition R to the set of premises, as in $P, P \rightarrow Q, R \vdash Q$. With the specification of the premise propositions as in (1a) to (1b), we do not find this revised sequent to be valid, but this is only because we no longer find the proposition "If the weather is fine, John runs for a mile" ($= P \rightarrow Q$) to be true when we reevaluate the inference. In contrast, if we force ourselves to assume that all the three premise propositions P, $P \rightarrow Q$, R are true in a model of interpretation, then in that model, we have to conclude Q. Some might find it difficult to think of such an interpretation model, because common-sense knowledge tells us that a person normally does not run with a broken leg, but we may sometimes make a claim such as, "If the weather is fine, John (always) runs for a mile. It does not matter if he gets injured. He always runs for a mile." Thus, it is not impossible to force ourselves to think of models in which the proposition, "If the weather is fine, John runs for a mile" is true despite the fact that John has broken his leg. In such models, whenever P is true, Q is also true: that is, the inference goes through.

As is clear from the foregoing discussion, typical interpretation data that allegedly show that monotonic logic cannot capture our propositional inferences include a change of models in which we evaluate the propositional sequents. The valid inference in (1) can be re-stated as "In each model in which P and $P \rightarrow Q$ are both true, Q is also true." The cancellation of the truth of Q arises because in the new model in which we "re-evaluate" the sequent, the premise $P \rightarrow Q$ is no longer true (or we find it more difficult to think of a model in which both $P \rightarrow Q$ and R are true). Because the initial valid inference in (1a) concludes R as a true proposition only on condition that $P \rightarrow Q$ and P are both true, this revision does not really involve the cancellation of the validity of the initial inference.

Given that specification of models in which logical formulas/sequents are evaluated is not part of either the syntax or the semantics of propositional logic languages, it is not clear that alleged non-monotonicity of reasoning, which arises because of the revision of models, requires formalising logical inference as non-monotonic or probabilistic. No doubt people hold beliefs with varying degrees of strength, and a result of reasoning is that these degrees of strength are changed. But one can agree with O&C on this point without thinking that the formal inference system itself is probabilistic.

Space here does not permit discussion of how some propositions are accepted and others rejected, but we are sympathetic to O&C's claim that heuristics that are sensitive to information gain must be involved, with the caveat that it cannot simply be information that is sought, but information that is important to the reasoner at a reasonable processing cost. This recalls discussion of relevance in Gricean pragmatics and Sperber and Wilson's relevance theory.

As Oakford & Chater note, for one's belief in the conditional in (2) it matters whether one discovers, for example, an unstarted car or is told that a car did not start.

(2) If the key is turned the car will start.

A pragmatic explanation in terms of the tendency of speakers to produce utterances relevant to their audience is natural. Effects of the order in which information is presented (see *BR*, pp. 157ff) also require such an explanation, we believe.

This raises a methodological point. To understand human reasoning, both classical logic and O&C's probabilistic account of conditionals and of inference must be supplemented by accounts of processing, and of the pragmatics of utterance interpretation. Thus, it is not obvious that the probabilistic account is more parsimonious.

Identifying the optimal response is not a necessary step toward explaining function

doi:10.1017/S0140525X09000302

Henry Brighton and Henrik Olsson
Center for Adaptive Behavior and Cognition, Max Planck Institute for Human Development, 14195 Berlin, Germany.
hbrighton@mpib-berlin.mpg.de
olsson@mpib-berlin.mpg.de

Abstract: Oaksford & Chater (O&C) argue that a rational analysis is required to explain why a functional process model is successful, and that, when a rational analysis is intractable, the prospects for understanding cognition from a functional perspective are gloomy. We discuss how functional explanations can be arrived at without seeking the optimal response function demanded by a rational analysis, and argue that explaining function does not require optimality.

Oaksford & Chater (O&C) argue in *Bayesian Rationality* (Oaksford & Chater 2007, henceforth *BR*) that a rational analysis is an essential step toward understanding process models from a functional perspective and that "[d]oing this requires developing an account of the optimal behaviour" (p. 268). We argue that relative uses of rational principles of inductive inference can be used to explain function without knowing the optimal response function, and propose that multiple forms of functional analysis are required to understand the cognitive system from a functional perspective.

Rational principles of inductive inference such as Bayesian and simplicity principles are perhaps most frequently used as criteria in relative statements of the form "P is a better response than Q for problem X" For instance, the use of rational principles as model selection criteria (Kearns et al. 1997; Pitt et al. 2002). In contrast, step four of Anderson's (1991a) rational analysis excludes anything but statements of the form "P is the optimal response to task X." Because rational principles can be used to compare the behavior of process models without knowing the optimal response, rational analysis is an instance of a broader class of methodologies adopting rational principles to understand function (e.g., Gigerenzer et al. 1999). Given this, does knowledge of the optimal response offer any intrinsic advantage when explaining why a process is successful; and what price do we pay by demanding knowledge of the optimal response function?

Knowledge of the optimal solution is one way of establishing *that* a particular process or organism is successful; but explaining *why* the process is successful is not implied by this finding. For example, knowing the optimality conditions of the naïve Bayes classifier does not by itself tell us why it is successful. One explanation for why naïve Bayes is successful might be that the independence assumption results in fewer parameters. In certain contexts which violate this independence assumption, the assumption nevertheless causes a reduction in the variance component of error relative to a learning algorithm that assumes that the features are dependent (Domingos & Pazzani 1997). This causal explanation for why naïve Bayes is successful does not require knowledge of the optimal response. It can be established using relative uses of rational principles when the optimal response is incalculable. Furthermore, for realistic contexts of sparse exposure the optimality conditions for naïve Bayes, despite its simplicity, are not fully known (Kuncheva 2006). Although typically unavailable, knowing the optimality conditions for an algorithm would undoubtedly provide a good starting point to understand its function; but optimality conditions are neither a required starting point, nor do they by themselves offer a causal explanation for *why* the algorithm is functional.

Functional explanations arising from Bayesian rational analyses also aim to tell us why a pattern of behavior is rationally justified given that the environment has a certain probabilistic structure. The optimal response function provides a reference point against which to measure behavior without committing to how that behavior is achieved. Because relative uses of rational principles are order relations over theories, or models, of problem solving, they cannot be used in this way. However, abstracting from the process level limits what causal explanation one can offer for why the organism does what it does. One can say *that* it is successful, but whether or not such a finding provides a satisfactory causal explanation for *why* it is successful is not so clear (Danks 2008).

Problems arising from the intractability of reliably identifying the optimal response can also make it less desirable. O&C contemplate the possibility that when a rational analysis is intractable, cognitive science may also be intractable (*BR*, p. 283). Alternatively, this makes complementary forms of functional analysis all the more necessary. The distinction between relative and absolute uses of rational principles mirrors the distinction between, and relative difficulty of, verification and search problems in complexity theory. For example, for an instance of the traveling salesperson problem, the comparative statement "Tour P is shorter than tour Q" is trivial to verify, but the absolute statement "Tour P is the shortest tour" will often be intractable to establish. Many problems take this form and are NP complete (Nondeterministic Polynomial time), including the computation of optimal Bayesian responses and approximations in many settings (e.g., Cooper 1990).

Functional analyses based on rational principles of induction rest on several idealizations: (a) Although different rational principles of induction point to a coherent theoretical picture of what makes a good inference, their practical implementations are often inconsistent, and point to different conclusions (Kearns 1997); (b) functional models will not capture all forms of uncertainty impacting on the problem, some of which may change the character of a functional response (Bookstaber & Langsam 1985); (c) functional models always consider local goals, which only partially inherit the properties of the global goal being examined; (d) rational principles of inductive inference are approximate models of function, which do not consider functional pressures arising from, for example, processing (Brighton & Gigerenzer 2008; Todd & Gigerenzer 2003). These are unavoidable realities of modeling, and apply to both relative and absolute uses of rational principles of induction.

An explanation requiring an optimal response function must also consider that: (e) for problems of inductive inference, the optimal response is often analytically intractable to determine with exact methods, and will not be unique; (f) behavioral responses are typically approximately optimal, revealing a tendency rather than a correspondence; (g) successfully optimizing a local goal does not necessarily take us toward the global optimal when other dependencies are known to be only approximately fulfilled (Lipsey & Lancaster 1956). These additional factors lead to increased flexibility in what behaviors we choose to label optimal. Our point is that these further assumptions are a choice rather than a necessity, and are only required to support certain forms of explanation. We agree with (O&C) on the importance of understanding the ecological function of cognitive processes. We also agree that rational analysis represents a powerful move in this direction. But functional analysis can also proceed without seeking to establish optimality with respect to the organism.

Explaining norms and norms explained

doi:10.1017/S0140525X09000314

David Danks[a] and Frederick Eberhardt[b]

[a]*Department of Philosophy, Carnegie Mellon University and Institute for Human and Machine Cognition, Pittsburgh, PA 15213;* [b]*Institute of Cognitive and Brain Sciences, University of California, Berkeley, Berkeley, CA 94720, and Department of Philosophy, Washington University in St. Louis, St. Louis, MO 63130.*

ddanks@cmu.edu

http://www.hss.cmu.edu/philosophy/faculty-danks.php

fde@berkeley.edu

http://www.phil.cmu.edu/~fde/

Abstract: Oaksford & Chater (O&C) aim to provide teleological explanations of behavior by giving an appropriate normative standard: Bayesian inference. We argue that there is no uncontroversial independent justification for the normativity of Bayesian inference, and that O&C fail to satisfy a necessary condition for teleological explanations: demonstration that the normative prescription played a causal role in the behavior's existence.

In *Bayesian Rationality* (Oaksford & Chater 2007, henceforth *BR*) we understand Oaksford & Chater (O&C) as providing the