THE REVIEW OF SYMBOLIC LOGIC
Volume 1, Number 4, December 2008

FORMAL EPISTEMOLOGY, CONTEXT AND CONTENT: INTRODUCTION TO SPECIAL ISSUE ON RECENT DEVELOPMENTS IN FORMAL EPISTEMOLOGY

HORACIO ARLÓ-COSTA

This special issue presents a series of articles focusing on recent work in formal epistemology and formal philosophy. The articles in the latter category elaborate on the notion of context and content and their relationships. This work is not unrelated to recent developments in formal epistemology. Logical models of context, when connected with the representation of epistemic context, are clearly relevant for many issues considered by formal epistemologists. For example, the semantic framework Joe Halpern uses in his article for this issue has been applied elsewhere to solve problems in interactive epistemology (which can be seen as an active branch of contemporary formal epistemology).

Formal epistemology is today a lively and active area of research in philosophy and neighborhood areas like artificial intelligence, psychology and the decision sciences. Various ideas and concepts used in traditional epistemology have found sophisticated formal counterparts in the development of a new generation of epistemic logics; and the tools recently used to represent belief and belief change have made possible various new philosophical insights. Moreover, formalisms originally designed to solve problems in artificial intelligence have been recently used to solve traditional problems in philosophical logic. This is the case with Halpern's article, which applies and adapts a semantic framework first presented in Fagin *et al.* (1995) to solve the Sorites paradox in particular and the problem of vagueness in general.

Worlds in this semantic framework have the form $(o, s_1, \ldots, s_n)$, where o is an objective state and the states s_i are subjective states of relevant agents. Agent i's subjective state s_i represents '[. . .] i's perception of the world and everything else about the agent's makeup that determines the agent's report.' So, possible worlds in Halpern's paper are structured complexes rather than primitive points. By the same token, the truth of formulas is relative to both the agent and the world. This and other distinctive features of the semantic apparatus first presented in Fagin *et al.* (1995), gives Halpern the semantic flexibility needed to tackle the subtle problem of vagueness. A new logic is presented both syntactically and semantically. Halpern claims in addition that this formalism permits a resolution of the Sorites paradox.

Haim Gaifman also presents a formalism that can be used to solve the problem of vagueness. He focuses on the formalization of the notion of context. In a precursor of the paper for this issue Gaifman (manuscript) uses a logic of context to solve the Sorites paradox. The logic of context used in this paper was nevertheless a particular case of this type of logic. Gaifman (manuscript) suggested a generalization capable of treating contexts of other kinds. This is the idea that is pursued in the paper published in this collection. Most of the applications considered in the paper published in this collection are concerned with issues in the philosophy of language like the treatment of indexicals, demonstratives and proper names of natural language.

doi:10.1017/S1755020308090011

The logic of context presented in the paper makes an important contribution to the literature devoted to understand and apply the notion of context. The proposal amounts to enriching first order logic by adding *context operators* which apply to sentences (more generally well-formed formulas, and sometimes smaller units) and which represent contexts of different types. The result of applying a contextual operator to a sentence is to produce the relevant context for this sentence.

Let me now focus on the papers related to formal epistemology. A group of contributed papers focus directly or indirectly on the so-called problem of logical omniscience (the papers of Sergei Artemov, Rohit Parikh and Giacomo Sillari). The problem of logical omniscience is related to the deficiencies of the standard formalisms used to represent belief and knowledge. For example, most epistemic logics appeal to semantics formulated in terms of Kripke structures and therefore assume that if an agent believes (knows) ψ and $\psi \rightarrow \phi$, then the agent should believe (know) ϕ. In addition these logics endorse axioms stating that rational agents believe (know) all tautologies; and they also assume that a logically inconsistent formula can be neither known nor believed.

The ideal that agents believe (know) all the consequences of believed (known) facts is for sure too demanding to impose on real agents. But the idea can be turned into a reasonable doxastic (epistemic) requirement by stipulating that logical closure should be imposed over sets representing the *logical commitments* of rational agents. By the same token rational agents should be committed to believe (know) all tautologies and an agent can never be committed to believe (know) a logically inconsistent formula. So, as Levi (1983) first proposed, if the issue is to represent the doxastic (epistemic) commitments of rational agents, the problem of logical omniscience has less force (or perhaps vanishes).

Yet many philosophers, logicians and psychologists are interested in representing actual belief rather than doxastic commitments. And in this case the problem of logical omniscience reappears. Perhaps this problem is too unruly or complicated to be modeled via logical or probabilistic tools. But many of the authors writing about this problem in this issue think that logical or probabilistic models of actual belief (knowledge) are indeed possible.

Let me start with Rohit Parikh's essay. It focuses not only on the problem of logical omniscience but also on developing models of actual belief and knowledge which are sensible to this problem. Parikh defines two types of belief (knowledge). His definitions draw on the work of Ramsey (1990), de Finetti (1937) and Savage (1954) among others, so they are inspired by a tradition that is somewhat orthogonal to the one that motivated the development of various doxastic and epistemic logics in the second half of the twentieth century. One of the definitions is consistent with the assumption of some forms of logical omniscience, while the other (the notion of i-belief/knowledge) does not obey any of the constraints usually associated with logical omniscience. The definitions start with behavior and the notion of planning rather than focusing on logical systems obeying weaker axioms than the ones imposed in normal modal systems. Many issues remain open here. For example, it would be interesting to explore how the notion of partial belief works in this setting. As Parikh indicates, one would expect departures from the notion of probability axiomatized by Kolmogorov.

Parikh's model of belief seems to be related to (and perhaps inspired by) the efforts of many theoretical psychologists who have recently tried to accommodate and explain robust departures from Bayesian and logical standards of rationality (Kahneman, 2003). In the case of decision, Kahneman & Tversky (1979) have proposed a theory which they claim is empirically sound: prospect theory. Nevertheless, there are many instances of documented behavior that seem to violate prospect theory. On the one hand, Parikh's

proposal has the relative advantage of being quite flexible and open to adjustments in comparison with theories like prospect theory, which is quite regimented. On the other hand, some of the alleged departures from Bayesian standards might reveal the adherence to alternative normative patterns rather than constituting evidence that Bayesian norms of rationality are violated systematically in everyday use (some of the examples of failures of rationality offered by Parikh, like Linda's case, seems to be of this sort). Perhaps we are not witnessing cognitive errors in these cases, but the operation of sound patterns of rationality not recognized by the Bayesian canon. Parikh seems to be sympathetic with views of this sort, but his emphasis is on how to model cognition reliably, including cognitive errors.

Artemov's paper also focuses on the problem of logical omniscience, but as in the case of Parikh's paper, the article addresses broader issues in epistemology. The paper presents a logical framework, Justification Logic, capable of reasoning about epistemic justification. Justification Logic is based on classical propositional logic augmented by justification assertions t:F that read *t is a justification for F*. The theory that thus arises fills an important gap that has contributed to separate work in traditional epistemology from recent work in epistemic logic. In fact, most epistemic logics lack the capacity for representing and reasoning about justification, while the informal notion of justification plays a central role in traditional epistemology. Artemov presents a general Correspondence Theorem showing that behind each epistemic modal logic, there is a robust system of justifications. Also the increased expressive power of the formalism permits the direct analysis of paradoxes presented informally in the literature like the Goldman-Kripke Red Barn example as well as the traditional Gettier example and the Knower paradox.

Many issues remain open here as well. Perhaps one of the central open issues is the following: contemporary epistemology seems to have undergone a transformation in recent years according to which the central issue is not how to justify static bodies of belief[1] but how to justify changes of view. Is it possible to make Justification Logic more dynamic in such a way that it can be used to reason about changes of view? This is a central challenge to this position. Artemov and associates are currently considering this and related problems via extensions of their logical framework.

The offending inference in the case of Goldman-Kripke Red Barn example (reconstructed in Artemov's paper) is the one from $\Box(B \wedge R)$ to $\Box(B)$. This inference is sanctioned in all *normal* epistemic modal logics. This shows that the family of normal modal logics might not be the most adequate logical tool to formalize belief and knowledge. In spite of this, epistemic logicians have tried to force the assumption of normality in most epistemic logics. Artemov's formalism offers a general framework that circumvents the assumption of normality and therefore permits a general treatment of the problem of logical omniscience. But not all traditional epistemic logics are normal. One way of circumventing the normality assumption is to change semantics and adopt the so-called *neighborhood semantics* first proposed by Dana Scott and Richard Montague in the 1960s. For years the research program based on this alternative semantic framework has been neglected in favor of the more popular Kripkean-style semantics. But as I argued elsewhere (Arló-Costa & Pacuit, 2006; Arló-Costa, 2002) neighborhood semantics can be interpreted epistemically. One can see the neighborhood associated with a world as the explicit representation of the propositions believed (known) by an agent at that world. The idea can be robustly and fruitfully extended to the first order case.

[1] This is the type of epistemology that van Fraassen (1989) calls 'defensive epistemology'.

https://doi.org/10.1017/S1755020308090011 Published online by Cambridge University Press

But perhaps neighborhood semantics is better suited to study epistemic operators that fail to obey some forms of logical omniscience but retain others, like logical omniscience over logical consequence. A typical example is the notion of high probability that abandons the axiom that requires that $\{\Box(A), \Box(B)\}$ entail $\Box(A \wedge B)$; but retains other forms of logical omniscience. In particular all *classical* systems of modal logic obey an axiom establishing that $\Box(A) \leftrightarrow \Box(B)$ is entailed by $A \leftrightarrow B$. Sillari has proposed in his paper pairing neighborhood structures with other approaches to logical omniscience with the goal of making neighborhood models more adequate to represent actual belief and knowledge. The resulting approaches manage to circumvent all forms of logical omniscience.

The paper by Sillari elaborates on two logical frameworks that have been proposed to deal with the problem of logical omniscience, interpreted both over Kripke semantics and over neighborhood semantics. The two approaches considered by Sillari are awareness models and impossible world structures. The first approach was initially proposed by Fagin and Halpern. The idea is to distinguish between implicit and explicit knowledge, and avoids logical omniscience for explicit knowledge. The second approach was pioneered by Rantala and Hintikka. It postulates logically impossible worlds to which the agent has epistemological access. The two approaches are known to be equally expressive in propositional systems interpreted over Kripke semantics. Sillari shows that the two approaches continue to be equally expressive in propositional systems interpreted over Montague-Scott (neighborhood) semantics. In addition, Sillari extends this result to the first order case, building on results first presented in Arló-Costa & Pacuit (2006).

Helzner's paper is in part inspired by the work of Stalnaker (1999) in his recent book *Context and Content*. One of the central ideas of this book derives from the work of Grice in philosophy of language. The idea is to situate the analysis of speech acts within a wider account of rational activity. Helzner points out that Savage had insights that in this respect seem quite relevant to the ideas elaborated by Grice and Stalnaker:

> Whatever an assertion may be, it is an act; and deciding what to assert is an instance of deciding how to act. (Savage, 1954, pp. 159–160).

Helzner wants to reconnect this basic idea with a decision theoretic approach and at the same time he generalizes the approach to *all* acts not only speech acts. The main idea of the proposed approach is that '[. . .] the content of an act is the decision makers expectation concerning the change that would take place if the act were to be performed.' Under this account, an act of arbitrary type is assigned a proposition in an interpreted Boolean algebra of sets, although the elements of such sets are not possible worlds or states. Rather, '[. . .] the points that underly this algebra of sets are dynamic objects that are constructed over a collection of what are essentially experiments.' The resulting notion of expected content is shown to have close connections with the classical (and better known) concept of expected utility that has played a crucial role in the development of the decision sciences.

Finally we have Tennant's essay, which focuses on recent work on belief revision and applications to the analysis of conditionals (via ideas first advanced by Frank Ramsey). Gärdenfors (1988) proposed the use of epistemic models for the study of conditionals. His proposal utilized a test for acceptance of conditionals inspired by the writings of Ramsey (1990):

$$\text{(GRT) } \phi > \psi \in K \text{ if and only if } \psi \in K * \phi$$

The test proposes that a conditional $\phi > \psi$ is accepted with respect to a corpus K if and only if ψ is in the suppositional state generated by revising K hypothetically with ϕ. Gärdenfors' original idea was to study the properties of conditionals induced by his test

when the notion '$*$' that appears in it obeys the axioms of the so-called AGM theory of belief change (Alchourrón *et al.*, 1985). Gärdenfors, who helped to develop AGM theory, considered it our best approximation to solve the problem of belief change. With some provisos his judgment continues to be true today. But in an important paper Gärdenfors (1986) showed that this task is impossible. He showed that GRT and the following important AGM postulate are in conflict with each other:

$$\text{(Preservation) If } \neg\phi \notin K, \text{ then } K \subseteq K * \phi$$

As a matter of fact, as it is shown in Arló-Costa & Levi (1996), GRT is in conflict with the weaker postulate that we can call Open Preservation:

$$\text{(Open-Preservation) If } \neg\phi, \phi \notin K, \text{ then } K \subseteq K * \phi$$

David Lewis proved in 1976 the probabilistic counterpart of Gärdenfors' impossibility result by showing that the probability of conditionals cannot be captured by the corresponding conditional probabilities (Lewis, 1976). Lewis' reaction to the impossibility was to preserve a probabilistic version of the test and to change the underlying notion of conditional probability to a new notion called 'imaging.' A similar result holds for the qualitative test proposed by Gärdenfors. In fact, in Arló-Costa & Levi (1996) it is shown that Lewis's system **VC** can be obtained by conjoining GRT with the axioms of a qualitative version of Lewis' imaging.

A different reaction to the impossibility is to give up GRT and to adopt a form of the test that is compatible with AGM. Arló-Costa & Levi (1996) and Levi (1996) argued that it is possible to modify GRT in such a way as to preserve some of the central insights provided by Frank Ramsey and that this modified version of GRT is compatible with AGM. So, there is a qualitative version of the Ramsey test that is indeed compatible with AGM (for a gentle introduction to this topic see Arló-Costa, 2007).

In his article Tennant revisits part of the literature mentioned above, particularly Gärdenfors' original impossibility proof. He claims that it is indeed possible to have the Ramsey test, and what he considers are weak and reasonable demands on a system of belief revision. But the Ramsey test that he has in mind is not GRT but a weakened version of it, and the weak and reasonable demands on belief revision are based not on AGM but on a modification of it that he proposed in Tennant (2006). So, Tennant's point is that a modified (weaker) version of the Ramsey test is compatible with a modified version of AGM. The main idea seems to be to rethink the role of the Ramsey test in the context of a different theory of belief change. The weakened version of GRT that Tennant proposes is:

$$(R^{\dagger}) \text{ If } \phi \nvdash \bot,\ \psi \nvdash \bot,\ K, \phi \vdash \bot, \text{ then } \phi > \psi \in K \text{ if and only if } \psi \in K * \phi$$

This test only provides acceptance conditions for conditionals that are *counter-doxastic*. GRT (as well as imaging) entail the following important monotonicity property (which is in tension with AGM): If $K \subseteq H$, then $K * \phi \subseteq H * \phi$. $R^{\dagger}$ entails a weakened version of monotonicity that requires that ϕ is incompatible with K.

Tennant includes in his essay excerpts of a discussion with me where I present some objections regarding the new proposal. So here I do not need to rehearse the main ideas of this exchange. Yet there is perhaps an issue that should be stressed. Tennant thinks that one should develop a theory of revision, where *expansion* is not a subspecies of revision – the expansion of K with a sentence ϕ is just defined as taking $K + \phi = Cn(K \cup \{\phi\})$. So, his theory has no room for a Preservation postulate or even for the weaker postulate that we called above Open-Preservation (these postulates get trivialized in the new account – see

https://doi.org/10.1017/S1755020308090011 Published online by Cambridge University Press

the corresponding objection in Tennant's article). By the same token his modification of the Ramsey test does not provide any guidance as to how to accept 'If p, then q' when one is in doubt about p, even when this is the central case explicitly considered by Ramsey.[2] It seems that any test that tries to capture Ramsey's ideas should preserve this crucial part of the test. Tennant seems to be motivated instead by the second part of the Ramsey test, which addresses the case where the supposition is counter-doxastic. In this case Ramsey does not propose any concrete procedure aside from a passing reference to the fact that in this case the conditional ceases to mean anything for the interpreter and the only guidance is to consider what follows from certain laws or hypotheses.

Open-preservation, nonetheless, provides guidance as to what to do when one considers 'If p, then q' but one is in doubt about p. In this case one should not abandon any piece of information that is held as background information. It is exactly this postulate that is challenged by the methods of change compatible with GRT (like imaging). So, Open-Preservation occupies both a central role in AGM and in the theories of conditionals based on it. The notions of change compatible with GRT adopt a weaker version of Preservation (Weak Preservation) that stipulates that if ϕ is already believed in K, then $K = K * \phi$. Unfortunately, as we explained above, Open-Preservation (and Weak Preservation) is not even expressible in the modified theory of belief revision proposed by Tennnat.

Nevertheless, the consistency result mentioned in Tennant's paper and presented in Tennant (2006) invites students of belief revision conditionals to axiomatize the conditionals that are validated by $R^{\dagger}$. It is unclear whether the theory can be used to propose epistemic models for truth-value bearing conditionals of the sort studied by Lewis, for example. Many formal and conceptual issues remain open regarding Tennant's intriguing proposal.

BIBLIOGRAPHY

Alchourrón, C., Gärdenfors, P., & Makinson, D. (1985). On the logic of theory change: partial meet contraction and revision functions. *Journal of Symbolic Logic*, **50**, 510–530.

Arló-Costa, H. (2002). First order extensions of classical systems of modal logic. *Studia Logica*, **71**, 87–118.

Arló-Costa, H. (2007). The logic of conditionals. *Stanford Encyclopedia of Philosophy*. Available from: http://plato.stanford.edu/entries/logic-conditionals/

Arló-Costa, H., & Pacuit, E. (2006). First-order classical modal logic. *Studia Logica*, **84**(2), 171–210.

Arló-Costa, H. L., & Levi, I. (1996). Two notions of epistemic validity. *Synthese*, **109**(2), 217–262.

de Finetti, B. (1937). Foresight: its logical laws, its subjective sources. *Annales de l'Institut Henri Poincare*, **7**, 1–68. Translation by Kyburg, H. (1980) in *Studies in Subjective Probability* (Kyburg and Smokler, editors). Krieger Publishing Company, pp. 53–118.

Fagin, R., Halpern, J. Y., Moses, Y., & Vardi, M. Y. (1995). *Reasoning About Knowledge*. Cambridge, MA: MIT Press (a slightly revised paperback version was published in 2003).

Gaifman, H. Vagueness, tolerance and contextual logic. Manuscript, Columbia University. Available from: http://www.columbia.edu/∼hg17/, 41 pp.

[2] Ramsey says about this case: 'If two people are arguing If p will q? and are both in doubt as to p, they are adding p hypothetically to their stock of knowledge and arguing on that basis about q, [. . .]'.

Gärdenfors, P. (1986). Belief revision and the Ramsey test for conditionals. *Philosophical Review*, **95**, 81–93.
Gärdenfors, P. (1988). *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. Cambridge, MA: MIT Press.
Kahneman, D. (2003). A perspective on judgment and choice: mapping bounded rationality. *American Psychologist*, **58**, 697–720.
Kahneman, D., & Tversky, A. (1979). Prospect theory: an analysis of decisions under risk. *Econometrica*, **47**, 313–327.
Levi, I. (1983). *The Enterprise of Knowledge: An Essay on Knowledge, Credal Probability, and Chance*. Cambridge, MA: MIT Press.
Levi, I. (1996). *For the Sake of the Argument: Ramsey Test Conditionals, Inductive Inference, and Nonmonotonic Change*. Cambridge: Cambridge University Press.
Lewis, D. (1976). Probabilities of conditionals and conditional probabilities. *Philosophical Review*, **85**, 297–315.
Ramsey, F. P. (1990). *Philosophical Papers* (Mellor, D. H., editor). Cambridge: Cambridge University Press.
Savage, L. J. (1954). *The Foundations of Statistics*. New York: Wiley.
Stalnaker, R. (1999). *Context and Content*. Oxford: Oxford University Press.
van Frassen, B. (1989). *Laws and Symmetry*, Oxford, : Oxford University Press.
Tennant, N. (2006). New foundations for a relational theory of theory-revision. *Journal of Philosophical Logic*, **35**(5), 489–528.

DEPARTMENT OF PHILOSOPHY
CARNEGIE MELLON UNIVERSITY
E-mail: hcosta@andrew.cmu.edu