# Review of Hans Rott's *Change, Choice and Inference*

Gregory Wheeler

Artificial Intelligence Center - CENTRIA
Department of Computer Science, Universidade Nova de Lisboa
2829-516 Caparica, Portugal
greg@di.fct.unl.pt

We are currently witnessing a revival in formal approaches to epistemology. Researchers from several fields—philosophy, artificial intelligence, statistics, economics, and other branches of applied logic & mathematics—seem to be, in growing numbers, taking up core questions in epistemology. This interest in formal epistemology from researchers in various fields appears to be more than a passing trend, evidenced by an increasingly interdisciplinary make-up of the editorial boards of some prominent, longstanding journals, such as *The Journal of Philosophical Logic* and the new special section of *Synthese*, *Knowledge, Rationality and Action*, edited by Wiebe van der Hoek, along with several relatively young interdisciplinary journals that are having a growing impact on the field, such as *The Journal of Applied Logic* and *Logic, Language and Information*. The last few years have also witnessed an increasing number of genuinely interdisciplinary workshops and conferences, including the *Computation and Philosophy* conferences in Europe and North America organized by the *International Association for Computing and Philosophy* (IACAP), *Computational Models of Scientific Reasoning and Applications* (CMSRA), which targets formal epistemology and computation, and *Progic* and the ISIPTA *Imprecise Probabilities* workshops, each addressing themselves to logic and probability.

The link between Artificial Intelligence and formal epistemology in this network is particularly strong and has, in a certain sense, matured. Research of high-level cognitive reasoning is currently being driven both by practical concerns and by purely theoretical concerns. From a practical point of view, distributed computing and autonomous robotic agents offer two examples of research areas where there is increasing interest in improving the capabilities of agents (or distributed processes) to reason about what each 'knows'. From a theoretical point of view, the natural emphasis on formal semantics *and* syntax that computational modeling demands is generating new formal studies of concepts and relations that have been traditionally studied by philosophical logics. The result is an exciting breakdown of the barrier between applied and philosophical logic—a form of 'computational naturalism'—that is opening promising new lines of research.

One fundamental question for formal approaches to epistemology is the extent to which traditional epistemology should be engaged if not abandoned altogether. One of the many interesting ideas in Hans Rott's ambitious and very important *Change, Choice and Inference* (Oxford: Oxford University Press, 2001)

is his methodological proposal to view the theory of belief change as a necessary component to the proper conduct of traditional epistemology. Rott points out a familiar shortcoming to traditional analytic epistemology, namely a reliance on key notions—in Rott's example, Lehrer's notions of *personal justification* and *undefeated justification*—that are not well defined. But rather than cite this as grounds for abandoning traditional epistemology, Rott suggests that progress may be made by turning to a well-founded theory of belief change to understand some of traditional epistemology's key notions, including *justification* and *defeasibility*. Rott's program then is an attempt to unify traditional and formal epistemology by providing a fundamentally new, decision theoretic foundation for belief change within which to study key epistemic notions and relations.

The twin aims of Rott's program—namely, the thesis that traditional epistemology and the theory of belief revision are interdependent topics of research, and Rott's new foundation for belief change—may be read as direct replies to Isaac Levi's pioneering work in formal epistemology dating since the 1960s. Levi has long counseled us to abandon traditional epistemology's interest in the problem of skepticism, calling it 'an exercise in paper doubts', and has warned that our profession's fixation on language has blinded us to the fundamental pragmatic nature of decision and belief fixation—a fixation he calls 'The Curse of Frege'. So, it is important to understand the scope of Rott's project: by attempting to formulate rational choice within essentially a Stalnaker-Lewis-Chellas approach to semantics for conditionals (p. 170), Rott is proposing at once to place both the theory of belief change *and* non-monotonic reason on decision theoretic foundations and to place this entire framework squarely within the bounds of traditional, analytic philosophy. Levi himself has addressed these two points his most recent book, *Mild Contractions* (Oxford: Oxford University Press, 2004). This book is both the latest refinement of Levi's proposal to formulate the aim belief contraction to be the minimization of loss of informational value and a critical study of Rott's approach, which views contraction ('severe withdrawal') as a choice among the most plausible options open to an agent.

At this point it should be pointed out to general readers intrigued by the scope and ambition of this project that Rott's book is not easy reading. For while the first three chapters provide an informal introduction to Rott's project, the richest ideas are found in the last three chapters, which presuppose a working understanding of the main results in the field of belief revision and a familiarity with Levi's program to appreciate the novelty of Rott's ideas.

Unfortunately, too, the easiest to read chapter is the book's weakest. Chapter 2, which is designed to motivate traditional epistemologists to explore the theory of belief change, is a bit too quick to be effective: here Rott conflates reliabilism with externalism, provides a very quick gloss on Gettier cases, and does not exercise caution against collapsing the distinction between theories of justification and theories of knowledge when using the overloaded terms 'foundationalism' and 'coherentism'. Such errors are not likely to win over many traditional epistemologists. However, it should be pointed out to those readers that it would

be their loss: these mistakes are very minor points when compared to the scope and aim of this book.

Chapter 3 includes a very nice description of the belief change problem. Rott asks us to think of belief change as a two dimensional process, each with its own maxims for coherence. Along one dimension, the 'vertical dimension', static changes in belief occur. Vertical change operations are typically very simple: basic set-theoretic operations on set of propositions that are in turn closed under logical consequence. The other dimension of belief change, the 'horizontal dimension', is extended through time. Processes along the horizontal dimension represent changes in belief states. In contrast to the simple operations on propositions in the vertical dimension, the horizontal dimension includes sophisticated state-change operations. Rott's thesis is that choices, to be rational, must satisfy coherence maxims and that rational choice is rooted in a preference relation that may be regarded as primitive: 'In many cases, the reasoner can be ascribed some context-independent preference orderings that determine his choices, decisions and actions. The existence of underlying preferences is exactly what guarantees that his choice behavior is coherent' (p. 71). The remainder of Rott's book is organized to deliver the theoretical underpinning to this claim.[1]

The fourth chapter presents an overview of the various sets of rationality postulates to have been presented in the literature, for both AGM-style belief change operators and postulates for non-monotonic inference operators. A convenient table summarizes the relations between the postulates for these two frameworks. One minor complaint at this point involves the discussion of the recovery postulate, which, when introduced (p. 102), is described as 'the most controversial AGM axiom' (p. 103) without providing a sense of what the controversy is. An uninitiated reader may be frustrated by not having even a brief description provided for him.

Here then is a sketch of the controversy. The AGM model of belief revision provides postulates for three belief-state operations, one operation for adding a consistent formula to a belief state (expansion), another operation for removing a belief from a belief state (contraction) and a third, called revision, which in AGM is equivalent to a successive application of contraction followed by expansion, called the *Levi Identity*.

The Recovery Postulate states that removing a target belief from a belief state and then adding the very same belief back by expansion returns you directly to the same belief state. The important point here is that Recovery is presumed to *not* lose information. However, it appears that some applications

[1] A brief remark on the difference between Rott's approach and Levi's. On Levi's view, there is 'no rift between theoretical and practical rationality' (Levi 2004: 18) because all justification for change in belief is pragmatic in the sense that justification for belief fixation and change are rooted in strategies for promoting the goals of a given inquiry, Setting these parameters for a particular inquiry fixed the theoretical constraints for the inquiring agent. The important point to stress here is that there is no conflict between theoretical and practical reasoning on Levi's approach, since the prescriptions of Levi's theory are *not* derived from minimal principles of rational consistency or coherence (Levi 2004: 77).

of Recovery *would* lose information. One reason for thinking that recovery may lose information is that not all evidential relationships between beliefs are deductive. Consider this example: Suppose that an agent believes that his home is infested with wood worms, that a well-regarded exterminator, Stan, treats the house and that by a subsequent test, his home tests negative for wood worms. Suppose the agent also believes that the test is extremely reliable in the sense that close to 100 percent of homes having wood worms and receiving Stan's treatment then test negative for wood worms. If we contract the agent's beliefs by 'Stan treats the house', the replacement belief set would omit 'Stan treats the house' *along with the belief that the house tested negative*, since it is Stan's treatment that is the basis for the belief in the test's outcome. However, restoring the judgment that Stan treated the house would not return the belief that the house tested negative since this is not a logical consequence of Stan's treatment. Hence, expansion after contraction does not equal the same set of beliefs. Hence, the Recovery Postulate is not satisfied.[2]

The fifth chapter studies in detail a notion introduced in Chapter 3, namely how to make the direct mode of belief revision more concrete. The *direct mode of belief revision* takes as an operand a belief base, which is a set of basic (underived) beliefs that does not necessarily satisfy any (logical) coherence conditions. This view of belief revision allows for very simple belief revision operations but calls for more complex inference operations in order to yield a coherent belief *set*, which may be thought of here as a deductively closed belief *base*. Constructing such inference operations is no trivial undertaking. Rott's novel proposal is, roughly, that his direct mode operators take *prioritized* bases (sequences of ordered sentences) as operands. While input bases may be inconsistent, inference operators are proposed to exploit (and change operations are designed to preserve) this additional structure to yield a menu of 'optimal' ways to restore consistency.

The sixth chapter presents a compact introduction to classical rational choice theory and marks the beginning of Rott's argument for viewing the theory of belief change as resting on the foundations of the theory of rational choice and revealed preferences.

Chapter seven is the most significant contribution of Rott's book. It is often a feature of belief change frameworks and non-monotonic formalisms that some extra-logical selection function is used: for example, the AGM revision and contraction operators do not guarantee a unique revised theory, and there may be more than one extension for a classical (i.e., Reiter) default theory. What Rott does in this chapter is provide a connection between selection functions defined at the level of models (such as is done in the AGM literature) with selection functions defined on sets of sentences, such as is done in both default logic and the belief base change literature. Rott's correspondence results then allow him to provide a new foundation for belief change and non-monotonic reasoning in terms of rational choice.

[2] This example is a compressed version of examples offered by Isaac Levi. See (Levi 2004: 62-63; Levi 1991: 134-35).

The final chapter reformulates the relation of *epistemic entrenchment* in terms of revealed preference. Rott provides correspondence results between the postulates for epistemic entrenchment and the conditions for his formulation of rational choice. The idea here is that an entrenchment ordering is what rationalizes the set of selection functions that are called upon for a change of belief. This last chapter ends abruptly, immediately after the final correspondence results between postulates for epistemic entrenchment and rational selection functions are presented. Concluding remarks summarizing his contributions would have made the book stronger.

With this book—which, it should be noted, includes important joint work with Maurice Pagnucco (Pagnucco and Rott 1999)—Rott offers us a new conception of what is required of a decision-theoretic rational for belief change. The result is a very significant contribution that will be required reading for anyone with an interest in belief revision. The book is also a very significant contribution to formal epistemology in general, since the clash between this program and Levi's exercises fundamental concepts in formal epistemology and draws direct attention to the central methodological questions concerning how formal concepts should be applied to philosophical problems. For this reason, it should be required reading for all of us working in formal epistemology.[3]

## References


[1] Horacio Arló-Costa and Isaac Levi, 2004. A complete characterization of a notion of contraction based on information-value, in *Proceedings of 10th International Workshop on Non-monotonic Reasoning (NMR-2004)*, James Delgrande and Torsten Schaub (eds.), Whistler, BC: 25–33.

[2] Isaac Levi. 1999. *The Fixation of Belief and Its Undoing*, Cambridge: Cambridge University Press.

[3] Isaac Levi. 2004. *Mild Contraction*. Oxford: Oxford University Press.

[4] Maurice Pagnucco and Hans Rott. 1999. "Severe Withdrawal—and Recovery", *Journal of Philosophical Logic*, 28: 507–547. Note: Reprinted with corrections in *JPL* vol. 29, 2000.


[3] Thanks to Horacio Arló-Costa for his generous and insightful correspondence on epistemic value.