Means-Ends Rationality and Categorical Imperatives in Empirical Inquiry

Review of *Elements of Scientific Inquiry*, E. Martin and D. Osherson (MIT Press 1998).

Kant taught us that there are two kinds of norms: Categorical imperatives that one ought to follow regardless of one's personal aims and circumstances, and hypothetical imperatives that direct us to employ the means towards our chosen ends. Kant's distinction separates two approaches to normative epistemology. On the one hand, we have principles of "inductive rationality", typically supported by considerations such as intuitive plausibility, conformity with exemplary practice, and internal consistency. On the other hand, we may assess rules for forming belief by how well they attain the objectives that motivate inquiry; in Levi's words, "the ends of inquiry control the legitimacy of inferences" [Levi 67, p. 241]. A doctrinaire attitude would ignore one of these perspectives in favour of the other; a balanced approach is to develop both and compare [cf. Helmann 97, Sec.2]. There are three possible relationships between hypothetical and categorical imperatives for empirical inquiry.

1. The categorical imperative will lead an inquirer to obtain his cognitive goals. In that case means-ends analysis vindicates the categorical imperative.

2. The categorical imperative may prevent an inquirer from achieving his aims. In that case the categorical imperative restricts the scope of inquiry.

3. Some methods meet both the categorical imperative and the goals of inquiry, and others don't. Then we may well prefer the former over the latter.

Which of these three cases obtains depends on details such as: the operative methodological norms, the questions under investigation, the background assumptions that the agent brings to inquiry, the observational means at her disposal, her cognitive capacities, and her epistemic aims. For a given constellation of these factors, two questions arise: (1) What are the optimal investigative methods for attaining the goals of inquiry? (2) What is the relationship between epistemic aims and norms of "inductive rationality"? *Elements of Scientific Discovery* presents a systematic set of answers to these questions that covers many aspects of empirical inquiry and scientific rationality.

A prominent norm for belief change is Bayes' rule – the principle that inductive inquiry should incorporate new evidence e by updating a subjective "prior" probability p with the probability measure p conditional on e. {footnote:

Formally, the principle is that the updated probability of an assertion h should be p(h and e)/p(e) assuming that p(e)>0.} Suppose that a Bayesian agent is investigating a hypothesis h. If the Bayesian successfully makes use of the incoming observations, his updated degrees of belief in h will approach or reach 1 if h is true, and 0 if h is false. If the Bayesian succeeds in every possible world consistent with given background knowledge, we may say that his prior probability p is objectively reliable [Juhl 97]. Reliability imposes an objective constraint on degrees of belief, of the sort that many methodologists have sought in answer to the charge that Bayesian inferences are subjective and hence unfit for scientific objectivity [Jeffreys 61, Fisher 93 pp.6-7, Howson 97, Maher 96, Sec. 3.3]. However, this happy marriage between reliability and Bayesian updating can take place only when objectively reliable prior probabilities exist. We cannot expect Bayesian updating to accomplish the imposssible; but in situations in which some inference rule is guaranteed to lead to the truth, is there always a Bayesian agent who can match this performance? Earman [92, Ch.9, Sec.6] conjectured that the answer is yes. Juhl [97] provided a partial confirmation of Earman's conjecture. He proved that if there are only two potential evidence items (e.g., "particle discovered" vs. "particle not yet found"), and there is some method guaranteed to settle on the truth about a given hypothesis h, then there is an objectively reliable prior for investigating h. It remains an important open question whether Juhl's theorem can be extended to situations in which there are more potential observations.

Martin and Osherson take a dimmer view of the potential of Bayesian inquiry. One way to describe their argument is in terms of Lewis' Principal Principle [Lewis 86]. The Principal Principle implies that if an agent knows the objective chance c of an event, then the agent's degree of belief in the event ought to be exactly c. Consider the following extension of the principle: Suppose that the agent knows that the objective chance of certain events is governed by one of a set C of chance distributions. Then the agent's subjective probability concerning the events should be a member of C. Martin and Osherson give an example of an hypothesis h such that (1) some rule for investigating h is guaranteed to eventually arrive at the truth, but (2) for every Bayesian agent, the world might be such that h is false, and yet the Bayesian will infinitely often assign almost maximal belief to h – if the Bayesian's prior probability measure incorporates background knowledge about certain probabilistic independencies between events [Sec.3.6.9]. This result shows that even when there is an objectively reliable subjective prior, there may not be one that satisfies the Extended Principal Principle.

Much of *Elements of Scientific Discovery* focuses on the maxim that an agent ought to accommodate new evidence with a minimal change in his beliefs. Epistemologists such as Quine [51], Levi [80] and Harman [86] have endorsed various versions of this norm. A recent branch of philosophical logic known as belief revision theory has endeavoured to give precise content to the notion of minimal change. This work has led to a set of axioms known as the AGM postulates [Gardenfors 88]. A characteristic AGM postulate is the "preservation principle": if an agent's beliefs are consistent with the evidence, the agent ought not to give up any of his beliefs. Putnam used Peirce's vivid term "tenacity" to refer to this precept.

At first glance, tenacity would seem to hamstring the quest for truth. Consider the following example. A particle physicist wants to investigate whether a certain type of particle exists (not a hidden particle). Suppose that her initial conjecture is that the particle exists. Now what will happen if there is no such particle and the inquirer revises her beliefs according to the AGM theory? Since she is tenacious, she will not give up her belief in the existence of the particle no matter how many experiments fail to turn it up.

Popper had a remedy for such inadequate performance: Start with hypotheses that are falsifiable. More precisely, if minimal belief change is to eventually eliminate false beliefs, it must adopt beliefs that, if false, will eventually be inconsistent with the evidence. {footnote For a discussion of this and other notions of falsifiability, cf. [Schulte and Juhl 97].} In our particle example, if the inquirer adopts as her initial hypothesis that the particle does not exist – in the spirit of Occam's Razor – then tenacity is compatible with the quest for truth. For if the particle exists, its eventual discovery will falsify the agent's initial hypothesis. If it does not exist, the inquirer's hypothesis will never be falsified, and being tenacious, she will not give it up. In either case, she settles on the correct belief about the existence of the particle.

In general, reliable belief revision via minimal changes looks like a Popperian conjectures-and-refutations scheme, in which the inquirer adopts falsifiable beliefs and hangs on to them until they are logically inconsistent with the data, in which case he adopts other falsifiable beliefs, etc. One of the major achievements of *Elements of Scientific Discovery* is a proof that whenever (just about) any kind of inductive method can reliably arrive at correct beliefs in the long run, then so can a method that satisfies the AGM postulates for minimal belief change [Sec. 4.2.4]. This is a remarkable theorem in formal epistemology that draws on powerful mathematics (such as well-orderings for high-cardinality sets). In addition, the book presents a rich set of results about various assumptions under which there are universal aspects of reliable belief revision that need not depend on the particular context of inquiry.

Several epistemologists have suggested that realistic accounts of rational scientific inquiry should be theories of "bounded rationality". *Elements of Scientific Discovery* examines the scientific potential of inquirers with two types of cognitive limitation: bounded memory and bounded logical powers.

An inquirer's memory of the evidence may be limited to some fixed constant amount; as an extreme case, an agent may take into account only the last datum. Martin and Osherson show that in a large class of inductive problems, if any agent reliably settles on a correct hypothesis, then so does an inquirer who revises his beliefs in a minimal way considering the last datum only [Sec. 4.6.2]. They also investigate the empirical potential of memory-bounded agents whose belief changes need not be minimal [Sec. 2.4].

A rich body of work known as computational learning theory studies inquirers whose cognitive powers are those of a Turing computer. *Elements of Scientific Discovery* provides a concise treatment of some of the basic results [Sec. 2.5, 3.5]. Often categorical imperatives that do not prevent logically omniscient agents from reliably finding the truth limit the scope of inquirers with bounded computational capacities. For example, consider the seemingly innocuous norm of consistency: Believe that h is false as soon as the evidence is logically inconsistent with h. The consistency principle is part of both Bayesian confirmation theory and AGM belief revision. Martin and Osherson describe an empirical question that can be reliably solved by a computer – but not if the computer satisfies the consistency principle. Kelly and Schulte [95] show that consistency prevents even agents with infinitely uncomputable cognitive powers from reliably assessing certain hypotheses. These examples show that if a theory is sufficiently complex, agents who are not logically omniscient may be unable to determine immediately whether a given piece of evidence is consistent with the theory, and need to collect more data to detect the inconsistency. But the consistency principle – and a fortiori, Bayesian updating and AGM belief revision – rule out this kind of scientific strategy.

Elements of Scientific Discovery addresses many other topics in formal epistemology. For example: Some empirical questions are so complex that no method of inquiry is powerful enough to provide a guarantee of converging on the right answer. What are necessary and sufficient conditions for reliable inquiry? [Sec. 2.2, 3.2] The authors define a precise sense in which inquiry arrives at the truth as quickly as possible. What inductive methods are as fast as possible? [Sec. 2.3, 3.4] Can minimal belief changes lead to the right answer as quickly as possible? (The answer is generally, yes.) [Sec. 4.4, 4.6.2] How can inquiry actively direct observations to find the truth faster than with passive observations? [Sec. 3.4.3] What is the relationship between methods guaranteed to arrive at the truth under given background assumption, and methods that have only a high probability of doing so? [Sec. 2.7, 3.6]

This book is both a research monograph and an introduction to formal learning theory, the authors' mathematical framework. The pace of the book is good for an introduction. The authors introduce each topic with helpful explanations, establish the basic facts, and give references to further reading. Those new to learning theory could use more informal motivation, and diagrams, to go with the formal definitions. The authors provide examples for almost every concept they introduce; however, these are mostly of an abstract mathematical sort. Most helpful would be examples that relate the mathematical apparatus to familiar kinds of evidence and empirical hypotheses. Philosophers of science will want to see how the theorems reported apply to questions that arise more directly from scientific practice than formal epistemology. Before tackling *Elements of Scientific Discovery*, readers may wish to consult literature that introduces the learning-theoretic approach to inductive inference at an accessible level of technicality, and that discusses applications of means-ends epistemology to questions pertaining to scientific practice. There is no shortage of such publications; a selective list is: [Glymour 91], [Kelly 96, Chs. 2,3] [Glymour and Kelly 92], [Kelly et al. 97], [Schulte and Juhl 97], [Glymour 94], [Bub 94].

Means-ends rationality and pure norms for inductive inference constitute two distinct perspectives on how scientific inquiry ought to proceed. The relationship between these two perspectives is a central issue for normative epistemology. *Elements of Scientific Discovery* is a mathematically sophisticated analysis of both approaches, rich in fundamental and often surprising insights. It sets a milestone in formal epistemology that is a fruitful starting point for further studies of rational empirical inquiry.

Oliver Schulte
Department of Philosophy
University of Alberta
4-89 Humanities Centre
Edmonton, AB T6G 2E5
Canada

References

- Bub, J. [1994]: `Testing Models of Cognition Through the Analysis of Brain-Damaged Performance', *British Journal for the Philosophy of Science*, 45, pp.837-55.
- Earman, J. [1992]: *Bayes or Bust?*. Cambridge, Mass.: MIT Press.
- Fisher, R.A. [1993]: *Statistical Methods, Experimental Design and Scientific Inference*, Oxford: Oxford University Press.
- Gardenfors, P. [1988]: *Knowledge In Flux: modeling the dynamics of epistemic states*, Cambridge, Mass.: MIT Press.
- Glymour, C. [1991]: `The Hierarchies of Knowledge and the Mathematics of Discovery', *Minds and Machines*, 1: pp. 75-95.
- Glymour, C. [1994]: `On the Methods of Cognitive Neuropsychology', *British Journal for the Philosophy of Science*, 45, pp. 815-35.
- Glymour, C. and Kelly, K. [1992]: `Thoroughly Modern Meno', in: *Inference, Explanation and Other Frustrations*, ed. John Earman, University of California Press.
- Harman, G. [1986]: *Change in View: Principles of Reasoning*, Cambridge, Mass.: MIT Press.
- Hellman, G. [1997]: `Responses to Maher, and to Kelly, Schulte and Juhl', *Philosophy of Science*, 64, pp. 317-322.
- Howson, C. [1997]: `A Logic of Induction', *Philosophy of Science*, 64,

pp. 268-90.

- Jeffreys, H. [1961]: *Theory of Probability*, 3rd edition, Oxford: Clarendon Press.
- Juhl, C. [1997]: `Objectively Reliable Subjective Probabilities', *Synthese* 109, pp. 293-309.
- Kelly, K. [1996]: *The Logic of Reliable Inquiry*, Oxford: Oxford University Press.
- Kelly, K. and Schulte, O. [1995]: `The Computable Testability of Theories Making Uncomputable Predictions', *Erkenntnis* 43, pp. 29-66.
- Kelly, K., Schulte, O. and Juhl, C. [1997]: `Learning Theory and the Philosophy of Science', *Philosophy of Science*, 64, 245-67.
- Levi, I. [1967]: *Gambling With Truth*, New York: Alfred Knopf.
- Levi, I. [1980]: *The Enterprise of Knowledge*, Cambridge: MIT Press.
- Lewis, D. [1986]: `A Subjectivist's Guide to Objective Chance', in: *Philosophical Papers*, Vol. 2, Oxford: Oxford University Press.
- Maher, P. [1996]: `Subjective and Objective Confirmation', *Philosophy of Science* 63, pp. 149-174.
- Putnam, H. [1963]: ``Degree of Confirmation' and Inductive Logic', in: *The Philosophy of Rudolf Carnap*, ed. A. Schilpp, La Salle, Ill: Open Court.
- Schulte, O. and Juhl, C. [1997]: `Topology as Epistemology', *The Monist*, 79, pp.141-48.
- Quine, W. [1951]: `Two Dogmas of Empiricism', *Philosophical Review* 60, pp. 20-43.