Machine Guessing – I

David Miller
Department of Philosophy
University of Warwick
COVENTRY CV4 7AL UK
e-mail: dwmiller57@yahoo.com

Abstract

According to Karl Popper, the evolution of science, logically, methodologically, and even psychologically, is an involved interplay of acute conjectures and blunt refutations. Like biological evolution, it is an endless round of blind variation and selective retention. But unlike biological evolution, it incorporates, at the stage of selection, the use of reason. Part I of this two-part paper begins by repudiating the common beliefs that Hume's problem of induction, which compellingly confutes the thesis that science is rational in the way that most people think that it is rational, can be solved by assuming that science is rational, or by assuming that Hume was irrational (that is, by ignoring his argument). The problem of induction can be solved only by a non-authoritarian theory of rationality. It is shown also that because hypotheses cannot be distilled directly from experience, all knowledge is eventually dependent on blind conjecture, and therefore itself conjectural. In particular, the use of rules of inference, or of good or bad rules for generating conjectures, is conjectural. Part II of the paper expounds a form of Popper's critical rationalism that locates the rationality of science entirely in the deductive processes by which conjectures are criticized and improved. But extreme forms of deductivism are rejected. The paper concludes with a sharp dismissal of the view that work in artificial intelligence, including the JSM method cultivated extensively by Victor Finn, does anything to upset critical rationalism. Machine learning sheds little light on either the problem of induction or on the role that logic plays in the development of scientific knowledge.

0 Introduction

There is an antiquated view about logic, originating perhaps with Aristotle, that reasoning in mathematics is bound by the canons of deduction, while reasoning in the empirical sciences is bound by the quite distinct canons of induction. For all their differences, however, deductive logic and inductive logic are both departments of the science of logic and are held to share three important characteristics: both are *formal* disciplines, in the sense that, for the most part, their principles do not vary with the subject matter, be it mathematical or empirical, to which they are applied; both are *creative* disciplines, in the sense that they furnish means by which we may extend our knowledge from what, if anything, is given in intuition, or in experience; and both are *regulative* disciplines, in the sense that they authorize or legitimize or justify their conclusions, endowing them with a seal of reliability and credibility, if not a guarantee of apodeictic certainty.

A paper presented at the conference J. S. Mill's Ideas on Induction and Logic of the Humanities and Social Sciences in the Cognitive Research and in Artificial Intelligence Systems, held at the Russian State University for the Humanities (URSS), Moscow, in June 2011. The central theme of the meeting was the work of Victor Finn in the theory of intelligent systems, which is briefly discussed in § 7 below. A Russian translation by Delir Lakhuti has been published in *Voprosy Filosofii* 7, pp. 110–119, and 8, pp. 117–126, in 2012. My thanks are due to Ali Paya, who drew my attention to Searle (1999), and to Alain Boyer for some useful hints.

The doctrine that deductive inferences are peculiar to mathematics is now wholly discredited. What was, in my judgement, Mill's most potent contribution to logic, his argument that the conclusion of a syllogism (and, by extension, any deductive inference) only repeats, partly or wholly, what is stated in the premises, and therefore cannot be better secured by the premises than it is secured by itself (1843, Book II, Chapter III, § 1), shows that deductive logic, though admirably formal, has no creative or regulative force. Even within mathematics, deductive inferences are not demonstrations, proofs of their conclusions, but only derivations, capable of transmitting truth but not of introducing it. The term *non-demonstrative inference* is nonetheless used exclusively for deductively invalid inferences, not only those inferences traditionally called inductive but also those called abductive. It defines an area in which the doctrines mentioned in the previous paragraph have become permanent possibilities of dissension. The difficulties exposed by Goodman (1954) have revealed a high cost in maintaining that inductive logic consists of purely formal rules. While allowing that the conclusion of an inference that generalizes from instances (induction by enumeration) may be credible, indeed compelling, Hume (1738), Book I, Part III, § xii, subverted the idea that this credibility is of other than psychological significance, even if the empirical premises are beyond dispute. Although some philosophers may have freed themselves from Hume's sceptical doubts, they have not identified clear errors in his premises or his argument, and the justification of induction remains an insubstantial dream. Yet there is growing support for the idea that non-demonstrative inferences may be both rational and creative, and that a working logic of discovery provides a way of vanquishing humean scepticism. This prospect has been popularized especially by workers in the domain of artificial intelligence.

The objective of this paper is to instil disenchantment. Non-demonstrative inference is either simple guesswork, untouched by reasoning, or the application of procedures that have themselves been guessed to be congenial to the task. The role of logic in scientific discovery is the purely negative one of abetting the elimination of bad guesses. The development of learning machines, in brief, may constitute an impressive engineering feat, and may provoke metaphysical speculation, but it is devoid of methodological significance. Hume's problem is not solved by circumvention.

My own view is that Hume's logical problem of induction is best solved by renouncing the prejudice that knowledge needs justification, and embracing wholeheartedly, but not slavishly, Popper's non-justificationist *critical rationalism*. In § 1, I shall say a little about why justificationist solutions to the problem flounder, and about why the problem is not wisely disregarded. The aim of § 2 is to put a brake on the popular doctrine, embraced by Hume, that although induction is logically unsound, it is a fact of psychological life, and must be properly acknowledged. In § 3, I note that it is guesswork that is a fact of life. In § 4, I provide a summary of my version of critical rationalism, and in § 5, I distinguish critical rationalism from other positions that might well be called deductivism. For more details, and especially more defence, Miller (1994), Chapters 1–3, (2006), Chapters 3–5, and (2007b) may be consulted. I shall then look more closely at the submissions made on behalf of machine learning. In § 6, I shall show what is mistaken in Gillies's thesis that the successes achieved in the development of artificial intelligence cast a dark shadow on the critical rationalist dismissal of inductive inference; and finally, in § 7, I shall look at Finn's claims on behalf of his JSM method, which is said to offer 'a synthesis of induction, analogy and abduction' (2002). I hope to show why it too is methodologically inert.

1 The problem of induction

What Carnap, writing in 1968 about inductive logic, called 'the grand old problem of justification, discussed so much since Hume's time', is now discussed not so much by philosophers, except in the classroom. There are several pretexts for this inattention. Carnap himself declared that 'the old problem is not the right way of asking the question' (p. 258). He proceeded (pp. 259f.):

> Many believe that the essence of inductive reasoning is to make inferences of a special kind [, so that] the demand for justification would naturally be directed at any alleged principles of inductive inference. I would not say that is wrong to regard inference making as the aim [b]ut ... it seems preferable to take as the essential point of inductive reasoning the determination of probability values. ... Inductive probability is related to degree of belief, as Ramsey explained .... [I]n inductive logic we are ... concerned with the *rational* degree of belief. Thus we have to consider the question: 'how can we find principles which guide us to reasonable degrees of belief?'

Austere personalist Bayesians regard even this question of Carnap's as asking too much; what matters are only reasonable principles for the coherence of degrees of belief (or credences) and for their modification in the light of evidence. Other personalists, such as Lewis (1980), think that some non-extreme credences can be rationally evaluated. Credence management, which consists entirely of deductive moves within the theory of probability, is often, tendentiously, called *inductive inference*. But since Carnap had little time, and more personalist thinkers have no time, for inductive inferences with factual conclusions, it is fair to say that Hume's problem has been put to bed unsolved (Miller 1994, Chapter 6, § 5). Despite much talk of the justification of belief, as in the title of Howson (2000), Bayesians have not told us how beliefs can be justified.

Popper too, of course, thought that the traditional philosophical problem of induction, that is to say, the question 'What is the justification for the belief that the future will be (largely) like the past? Or, perhaps, What is the justification for inductive inferences?' is 'wrongly put' (proem to 1971). His objection to these formulations was that they appear to assume first that the future will be (largely) like the past, which is a big assumption, and secondly that there exist inductive inferences that are in need of justification. Popper's well known view, here endorsed, is that there exist no such inferences, and hence there is no need to agonize about justifying them. What there are instead are conjectures, or guesses, or hypotheses, which are not generated by any cognizable procedure, and refutations, which are generated deductively. Neither conjectures nor refutations are justified, or justifiable, but rationality is not the same thing as justification.

Some authors who push aside the traditional problem of induction do so even though, unlike Carnap and Popper, they seem not to know where to go next. Papineau (1995), pp. 4f., writes:

> It is true that induction presents an abstract philosophical puzzle. Inductive inferences are not logically compelling, and because of this their ultimate authority is an issue of philosophical controversy. But this is a puzzle, not the start of a philosophical system. It is akin to the question, 'How do I know there is a table in front of me?' This is a good issue for first-year philosophy students to cut their teeth on. But outside the classroom nobody seriously doubts that we do know about tables.

Similarly scornful of the value of abstract philosophical inquiry, Searle (1999), pp. 2079f., writes:

> I believe that epistemic problems, [such as] 'How is it possible that we can have knowledge at all in the light of the various sceptical paradoxes?', should be regarded in the same way as other such paradoxes have been regarded in the history of philosophy. Zeno's paradoxes about space and time, for example, pose interesting puzzles, but no one supposes that we cannot seriously attempt to cross a room until we have first answered Zeno's scepticism about the possibility of moving through space. Analogously, I believe, we should have the same attitude towards the paradoxes about the possibility of knowledge that were advanced by sceptical philosophers. That is,

> these are interesting puzzles, and they provide good five-finger exercises for training young philosophers, but we should not suppose that the possibility of knowledge and understanding rests on our first being able to refute Hume's scepticism.

It seems to have occurred to neither of these writers that there might be something valuable to be learnt from an attempt to resolve the problem of induction, just as much was learnt from attempts to resolve Zeno's paradoxes. Both Papineau ('science can never yield any positive findings') and Searle ('the scientist does not arrive at truths about nature') attribute to Popper a depiction of science as an enterprise with no positive features. Had they interested themselves in his solution to Hume's problem, they might have learnt how wrong they are: that science consists of both conjectures and refutations, and whereas the latter are negative, the former are affirmative — not in the authoritarian sense of being justified, which both Papineau and Searle evidently see as the only standard of knowledge, but in the sense of being claimants to the truth.

According to Russell, Hume's 'destruction of empiricism' rested 'entirely on his rejection of the principle of induction', the principle, in short, that exceptionless regularities observed profusely in the past will probably be abode by on the next occasion (1946, p. 699). Moreover,

> [i]f this principle is not true, every attempt to arrive at general scientific laws from particular observations is fallacious, and Hume's scepticism is inescapable for an empiricist. The principle itself cannot, of course, without circularity, be inferred from observed uniformities, since it is required to justify any such inference. It must therefore be, or be deduced from, an independent principle not based on experience. To this extent Hume has proved that pure empiricism is not a sufficient basis for science.

But Russell was mistaken here. It was not the weakness of empiricism that Hume exposed, but the presumption of justificationism, the idea that the only respectable way to 'arrive at general scientific laws', or other general statements, is by inference, rather than by guesswork. True or not, empirical or metaphysical, the principle of induction, and any 'independent principle' from which it is deduced, remain unjustified, and cannot provide justification in their turn. We must look elsewhere for the rational element in science, since nothing can give it 'a sufficient basis'. A non-justificationist theory of rationality is needed, and this is what critical rationalism supplies.

A recent example of the strategy recommended by Russell is Zahar (2007), Chapter II, who introduces a (metalinguistic) principle of induction, admitted to be both 'unjustifiable and uncriticizable', that, he says, 'allows us to assess science' (p. 19). No doubt fundamentalist religious zealots can suggest alternative principles of assessment. Zahar reports Popper's position, as others do, through the eyes of someone who does not know that critical rationalism has no use for justification. Discussing Popper's approach in (1971), § 9, to the problem of how scientific theories are used in practice, but without reference to the improved treatment in (1974b), § 14, Zahar infers from Popper's declaration that '*there can be no good reasons* ... for expecting that it [a theory] will in practice be a successful choice' (§ 9) that, according to Popper, 'non-rationality has to be admitted into the sciences ... at the level of technology' (p. 10). It is remarkable that Popper's appeal to the non-justificationist theory of rationality sketched in Chapter 24 of his (1945), could be admonished in 2007, as it is by Zahar (p. 15), for being 'somewhat disingenuous'.

2 Why learning by repetition is impossible

What is the source of the general ideas, the universal hypotheses, that individual agents are in possession of? There seem to me to be four possible answers: (i) they are prompted directly by perceptual experience; (ii) they are inherited genetically, or suggested by teachers; (iii) they are the product of pure guesswork; and (iv) they are inferred from more primitive ideas (perhaps themselves originated by processes (i)–(iii)). It should be evident that (ii) cannot be the whole answer, and I shall not consider it apart from (i). My purpose in this section is to confound (i).

The empiricist doctrine that we learn wholly from experience has to be rejected; not because of its emphasis on experience, about which I shall say something more in a moment, but because of the direction implicit in the phrase 'from experience'. Experience is not given in experience; it is not another name for interaction with something else. That much is familiar, and is by no means controversial. Kant decided that at least the form of our experience, though not its content, is given prior to experience. A more modern version of this doctrine has it that our sense organs, being largely pre-formed, impose not only some form, but also much content, on what is experienced. The frog, who sees only moving objects, is frequently cited as an example.

A more biologically informed empiricism may therefore hold that although our knowledge is not fully explicable in terms of our encounters with the external world (even supposing that there is a sharp distinction between what is internal and what is external), it is fully explicable by the joint action of our experiences and our inheritance. Inherited knowledge delivers expectations, universal hypotheses, generalizations, statements of regularities (including statistical regularities); experience delivers singular statements, interpretations of experience in accordance with those expectations. Working together, these sources could perhaps provide enough knowledge for survival and reproduction, enough knowledge to maintain life at a low and almost mechanical level. But it is incredible to suppose that all human knowledge of regularities is inherited knowledge. Even when experience is recognized not to be the unstructured starting point once imagined, empiricism is driven to invent some mode of generalization from perceptual experience.

This is induction, still popularly thought to be an essential factor in our learning about the world. It is admitted, of course, that induction, whatever psychological process it is, is a fallible process, and that what is learnt by induction may be wholly incorrect. Logically the situation is fairly clear: induction, whatever it is, is invalid. What is much less clear is what induction is.

The difficulty in making sense of induction as a mode of generalization from experience is a difficulty in working out what exactly it is that is being generalized. An experience may be a singular event in space-time, but that is not something simple; on the contrary, it is so complex that there is no saying what might be a repetition of it (Popper 1957a, § IV). There are any number of ways of generalizing from a particular experience or set of experiences, each yielding different results, and it is therefore absurdly tendentious to represent induction as a determinate method of generalizing from experience. Induction, that is to say, is not a specific process, but at best a label for innumerable unrelated acts of generalization (as the problem of curve fitting, and Goodman *ibidem*, suggest). It is apparent too that even when the matter of its validity is not an issue, the usual problem concerning the status of induction does not entirely go away. Induction cannot be a skill learnt from experience, which means that it is an inherited skill, which means that there must at least exist in most organisms an inherited disposition to generalize, but not to generalize in any particular way (Popper 1974a, § 10). Organisms, that is, inherit not only expectations that express regularities; they inherit also an expectation of unspecified regularity.

If this is so, then induction, thought of as a mode of learning rather than as a mode of inference, is arbitrary not because it moves from the singular to the universal but because it moves from the existential to the singular. Let a be some action that an agent performs that is followed by some advantageous outcome c. How is the agent to learn from this success? If the position presented here (which has something in common with the 'fourth [metaphysical] stage of the problem of induction' in Popper 1983, Part I, § 5) is correct, then what is required is a transition from the expectation of regularity (which may be expressed by 'there exists a regularity here') and the experience of action a's being followed by outcome c, to the expectation that a's being followed by c is an instance of a particular regularity; in the formal mode of speech, 'actions of type X [such as a] are invariably followed by outcomes of type Z [such as c]'. The hard part of this process is the categorization of the action a and the outcome c as instances of the types X and Z respectively. What should be obvious is that the appropriate types, if there are any, are not learnt from the experience of a and c. Of course, some anticipation (correct or incorrect) of which types are appropriate may be inherited, and may be applied automatically. In the limiting case, there is only a response to experience. The phenomenon of imprinting, discovered by Lorenz, is a perfect example (though its irreversibility may not be characteristic).

The logical problem of induction is a problem that arises when the events under investigation are correctly (if not appropriately) sorted into types, but there is no premise that implies the existential statement 'there exists a regularity here'. There is no valid move from 'an action of type X here was followed by an outcome of type Z' to 'actions of type X are always followed by outcomes of type Z'. Induction is logically bankrupt. The psychological (or, better, the biological) problem of induction, for its part, is best seen as a problem that arises when the initial situation is the other way round. The premise 'there exists a regularity here' is in place, and what is missing is an appropriate categorization into types of the events experienced. In situations crucial to survival and reproduction, inheritance may provide clues, but if it does not, then no profitable expectation is formed. Induction is biologically and psychologically bankrupt.

The inference from an existential statement $\exists yFy$ to an instance Fb is of course logically equivalent to the inference from the instance $\neg Fb$ to the universal statement $\forall y \neg Fy$. To avoid misunderstanding, I should stress that the inferences *existential* $\rightarrow$ *singular* and *singular* $\rightarrow$ *universal* considered in the last two paragraphs do not stand in the relation of equivalence to each other.

Having said all this, I must admit that the previous two paragraphs present the shortcomings of induction too ingenuously. For in both ways of looking at the problem, logical and biological, the invalid move or inference in question is not validated by a premise expressing, or an expectation of, the existence of regularities. What is needed is at least a premise to the effect that there exists a regularity connecting events of one of the mutually exclusive and effectively discriminable types $X_0, \ldots$, of which X is one, with outcomes of one of the mutually exclusive and effectively discriminable types $Z_0, \ldots$, of which Z is one. (For a more sensitive discussion of what is the weakest extra premise needed to validate an inductive inference, see my 1995.) The most that can be learnt from experience when the action a belongs to some particular X_i is that the outcome b belongs to some particular Z_k. What is learnt from experience is singular. Of course we learn, and to deny that experience plays a role in our learning would be fantastic. We need only consider how children learn their native language. It is uncontroversial that many of the abilities needed to learn a language are inherited, but that it is the child's early experiences that determine the particular language or languages learnt. Another explanation is therefore required of the role that experience plays in the creation of our general knowledge of the world.

It was this problem that Popper solved with what was initially a somewhat modest proposal concerning the logical problem of induction. Universal hypotheses, said Hume, cannot be derived from reports of experience. They must therefore, said Popper, precede experience, rather than follow it. But experience still has a service to perform. Universal hypotheses can be contradicted by reports of experience even if they cannot be derived from them. The role of experience is accordingly never to suggest hypotheses, which, as we have seen, it has no power to do, but to exclude them. This idea, which is logically trite, is of considerable significance. For by putting the hypotheses ahead of the experiences, we obtain a simple solution, though not effective one, to the biological problem of induction: the aspects of our experiences that are appropriately generalized are those that occur in hypotheses that will survive the later exposure to experience. This is the fundamental idea of the epistemology of trial and error, of conjectures and refutations.

When I say that the role of experience in learning is to exclude, I do not mean that experience can teach an organism to avoid mistakes, but not how to get things right. The contrast is not between the stick and the carrot, between the effectiveness of punishment and the ineffectiveness of reward. All that experience can impart is that a mistake has been made. If the organism was acting under the conjecture that such actions would be universally successful, then that conjecture may be discarded. But that is all that the organism can learn from its experience. If it is lucky, a different conjecture, a better one, may suggest itself. But only if the organism can recognize that the new conjecture is a different conjecture will it be able to improve its knowledge. Otherwise it may repeat the mistaken action, learning nothing. We see the great advantage for a species, if not for the individual, of the copying mechanism of the genetic code.

3 Invention

There is a familiar distinction, often attributed to Reichenbach, but essentially going back to Whewell (Hoyningen-Huene 1987, § II) between the discovery (better: the invention) of a hypothesis and its justification (better: its evaluation). The distinction is made, for example, in Whewell (1847), Book XI, Chapter VI, § 7. We can all agree that a hypothesis cannot be evaluated before it has been invented, but it is possible that the two processes, of invention and of evaluation, might be amalgamated in self-regulating rules of inference. Mill's 'methods of experimental inquiry' are sometimes presented as rules of inductive inference of this kind, leading unerringly from carefully assembled data to true causal laws. Mill himself originally described the method of agreement as a 'mode of discovering and proving laws of nature' (1843, Book III, Chapter VIII, § 1). Whewell, who disagreed often with Mill, wrote that, on the contrary, there are no such rules of discovery (*ibidem.*, Chapter II, § IV): 'Scientific discovery must ever depend upon some happy thought, of which we cannot trace the origin; — some fortunate cast of intellect, rising above all rules. No maxims can be given which inevitably lead to discovery.'

In Lecture VI.2 of *Lectures on Pragmatism*, Peirce (1903) distinguished three modes of reasoning: deduction, induction (confirmation), abduction (guesswork). In Lecture VI.4 he elaborated:

> Abduction is the process of forming an explanatory hypothesis. It is the only logical operation which introduces any new idea; for induction does nothing but determine a value, and deduction merely evolves the necessary consequences of a pure hypothesis. . . . Deduction proves that something *must* be. Induction shows that something *actually is* operative. Abduction merely suggests that something *may* be.

He notes in the same section that the 'only justification' for using abduction is that 'from its suggestion deduction can draw a prediction which can be tested by induction', and that it is only by abduction that we shall 'understand phenomena at all'. But this explains only why abduction is useful, not why it is successful when it is successful. 'No reason whatsoever can be given for it . . . , as far as I can discover; and it needs no reason since it merely offers suggestions.'

Popper too thought that 'the act of conceiving or inventing a theory . . . seems . . . neither to call for logical analysis nor to be susceptible to it'. Empirical psychology may be interested in 'the question how it happens that a new idea occurs to a man', but it is 'irrelevant to the logical analysis of scientific knowledge' (1934, § 2). He had previously identified 'the logic of scientific discovery' and 'the logic of knowledge' (*ibidem*, proem to Chapter 1), and went on to suggest that its task — 'in contradistinction to [that of] the psychology of knowledge — consists solely in investigating the methods employed in those systematic tests to which every new idea must be subjected . . . '. Despite this aversion to psychologism, in his early work *The Two Fundamental Problems of the Theory of Knowledge*, Popper, influenced by the biologically oriented approach of

Otto Selz (ter Hark 2004, 2006), had tentatively floated what he called 'a deductivist psychology of knowledge' (1930–1932, Chapter II, § 4). It was because it is both deductively invalid and superfluous that, like Hume, Popper excluded inductive inference from the logic of knowledge. It was because it is logically impossible (see § 2 above), that, unlike Hume, he excluded inductive learning (learning by repetition) from the psychology of knowledge (Popper, 1971, § 10; 1974a, *ibidem*). But, beginning with Whewell, many philosophers, contrary to both Hume and Popper, have denied any role to induction (generalizing from instances) in the psychology of discovery but have insisted on the indispensability of confirmation by instances in the logic of evaluation.

The view that induction plays a part in the formulation of scientific hypotheses was so outmoded 30 years ago that we read in a text for workers in AI (Mortimer, 1982, pp. 90f.) that

> ... inductionists do not question the hypothetico-deductive method propagated by Popper, they only express the opinion that induction is the essential element of this method. ... Popper's arguments against inductionism usually seem to be aimed at this conception of induction, whose adherents would be hard to find today, namely, the view that induction is the process of getting to hypotheses in the heuristic sense. ... This kind of criticism can be considered as relevant only with reference to some grotesque version of inductionism. It would be difficult to quote an example of a contemporary inductionist who could be accused of psychologism. ... All known attempts at creating an inductive logic refer to the problem of justification, to the problem of the criteria for correct acceptance of hypotheses and not to the problem of where the ideas of hypotheses in the mind of a research worker come from.

This passage is an odd way of opening a chapter entitled 'Popper's anti-inductionism and anti-probabilism', since the thesis that an important feature of a hypothesis is its probability relative to the available evidence belongs indubitably (though erroneously) to the stage of its evaluation, not to the moment of its creation. Mortimer's judgement that attention to the empirical psychology of knowledge is firmly out of favour (except amongst psychologists, such as the contributors to Oaksford & Chater 2008) may, however, still be sound. For it is not to the psychology of science that today's advocates of abduction and heuristics aspire to contribute, but to its logic.

The principal point of disagreement, I take it, between authors such as Aliseda (2006), Catton (2004), Finn (2002), (2011), Josephson & Josephson (1994), Simon (1973), and Zahar (2007), who set out to formulate some logic of discovery, and critical rationalists, who do not, is whether new hypotheses are free conjectures, or whether they are formed by some process, however dimly articulated, of inference or reasoning. These are the alternatives that were numbered (iii) and (iv) at the beginning of § 2. There are plenty of opportunities for confusion here, as we have seen, since Peirce called abduction a process of reasoning, and later in the same section used the term *guesses* to refer to scientific hypotheses ('man ... cannot give any exact reason for his best guesses'). Let us therefore adopt the convention that inferences are constrained by rules, and guesses are not. But there is extant another serious ambiguity concerning the term *abduction*, which is sometimes used not for the process of invention, but for a process of selection amongst hypotheses that have already been formulated and scrutinized. Abduction in this sense, more often known as *inference to the best explanation*, is subsequent to empirical evaluation, and does not properly belong to any logic of discovery. Understood as inference, indeed, *abduction is abductively unsound* — just as induction is inductively unsound (Popper 1971, § 6) — since there exists a much better explanation of how our knowledge grows. This is critical rationalism.

References for Part I

Aliseda, A. (2006). *Abductive Reasoning. Logical Investigations into Discovery and Explanation.* Dordrecht: Springer.

Carnap, R. (1968). 'Inductive Logic and Inductive Intuition'. In I. Lakatos, editor (1968), pp. 258–267. *The Problem of Inductive Logic.* Amsterdam: North-Holland Publishing Company.

Catton, P. (2004). 'Constructive Criticism'. In P. Catton & G Macdonald, editors (2004), pp. 50–57. *Karl Popper. Critical Appraisals.* London & New York: Routledge.

Finn, V. K.(2002). 'The Use of Induction in Plausible Reasoning in Intelligent Systems'. Abstracts, Karl Popper 2002, p. 19. Vienna: Karl Popper Institute. Reprinted as Appendix 4 of Finn (2011).

——— (2011). *Artificial Intelligence. Methodology, Applications, Philosophy.* Moscow: Krasand/URSS.

Goodman, N. (1954). *Fact, Fiction, and Forecast.* London: Athlone Press.

ter Hark (2004). *Popper, Selz and the Rise of Evolutionary Epistemology.* Cambridge: Cambridge University Press.

——— (2006). 'The Historical Roots of Popper's Theory of the Searchlight'. In I. C. Jarvie, K. M. Milford, & D. W. Miller, editors (2006a), pp. 37–56. *Karl Popper: A Centenary Assessment.* Volume I: *Life and Times, and Values in a World of Facts.* Aldershot & Burlington VT: Ashgate.

Howson, C. (2000). *Hume's Problem. Induction and the Justification of Belief.* Oxford: Clarendon Press.

Hoyningen-Huene P. (1987). 'Context of Discovery and Context of Justification'. *Studies in the History and Philosophy of Science* **18**, 4, pp. 501–505.

Hume, D. (1738). *A Treatise of Human Nature.* Book I, *Of the Understanding.* London: John Noon.

Josephson, J. R. (1994). 'Conceptual Analysis of Abduction'. In J. R. Josephson & S. G. Josephson, editors (1994), pp. 5–29. *Abductive Inference: Computation, Philosophy, Technology.* Cambridge & elsewhere: Cambridge University Press.

Lewis, D. K. (1980). 'A Subjectivist's Guide to Objective Chance'. In R. C. Jeffrey, editor (1980), pp. 263–293. *Studies in Inductive Logic and Probability*, Volume II. Berkeley & Los Angeles CA: University of California Press. Reprinted as Chapter 19 of Lewis (1987). *Philosophical Papers*, Volume II. Oxford: Oxford University Press.

Mill, J. S. (1843). *A System of Logic, Ratiocinative and Inductive.* London, New York, & Bombay: Longmans, Green & Co. Ltd.

Miller, D. W. (1994). *Critical Rationalism. A Restatement and Defence.* Chicago & La Salle IL: Open Court.

——— (1995). 'How Little Uniformity Need an Inductive Inference Presuppose?'. In I. C. Jarvie & N. Laor, editors (1995), pp. 143–148. *Critical Rationalism, Metaphysics and Science. Essays for Joseph Agassi.* Volume I. Dordrecht: Kluwer. Reprinted as Chapter 8 of Miller (2006).

——— (2006). *Out of Error. Further Essays on Critical Rationalism.* Aldershot & Burlington VT: Ashgate.

——— (2007a). 'El modo único para aprender'. *Estudios de Filosofía* (Antioquia), 36, agosto 2007, pp. 25–39. On line at `http://www.scielo.org.co/pdf/ef/n36/n36a03.pdf/`.

——— (2007b). 'The Objectives of Science'. *Philosophia Scientiæ* **11**, 1, pp. 21–43. Revised version of §0, §1, §5, and §6 of Chapter 4 of Miller (2006). On line at `http://go.warwick.ac.uk/dwmiller/poincare.pdf/`.

Mortimer, H. (1982). *Logika indukcji.* Warsaw: Panstwowe Wydawnictwo Naukowe. References are to the English translation 1988. *The Logic of Induction.* Chichester: Ellis Horwood Ltd.

Oaksford, M. & Chater, N., editors (2008). *The Probabilistic Mind.* Oxford: Oxford University Press.

Papineau, D. A. (1995). 'Open Society, Closed Thinker'. *The Times Literary Supplement*, 23.vi.1995, pp. 4f.

Peirce, C. S. S. (1903). 'Lectures on Pragmatism'. In C. Hartshorne & P. Weiss, editors (1934–1935), pp. 11–131. *Collected Papers of Charles Sanders Peirce*, Volume V. Cambridge MA: Harvard University Press.

Popper, K. R. (1930–1932). *Die beiden Grundprobleme der Erkenntnistheorie*. Tübingen: J. C. B. Mohr (Paul Siebeck). Published in 1979. English translation 2009: *The Two Fundamental Problems of the Theory of Knowledge*. Abingdon: Routledge.

——— (1934). *Logik der Forschung*. Vienna: Julius Springer Verlag. Expanded English translation 1959. *The Logic of Scientific Discovery*. London: Hutchinson Educational. Repaginated Classics edition 2002. Abingdon: Routledge.

——— (1945). *The Open Society and Its Enemies*. London: George Routledge & Sons. 5th edition 1966. London: Routledge & Kegan Paul. Repaginated Classics edition 2003. Abingdon: Routledge.

——— (1957a). 'Philosophy of Science: A Personal Report'. In C. A. Mace, editor (1957), pp. 155–191. *British Philosophy in the Mid-Century*. London: George Allen & Unwin Ltd. Reprinted as Chapter 1 of Popper (1963).

——— (1963). *Conjectures and Refutations*. London: Routledge & Kegan Paul. 5th edition 1989. Repaginated Classics edition 2002. Abingdon: Routledge.

——— (1971). 'Conjectural Knowledge. My Solution to the Problem of Induction'. *Revue internationale de philosophie* **25**, 95–96, fasc. 1–2, pp. 167–197. References are to the reprint as Chapter 1 of Popper (1972).

——— (1972). *Objective Knowledge. An Evolutionary Approach*. Oxford: Clarendon Press. 2nd edition 1979.

——— (1974a). 'Intellectual Autobiography'. In Schilpp (1974), pp. 1–181. Reprinted separately as K. R. Popper (1976). *Unended Quest*. London: Fontana. Repaginated Classics edition 2002. Abingdon: Routledge.

——— (1974b). 'Replies to My Critics'. In Schilpp (1974), pp. 961–1197.

——— (1983). *Realism and the Aim of Science*. Edited by W. W. Bartley III. London: Hutchinson Educational. Current edition. Abingdon: Routledge.

Russell, B. A. W. (1946). *History of Western Philosophy and its Connection with Political and Social Circumstances from the Earliest Times to the Present Day*. London: George Allen & Unwin Ltd. Repaginated Classics edition 2004. Abingdon: Routledge.

Schilpp, P. A., editor (1974). *The Philosophy of Karl Popper*. The Library of Living Philosophers, Volume 14. La Salle IL: Open Court.

Searle, J. R. (1999). 'The Future of Philosophy'. *Philosophical Transactions: Biological Sciences* **354**, 1392, 29.xii.1999, pp. 2069–2080.

Simon, H. A. (1973). 'Does Scientific Discovery Have a Logic?'. *Philosophy of Science* **40**, 4, pp. 471–480. Reprinted as Chapter 11 of H. Keuth, editor (1998), pp. 235–248. *Karl Popper. Logik der Forschung*. Berlin: Akademie Verlag.

Whewell, W. (1847). *Philosophy of the Inductive Sciences, Founded upon their History*. 2nd edition. London: J. W. Parker.

Zahar, E. G. (2007). *Why Science Needs Metaphysics*. Chicago IL: Open Court.