1

Introduction

> It is a curiosity of the philosophical temperament, this passion for radical solutions. Do you feel a little twinge in your epistemology? Absolute skepticism is the thing to try . . . Apparently the rule is this: if aspirin doesn't work, try cutting of your head.
>
> Jerry Fodor (1985)

Humans are in pursuit of knowledge. It plays a significant role in deliberation, decision and action in all walks of everyday and scientific life. The systematic and detailed study of knowledge, its criteria of acquisition and its limits and modes of justification is known as epistemology.

Despite the admirable epistemic aim of acquiring knowledge, humans are cognitively accident-prone and make mistakes perceptually, inferentially, experimentally, theoretically or otherwise. Epistemology is the study of the possibility of knowledge and how prone we are to making mistakes. Error is the starting point of skepticism. Skepticism asks how knowledge is possible given the possibility of error. Skeptics have for centuries cited *prima facie* possibilities of error as the most substantial arguments against knowledge claims. From this perspective, epistemology may be viewed as a reply to skepticism and skeptical challenges. Skepticism is the bane of epistemology, but apparently also a blessing, according to Santayana (1955): "Skepticism is the chastity of the intellect, and it is shameful to surrender it too soon or to the first comer" (p. 50).

Skepticism is a tough challenge and requires strong countermeasures. In set theory, a powerful combinatorial technique for proving statements consistent with the axioms of set theory was invented by P. Cohen in the

1960s. The technique is called *forcing*. In particular, Cohen developed forcing in order to prove that the negation of the Axiom of Choice and the negation of the Continuum Hypothesis are consistent with the axioms of set theory. Today, there are various ways of using the forcing technique. One way is to construct an object with certain properties or to construct a model in which there are no objects with certain properties, thus forcing what you want directly – either constructing the object or iteratively destroying any such object.

Contemporary epistemologies have developed a family of countermeasures for standing up to the skeptical challenge; these exhibit a type of 'bluntness' similar to that of set-theoretical forcing.[1] The idea of epistemological forcing is as follows: whenever skeptics cite possibilities of error as arguments against knowledge claims, the strategy is to show that, although they are possibilities of error, they fail to be *relevant* possibilities of error. Some possibilities of error are simply not genuine – they are too remote, too speculative, or too much. These possibilities may accordingly be *forced* out and are henceforth not to be considered during the knowledge acquisition process. If the agent can succeed over the possibilities deemed relevant, then that is good enough for knowledge – knowledge will, or should, exhibit all the classical characteristics under forcing.

The influential *epistemic reliabilism* of Goldman, Nozick's elegant formulation of the *counterfactual epistemology* and Lewis's new *contextual epistemology* are all informal epistemological proposals observing the forcing relation.

Epistemic reliabilism (Goldman 1979, 1986) and especially the recent versions outlined in Goldman (1992, 1996) acknowledge the agent's limited cognitive abilities and accordingly deflate the agent's epistemic responsibilities. The idea is to replace the rather demanding requirements typically proposed by skepticism for justified knowledge possession with more lenient conditions. In principle, a particular justified belief may be false; however, its method or mode of acquisition must in general lead to true convictions. For knowledge to come about, besides the truth of the belief in question, its method of acquisition must rule out all relevant possibilities of error. The forcing technique is included in the method of acquisition. The method may not be able to exclude the possibility that Descartes' devious demon is feeding the agent systematically misleading sensations. Then again, this is not a relevant possibility of error. Or so

[1] Otherwise set-theoretical and epistemological forcing bear little resemblance to one another. In a certain sense one may even call them opposites. The term 'forcing epistemology' was coined in Hendricks 2001.

it is claimed. According to epistemic reliabilists, infallible methods are not required for knowledge. The development of epistemic reliabilism up to the current versions is scrutinized from the forcing perspective in Chapter 3.

Nozick's (1981) counterfactual reliabilistic knowledge definition, an adapted and supplemented version of a proposal put forth by Dretske (1970), is likewise a forcing proposal. The goal is to show that knowledge is in fact possible. The inherent decision procedure in Nozick's definition of knowledge, together with the counterfactual semantics, requires the agent to succeed in all possible worlds sufficiently close to the actual world. The agent may not know that he is not a brain in a vat – a famous thought experiment suggested by Putnam (1981) – but that possibility of error is so remote, and the semantics governing the counterfactual conditional guarantees the long distance. This counterfactual epistemology is the topic of Chapter 4.

Whereas both epistemic reliabilism and Nozickian counterfactual epistemology begin by confronting the skeptical challenge, Lewis's (1996) contextual epistemology, in contrast, assumes knowledge of a great many things in a variety of different contexts, particularly conversational contexts. 'Contextualists' hold the view that the standards for knowledge acquisition, possession and maintenance fluctuate with what is at issue – and at stake – in the particular linguistic context. The current interlocutors determine which possible worlds are real or relevant and also why and when. The knowledge that you are currently wearing sneakers may evaporate into thin air once you set foot in an epistemology class because in this new context you may doubt whether you even have feet to put your sneakers on. Be that as it may, we have knowledge, and epistemology starts from there – not from ignorance or demonstrations of the mere possibility of knowledge. Considering brains in vats and Cartesian demons is to 'epistemologize', which may make knowledge 'elusive' especially in an epistemology class. What is needed for obtaining knowledge are regulatives to rule out possible worlds dictated by the current (conversational) context and then describe how we avoid error and gain truth in the ones that are left. Contextual epistemology is discussed in Chapter 5.

It turns out that a host of *formal* epistemological proposals also share the forcing heuristics. Knowledge claims may be restricted by algebraic constraints defined for the accessibility relation between possible worlds, which is the forcing foundation for *epistemic logic* or *logical epistemology*. Logical epistemology originates with Von Wright (1951) and was propounded most notably by Hintikka (1962). The algebraic properties of the accessibility relation between possible worlds may sometimes be

defined in such a way that the skeptic has nowhere to go. The forcing characteristics and the (often neglected) epistemological significance of epistemic logic are the topics of Chapter 6.

Formal learning theory, also dubbed *computational epistemology* by Kelly (1996), focuses on the intrinsic solvability of inductive epistemic problems for both ideal and computationally bounded agents (Kelly 2000). The basic idea is that when an agent is faced with an epistemic learning problem, the problem determines a set of possible worlds in each of which the agent has to succeed to solve the problem and acquire knowledge. This is also forcing. Brains in vats sever the connection between knowledge acquisition and reliable inquiry, but short of that, agents may have quite a bit of reliable inductive knowledge. Although it is a logical paradigm, in that it utilizes tools from mathematical logic, it is also a procedural or effective paradigm, as it concentrates on learning and knowledge acquisition issues rather than modal operators, axiomatics and validity, as logical epistemology does. Computational epistemology is the topic of Chapter 7.

The last epistemological proposal to be considered is called *modal operator epistemology*.[2] Modal operator epistemology is a mixture of epistemic, tense and alethic logic and a few concepts drawn from computational epistemology. It was developed in order to study the validity of limiting convergent knowledge (Hendricks 2001). To obtain limiting convergent knowledge, the agent has to converge to the true hypothesis only in the possible worlds consistent with what has been observed so far. This approach also pays homage to the forcing relation. Brains in vats are as devastating here as elsewhere, but if blocked, knowledge may in the limiting end have a certain strength measurable by a yardstick devised by logical epistemology. An outline of the modal operator theory of knowledge, together with an analysis of its epistemological importance, is provided in Chapter 8.

Epistemology may be pursued in different ways:

- 'Mainstream' epistemology (which encompasses epistemic reliabilism, counterfactual epistemology and contextual epistemology) seeks necessary and sufficient conditions for the possession of knowledge using largely common-sense considerations and folksy examples and counterexamples (see Fig. 1.1).

[2] Elsewhere the paradigm is also known as *modal operator theory*, since the paradigm is flexible enough to study other modalitites than knowledge.

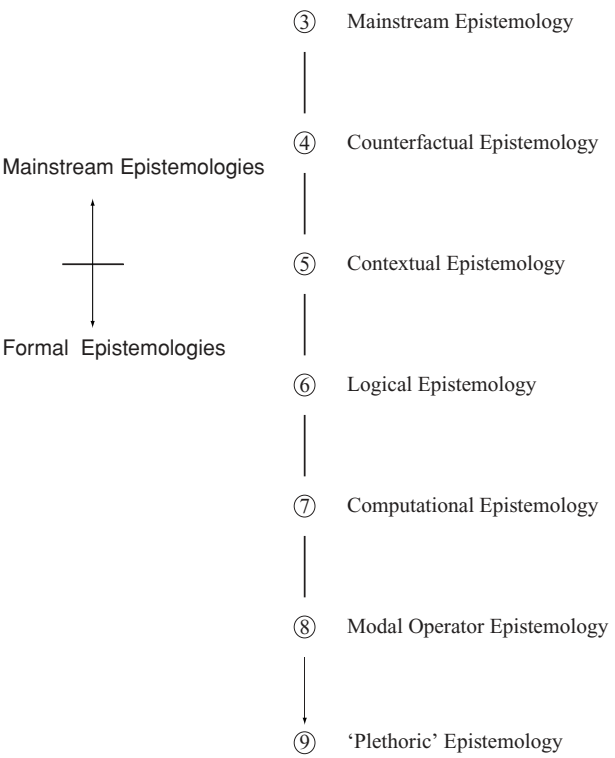

FIGURE 1.1. Epistemologies covered in this book and the chapters in which they are discussed.

- 'Formal' approaches to epistemology (which include logical epistemology, computational epistemology and modal operator epistemology) either proceed axiomatically or concentrate on learning and knowledge acquisition using toolboxes from logic and computability theory.

The two traditions have regrettably not paid much attention to each other. But both approaches, or rather their current exponents, employ the regulative forcing principle to combat skepticism. Based on this common denominator, the fundamental epistemological similarities and differences of the six paradigms may be cashed out in terms of how they each determine the set of possible worlds required for successful knowledge possession or acquisition.

The two approaches to the theory of knowledge share something else as well. One of the primary debates in contemporary epistemology

concerns the justification condition of the standard tripartite definition of knowledge as justified true belief. Time and time again, philosophers attempt to remedy the justification condition in order to avoid 'Gettierization' (Gettier 1963) and other epistemic unpleasantries. The justification condition is supposed to ensure that the belief and the truth conditions of the tripartite definition are 'adequately connected', that is, that the reasons for believing are truth-conducive and, insofar as they are, indicate what is meant by rational inquiry. Philosophy of science includes a subdiscipline concerned with exactly the same thing: methodology.

Methodology may crudely be characterized as the study of the methods by which science arrives at its posited truths. Methodologists and formally minded philosophers have a large technical toolbox available for analyzing and hopefully ensuring the truth-conduciveness of the methods of science. These techniques range from various inductive and nonmonotonic logics to Bayesianism, game theory and belief revision theory to formal learning theory, and so forth. When mainstream philosophers talk about justification, formalists speak of methodology. A philosopher may choose to invoke reliability; the formalist then asks how reliability is to be defined and what it can do for you methodologically. The mainstream epistemologist calls for a defeasibility condition, and the philosophical logician starts to think about default rules and nonmonotonic logics; the mainstreamer wants to get to the truth sooner or later, the computational epistemologist, say, begins to consider solvability and criteria of successful convergence; accumulating evidential support the mainstream community decides for and the Bayesian community will start conditionalizing; minimum mutilation of the web of beliefs and the belief revision theorists will work on revision functions and entrenchment relations; an epistemologist may worry about rationality, the game-theorist will start to consider, say, strategies for winning noncooperative games of perfect information. What the mainstream epistemologists are looking for may to some extent be what the formal epistemologists have to offer. But what the formal epistemologists have to offer the mainstream community, and vice versa being a two-way street, may also be quite sensitive to the perspectives on inquiry that the different approaches adopt.

The general prerequisites for studying these epistemo-methodological affinities are outlined in Chapter 2, then applied systematically in the subsequent chapters. Finally in Chapter 9, they are used for the purpose of outlining a program of 'plethoric' epistemology.

2

Priming the Pump

The epistemo-methodological prerequisites for comparing mainstream and formal epistemologies concentrate on the following items: the modality of knowledge, infallibility, forcing and the reply to skepticism; the interaction between epistemology and methodology; the strength and validity of knowledge; reliability; and the distinction between a first-person perspective and a third-person perspective on inquiry.

> If knowledge can create problems, it is not through ignorance we can solve them.
>
> Isaac Asimov

2.1 Modal Knowledge, Infallibility and Forcing

Agents inquire to replace ignorance with knowledge. Knowledge is a kind of epistemic commitment or attitude held toward propositions or hypotheses describing some aspect of the world under consideration.[1] Agents may in general hold a host of different propositional attitudes, such as belief, hope, wish, desire etc. But there is a special property that knowledge enjoys over and above the other commitments. As Plato pointed out, a distinct property of knowledge is truth. Whatever is known must be true; otherwise it is not knowledge, even though it very well may qualify as belief or some other propositional attitude.

Contemporary notions of knowledge are often *modal* in nature. Knowledge is defined with respect to other possible states of affairs besides the actual state of affairs (Fig. 2.1). The possibility of knowledge seems ruled

[1] The terms 'hypothesis' and 'proposition' are used interchangeably unless otherwise stated.

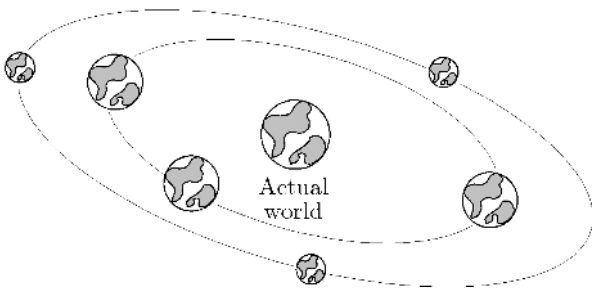

FIGURE 2.1. Modal knowledge is defined with respect to other possible worlds.

out when it is possible that we err. Introducing other possible state of affairs is an attempt to preclude exactly these error possibilities. Knowledge must be *infallible* by definition. As Lewis (1996) puts it, "To speak of fallible knowledge, of knowledge despite uneliminated possibilities of error, just *sounds* like a contradiction" (p. 367). A fallible notion of knowledge is not much different from a concept of belief potentially allowing the agent to 'know' a falsehood, severing the connection between knowledge and truth.

Plato also observed that knowledge, as opposed to mere true belief, is stable in nature. Knowledge has steadfastness and indefeasibility attached to it. True belief is quite useful as far as it goes, but in the light of true evidence, it may vanish. In the light of true evidence, knowledge will not evaporate. Inevaporability makes for the robust usefulness of knowledge compared with beliefs that are simply true. True belief in the actual world is not necessarily preserved if circumstances were to change but slightly. On the other hand, knowledge in the actual world is assumed to be stable across quite radically varying circumstances. Thus, among both informally and formally minded epistemologists, there is an agreement that knowledge is defined with respect to other 'possible worlds'. As Hintikka (2003a) notes,

> In order to speak of what a certain person *a* knows and does not know, we have to assume a class ('space') of possibilities. These possibilities will be called scenarios. Philosophers typically call them possible worlds. This usage is a symptom of intellectual megalomania. (p. 19)

There is an immediate difference between a philosophical logician and a philosopher. The logician typically remains agnostic about the ontological significance of the possible worlds and may refer to them as scenarios,

situations, states, contexts or conceptual constructions. The philosopher is usually quite concerned with the metaphysical baggage that comes along with the notion.[2]

Be that as it may, the stability and robustness of knowledge over other possible worlds leaves open the question of which space of worlds should be considered relevant for epistemic success. The classical conception of infallibilism is taken to require that an agent, in order to have knowledge of some hypothesis, must be able to eliminate *all* the possibilities of error associated with the hypothesis in question. The set of *all* worlds is considered. This set of possible worlds is too big for knowledge to have scope over. The set includes some rather bizarre worlds inhabited by odd beasts ranging from demons to mad and malicious scientists who have decided to stick your brain in a tank of nutritious fluids to systematically fool you. Or worlds in which contradictions can be true and tautologies can be false, like 'impossible, possible worlds' (Hintikka 1975). If these worlds were to be considered relevant all of the time, skepticism would have the upper hand all of the time. Epistemology may just end up with a fallibilistic notion of knowledge after all: there may not be a way for an agent to determine that he is not in the world of the beast or the brain. But then again, a fallibilistic notion of knowledge hardly qualifies as knowledge at all. At most, it amounts to a description of knowledge-seeking practices. Consequently, *if infallibilism is to be a viable reply to the skeptic, then infallibilism cannot be defined with respect to all possible worlds.* This is where epistemological forcing comes in.

The bizarre and extravagant possibilities of error may, under the right circumstances, be ignored even though they are logically possible, for all the agent knows. Knowledge may accordingly remain infallible but with world restrictions imposed. Forcing is more of an heuristic principle than an epistemological thesis proper:

> Whenever knowledge claims are challenged by alleged possibilities of error, the strategy is to show that the possibilities of error fail to be genuine in the relevant sense.

[2] To stay with currently adopted jargon, other scenarios, situations, states or contexts will be referred to as 'possible worlds' but nothing metaphysical is necessarily implied by the usage. Possible worlds are not to be understood as ontological or semantical totalities complete in their spatiotemporal history. Later it will become apparent that possible worlds may be endowed with enough formal structure to actually facilitate the achievement of important epistemological results.

Contemporary epistemologists choose to speak of the *relevant* possible worlds as a subset of the set of all possible worlds.[3] The philosophical logicians and other formal epistemologists consider an *accessibility* relation between worlds in a designated class within the entire universe of possible worlds. It will become apparent that there is no principled difference between relevance and accessibility. Informal epistemologies differ by the way in which relevance is forced given, say, perceptual equivalence conditions, counterfactual proximities or conversational contexts circumscribing the possible worlds. Formal epistemologies differ by the way in which the accessibility relation is defined over possible worlds. For example, philosophical logicians obtain different epistemic modal systems valid for a knowledge operator by varying (adding, dropping or relativizing) the properties of the accessibility relation, which might change from being reflexive and transitive to being reflexive, symmetric and transitive, for example.

Computational epistemology also forces as inductive epistemic problems to be checked for solvability specify a set of possible worlds for the agent to succeed over. Modal operator epistemology assumes that limiting convergent knowledge is restricted to possible worlds that are consistent with what has been so far observed.

> An informal epistemological paradigm may be a forcing strategy, and a formal one may be too. The task is then to find out how they each force so they can be put on a par for comparison.

Following Lewis (1996), one may say that it is a basic epistemic condition for agents to force for knowledge. The technically minded theorists of knowledge choose to formalize this basic epistemological condition.

2.2 Skepticism

The unstated premise of epistemology is, of course, that agents are seekers of knowledge or information – that is the whole point of conducting inquiry. Skepticism argues that even though gaining truth and avoiding error is the point of inquiry, the acquisition of knowledge is impossible given the standing possibility of error. In the end, we are left with ignorance rather than insight. Skeptics often cite two lines of argument in favor of their pessimistic conclusion.

[3] Explicit forcing proposals in the epistemological literature are sometimes referred to as 'relevant alternatives proposals'. Cf. Bernecker and Dretske 2000.