

Synthetic Epistemology

Roger Bishop Jones

2022/03/13

Contents

Preface	2
1 Introduction	3
2 Epistemic Futures	5
2.1 As things are	6
2.2 Reality Check	7
2.3 Motivations and Economics	7
2.4 Proliferate What?	8
3 The Logical Structure of Knowledge	9
4 Matters of Logic	12
4.1 Temporary Sketch	16
5 Matters of Fact	17
6 Values and Morals	17
Bibliography	17

Preface

This is intended to be a sketch of a larger, book sized, project and a contender for the introduction to the book, if ultimately this kind of introduction seems a good idea. Otherwise, its still a stage in working out the ideas to be addressed in the larger format.

[1](#)

¹There may be “hyperlinks” in the PDF version of this document which either link to another point in the document (if coloured blue) or to an internet resource (if coloured red) giving direct access to the materials referred to (e.g. a Youtube video) if the document is read using some internet connected device. Important links also appear explicitly in the bibliography.

1 Introduction

Epistemology is the theory of knowledge, a branch of philosophy. The word “episteme” has roots at least as far back as Aristotle, who used it’s Greek predecessor as a term for scientific knowledge, as it was then known. This he conceived of as theoretical knowledge rather than the less distinguished “techne” concerned with practice.

Both the contemporary terms *epistemology* (the theory of knowledge) and *episteme* (the standards for what counts as knowledge, or for judging the truth of conjectures) derive from those roots. In this work the word *epistemology* will be used broadly, concerning all kinds of knowledge. When I speak of epistemes, by contrast, I am concerned with standards for theoretical knowledge rather than practical. The line between theoretical and practical for our purposes is similar to that between *knowing that* and *knowing how*, that between knowledge which can be articulated and communicated as propositions, and knowledge which consists in the possession of some skill which one acquires through observation, practice and coaching.

In the hands of postmodern philosopher Michel Foucault the term “episteme” was used to speak of the standards imposed on knowledge more generally, treated as a power-play by culturally dominant groups. This cynical view of epistemic standards later provided a justification for the weaponisation of radical epistemic innovations, ostensibly for the emancipation of the underprivileged, by political activists dedicated to the complete overthrow of the western cultural hegemony, which was held to sustain global capitalism by inculcating false consciousness into its victims.

I will use the term here to refer to those aspects of a culture or subculture which determine or influence what is accepted as knowledge. The term “synthetic epistemology” I have coined to talk about the philosophical activity of synthesising epistemes, or reasoning about such syntheses. In the context of Michel Foucault’s philosophy this constitutes a power play, in which reading I acquiesce.

By way of mitigating that acquiescence, my inclination is to re-

gard evolution as cultivating instrumental epistemes, those which are conducive to the proliferation of humanity or its progeny. A distinction between domains in which knowledge is objective and those in which standards are discretionary is central to my synthesis. In the domain of the objective, instrumentality abhors the intervention of political motivations into epistemic standards. They simply are not effective or relevant. The syntheses undertaken in this context, are therefore not intended to align with the interests of any social grouping, but with considerations which are of general utility. This mitigation bears only upon the realms, should there be any, of objective knowledge, and does not extend to the establishment of value systems. The distinction between these two may not be straightforward to make in an objective way, and is one of the first to have been challenged by the post-Marxist ideologies to which Foucault has (perhaps unwittingly) contributed. I don't expect to convince those wed to the radical epistemological innovations of our time.

The discussion is divided into the following parts. First it is necessary to speculate about how the proliferation of humanity and its progeny or product will proceed. This is undertaken in the context of a conception of evolution which is post-Darwinian, indeed post-Modern Synthesis, and so as brief as possible an account of the kinds of evolution which will shape our future will be necessary. Then on the basis of that speculation, I will outline some of the broad epistemological features likely to emerge. An evolutionary perspective will prove important in these speculations. On the basis of those anticipated features I will then begin to synthesise an episteme in three stages, addressing logical, empirical and evaluative knowledge following that classification derived from the philosophy of David Hume substantially as as refined in the works of Rudolf Carnap.

Hume classification came as two distinctions both of which have sometimes been called "Hume's fork". The first is the distinction between, as Hume put it, relations between ideas, and matters of fact. In the more recent terminology this is the distinction between logical and empirical propositions. The second is that suggested by

Hume's dictum, "you cannot derive and ought from an is", which presumes a clear distinction moral judgements and factual or logical claims. This latter distinction can be extended to include with moral judgements other (though not all) evaluative claims.

These distinctions have a long history, perhaps not of precise articulation, but of faltering approaches, gradually homing in on the fundamental distinctions behind those tenuous descriptions, some of which I will sketch below.

It is not my intention here to discover truths of epistemology. It is for us to choose how to discover, evaluate, organise and apply knowledge, and though common standards are highly desirable, there is always room for debate about what they should be. This is written as a contribution to the continuous evolution of epistemological standards.

In that description of my enterprise you may immediately sense my divergence from at least two other attitudes to epistemology. The first is any kind of epistemological absolutism. It is up to us where we place the boundary between conjecture and fact, or indeed, whether we have such a dichotomy rather than some more complex characterisation.² The second is the radical relativism which we might see in Michel Foucault, and the weaponisation of epistemology through standpoint epistemology which was enabled by it. Knowledge is important, it enables those who possess it to flourish, and facilitates their fulfillment. Epistemes facilitate the acquisition of knowledge and guard against the degradation of our understanding of the world. They are not all equally effective, they evolve with us, and this synthesis is my attempt to contribute to that evolutionary process.

2 Epistemic Futures

What does the future hold, what knowledge would facilitate that future and what kind of episteme would best enable its cultivation?

²Just as hot and cold were displaced for many purposes by quantitative measures, some philosophers have sought quantitative evaluation of the evidential support for a scientific hypothesis.

tion? These questions are central to the epistemological synthesis at hand.

The episteme I sketch below is my idea of what might serve us well in the future, and so I intend here to give my take on how the future will be, focussing on those aspects of the future which I think most significant epistemologically.

I am thinking well into the future, expecting the proliferation of humanity and its progeny (not necessarily biological) across the Galaxy. I'm going to assume that something along those lines is possible, and will happen, despite any imminent risk of self-annihilation.

2.1 As things are

As I write, we have SpaceX, lead by Elon Musk, set up in 2002 with the aim of establishing on Mars an independent colony which could survive if humanity on earth should become extinct or too far degraded to lend further support. After achieving radical reductions in the cost of launch into orbit by developing reusable launch technology, SpaceX are on the verge of testing orbital launch for a system designed to be capable of carrying a payload of 100 tons and one hundred passengers to Mars.

Musk believes that to achieve independence a colony a million people would be needed, and has built a re-usable rocket system which he hopes will be up to transporting the one million tons of materials which would be necessary to establish the colony, at a cost of up to \$10Trillion.

The establishment of a colony would have to be undertaken primarily by automated systems rather than human astronauts since otherwise human life support systems would have to be in place. Musk is developing intelligent android robots in part for that purpose.

The motivation for colonisation of another planet is for Musk, the resilience given to homo sapiens against earthly catastrophe. We know however, that the sun will in due course render the solar system increasingly hostile to human habitation, and so it a natural

next step to attempt interstellar migration.

There is an initiative in progress to send a mission to our nearest neighbouring star, Alpha Centauri. The aim is a 1 gram payload.

2.2 Reality Check

Its pretty clear that Elon Musk, the only person trying to colonise Mars, doesn't know how to do it. He does have a good track record for seeing that something can be done in principle and then, by hook or by crook, trial and error, eventually figuring out how to do it. He now seems close to having the technology to ship materials to Mars on a much larger scale than has so far been possible. He has made some back-of-the-envelope calculations of how much resource would be needed for an independent settlement on Mars, and this seems to be within the scope of the technologies he is close to establishing, but subject to a very large price tag. To succeed it is likely that he would have to automate so much of the enterprise that he could have an independent settlement lacking human's (which is not his goal!). There will of course be a huge body of difficult problems to be solved, just to establish a working environment before independence is realised, and there clearly are questions about the scale of the enterprise and how it could possibly be funded.

2.3 Motivations and Economics

For my epistemological purposes I look further forward, anticipating interstellar migration.

What would be the point?

There will undoubtedly be enthusiasts for such an enterprise, but for it to succeed there would have to be very large outlays and success will probably be contingent on making the economics work. The kinds of interstellar proliferation which will dominate, will be those which successfully stimulate the required investments. It will be a great advantage if the project actually generates earthly revenue, so its worth asking how it might do that, and the answer to that question will not be unrelated to the question what value

could it have to the population, the first instance, of the earth, and later, to successive generations of replicating intelligent systems.

The distances involved mitigate against material transfers, and the transmission delays incurred by communications mitigate against the provision of remote computing facilities, unless of results obtained by large scale computations over long periods.

In the most distant regions of intelligent self-replicating systems, the imperative to proliferate will place a premium on the kinds of research and development which would advance the rate of proliferation. Each new generation must first travel to a location suitable for it to “germinate” (though not biological, which may allow for a wider range of environments to be suitable). It will then undertake a process of development, during which, using the energy and mass available at that location it grows key capabilities in a manner analogous to the development of embryos. An important aspect of this is the growth of computing infrastructure.

2.4 Proliferate What?

There has long been controversy about whether space exploration should be primarily by manned missions, or whether unmanned probes could achieve more at lower cost. This does depend of course on the purpose. If, as in the thinking of Elon Musk, the purpose is to make homo sapiens a multiplanetary species so that we will survive the worst catastrophes which might occur on a single planet, then manned missions will eventually be necessary. It might still be expected that a large proportion of the work involved in establishing an independent habitable environment will be undertaken by autonomous machinery.

A spectrum of possibilities for the kinds of system which might be proliferated across the Galaxy (or the Cosmos) is:

- Homo Sapiens
- Homo Post-sapiens

By this is intended biological species descended from homo-sapiens in ways which are adaptive for interstellar travel and

synthetic environments, possibly evolved by design using genome editing or synthesis.

- Synthetic self-reproducing intelligent biosystems

This refers to intelligent organisms designed for interstellar migration.

- Non-biological intelligent self-reproducing systems

Non-biological self replicating intelligent systems might achieve more rapid rates of proliferation, benefitting from the possibility of miniaturisation. This reduces transport cost, increases velocities, and simplifies environmental requirements for proliferation.

- Intelligent exploratory probes

These are not exclusive, it is not unlikely that several of these will progress in tandem. In that case, it seems likely that the earlier options in the list will be those least widely progressed from the epicentre (earth). I would guess that by the time the indirect reach of homo sapiens over the Galaxy reaches 100,000 light years, the largest bulk of that expansion will be of synthetic intelligence, with a penumbra of non proliferating probes.

3 The Logical Structure of Knowledge

[The topic of this section is now being addressed in a separate document and so if work on this document continues this section is likely to be reduced by reference to that other document [1]]

We learn about the world around us, and about how to manipulate it for our benefit, by observation, experiment and reason. These processes are sometimes tentative and unreliable, but over time they yield substantive results, and are responsible for the technologies which have made us healthy and prosperous.

Alongside the development of science and technology, playing a vital role in enabling that development, there have been advances

in the manner in which knowledge has been gathered. Some of these may be considered philosophical, even though perhaps more often effected by scientists than philosophers.

The records we have of the development of science and technology extend back only a few thousand years, out of the 300,000 years that homo sapiens may have lived on earth. Fairly early in that development the ancient Greeks turned the mathematical techniques which they inherited, mostly from Babylonians and Egyptians, into a *theoretical* discipline. By this is meant that instead of aggregating mathematical techniques for solving practical problems, they undertook the establishment of general principles governing those techniques, which they did by constructing deductive proofs or demonstrations.

Mathematics as a theoretical discipline was very successful, and the methods used were eventually codified as the 'axiomatic method' (a way of organising deductive reasoning), documented by Euclid, and best known in the development of Euclidian Geometry which was a high watermark in standards of logical rigour until late in the 19th Century. In the abstract domain of mathematical truth, the deductive method proved reliable from its very beginnings, but the faculty of reason which is employed proved much less satisfactory when used to reason about the concrete world about us, which the early Greek philosophers sought to do. To arrive at satisfactory general principles about how the world works required different methods, the methods which would eventually characterise empirical science, but which at their best conferred a lesser degree of confidence than the logical reasoning which sufficed for mathematics.

We may therefore see from the very earliest philosophical and scientific investigations, two distinct kinds of subject matter for the investigation of which quite different methods were appropriate, those of logical truth and empirical science. The relationship between these two and the precise location of the line which set them apart would be a continuing matter of controversy until the present day.

The construction of the 'episteme' which I undertake here, re-

flects a particular view on this issue.

In due course a further domain was to be identified, marking the distinction between factual claims of a logical or empirical nature from moral judgements.

Though we have here a separation between three domains, these are not islands. Though deductive reasoning is sufficient by itself in establishing only logical claims, together with empirical premises it can lead us to further empirical results, and aided both by moral and empirical principles it is a means to moral judgements.

This informal delineation of three domains of knowledge and the relationship between them forms the basis for an architecture for the representation and aggregation of knowledge which forms the backbone of the episteme which I propose.

The relationship between these three domains has the following characteristic. Each domain is deductively closed. That means that from true premises which belong to that domain we obtain only true premises which also belong to that domain. In each domain we seek premises which have as consequences the answers to the various practical problems which face us in the relevant domain.

In the logical domain this is a small number of logical truths from which all or most logical truths can be derived, allowing that various definitions of concepts are taken into account in the demonstration of propositions which make use of those concepts. In the empirical domain, the principles required are the laws of empirical science which are discovered by observation and experiment, and are then exploited by demonstration of practical results of interest.

Within the moral domain, reasoning also depends in part on moral principles within the deductive closure of which applicable results are to be found. The aspiration to embrace all moral truth in such a deductive closure is common, and is exhibited particularly in the high noon of rationalist philosophy by Spinoza's work on ethics.

In all three domains the existence of a set of principles in whose deductive closure all the truths of the domain can be found seems unrealisable. There are fundamental results in logic which show that even in the domain of logic this incompleteness is inescapable,

though in the case of logic the shortfall may not be significant. In the other two domains, incompleteness of the deductive closure follows from incompleteness in the logical domain, but it is likely that much more serious gaps arise from shortfalls in the principles of that domain from which deduction begins. In the empirical domain the shortfall appears both in the scientific laws and in the description of the concrete details of the world in which we seek to apply the laws. In relation to the laws scientists may believe that ultimately all the laws governing the universe may be known, but the past tells us only that our knowledge however much it progresses remains incomplete. As far as the description of the universe as it is, rather than how it progresses, it is easy to see how large a problem that is, and even when confined to some region, can only be given to a certain degree of precision which is inevitably imperfect.

In the moral domain, matters are even more difficult, for there is no known way to settle the principles from which reasoning might begin, and it is very likely that those principles will refer to matters for the differentiation of which intuitive rather than deductive reasoning is essential. It is therefore likely that in the domain of moral truths the application of deductive methods will be more problematic than in the other two domains.

The above account of the relationship is I hope a good place to start, but we need immediately to refine this sketch. So first I'll point out some weaknesses in that sketch. But before doing that I would like to say a few words about the aspiration, which may help to motivate the refinements which I will then propose.

I have been talking about deductive closure. I'd like to say why that is so important in this proposed architecture of knowledge, and then look closer at what it takes to make it work.

4 Matters of Logic

The distinction between logical and empirical truth is fundamental to the epistemic developments discussed in this work, and it will be helpful in clarifying that distinction to talk about its history.

This account of the history of the development of logic is focussed on making intelligible a particular view of the nature of logical truth and the purposes for which and ways in which a clear understanding of this conception of logical truth can be instrumental. The notion of logical truth has been controversial, as indeed most philosophical concepts can be from time to time, particularly because of the centrality of that notion in the doctrines of Rudolf Carnap and the logical positivists, and (I would say) the importance to American philosophers of finding their own way and asserting their strength in mid 20th Century philosophical politics. There is no intention here to revive the debate about what is “logical truth”, my use of that term, even though it is very close to the preferred usage of Rudolf Carnap, is not a claim about the meaning of the term, it is a choice of terminology which is convenient for describing the *episteme* to the construction of which this work is devoted.

The distinction between logical and empirical truth is rooted in the nature of language, and some practical application of logical truth is implicit in the skills which constitute linguistic competence. The part of language which we can call ‘propositional’ and which consists in the use of indicative sentences, depends for its use upon an understanding of the meaning of the concepts which can be asserted of the world or of its constituents. One aspect of that knowledge of language is the understanding of various conceptual inclusions. A rich source of these is taxonomic hierarchies, and the knowledge of the conceptual inclusion in such hierarchies yields logical truths of which mere competence in language assures us, without need to inspect the world for corroboration. The competence spoken of here is operational competence in making natural inferences, or in actions which involve implicitly such inference, rather than in being able to articulate conceptual inclusions or in understanding what a ‘conceptual inclusion’ is.

For signs of systematic use of logical inference we have to wait quite a long time, from the birth of language, maybe coeval with homo sapiens some two or three hundred thousand years ago to the advent of mathematics as a theoretical discipline in ancient

Greece. Prior mathematics, Babylonian arithmetic and Egyptian geometry, consisted in techniques for achieving certain practical ends, which Greek antiquity would have called logistics rather than mathematics. From about 600 BC the development of mathematics as a theoretical discipline began, seeking theoretical principles established not by successful application, but by deductive proof.

The deductive method was applied throughout the development of Greek mathematics, but was best exhibited in the development of Euclidean geometry, which showed it to be a highly reliable method for establishing geometric truths. These early Greek philosophers also sought to understand the world and the cosmos through reason, and in these broader realms were rather less successful. Theories advanced in these matters were not supported by deductive proof, were frequently contradicted, and eluded consensus.

In this way, the axiomatic method deployed for geometry and later systematised by Aristotle and Euclid came to be envied by all who wanted their opinions to be taken as gospel, and would often be misapplied or its merits fraudulently claimed for less reliable methods. In clarifying the nature and boundaries of *logical truth* the scope of deductive methods can be made clear. This makes a proper evaluation of other kinds of claims to knowledge more readily achievable.

After the period of pre-Socratic philosophy, in which the disparate achievements of reason were laid bare, resolution of the disparities was sought in the two great classical philosophical systems, those of Plato and Aristotle. It is here that much of the terminology connected with the distinction between logical and empirical truth discussed, as well as the ideas behind later terminological developments.

Plato was influenced by the philosophers Heraclitus and Parmenides. Heraclitus held that everything was in constant flux, Parmenides that change was impossible. Parmenides arrived at his position by reason, Heraclitus was acquainted with the chaos of the world through his senses. Plato resolved these apparently

contradictory philosophies by identifying two distinct domains, a realm of ideal forms, immutable and accessible to reason, of which certain knowledge could therefore be had, and a realm of appearances, fleeting and ephemeral, brought to us by our senses, of which we could have only opinion, not knowledge. The realm of Platonic “forms” or universals, and one imperfectly revealed by our senses. A classical predecessor to what might now be rendered as a realm of abstract objects, and one of material entities.

David Hume offered an intermediate position between this classical philosophies of Plato and Aristotle and more contemporary philosophy of Rudolf Carnap.

Hume makes the distinction thus:

“ALL the objects of human reason or enquiry may naturally be divided into two kinds, to wit, Relations of Ideas, and Matters of Fact.”

Which we may characterise as sharing with Plato a classification by subject matter with multiple dimensions of significance.

He gives the following description of “Relations of ideas”:

“Of the first kind are the sciences of Geometry, Algebra, and Arithmetic; and in short, every affirmation which is either intuitively or demonstratively certain. That the square of the hypotenuse is equal to the square of the two sides, is a proposition which expresses a relation between these figures. That three times five is equal to the half of thirty, expresses a relation between these numbers. Propositions of this kind are discoverable by the mere operation of thought, without dependence on what is anywhere existent in the universe. Though there never were a circle or triangle in nature, the truths demonstrated by Euclid would for ever retain their certainty and evidence.”

I’ll go through this in some detail. This connects with Plato in that the primary distinction is about subject matter, and that

the distinction is connected with certain other distinctions. The other distinctions found in Plato are combined into one as the distinction that we obtain knowledge about the forms by reason, but only have opinion about the world of appearances derived from our senses. This distinguishes two ways of coming to knowledge or opinion, and two corresponding results of such enquiries. What we have here might be called a “triple dichotomy” in which three different characterisations are held to yield closely related distinctions, and we will see that progress in clarifying the notion of logical truth has advanced through the refinement of similar dichotomies.

The things being classified are usually sentences, judgements or propositions. The distinctions made of of the following kinds:

Semantic Concerned with the meaning of a sentence, or its subject matter, what the sentence is about or what it says about its subject.

Epistemological Concerned with how knowledge or opinion about the truth of the sentence or proposition is to be obtained, and/or justified.

Confidence Whether one can have certainty on the matter or not.

Modal Status Whether the sentence or proposition is necessary or contingent.

4.1 Temporary Sketch

The stages:

- the origin of language
- the beginnings of deductive mathematics
- the Platonic dichotomies
- Aristotelian logic and metaphysics
 - The Logic of Syllogism

- Category Theory and the Necessary/Contingent dichotomy
 - Demonstrative Science
- Leibniz and the mechanisation of logic The mechanisability of Logic
- The Humean dichotomies
- Kant’s retreat
 - The concept of analyticity
 - The synthetic a priori
 -
- Bolzano
 - The break with Aristotle
 - Logic as Theory of Science
 - Proposition as meaning of a sentence

5 Matters of Fact

6 Values and Morals

References

- [1] Roger Bishop Jones. An Architecture for Propositional Knowledge. 2022.

Started 2022/10/09

<http://www.rbjones.com/rbjpub/www/papers/p049.pdf>