LIBRARIES AND FALLIBLE KNOWLEDGE
Stephen Hetherington

Recently, I wandered into a substantial research library in downtown Sydney – the State Library of New South Wales. A moment prior, I had been evading cars, listening unwillingly to city noises. The next moment, I was inside the library: elegance, peace, stillness – books enveloping me. The library's reading room: an imposing collection, especially of non-fiction. So I sat; I pondered; and I raised my eyes to the surrounding walls of books. Calmness for a while; until suddenly I was troubled. A fundamental *philosophical* problem about the nature of that room had crept into my mind. Here is the story of my collision with that problem – and of how I tried to resolve it.

Questions

Silently, I allowed an intellectual challenge to unfold.

Might there be no knowledge – none at all, anywhere? Seemingly, this is a possibility I should not reject without reflection. But *can* I seriously engage with that idea? Here, of all places (I thought), that sounds bizarre. Look at these walls (I said, to myself), inside one of Australia's greatest libraries. Behold these non-fiction books. They encapsulate people's attempts to understand and learn about our world – along one wall, tomes on the First Fleet from England, travelling so far so long ago; along another wall, books about the later Federation of the States (New South Wales, Victoria, etc.) into the Commonwealth of Australia; everywhere in the room, something serious. And part of the aim in the library's storing such books is to enable people likewise to understand and learn – not

merely what the authors thought, but truths reflected in those thoughts.

Unsurprisingly, therefore, as I gazed upon the room I wondered how there *could* be no knowledge: 'We know nothing about the First Fleet? Not even that it arrived in 1788? And Federation in 1901: no knowledge of that, either? Surely not.' A supposition of knowledge's nonexistence sounds like a circuitous way of denying that a real library ever *exists* around me – or at least that the library, even if it does exist, can be storing and transmitting *knowledge*. Yet knowledge's preservation and propagation seems essential to the library's aim.

Those thoughts brought me face to face with a *conceptual* challenge: Is even the *idea* of a research library confused? To understand and to learn is partly to come to know. So if there is no knowledge there is no understanding and no learning. But if there is no knowledge, what is accomplished by the research stored here? What are those books about Federation and about the First Fleet achieving *if* never storing and imparting knowledge? What *does* the library accomplish with its non-fiction holdings, if not the storing and imparting of knowledge?

Knowledge-Scepticism

Well, *are* those non-fiction books never storing and imparting knowledge? This depends upon whether there are good *arguments* that could sensibly lead me to regard those walls as empty of knowledge.

That dire view of the walls is what epistemologists (philosophers of knowledge) call a *sceptical* idea. It's a long-standing philosophical approach, receiving over the centuries several clever formulations and defences. It's also surprisingly difficult to shift from philosophy. It has a seductive allure. Even among philosophers who would spurn it, there is no generally welcomed way to do so.

Perhaps the sort of sceptical argument most pertinent to this setting – the library – is a philosophically ancient

sceptical path. It concerns kinds of *disagreement*, and it comes from the Greek philosopher Pyrrho of Elis (about 360 to 270 BC). He was not talking directly about knowledge; belief was his target. Nonetheless, his account had implications for knowledge. If one is to know something, one needs a belief or acceptance that it is true. (More is needed, too. I'll comment on that soon.) Pyrrho found himself noticing how our beliefs about what the world is really like rest on appearances – which always clash or disagree, it seems to transpire, with some other appearances. How do we ever sensibly trust one of these over others? We do not. We only appear to do so. The world doesn't strike me quite as it does – or so it seems – my dog or my neighbour, even a past or future neighbour. There is me-today versus me-a-week-ago. And what of me-well and me-ill, me-here and me-over-there, my eyes as against my hearing? There's what I am used to ('The First Fleet landed here in 1788') and what is surprising ('No, it didn't'). Nor must I forget a current favourite (yet also discussed, so long ago, by Pyrrho) – different cultural ideas about what should be done and what is morally right. Pyrrho offered more besides (as his views are portrayed in our chief source for them – Sextus Empiricus' *Outlines of Pyrrhonism*); but it's apparent where this leads. There's no way of choosing between conflicting appearances except by consulting appearances – possibly an impossibly unending process, and hardly an unbiased method. To Pyrrho, the moral seemed clear:

> Don't insist on the world's really being as it appears to be; which is to say, don't have beliefs about how the world really is. Simply notice appearances. Feel a resulting tranquillity.

But a lack of belief as to what is true is in turn a lack of knowledge. So, Pyrrhonian scepticism implies our never knowing how the world is.

That form of scepticism is still with us, in a less careful form. Instead of describing how a person may allow beliefs

to fade away, it concludes that no truths are available to *be* believed. (And if there is no truth, there is no knowledge, since knowledge is always of truth.) One cannot travel far within our society before hearing this sort of well-meaning but loose talk:

> Because people disagree so often as to what is true, no one can decide authoritatively what's true. ('Who's to say?') Better to walk away from truth-seeking. Conclude that there is no truth (only sincere opinions).

Thankfully, that is fallacious reasoning. Disagreement doesn't entail there being no truth. It entails at most that not everyone involved in the dispute is correct – a result compatible with some of them being right. An inability to agree on facts does not ensure there being no facts. (Ironically, that would follow only if the world forever awaits our collective views on it before taking shape. Social facts such as something's being money are like that. But is all reality social? Language is; not everything being described is. The word 'mountain' is; the existence of that rocky mass is not.)

Is there a better argument for a sceptical result? Here is another popular one:

> There is no certainty. Doubt is forever possible, because there cannot be an absolute or perfect viewpoint on reality, a rationally final perspective on truth. Hence, knowledge isn't possible.

In the seventeenth century, René Descartes was a philosophical champion of some knowledge's requiring certainty. Arguably, he required only metaphysical knowledge, not scientific knowledge, to be certain. In any case, often people do treat all knowledge in that demanding way. They infer that, because forever we lack God-like access to truths, we have no knowledge. Is it absolutely certain,

rationally unquestionable, that the First Fleet landed in Port Jackson (what is now Sydney Harbour) in 1788? We cite what we take to be excellent archival evidence for this; not fundamentally and finally certain evidence, though. Deception and delusion – small or large – have always been possible.

Fallible Knowledge

And so (I inferred, when musing upon all of this) there are ancient and not-so-ancient sceptical pressures upon us. We may absorb *this* much from Pyrrho (and from other philosophers since): Bow to our limits of awareness, of perception, of reason. Should I therefore have decided, pessimistically, to be persuaded by the sceptical thinking? Should I have concluded that there is no knowledge anywhere – returning in that way to my initial conceptual challenge about how to understand the library's role if there is no knowledge for it to store and impart.

Not yet, I decided; for it seemed to me that epistemology permits an alternative interpretation of the situation. I realised that I could regard anew the grand library in which I was seated: I could view it partly through an idea of *fallible knowledge*.

This is the idea of knowing a truth in a way which *allows the possibility* of being mistaken. To know fallibly is to be correct, using evidence which is good *yet not perfect*. Evidence can be like that even when one regards it as perfect. Or, open-mindedly, one might realise that one's evidence is not perfect. (Isn't this essential to being open-minded?) Even then, is knowledge still possible? (It would be fallible, not having required perfect certainty. And it would be knowledge of a truth, with the fallibility remaining just a potential for mistake.) Does that idea make sense?

I hope so, because I also wished to take seriously this suggestion: *If* we know, that's *how* we know – fallibly, that is. To know even a simple truth would be to do so fallibly.

The suggestion counsels humility, maybe realism, about how we know: never, anywhere, do we have even some infallible knowledge.

Should I accept that view of myself – as only ever knowing fallibly if at all? Can I conceive comfortably of myself as vulnerable, even when knowing, to being mistaken? The idea is abstract, but I believed that I could render it less so. I did this by looking around again at my august surroundings – but now with the potentially reassuring thought that *the library itself* is a model of the idea of fallible knowledge. The library manifests both elements of that idea – fallibility and knowledge.

How so? Such a large long-lived library is full of conflicting books, competing ones – books locked in disagreement forever, staring disapprovingly at each other across the shelves. But I accepted earlier that conflict doesn't entail all parties being mistaken: truth can coexist with disagreements about its location. Some books could be correct even amidst such disagreement. Those books are like large and complex correct beliefs. The well-researched ones also contain good evidence for their correct conclusions. Hence, those books contain knowledge, such as about the First Fleet or Federation. In theory, they could be so knowledgeable as to include only knowledge – saying nothing false, conveying it learnedly.

Yet would that knowledge be certain? Infallible? Perfect? No. It remains fallible even while knowledge (and even within a book saying nothing false). That fallibility is recognised in practice by how the library functions. For a start, the library *keeps* many competing books. I can imagine everyone's accepting a particular book, to the exclusion of a conflicting one, regarding how Federation unfurled. Even so, the book with no advocates might well not be discarded by the library. Why not? Partly, because the book with no current adherents could still be correct. Although (in that imagined circumstance) we do not *think* of that book as being correct, we may acknowledge that our attitude might

subsequently change. So, by retaining the disfavoured book, we would be accepting the favoured book along *with* a concession of our present fallibility. We might – who knows when? – come to decide – who knows how? – that there is truth in the currently disfavoured book's alternative story about Federation. New times, new thoughts: we accept this possibility; the library is an institutionalised way of accepting it, too.

I also tried to imagine a library that *aims* to contain only wholly true books. Would a good library really strive to discard relentlessly whatever seems false? Surely not. There is rightness in keeping some, even many, books we believe to be mistaken. And the explanatory model behind this assessment isn't one of slavishly gripping old unloved drafts of a book. The model is more one of retaining examples of what we think we should not believe – accepting that these could turn out, after all, to be what we should believe. We would be keeping examples of what we might one day aptly come to believe – because we realise we don't always believe what, we later believe, we should have believed.

Libraries include knowledge, we hope. *If* they do, it is partly because they incorporate fallibility. By storing so much, they are equally our chance to check – fallibly – on whether we know fallibly. We would know by remaining vigilant about the real possibility of not knowing.

In that sense, libraries *aim* to include not only knowledge but also fallibility. This makes them useful and human. They model us in this respect. They *are* us writ large in this way. With effort, we accurately see ourselves in them. The result is some fundamental self-knowledge – all the while still fallible, of course.

Relief

. . . and so I arose, relieved; and I rejoined the world outside the library. But I was changed. Now I understood – I understand – better how the library is *like* the world

outside it. It is like *us*, at least as we may be when knowing that world around us. As we can know the world, so can the library. And as we can know the library, so can we know the world. All is fallible; there might be knowledge, even so.

Stephen Hetherington is Professor of Philosophy at the School of Humanities, University of New South Wales.