A Russellian Philosophy of Philosophy: Systematicity, Entitlement, and the End of Days

Samuel Lebens

My project, in this chapter, is to excavate a number of Bertrand Russell's various statements regarding philosophical method, and the epistemology of philosophy. What, for Russell, constitutes good philosophical practice? What does philosophy aim at? How are its theories to be justified? In short, what was Russell's philosophy of philosophy?

This excavation will yield something like a rough and ready ore. From that ore, a philosophy of philosophy can be extracted. The raw materials that we'll find in the primary sources contain internal tensions; imperfections. My aim is to smelt away the imperfections and produce something out of these raw materials; to present the final, polished result, as a plausible philosophy of philosophy, worthy of consideration in its own right; irrespective of historical pedigree, or fidelity to the original textual terrain from which it was mined.

In §1, I draw out an abductive account of philosophical methodology from Russell's *Problems of Philosophy*. In §2, I raise and dismiss a problem that emerges from attributing this account to Russell. It seems to (but it doesn't) misconstrue Russell's attitudes towards common sense. In §3, I raise a more difficult problem: the account I draw in §1 conflicts with Russell's oft-stated foundationalist aspirations. I sketch, and dismiss, two possible solutions to this problem, before concluding that there is, indeed a tension in Russell's philosophy of philosophy. It has to be smelted away to give rise to something that could be called "foundherentist" (borrowing a label that Susan Haack (1997) applies to her own account of epistemic justification).

In §4, I interrogate our newly purified metal on various fronts: what role should philosophers give to their intuitions, what is the epistemic status of our metaphysical convictions, and will there be an epistemic end of days?

§1. An Abductive Methodology

Russell's *Problems of Philosophy* was dismissed by Wittgenstein as a "shilling shocker."[1] I reject Wittgenstein's attitude. But, whatever your attitude to the book as a whole, its conception of the philosophy of philosophy, I shall argue, is worthy of serious consideration.

Having presented Berkeley's denial of a mind-independent reality, Russell's main line of resistance to it stems, not from argument, but from the existence of his doggedly persistent belief in exactly that which Berkeley denies: "We find this belief [in a mind-independent reality] ready in ourselves as soon as we begin to reflect: it is what may be called an instinctive belief" (Russell, 1912, p. 11).

Originally, and especially in the modality of sight, it seems "as if the sense-datum itself were instinctively believed to be the independent object" that we're looking at (Ibid.). In other words: as we look at the table, we think that we're seeing the table itself. But the sensible table changes colour and shape as we move around it, which isn't something we imagine to be true of the *table*, and thus "argument shows that the object cannot be identical with the sense-datum"(Ibid.). Accordingly, Russell is willing to depart from his naïve starting point. But, to go as far as Berkeley would be to go too far.

The discovery that our sense-datum must be distinct from its mind-independent cause needn't lead us to deny altogether that its mind-independent cause *exists*. In fact:

> Since this belief does not lead to any difficulties, but on the contrary tends to simplify and systematize our account of our experiences, there seems no good reason for rejecting it. We may therefore admit — though with a slight doubt derived from dreams — that the external world does really exist, and is not wholly dependent for its existence upon our continuing to perceive it.
>
> (Ibid.)

In short, Russell comes to the world with an instinctive belief in a mind-independent reality. Perhaps this is a function of nature, or nurture, or both. That doesn't seem to matter. All that matters is that the belief is there.

[1] (Russell, 1913, p. xix)

To be fair, Berkeley had tried to demonstrate that the belief in mind-independent matter was *incoherent*. Had Russell accepted these arguments, he would have given up his "instinctive belief". But Russell saw no good argument from Berkeley. According to Russell, Berkeley had shown us nothing more than the fact that "the existence of matter is capable of being denied without absurdity" (Ibid., p. 4). That's not a good enough reason for Russell to give up on an instinctive belief.

We needn't give up an instinctive belief for "no good reason" (Ibid., p. 11). Especially if it can play a role in "simplify[ing] and "systematiz[ing] our account of our experiences" (Ibid.) Russell goes further:

> All knowledge, we find, must be built up upon our instinctive beliefs, and if these are rejected, nothing is left. But among our instinctive beliefs some are much stronger than others, while many have, by habit and association, become entangled with other beliefs, not really instinctive, but falsely supposed to be part of what is believed instinctively.
>
> Philosophy should show us the hierarchy of our instinctive beliefs, beginning with those we hold most strongly, and presenting each as much isolated and as free from irrelevant additions as possible. It should take care to show that, in the form in which they are finally set forth, our instinctive beliefs do not clash, but form a harmonious system.
>
> Hence, by organizing our instinctive beliefs and their consequences, by considering which among them is most possible, if necessary, to modify or abandon, we can arrive, on the basis of accepting as our sole data what we instinctively believe, at an orderly systematic organization of our knowledge, in which, though the *possibility* of error remains, its likelihood is diminished by the interrelation of the parts and by the critical scrutiny which has preceded acquiescence.
>
> This function, at least, philosophy can perform.
>
> (Ibid., pp. 11-12)

What has emerged from these quotes is that philosophy is, ultimately, about marshalling the data – empirical data, sense-data, instinctive beliefs – into an harmonious whole. We start with the dictates of common sense, centre stage, and those beliefs that survive the system-building and harmonizing scrutiny of philosophy, will have earned their keep. Our end result

may be quite far from our starting point, but our starting point, at least, is common sense. As Russell put it elsewhere:

> The point of philosophy is to start with something so simple as not to seem worth stating, and to end with something so paradoxical that no one will believe it.
>
> (Russell, 1918, p. 193)

Alternatively, and much to the same affect:

> The process of sound philosophizing, to my mind, consists mainly in passing from … obvious, vague, ambiguous things, that we feel quite sure of, to something precise, clear, definite, which by reflection and analysis we find is involved in the vague thing that we started with, and is, so to speak, the real truth of which that vague thing is a sort of shadow.
>
> (Ibid., pp. 179-180)

Very often, Russell's philosophizing would find philosophically perspicuous paraphrases for our ontologically cumbersome colloquialisms. Post logical analysis, for example, he would allow us to carry on talking about Hamlet and Zeus, as *façons de parler*, without incurring ontological commitment to fictional beings, or Greek gods. This fact perhaps explains, in part, how Russell could say, "I am prepared to admit the ordinary beliefs of common sense, in practice if not in theory" (Russell, 1928, p. 2). He accepts them in practice, because he's found theoretical paraphrases that render the sentences we use to assert them acceptable, as figures of speech. Perhaps this explains how Russell, in his popular writings, could champion himself, despite the logical complexity of his formal work, as a warrior for common sense (see, for example, Russell, 1959).

Not only are we counselled to start from our "instinctive beliefs"; our passage from the *Problems of Philosophy* (Russell, 1912, pp. 11-12) contends that we should set out to *save* as many of them as we can. The system that we're able to build that is most simple, most explanatory, *and* that holds onto the most of our most instinctively held instinctive beliefs, is the system that we're counselled to construct.

§2. Common Sense and Savages

A problem: Russell had a much more complicated relationship to intuition and common sense than the relatively uncritical attitude described thus far. But this problem evaporates on further reflection.

Of course, according to Russell, we can *misunderstand* the content of our own instinctive beliefs; we can – for example – misconstrue their logical form. Revealing the precise logical contours of our vague and ambiguous instinctive beliefs, to uncover something "precise, clear, [and] definite," is part of the process that Russell endorses (Russell, 1918, p. 180). Another part of the process is to organise all of our beliefs into a coherent *system*. Both activities (the logical analysis of our beliefs, *and* the effort to render them consistent, coherent, and explanatory) will require us, in time, to jettison all sorts of prejudices.

And thus, at the same time as being an advocate of common sense, Russell can coherently condemn the *unpolished* deliverances of common-sense as the "metaphysics of savages." Russell didn't actually coin this (somewhat indelicate) phrase, commonly misattributed to him. This slur on common sense was *actually* coined by Douglas Clyde MacIntosh (1915, p. 243), in an effort to *encapsulate* Russell's attitude.[2] MacIntosh was probably intending to crystallise into slogan form the following words of Russell:

> Physics started from the common-sense belief in fairly permanent and fairly rigid bodies – tables and chairs, stones, mountains, the earth and moon and sun. This common-sense belief, it should be noticed, is a piece of audacious metaphysical theorizing; objects are not continually present to sensation, and it may be doubted whether they are there when they are not seen or felt. This problem, which has been acute since the time of Berkeley, is ignored by common sense, and has therefore hitherto been ignored by physicists. We have thus here a first departure from the immediate data of sensation, though it is a departure merely by way of extension, and was probably made by our savage ancestors in some very remote prehistoric epoch.
>
> (Russell, 1914, p. 107)

[2] For this correction to the historical record, I'm indebted to Stephen Leach (2020).

Russell is here sensitive to what later cognitive scientists would strive to uncover with more precision, namely, the existence of various theories of the world, somehow hard-wired into our cognitive architecture as an evolutionary endowment. The fact that these theories were produced in order to serve our Palaeolithic ancestors should give us pause for thought. Their environment, and their immediate interests, were very different to ours. But, this disdain for *unreflective* common sense doesn't undermine the philosophical methodology excavated in §1. It should merely encourage us to scrutinise thoroughly and with alacrity the deliverances of common sense. Those instinctive beliefs that survive analysis and system building, we have "no good reason for rejecting" (Russell, 1912, p. 11); whatever their biological provenance.

Instinctive beliefs have no right to automatic adoption. Rather, their instinctiveness generates something like a legitimate aspiration, on our part, to create a philosophical system that can salvage them, or, at least, some post-analysis precisification of them. Why? Perhaps because the most we can hope to do, as philosophers, is to render our beliefs into a coherent, consistent, elegant, explanatory, whole. If this can be done in such a way as to preserve our most instinctive beliefs, then – once again – we have "no good reason for rejecting" them (Ibid.). This could be motivated by some sort of epistemic conservativism, according to which, simply having a belief is *prima facie* reason for believing it. Alternatively, it might be motivated by a sort of fallibilism, according to which the philosopher knows that we cannot be certain that the final deliverances of our philosophising will be free of error. Coherence and explanatory power are perhaps the most that we can hope for, and if we can have these desiderata without getting rid of beliefs that we somehow cherish, then why shouldn't we?

As we shall see, in §4.2, I think Russell grants us permission to hold on to our surviving instinctive beliefs neither in recognition of our fallibilism nor in virtue of a crude conservativism. But we'll cross that bridge when we come to it.

Either way, the vision for philosophy that emerges is one that gives centre stage to common-sense as one of the many inputs to a reflective equilibrium that we hope for philosophy someday to achieve. This chimes especially well with Fraser MacBride's account of the "eschatological burden" and "synoptic commission" of philosophy, namely: "to establish whether, or to what extent, the final outputs of physics, mathematics, psychology, political science and all the other special disciplines can be fused with common sense into one single periscopic vision of the Universe" (MacBride, 2014, p. 228).

§3. A Tension with Foundationalism

Another problem with attributing to Russell the model of philosophy that we've been developing so far, is its incompatibility with the foundationalism that Russell seems to have held at exactly the same time. For example, witness another excerpt from *Problems of Philosophy*, in which Russell outlines how we are supposed to arrive at "derivative knowledge" from "intuitive" foundations, and here, "intuitive" doesn't mean *instinctive*, it means *self-evident*:

> Our derivative knowledge of things, which we call knowledge by *description,* always involves both acquaintance with something and knowledge of truths. Our immediate knowledge of *truths* may be called *intuitive* knowledge, and the truths so known may be called *self-evident* truths. Among such truths are included those which merely state what is given in sense, and also certain abstract logical and arithmetical principles, and (though with less certainty) some ethical propositions. Our *derivative* knowledge of truths consists of everything that we can deduce from self-evident truths by the use of self-evident principles of deduction.
>
> (Russell, 1912, p. 63)

Either (1) we start our philosophy with the aspiration of organising our hodgepodge of instinctive beliefs, alongside other inputs, into something more harmonious, simple, and explanatory, *or* (2) we clear away all of the clutter of our instinctive beliefs and start our philosophy, Descartes-style, in search of self-evident axioms from which we can arrive at derivative knowledge through a process of truth-preserving inferences. Which is it to be? It can't be both.

There are three ways in which we can relate to Russell's foundationalist aspirations, so as to resolve the conflict with the model we've been building up until now:

- **Option 1.** We could distinguish two projects – a coherenetist project for metaphysics (for regulating *opinion*), and a foundationalist project for science (for generating *knowledge*).
- **Option 2.** We could accuse Russell of getting carried away with his own rhetoric, in his shilling shocker. Accordingly, we should either:

a. take his talk of instinctive beliefs with a pinch of salt and recognise him for the foundationalist that he really was, or
b. take his talk of self-evidence and foundationalism with a pinch of salt, and recognise him for the coherentist that he really was.

Option 3. Recognise that Russell endorsed neither a pure form of foundationalism nor a pure form of coherentism but some sort of hybrid of the two.

Let's investigate the evidence for these options in turn.

§3.1. Option 1: Two Projects

According to option 1, we need to distinguish between metaphysics[3], on the one hand, and science (and logic) on the other. Metaphysics can only give rise to *opinion*. Science, by contrast, aided by logic, can give rise to knowledge. Russell gives voice to this distinction explicitly:

> I believe the only difference between science and philosophy is, that science is what you more or less know and philosophy is what you do not know. Philosophy is that part of science which at present people choose to have opinions about, but which they have no knowledge about. Therefore every advance in knowledge robs philosophy of some problems which formerly it had, and if there is any truth, if there is any value in the kind of procedure of mathematical logic, it will follow that a number of problems which had belonged to philosophy will have ceased to belong to philosophy and will belong to science.
>
> (Russell, 1918, p. 281)

In the epistemic end of days, our philosophy will be the set of mere opinions that survive from among our instinctive beliefs in the furnace of system building. Our *science*, however, guided by mathematical logic, will have refined all of our other beliefs, basing them all, perhaps, upon a foundation of self-evident axioms, such that they shall be worthy of being called *knowledge*, and *science,* rather than *opinion*, and *philosophy*.

[3] I shall use 'metaphysics' and 'philosophy' interchangeably, not because I have a narrow view of the latter, but because I have a broad view of the former.

The role that logic plays, in aiding movement from opinion to knowledge, is that logic can help us to uncover the objective contours of reality. As Russell puts it:

> I think one might describe philosophical logic… as an inventory, or if you like a more humble word, a "zoo" containing all the different forms that facts may have.
>
> (Russell, 1918, p. 216)

We can use logic to analyse the content of our beliefs and, in so doing, we can reveal something about the world beyond, namely: the different forms that facts may have (since beliefs, to be truth-apt, must share some sort of structure with the facts that they seek to represent). This process can be one of the routes by which opinion gives way to knowledge: by scrutinizing the logical forms of our opinions, we can learn something about the contours of reality. Moreover, revealing the logical form of our beliefs can help us to uncover their inferential relations. This in turn can help us to discover which of our beliefs follow from self-evident axioms (and which don't).

Accordingly, when Russell describes an abductive method of building a worldview, by marshalling instinctive beliefs into an harmonious whole, we could read him as describing a methodology for *philosophy*; a method that regulates (instinctive) *opinion*. When he describes a more foundationalist project that moves in truth-preserving steps only from self-evident axioms, then he's describing a methodology for the *sciences*; which generates *knowledge*. Further evidence that option 1 is on the right track is that Russell actually distinguishes, in the *Problems of Philosophy*, between two forms of self-evidence:

> [T]wo different notions are combined in 'self- evidence' … one of them, which corresponds to the highest degree of self-evidence, is really an infallible guarantee of truth, while the other, which corresponds to all the other degrees, does not give an infallible guarantee, but only a greater or less presumption…
>
> (Russell, 1912, p. 68)

Later, he clarifies that:

> We may say that a truth is self-evident, in the first and most absolute sense, when we have acquaintance with the fact which corresponds to the truth…
>
> (Ibid., p. 79)

In some loose-sense of the word, Russell would say that it's self-evidently the case that there exists a mind-independent world. But all that that means is that it's intuitively the case; or that he finds himself with an instinctive belief that it's the case. If our starting point, in philosophy, is just a set of propositions with this *loose* form of self-evidence, then we're actually a long way from classical foundationalism. The ideal would be to start with *absolutely* self-evident propositions. This absolute form of self-evidence occurs when a person is acquainted with "the fact which corresponds" to the truth in question (Ibid.).

The picture that emerges is as follows. There are two projects in which we could engage. One starts from *loosely* self-evident propositions, and ends with an harmonious whole. Its contents have survived scrutiny, and thus we're entitled to believe them, albeit with an attendant sense of doubt (Ibid., p. 11-12). The second project starts from *absolutely* self-evident propositions, moving in truth-preserving steps, towards omniscience. In fact, Russell seems to distinguish these two projects explicitly. Having laid out the contours of the first (and less ambitious) project, he writes:

> This function, at least, philosophy can perform. Most philosophers, rightly or wrongly, believe that philosophy can do much more than this — that it can give us knowledge, not otherwise attainable, concerning the universe as a whole, and concerning the nature of ultimate reality. Whether this be the case or not, the more modest function we have spoken of can certainly be performed by philosophy, and certainly suffices, for those who have once begun to doubt the adequacy of common sense, to justify the arduous and difficult labours that philosophical problems involve.
>
> (Ibid., p. 12)

In other words, if philosophy can achieve nothing more than the production of a coherent, simple, elegant, and explanatory set of *opinions*, then that's already quite an achievement. But, that doesn't rule out a more ambitious project worth striving for, the ambition outlined in the later, foundationalist passages of the same book; a project which starts with absolute self-evidence and ends with knowledge.

There are two prizes in wait for philosophers, when we reach the epistemic end of days. The greatest prize would go to us were we to construct a "periscopic vision of the Universe," to

borrow MacBride's turn of phrase, that integrates all of the findings of every science, and also manages to derive them all from a set of absolutely self-evident axioms. If it does *that*, then, in a sense, philosophy will have been a victim of its own success. If philosophy becomes science when opinion becomes knowledge, then philosophy's grand prize in the eschaton is its own happy suicide, when all opinion is upgraded to knowledge; no distinctively philosophical topic will be left unconverted into science.[4]

Failing the top prize, we can, at least, hope to construct a "periscopic vision of the Universe" that integrates the findings of all of the sciences, with the dictates of common sense, without a complete axiomatization from *absolutely* self-evident propositions. This integrated whole will give appropriate weight to our most instinctive beliefs, and jettison the opinions that don't find a place in the whole. Some subset of this system might count as knowledge; *i.e.,* that subset grounded upon absolutely self-evident premises, but some of it will remain mere opinion – opinion that we're somehow entitled to hold; given its survival. Philosophy will be complete, but it will not pass away.

The first prize is a foundationalist periscopic vision. The second prize is a coherentist periscopic vision. The first prize is knowledge and science. The second prize is opinion and philosophy. The first prize is an infallible worldview. The second prize is aware of its fallibility.

This picture, I call option 1. I have little doubt that it captures an historically accurate thread of Russell's philosophy of philosophy. But there's a fatal problem at its heart. Here's the basic issue. Some of the propositions that Russell had at one point been willing to accept as maximally self-evident, he would later be willing to deny. So, if maximal self-evidence is acquaintance with a fact, then it turns out that we can be mistaken as to whether we are in such an epistemic state. Accordingly, it turns out that even maximal self-evidence isn't all that transparent; and so, in a crucial sense, it isn't self-evidence! Accordingly, we can't make the distinction between infallible science and fallible philosophy that Russell evidently hoped to make. And indeed, in other writings, Russell was well aware of this problem. The ore from which we're trying to mine a consistent philosophy of philosophy is simply conflicted.

[4] Of course, I borrow the notion of philosophy's happy suicide from Bradley's quite different but structurally analogous notion of thought's happy suicide (Bradley, 1897, p. 173).

One way to see the problem is to recognise that Russell's self-evident propositions aren't supposed to be *obviously* true. This is where the fallibility enters the picture, even for the foundationalist project. For Russell, it takes a lot of work to arrive at self-evident axioms. The axioms that underlie the proposition that 2+2=4, for example, are much more recondite than the proposition itself, and Russell was well aware of this fact (Russell, 1907, p. 272). We start, remember, from simple beginnings, and dig down until we find something that would strike the ordinary person as paradoxical.

Indeed, an entire chapter of Russell's aborted *Theory of Knowledge* manuscripts, of 1913, is dedicated to an analysis of self-evidence (or what, in 1912, he would have called maximal or *absolute* self-evidence). Self-evidence, he says:

> … is a property of judgements, consisting in the fact that, in the same experience with themselves, they are accompanied by acquaintance with their truth.
>
> (Russell, 1913, p. 166)

On this view, the proposition that 2+2=4 might be much easier to grasp than the axioms of logic. But, a person who is in the position to *assert* an axiom of logic, will, in the very same experience as asserting it, become acquainted with its truth. It takes a lot of work to get to that position. The axioms aren't obvious to those who haven't toiled to arrive at them. But they are, Russell would insist, self-evident to those who are in a position to assert them.

That's all well and good, but if self-evident propositions *aren't* obvious, then it's difficult to accept that we can't be mistaken about them. Surely Russell would have been willing to call the axioms of naïve set-theory maximally self-evident, before he discovered the paradox that bears his name. He may even have thought that he was acquainted with their truth in the very act of asserting them. But, by his own lights, he turned out to be wrong. Russellian self-evidence isn't self-evident. His foundationalist project is therefore built upon (potentially) sinking sand. How are we ever to know, infallibly, that our seemingly self-evident foundations are actually self-evident?

Another problem with this model is that Russell was aware that logical analysis itself – the very tool that would help us move from obvious truths to their self-evident axiomatic base – would have to be regulated by what he called a vivid sense of reality. For instance, this is how

Russell motivated his logical analysis of definite descriptions, which helped him avoid ontological commitment to Hamlet, unicorns, and Zeus:

> The sense of reality is vital in logic, and whoever juggles with it by pretending that Hamlet has another kind of reality is doing a disservice to thought. A robust sense of reality is very necessary in framing a correct analysis of propositions about unicorns, golden mountains, round squares, and other such pseudo-objects.
>
> (Russell, 1919, p. 170)

Witness also how he justified his logical analysis of assertion, which helped him avoid ontological commitment to the existence of false propositions, or to non-obtaining states of affairs. About false propositions, he wrote:

> I cannot believe they go about the real world. It is more than one can manage to believe, and I do think no person with a vivid sense of reality can imagine it. One of the difficulties of the study of logic is that it is an exceedingly abstract study dealing with the most abstract things imaginable, and yet you cannot pursue it properly unless you have a vivid instinct as to what is real … I think Meinong is rather deficient in just that instinct for reality.
>
> (Russell, 1918, p. 223)

Presumably, a vivid sense of reality is something of a regulatory device. Unlike our instinctive *beliefs*, this sense can't be reduced, at any given time, to a set of propositions. It is, rather, a *sense*. Just as our instinctive beliefs should be subjected to scrutiny, especially in light of their biological origins in our Palaeolithic ancestors, so too, I assume, should a thorough going Russellian be willing to subject our *sense of reality* to scrutiny – in the light of contemporary cognitive science. How does this *sense* feature in our cognitive architecture, and how did it get there? Nevertheless, it would be hard to proceed in our logical investigations *without* such a sense.

To summarise: Russell's self-evident axioms are not self-evidently self-evident, and his logical analyses, which are supposed to help us transform opinion into science, were to be regulated by nothing more than a *sense* of reality, a *common sense*. And thus, I'm happy to accept, with option 1, that a foundational project co-exists alongside a less ambitious abductive project, in the ore that we've mined from Russell's corpus. But, I would argue that, for philosophical

reasons, rather than historical ones, that foundational project needs to be expunged; it is bound to collapse into something abductive and fallible.

§3.2. Option 2: Regression

Option 2 denies that the aspiration for a foundationalist first prize, backed up by the promise of a coherentist second prize, was ever a serious aspiration in Russell's more considered work. *The Theory of Knowledge*, remember, was never published, and his shilling shocker was just a shilling shocker. Russell's more considered work, so option 2 would have it, was unambiguously committed either to (a) a thoroughgoing foundationalism that moves only from self-evident axioms, or (b) an abductive method that eschews any appeal to self-evidence, but never to the idea that we might strive somehow towards both, in the hope of achieving one or the other.

Option 2 (a) is probably the picture of Russell that captures the popular imagination given his commitment to logicism. Surely, the whole idea of trying to found mathematics upon the axioms of logic is to place mathematics on solid foundations. Beyond the mathematical realm, Russell's attempts to ground our knowledge of the mind-external world in logical constructions out of the sense-data with which we're directly acquainted also betrays his foundationalist aspirations. Forget *Problems of Philosophy*, and his unpublished work – we all know that Russell was a champion of foundationslim. This is the picture of Russell that emerges from numerous portraits: *e.g.*, Pears (1967, ch. 3), Eames (1969, p. 83), Ayer (1971, p. 12), and Sainsbury (1979, pp. 193-200), to name a few.

Swimming against this tide, option 2 (b) has been most thoroughly developed by Andrew David Irvine (1989). Based on a close reading of multiple sources, Irvine seeks to demonstrate that Russell's epistemology was never foundationalist at all, rather: "mathematical knowledge begins in the first instance from particular observations, e.g., the observation that two objects together with two distinct objects are four objects" (Ibid., p. 312). From observations such as those, we arrive at generalizations, such as the generalization that 2 + 2 = 4. As Russell readily admits, these generalizations are "sufficiently obvious" to be taken, by most people, as assumed premises in deductive proofs (Russell, 1907, p. 275). But, in addition

to the deductive proofs that move outwards from these generalisations, Russell advocates a body of mathematical knowledge that we can only reach "regressively" (Ibid.).

Deductive mathematical knowledge proceeds from obvious premises, such as that 2+2=4, to their logical consequences. But, in order to arrive at the *axioms* of mathematics, we have to engage in the regressive method, which Russell also calls "induction" – although abduction might have been a better term. This regressive method infers "premises from consequence… thus the method in investigating the principles of mathematics is really an inductive [or, better, abductive] method, and is substantially the same as the method of discovering general laws in any other science" (Ibid., p. 273f).

Far from being self-evident, the axioms of logic are discovered much as we would discover the general laws of a science. We look at the body of knowledge that we've gathered from observation and try to find the smallest set of principles from which those observations could be derived. The law of non-contradiction, for example: "[M]ust have been originally discovered by generalizing from instances, though, once discovered, it was found to be quite as indubitable as the instances" (Ibid., p. 274).

Unlike the distinctions that Russell draws in his more sloppy moments, between (1) science and (2) philosophy, (1) infallible knowledge and (2) fallible opinion, (1) absolute self-evidence and (2) relative self-evidence, Russell, in his more considered writings, recognises that no such distinctions can be drawn:

> Infallibility is never attainable, and therefore some element of doubt should always attach to every axiom and to all its consequences. In formal logic, the element of doubt is less than in most sciences, but it is not absent, as appears from the fact that the paradoxes followed from premisses which were not previously known to require limitation.
>
> (Russell & Whitehead, 1910-13, vol. 1, p. 59 (p. 62 of the first edition))

In fact, immediately prior to the excerpt quoted above, Russell and Whitehead wrote that "self-evidence is never more than a part of the reason for accepting an axiom, and is never indispensable." Their logicism was, therefore, never intended as an "epistemic logicism" that would demonstrate all of mathematics to follow from self-evident axioms. In fact, as is

evident from the previous quote, self-evidence was neither thought to be necessary nor sufficient for attaining an axiomatic status in the logic of *Principia*. Moreover, they wrote:

> If [an] axiom is apparently self-evident, that only means, practically, that it is nearly indubitable; for things have been thought to be self-evident and have yet turned out to be false.
>
> (Ibid.)

It's true that these quotes are resolutely opposed to the passages we saw in previous sections, which tended towards a form of foundationalism. These quotes also conflict with Russell's popular auto-biographical works which present his work in mathematics as a constant search for *certainty.*[5] But why should we accept Irvine's conclusion that these quotes reflect the real Russell, and that the other quotes reflect his more considered position?

If that was the case, and if there wasn't really a burning desire for the sort of knowledge that a foundationalist could endorse, then why did Russell struggle, throughout the summer of 1913, to make sense of self-evidence? Moreover, in the second edition of *Principia,* Russell rejects the axiom of reducibility, because it simply seemed too *ad hoc* to deserve to be called an axiom. This implies that Russell was never totally comfortable with non-self-evident axioms.

To be fair to Irvine, in a footnote, he concedes: "Russell apparently never gave up the hope of deducing such axioms from other, more self-evident logical truths. For example, see the Introduction to the second edition of *Principia*, Vol. 1, p. xiv" (Irvine, 1989, p. 323, f. 42).

Perhaps we'd get closer to the heart of the matter if we move onto option 3 and recognise that Russell was trying to create a hybrid between foundationalism and coherentism.

§3.3. Option 3: Russell's Bridge

The idea that Russell's actual philosophical method comprises some sort of hybrid between foundationalism and coherentism has been defended by Dustin Olson and Nicholas Griffin (2018). To this end, they focus on a key passage from Russell's 1948 book, *Human Knowledge:*

[5] See, for example (Russell, 1956, p. 53)

Its Scope and Limits which coheres beautifully with both the foundationalist and coherentist threads of his earlier work. The passage in questions reads:

> Given a number of propositions, each having a fairly high degree of intrinsic credibility, and given a system of inferences by virtue of which these various propositions increase each other's credibility, it may be possible in the end to arrive at a body of interconnected propositions having, as a whole, a very high degree of credibility. Within this body, some are only inferred, but none are only premises, for those which are premises are also conclusions. The edifice of knowledge may be compared to a bridge resting on many piers, each of which not only supports the roadway but helps the other piers to stand firm owing to interconnecting girders. The piers are the analogues to the propositions having some intrinsic credibility, while the upper portions of the bridge are the analogues of what is only inferred. But although each pier may be strengthened by the other piers, it is the solid ground that supports the whole, and in like manner, it is intrinsic credibility that supports the whole edifice of knowledge.
>
> (Russell, 1948, p. 413)

This is what Olson and Griffin call "Russell's bridge." If other readings of Russell force us to choose between his foundationalist and coherentist moments, as if these two threads of his thought were never purposefully embraced together, this excerpt should lead us to believe that there is no such choice confronting us. Russell's bridge is neither foundationalist nor coherentist. Instead, it borrows elements from both. Moreover, as far as Olson and Griffin are concerned, its evasion of easy characterization "is one of its most significant merits" (Olson & Griffin, 2018, p. 293).

Olson and Griffin categorise Russell's bridge among a number of contemporary theories of epistemic justification that bring together elements of both foundationalism and coherentism, including, "Susan Haack's *foundherentism* (1993; 1997), Earl Conee's *foundational coherentism* (1988), Kevin McCain's *explanationist evidentialism* (2014), Catherine Elgin's *holism* (1996; 2005), and William Lycan's *explanationist coherentism* (1988; 2012)" (Ibid., p. 294). These theories differ in detail, but the commonalities they share with each other and with Russell's bridge are numerous:

> First, they each accept a fallibilist conception of knowledge: we may never have certainty, but we can have knowledge. Furthermore, they each recognize that the two traditional theories [of epistemic justification, namely: coherentism and foundationalism] needn't be at odds concerning their core contributions to justification. On the one hand, they all adopt the foundationalist's contentions that justified beliefs must be suitably attuned to one's experiences and that a coherent set of beliefs cannot be justified if that set doesn't cohere with one's experiences; on the other side, they adopt the coherentist's contentions that coherence contributes to a belief's justification and that no autonomously justified belief is required for coherence to provide this justificatory work.
>
> (Ibid.)

Olson and Griffin would argue that Russell's foundherentist espitemology (to borrow the term from Susan Haack) was a relatively constant feature of his otherwise fast changing philosophy over the years. It explains why at times, within in one work, he sounds like a foundationalist, and other times he sounds like a coherentist.

Russell endorsed a form of foundherentism in the 1940s. I don't think I can be as charitable as Olson and Griffin, so as to think that this foundherentism was there from the start. If it was, I don't think that he would have worked so hard to give an account of absolute self-evidence, or to make his seemingly hard-and-fast distinction between philosophical-opinion (without absolutely self-evident foundations), on the one hand, and scientific-knowledge (built entirely upon self-evident foundations), on the other.

I think we're on safer ground simply to concede that, before the 1940s, Russell was torn between a regressive (or abductive) conception of logic, and a foundationalist conception of logic; between the abductive vision of a philosophy that never succeeds in eliminating itself, and the foundationalist conception of a philosophy that one day commits happy suicide, as opinion gives way to knowledge. Foundherentism on this view, was the wise and happy compromise that Russell arrived at only after years of internal tension.

§4. Conclusions and Eschatology

Once we refine the metal that we've smelted from our Russellian ore, we arrive – as it seems did Russell, in the 1940s – at a view that we could call foundherentist. We start with instinctive beliefs. We subject them to scrutiny. We try to marshal them together, with the findings of all of the sciences, into a coherent and harmonious, explanatory system. We use logical analyses, regulated by a vivid sense of reality, in order to impose an inferential order upon this system (even if that sense of reality has to be scrutinised too).

The emerging order of the system we build will help us to arrive at a set of axioms for much, if not for all of our knowledge, although the axioms themselves will often be, like the general laws of any science, less obvious than the observations they underpin and explain. Moreover, nothing much, if anything, is worthy of being called *self-evident*.

There are three aspects of this purified Russellian philosophy of philosophy that I'd like to explore. The first is its relationship to intuition. The second is its epistemic safety, an account of which will have to appeal to notions of systematicity and entitlement. The third is its answer to the following question: will there be an epistemic end of days; will philosophy ever arrive at consensus?

§4.1. The Role of Intuition

There is a tendency among contemporary analytic metaphysicians to place great weight upon intuition. Kripke tries to imagine a table retaining its identity without being composed of molecules, and he can't, so he concludes that tables are *essentially* made of molecules (Kripke, 1972, p. 47). He tries to imagine a human being retaining its identity whilst emerging from a different sperm-cell and/or egg (Ibid., p. 113). He can't, and so he arrives at his doctrine of origin essentialism. One might think that the centrality of intuitions like these would be justified by the Russellian philosophy of philosophy outlined in this paper. But that wouldn't be fair. As MacBride makes clear:

> Russell and Quine already understood that even our most assured and unreflective beliefs cannot be taken for granted. They recognized that thought and language aren't transparent so even here there is a pressing question of appearance and

> reality; but a significant minority of subsequent philosophers following Kripke have evidently forgotten.
>
> (MacBride, 2014, p. 235)

Russell's philosophy of philosophy grants us no license to hold onto our unreflective beliefs without question. Rather, we have to subject every belief to tremendous scrutiny, subjecting them all to logical analysis in order to discover their underlying form, and try to find a place for them within a system of thought. The only weight that's given to the instinctiveness or intuitiveness of a belief is that it creates a *prima facie* reason to try harder to find a place for it in our final system. That's all.

§4.2. Systematicity and entitlement

The second issue I want to address is the account of epistemic safety we can marshal for this Russellian philosophy of philosophy. At first flush, our Russellian meta-philosophy seems committed to some form of doxastic conservativism (as I mentioned in §2), according to which the mere fact that you find yourself believing some proposition counts as *prima facie* justification for believing that proposition. How safe can such a procedure be?

If I find myself believing that there are an even number of craters on the moon, and if that proposition survives my best attempts to scrutinise and system build, then it seems as if Russell would allow me to believe it. But why? Merely finding myself to be believing something, without even remembering *why* I believe it, seems to offer me no good reason for *continuing* to believe it!

A better suggestion would be to align our philosophy of philosophy with a *phenomenal* conservativism. According to this epistemology, although a belief receives no justification merely by being *held*, its belief *does* receive some justification from a *phenomenal seeming* (cf. eg., Huemer, 2001, ch. 5).

If it *seems* to you that the moon has an even number of craters, and if you've done your best to scrutinise that seeming, and you've arrived at a comprehensive theory of the universe, informed by all of the best observations and scientific theories available to you, and that belief manages to survive, *then* you're entitled to hold on to it. Not merely because you find yourself

believing it, but because the universe somehow *strikes* you that way. The moon *seems* to you to have a certain property, and that seeming has remained, after much work, unscathed. In that case, you have reason to believe it – at least until further evidence comes your way and undermines or defeats the seeming in question.

One problem with phenomenal conservativism is that it renders justified-belief too easy to come by. This is the safety concern. A second problem is that it's hard to believe that anything at all Russellian in spirit could depend upon any form of epistemological conservativism. Surely Russell was an evidentialist. "Perfect rationality," Russell wrote, "consists ... in attaching to every proposition a degree of belief corresponding to its degree of credibility" (Russell, 1948, p. 349).

In response to the second worry, we should note that phenomenal conservativism isn't in *direct* tension with Russell's evidentialism. Rather, the idea is that a seeming *is* evidence, albeit, defeasible evidence. Nobody is suggesting that evidence should be ignored, or that our confidence should in any way *outstrip* the evidence. But, even if phenomenal conservativism *is* compatible with, so to speak, the spirit of Russell, we're still left with our first worry, namely, that it makes justified belief too easy to come by.

Imagine that it *seems* to me that my religious elder is telling the truth. It just *seems* to me that he is. That's how things *strike* me. Imagine that he tells me that the world is flat. He also tells me not to listen to other people. He tells me that they're all liars. Well, it *seems* to me that he's telling the truth. The world sure *seems* flat. And he *seems* like a reliable person. That's how things *seem* to me. Then why am I not justified in believing that the world is flat, and justified in refusing to allow further evidence into consideration, since the seemingly-trustworthy authority has told me that all counter streams of evidence are liable to mislead me?

Thankfully, I don't have to answer this concern on behalf of phenomenal conservatism. Instead, it can be shown that our Russellian philosophy of philosophy requires no such doctrine. For the phenomenal conservative, a mere seeming is to be counted as *evidence*. But that *isn't* what Russell was arguing for the in the *Problems of Philosophy*. Rather, his claim, as I've already explained, was this: if it *seems* to you that P, then that seeming creates a *caeteris paribus* reason for trying to find room for P in your final theory. If, after sufficient intellectual

labour to construct such a theory, you find that P survives, you are then entitled, provisionally, to hold on to P – but, as Russell makes clear – you must do so with some degree of doubt.

Note: the *seeming* isn't evidence. Rather, it generates entitlement iff the belief in question survives a sufficient degree of scrutiny and system-building. Note also: the epistemic entitlement is always provisional, at least until we arrive at an epistemic end of days. Accordingly, one has to remain epistemically *open*, at least until the eschaton.

On what grounds are we "entitled" to hold the beliefs that survive? Why should we think that mere survival of the belief that P is any indication of its truth, if we still lack direct, positive evidence?

Russell's own answer to this question is that "though the *possibility* of error remains, its likelihood is diminished by [1] the interrelation of the parts" of the system one has built, "and by [2] the critical scrutiny which has preceded acquiescence" (Russell, 1912, p. 11). The second conjunct is easier to understand. The idea is that logical analysis often reveals the underlying ontological commitments of our beliefs. Many beliefs will fall at that hurdle, as we uncover their unsightly consequences. A belief that survives will have proved itself free from various forms of defects. This should increase its subjective likelihood. The first conjunct, however, is harder to fathom. Why should the interrelation of the parts of a system of thought offer epistemic grounds for believing in those parts?

Aaron Segal argues, in his contribution to this book, that the early analytic philosophers eschewed a certain sort of systematicity. Of course, they may have been systematic if that only means that they were (1) methodical. They may have been systematic in the more robust sense that they (2) sought to be comprehensive; taking a stance on every substantive philosophical issue. Some of them may have been systematic in a yet more robust sense, that they (3) sought to derive all of their philosophical conclusions from some single underlying principle or doctrine. And yet, Segal points to a further sense of systematicity that they lacked.

Although Segal provides a more precise definition, as his paper progresses, at a first pass, a philosophy is systematic in the sense that Segal's getting at if it, (4) suggests that all of the substantive philosophical issues are "highly interconnected." Call this fourth species of systematicity, Segal-systematicity. A philosophy *P* is Segal-systematic if, according to *P*, "One can't sensibly take a position on *any* metaphysical issue *in isolation*." Segal-systematic

philosophies contend that one's positions on substantive philosophical issues must come as a package deal.

Segal has good reason to suggest that Russell was not Segal-systematic. Indeed, Russell wrote:

> The essence of philosophy as thus conceived is analysis, not synthesis…What is feasible is…the division of traditional problems into a number of separate and less baffling questions. "Divide and conquer" is the maxim of success here as elsewhere.
>
> (Russell, 1917, p. 113)

But, in actual fact, this quote is incompatible with the Russellian philosophy of philosophy that we've extracted. Indeed, this quote has to be one of the imperfections that we smelt away. I say this because, according to the refined Russellian philosophy of philosophy, our provisional epistemic entitlement to hold onto our instinctive metaphysical beliefs, stems – in part – from the "interrelation of the parts" of the system of propositions that we manage to build, in which the instinctive belief survives, and finds a home.

On the one hand, and to be fair to Segal, it's clear that Russell's philosophy was far from Segal-systematic. On the other hand, it aspired to a degree of systematicity more robust than the three species that Segal notes in his prelude to Segal-systematicity. We could call it Russell-systematicity. And the idea is, roughly, that a philosophical thesis cannot receive (much) justification in isolation; it can only receive (significant) justification by surviving in a system of a certain sort. We're talking here about an epistemic notion of systematicity.

Not just any sort of coherence will do. The propositions that form the basis of a novel might be coherent, in the sense of being compossible, but *their* coherence doesn't give us any reason to believe that they're true. Russell writes:

> Assuming the usual laws of deduction, two obvious propositions of which one can be deduced from the other become more nearly certain than either would be in isolation; and thus, in a complicated deductive system, many parts of which are obvious, the total probability may become all but absolute certainty.
>
> (Russell, 1907, p. 279)

The idea is that if some members of the set are self-evidently true, or have some independent source of epistemic warrant, then (and only then) does the set hanging together in a certain way add credibility to all of its parts (even those that are not self-evident, or don't come with an independent source of warrant). In the words of Conor Mayo-Wilson: for Russell, the "coherence of a set of propositions is a mark of truth only when at least some of the propositions are independently supported by available evidence" (Mayo-Wilson, 2011).

At times, Russell talks as if this species of coherence increases the *probability* of each proposition in the set. But when "Russell wishes to be careful in distinguishing the mathematical concept of probability from the concept of evidential support, he speaks of a coherence theory of "credibility" rather than probability" (Ibid.):

> Given a number of propositions, each having a fairly high degree of intrinsic credibility, and given a system of inferences by virtue of which these various propositions increase each other's credibility… [we] arrive at a body of interconnected propositions having, as a whole, a very high degree of credibility.
>
> (Russell, 1948, p. 395)

Lots of questions remain. What relation needs to hold between the members of a set of propositions so as to generate Russell-systematicity? Does it come in degrees? Why should Russell-systematicity lend even a provisional epistemic entitlement to hold onto instinctive beliefs? Is it because the more explanatory, simple, elegant, and comprehensive a theory of the world becomes the more likely it is to be true? But why should we think that that's true?

Should we think that a system is more likely to be true the more it exhibits Russell-systematicity because it instinctively *seems* to us that such a system is more likely to be true? Wouldn't that be circular? Wouldn't that be to justify our way of justifying instinctive beliefs by appeal to another instinctive belief? Is circularity something we can hope to escape once we've given up the hope of classical foundationalism? Are all forms of circularity to be avoided? Sadly, these questions are beyond the scope of this paper – as is the important question as to how much work has to be done, in the process of system building, and logical analysis, before a person wins provisional epistemic entitlement for their surviving beliefs.

Despite these unanswered questions, there seems to me to be something deeply attractive about an epistemology of provisional epistemic entitlement won on the basis of constructing

a Russell-systematic theory. Instinctive beliefs that survive our Russell-system-building are (provisionally) warranted, and not merely tolerated in recognition of our fallibility. This epistemology would mean that we really can claim to hold (provisionally) justified metaphysical beliefs, without appeal to any crude form of conservativism.

§4.3. The Epistemic End of Days

Finally, I turn to a question. Will philosophy ever arrive at consensus? Will it ever arrive at an epistemic end of days, in which we finally coalesce around a single periscopic vision of the universe? Our Russellian philosophy of philosophy can remain doggedly neutral on this issue.

Peter van Inwagen, in his contribution to this volume, endorses a meta-metaphysics that stays broadly true to the spirit of our refined Russell. And yet, he takes a negative stance regarding the prospects for eschatological agreement. Philosophy can never arrive at consensus. The most we can each hope to do is to find a system that's coherent, comprehensive, systematic, elegant, and satisfies as many of our own, idiosyncratic, instinctive beliefs as possible. Each to their own. After all, we each start with different instinctive beliefs. If we're starting from different places, why should we think we'd end up at the same destination?

Fraser MacBride (2014), by contrast, paints his equally Russellian vision of philosophy in terms of an eschatological vision. MacBride is, in that sense, a messianist, living in the hope that we may, one day, arrive at consensus; an epistemic end of days.

As Russell repeatedly makes clear in his lectures on the philosophy of logical atomism, it isn't too hard to find propositions that all human beings take to be intuitive (or, at least, all of the human beings in his lecture theatres) – this should come as no surprise, since we share a basic cognitive architecture. To move towards consensus, we can first of all be careful to start from this inter-subjective basis.

Moreover, the more progress that we make in the natural and social sciences, and in normative philosophy, the more constraints are laid down on the possible contours of a final theory. Perhaps, in a final theory, there will still be some distinctively philosophical theses that will not be justified by any particular science but will survive as a mere vestige of the instinctive opinions held firmly from the outset. But, the more constraints in place, the less

room there will be for a diversity of such opinions. So, we've found two reasons for messianic hope. First: despite our different starting places, there's a large inter-subjective overlap between many of our instinctive pre-philosophical beliefs. Secondly: the hefty constraints upon a final theory, emerging from the special sciences, might encourage a large degree of convergence.

And yet, this hope leads us to a couple of daunting questions that I'll leave unanswered. First: if a given theory turns out to be the only available "grand theory of everything", would that entail that that theory must be *true* (or even *likely* to be true)? Second: would the conformity of opinion that would emerge in an epistemic end of days be a good thing, or a bad thing, for human flourishing? Are philosophical messianists hoping for something we might come to regret?

Bibliography

Ayer, A. J., 1971. *Russell and Moore: The Analytical Heritage.* London: Macmillan.

Bradley, F. H., 1897. *Appearance and Reality.* 2nd ed. Oxford: Calrendon Press in 1968.

Conee, E., 1988. The Basic Nature of Epistemic Justification. In: E. Conee & R. Feldman, eds. *Evidentialism: Essays in Epistemology.* Oxford: Oxford University Press, pp. 37-52.

Eames, E. R., 1969. *Bertrand Russell's Theory of Knowledge.* London: Allen and Unwin.

Elgin, C. Z., 1996. *Considered Judgements.* Princeton: Princeton University Press.

Elgin, C. Z., 2005. Non-foundationalist Epistemology: Holism, Coherence, and Tenability. In: M. Steup & E. Sosa, eds. *Contemporary Debates in Epistemology.* Malden Ma.: Blackwell, pp. 156-167.

Haack, S., 1993. Double-Aspect Foundherentism: A New Theory of Epistemic Justification. *Philosophy and Phenomenological Research,* Volume 53, pp. 113-128.

Haack, S., 1997. A Foundherentist Theory of Epistemic Justification. In: L. Pojman, ed. *The Theory of Knowledge: Classic and Contemporary Readings.* Belmont Ca.: Wadsworth, pp. 237-247.

Huemer, M., 2001. *Skepticism and the Veil of Perception.* Lanham: Rowman & Littlefield.

Irvine, A. D., 1989. Epistemic Logicism & Russell's Regressive Method. *Philosophical Studies,* 55(3), pp. 303-327.

Kripke, S., 1972. *Naming and Necessity.* Blackwell Publishing, in 1981.

Leach, S., 2020. Bertrand Russell & Common Sense for Savages. *Philosophy Now*, January.

Lycan, W. G., 1988. *Judgement and Justification.* Cambridge: Cambridge University Press.

Lycan, W. G., 2012. Explanationist Rebuttals (Coherence Defended Again). *The Southern Journal of Philosophy,* 50(1), pp. 5-20.

MacBride, F., 2014. Analytic Philosophy and its Synoptic Commission: Towards the Epistemic End of Days. *Royal Institute of Philosophy Supplement,* Volume 74, pp. 221-236.

MacIntosh, D. C. M., 1915. *The Problem of Knowledge.* New York: The Macmillan Company.

Mayo-Wilson, C., 2011. Russell on Logicism and Coherence. *Russell: The Journal of Bertrand Russell Studies,* 31(1), pp. 89-106.

McCain, K., 2014. *Evidentialism and Epistemic Justification.* New York: Routledge.

Olson, D. & Griffin, N., 2018. Russell's Bridge. In: R. Wahl, ed. *The Bloomsbury Companion to Bertrand Russell.* London: Bloomsbury, pp. 286 - 311.

Pears, D., 1967. *Bertrand Russell and the British Tradition in Philosophy.* London: Fontana.

Quine, W. V. O., 1953. On What There Is. In: *From a Logical Point of View.* New York: Harper and Ross, pp. 1-19.

Quine, W. V. O., 1966. Russell's Ontological Development. *Journal of Philosophy,* 63(21), pp. 657-667.

Russell, B., 1907. The Regressive Method of Discovering the Premises of Mathematics. In: *Essays in Analysis.* London: George Allen & Unwin, in 1973, pp. 272-283.

Russell, B., 1912. *Problems of Philosophy.* Oxford: OUP 1998 Edition.

Russell, B., 1913. *Theory of Knowledge: The 1913 Manuscript.* Abingdon: Routledge, in 1992.

Russell, B., 1914. *Our Knowledge of the External World.* London: Allen and Unwin.

Russell, B., 1917. *Mysticism and Logic.* London: George Allen & Unwin.

Russell, B., 1918. Lectures on the Philosophy of Logical Atomism. In: R. Marsh, ed. *Logic and Knowledge.* Nottingham: Spokesman Books in 2007, origionally published in 1956, pp. 175-282.

Russell, B., 1919. *Introduction to Mathematical Philosophy.* 3rd ed. London: George Allen and Unwin, published in 1924.

Russell, B., 1928. *Skeptical Essays.* Reissuued in 2004 by Routledge ed. London: George Allen & Unwin.

Russell, B., 1948. *Human Knowledge its Scope and Limits.* London: George, Allen and Unwin Ltd, Reissued by Routledge in 2009.

Russell, B., 1956. Reflections on my Eightieth Birthday. In: *Portraits From Memory.* London: George Allen & Unwin.

Russell, B., 1959. *Common Sense and Nuclear Warfare.* London: George Allen & Unwin.

Russell, B. & Whitehead, A. N., 1910-13. *Principia Mathematica.* Cambridge: Cambridge University Press, 2nd edtion, in 1925-27.

Sainsbury, R. M., 1979. *Russell.* London: Routledge and Kegan Paul.