Four Kinds of Russellian Monism*

Daniel Stoljar, ANU

"Russellian Monism" is a name given[1] to a family of views in philosophy of mind. The family is exciting because it seems to present an alternative both to materialism and to dualism. After briefly setting out the need for this alternative, I will in what follows distinguish four different kinds of Russellian Monism (RM), and assess their pros and cons. My own feeling, as will emerge in the final section of the paper, is that only the fourth of these represents a viable version of the view. But my main aim is less to state my feelings than to get clear on the different versions of the view and on what is involved in choosing among them.

Motivating RM

It is not hard to motivate the thought that neither materialism nor dualism is very attractive and that it would be better in principle if we could make out some alternative.

The problem for materialism is that it is on the face of it incomplete, for it seems to leave out properties constitutive of consciousness. A good way to bring this out is to operate with the following simple definition:[2]

M1. Materialism is true at a possible world *w* if and only if for every property *G* instantiated at *w*, there is some physical property (or some complex of physical properties) *F* instantiated at *w* such that *F* (metaphysically) necessitates *G*.

This provides an account of what it is for materialism to be true *at* some possible world or other; philosophers who are materialists believe that materialism is true at one world in particular, viz., the actual world.

* I am very much indebted to the following for their help: Torin Alter, David Chalmers, Uriah Kriegel, Leon Leontyev, Derk Pereboom, and Jon Simon.

[1] As far as I can make out, the first explicit occurrence of the name 'Russellian Monism' is in Chalmers 1999. It is now used widely—see, e.g., Chalmers 2010, Pereboom 2010, Alter and Nagasawa 2012, Alter and Nagasawa, forthcoming—and has its own entry on PhilPapers: http://philpapers.org/browse/russellian-monism. Similar names include (at least) 'The Russellian Identity Theory' (Lockwood 1989, ch.10), ''the Russellian View' (Chalmers 1996), "The Russellian theory of Mind" (Holman 2008), "Russellian Physicalism" (Montero, 2010), and the considerably more prosaic "o-physicalism" (Stoljar 2001) and "the Russellian version of the epistemic view" (Stoljar 2006).

[2] For statements of materialism (aka physicalism) of this sort, see Stoljar 2010. The definition used in the text is a simple one; for example, it takes for granted exactly what metaphysical necessitation is. I will set aside such complications here.

Why think that materialism so defined is incomplete? Well, at the moment (let us suppose) I am in some sort of total experiential state that might at least partially be captured in the following way. I am sitting in my office; I have a cup of coffee to my left; some people are talking softly in the corridor outside; I smell slightly of chlorine from the pool earlier; the light is coming in through the wooden venetian blinds; and so on. It is conceivable[3] that there could be someone identical to me in respect of whatever physical properties I instantiate and yet who is not in exactly the same experiential state I am in. Maybe, for example, the taste of his coffee is ever so slightly more bitter than the taste of mine. If what is conceivable is possible, there could be someone identical to me in respect of all physical properties and yet who does not taste the coffee as I do. But then it follows that this property—tasting the coffee in the precise way that I do—is both instantiated and yet is not necessitated by any physical property. By M1, however, this is inconsistent with materialism; hence materialism is false.

This argument against materialism—the conceivability argument, as it is usually called—looks simple enough, but in fact it is quite complex. As a consequence, the literature on it has become increasingly involved.[4] But suppose we provisionally agree that the conceivability argument is successful; what are our options then? The standard option is to adopt the kind of dualism according to which most, but not quite all, instantiated properties are necessitated by the physical. The exceptions are properties associated with consciousness such as tasting coffee—'experiential properties', as we can call them. According to the dualist, the relation between experiential and physical properties is not metaphysically necessary but is naturally necessary, i.e. necessary given the way the laws of nature in fact are. More generally, on this picture, there are contingent psychophysical laws that tell us that if such and such a complex physical property is instantiated, then such and such an experiential property is instantiated.

Even if it avoids the conceivability argument, however, dualism is on the face of it inelegant. While the definition above left this implicit, presumably the physical properties invoked in M1 are governed by a relatively small stock of laws that describe regularities in how these properties are instantiated and how they interact with each other. It is not that the materialist cannot countenance laws among complex non-physical properties as well, but these hold in virtue of the regularities that obtain among physical properties and the laws that govern them. By contrast, dualism seems to require a quite distinct set of psychophysical laws that connect complex physical properties with

[3] There is a very big literature on what 'conceivability' means; we will not go into the details here. (You will not go too far wrong if you read 'it is conceivable that' as 'it seems possible that'.) For some recent discussion see Chalmers 2010 and Pereboom 2011.

[4] See, e.g., Chalmers 2010 and Pereboom 2011

experiential properties (assumed now to be fundamental); it requires "nomological danglers", as it is often put (Feigl 1967, Smart 1959). But this, as the phrase 'dangler' suggests, is inelegant.

At this point, a natural line of thought is this. Could there not be a theory of the world that is (a) as elegant as materialism but (b) as comprehensive as dualism? Alternatively, could the world not be such that a theory of this sort is true? It is this that provides one part of the motivation for RM. The Russellian Monist, as I understand matters, is someone who agrees that materialism (at least of the usual sort) is defeated by the conceivability argument but holds out the hope for some alternative account that preserves its spirit and structure, and so avoids the inelegance of dualism.[5]

The desire for a theory that is at once elegant and comprehensive is part of the motivation for RM. The other part involves a suggestion about where to look when thinking about the shape of such a theory. This suggestion is to focus, not on the nature of consciousness but—to borrow a phrase from Leopold Stubenberg (1997)—on the "south-pole" of the mind-body relation, i.e., on the nature of the physical. For in fact discussions in philosophy of mind *do* seem to unfairly neglect this side of things. Notice, for example, that in my presentation of M1 and the conceivability argument above *I failed completely* to say what physical properties are supposed to be; in doing so I was being quite faithful to the tradition in which I work. But—Russellian Monists quite reasonably insist—this will not do; moreover, when we do attend to the nature of physical properties, we will (they say) be able to formulate a theory that is an alternative to both materialism and dualism.

But how should this suggestion be developed? The next four sections consider four different ways.

RM1

The first version of Russellian Monism (RM1) starts off from some ideas from John Locke's *Essay Concerning Human Understanding* (see Locke 1975). There are two ideas in particular that we need.[6] The first is a list of what Locke famously called the "primary qualities" of physical objects. These properties are: size, shape, position, duration, movability, divisibility and solidity. The second is an assumption about the limits of human understanding at least as regards physical objects,

[5] Of course there is no a priori guarantee that the world is such that a theory of this sort is true, or that we can formulate the theory even if it is. The hope rather is that there is no philosophical objection to there being such a theory, where a 'philosophical objection' in this context is, roughly, an objection to the existence of such a theory that can be made in advance of formulating it.

[6] Since Locke's essay was published in 1689, the historical issues are difficult, and I won't address them here. The point is to take from Locke (and to take from those who take from Locke) ideas that can then be used to formulate a version of RM. Among those who take from Locke I have in mind mainly Armstrong 1961 (esp. chap 15) here, but see also Armstrong 1968, Foster 1982 and Blackburn 1990. Langton 1998 presents a somewhat similar position, but suggests it constitutes the best interpretation of Kant.

namely, that the properties of physical objects that we may understand are the primary qualities and perhaps logical combinations thereof.

RM1 Formulated[7]

With these ideas in place, our characterization of the first version of RM proceeds in two stages. The first stage formulates a version of materialism by combining the definition above—i.e. M1—with a view about physical properties that is suggested by Locke's list of primary qualities; according to this view, physical properties just are primary qualities. We may call the resulting version of materialism 'primary quality materialism' or 'PQ-materialism' for short. PQ-materialism is true at a possible world *w* just in case for every property *G* instantiated at *w*, there is some primary quality (or complex of primary qualities) *F* instantiated at *w* such that *F* metaphysically necessitates *G*.

The second stage presents RM1 as the conjunction of three claims about PQ-materialism. The first is this:

RM1.a. PQ-materialism is false, and false for reasons quite distinct from those involved in the conceivability argument.

The reason for RM1.a begins from a consideration about one primary quality in particular, viz., solidity. To be solid, Locke said, an object must be such as to resist penetration from other objects. But this entails that solidity is a *dispositional* property of an object, where, to a first approximation, a dispositional property is a property that constitutively involves various claims about what things *would* or *might* be like in certain circumstances, rather than merely about what things *are* like. On the face of it, however, when an object has a dispositional property, there must be some further property it has in virtue of which it has the dispositional property in question. For example, when a chair is uncomfortable there must be some *other* property of the chair— its shape, say—in virtue of which it is it is uncomfortable. By analogy, if some object is solid, there must be some other property of it in virtue of which it is solid.

Why does this show that PQ-materialism is false? Well suppose some object α is solid. By the considerations just reviewed, α must have some further property—which for the moment we can call 'Stuff'—in virtue of which it is solid. If PQ-materialism true, Stuff must either be a primary quality or be necessitated by some primary quality (or some complex of primary qualities).

[7] *Nota bene:* in this section my aim is to provide an *exposition* of RM1 and the arguments for it (ditto for parallel sections about RM2-4). Assessment will come later.

But neither appears to be the case. First, Stuff is not identical with any primary quality on Locke's list, or any complex of such properties. It is not itself solidity, because it is by assumption distinct from solidity; nor is it size, shape, position, duration, divisibility, or moveability—or any complex of these—since these can be properties of empty regions of space, i.e., regions of space that contain no matter and so do not contain α in particular. Second, Stuff is not necessitated by any primary quality on Locke's list. It is not necessitated by solidity—indeed, the opposite is closer to the truth, since things are solid in virtue of instantiating Stuff; nor is it necessitated by size, shape etc.—or any combination of size, shape etc.—for the reason again that regions of space can have these properties and yet are not solid. Conclusion: Stuff is instantiated and yet is not necessitated by any primary quality. Hence PQ-materialism is false.

The second claim of RM1 is this:

RM1.b While PQ-materialism is false for the reason mentioned in RM1.a, there is a substitute thesis which is not false for that reason, and which preserves the spirit and structure of PQ-materialism.

The substitute thesis mentioned here—we may call it 'PQ-materialism+'—is exactly like PQ-materialism but with this difference: where PQ-materialism says that the physical properties are the primary qualities, PQ-materialism+ says that the physical properties are *either* the primary qualities *or* whatever properties objects must have in virtue of having primary qualities. So, for example, while PQ-materialism is refuted if Stuff is instantiated, PQ-materialism+ is not refuted, since Stuff counts as a physical property according to it. Does PQ-materialism+ preserve the spirit and structure of PQ-materialism? It would seem so. First, it may be formulated within the framework provided by M1. Second, it seems reasonable to include Stuff as a physical property since it (or a property like it) will be instantiated even at the most obvious possible worlds at which materialism is true. For example, take a world that contains only one lonely atom located in space. This atom is solid; hence, given our assumptions it will instantiate Stuff. And yet surely materialism is true at this world if it is true anywhere. So the existence and instantiation of Stuff is not anathema to materialism.

The third claim of RM1 is this:

RM1.c. While PQ-materialism is false for the reason mentioned in RM1.a, it is also false for another reason, viz., the conceivability argument; by contrast, PQ-materialism+ escapes this argument.

In part RM1.c says that PQ-materialism is defeated by the conceivability argument—we saw the reason for this when motivating RM at the start. But RM1.c also says that PQ-materialism+ is *not* defeated by the conceivability argument. The reason for this concerns the second idea we took from Locke above, the idea about human understanding. If sound, the line of reasoning we just considered tells us that α instantiates Stuff. But it does not tell us what Stuff is, or whether we can understand what it is. Can we understand it? Not if the second Lockean idea is right. That idea tells us that our understanding of the properties of physical objects is limited to the primary qualities. But Stuff is not a primary quality. Hence we cannot understand what it is. Of course, we *can* name it, refer to it, (partially) describe it, and so on; we called it 'Stuff' after all, and described is the property in virtue of which α is solid. But from the fact that we can name and describe a property it does not follow that we can understand it. Indeed, that is the position we are in with respect to Stuff if RM1 is right.

But why does this point about Stuff mean that PQ-materialism+ escapes the conceivability argument? Well, that argument began with the idea that it is conceivable that there is someone identical to me in respect of physical properties but who does not taste the coffee in quite the way that I do. But is Stuff included here as a physical property or not? This question poses a destructive dilemma. For suppose that Stuff *is* included as a physical property. Now the conceivability claim loses its force. For it is now not clear that it is conceivable that there is someone identical to me in respect of Stuff who is different from me experientially. Since I have no idea what Stuff is, to say that this is conceivable is to say something for which I have no rational grounds. On the other hand, suppose Stuff is *not* included as a physical property. Now the argument may succeed but only in a diminished form, for now it targets, not PQ-materialism+, but only PQ-materialism. In sum, the lesson of Stuff is that the conceivability argument may refute PQ-materialism, but it does not refute PQ-materialism+.

The following analogy illustrates the situation for PQ-materialism+. Suppose we have a very large box of fruit, perhaps an infinitely large one. And suppose we have some evidence (never mind why) that the box contains a certain kind of fruit—say citrus fruits. Suppose now someone argues as follows. "It is conceivable, and so possible, that there is a box which is identical to this box in every respect but which contains no apples." Is this persuasive? It is correct to answer "no" with something like this commentary: "You have not conceived what you say you have. You might have conceived a box that contains only citrus fruit but which has no apples, but you have not conceived a box that is identical in all respects to this box but which contains no apples. You are ignorant of the kinds of fruits there are in this box so you can't have conceived that." Just as the

conceivability argument fails in the case of the box and apples, it fails with respect to the physical world and consciousness, according to RM1.c

RM1 Assessed

There are clearly elements of RM1 that are not for everybody. For example, some philosophers insist that what is dispositional here is the *word* 'solidity' rather than any property; and others say that while solidity is dispositional there is no philosophical reason to assume that must be a further property that grounds it.[8]

But I don't want to focus on these issues, for there is a simpler and more serious problem with RM1. This is that the epistemological idea that we took from Locke is palpably false; we are *not* limited in our understanding of physical objects to the primary qualities on Locke's list. On the contrary, if we take scientific knowledge more or less at face value our knowledge of, and so understanding of, physical objects goes way beyond this. After all, consider all of the remarkable things that science has thrown up over the last few hundred years—the identification of matter and energy, quantum wave-functions states, fields, strings in n-dimensional space and so on and so forth. It is impossible to say that humans as such do not understand these things (though *some* humans don't of course). But then the epistemological idea at the heart of RM1 should be rejected.

Of course that RM1 makes implausible epistemological assumptions is hardly surprising; isn't that just what you get if you formulate a theory by adopting ideas that were *au courant* circa 1689? So it would seem. But it is also true that ideas very close to RM1 have been formulated and discussed sympathetically by many contemporary philosophers, e.g., those mentioned fn.6. How do they react to the point that Locke's epistemological assumptions are out of date? Their reaction (I think) is roughly this: true but not relevant to the substance of the issue, for a position rather like Locke's can be formulated no matter what we assume about our understanding of the world. As Armstrong puts it: "Modern theory…[is in a]… "no better position than…the Lockean Theory", for the distinctive postulates of modern theory themselves "dissolve into relations, or dispositions to have relations, that one particular has to another" (1961, 188).[9]

Adapting this to our own discussion, the suggestion is that while RM1 may be wrong for the reason given, there is a version of RM that is not wrong, a version that abstracts away from the

[8] See, e.g., the papers in Handfield 2009

[9] In the passage quoted, Armstrong apparently treats the possibility of relations and dispositions in the same way. I think to the contrary they are very different, since relations are properties that are instantiated by more than one thing, while dispositions are a particular sort of property or relation, i.e. ones that involve modality in a very distinctive way. But I will set this issue aside here.

particular claims about human understanding that Locke makes. To assess this suggestion I now turn to a second version of RM, a version that allows us all the knowledge and understanding it would be possible for us to have.[10]

RM2

The second version of Russellian Monism (RM2) starts off from some ideas set out in David Lewis's paper 'Ramseyan Humility' (see Lewis 2009).[11] There are two ideas we need. The first is that there is what Lewis calls "a final theory", a theory that delivers "a true and complete inventory of those fundamental properties that play an active role in the actual workings of nature" (2009, 205). We may never formulate the theory in question, Lewis says, but it "nevertheless exists…in the way never-to-be-written poems do" (p.219). In order to spell out RM2, we will assume something Lewis does not, viz., that the final theory of the world is a physical theory.

The second idea is that the final theory may be divided into (what Lewis calls) 't-terms' and 'o-terms'. The t-terms name the fundamental properties of the world, properties as Lewis says which are "not at all disjunctive, or determinable, or negative. They render their instances perfectly similar in some respect. They are intrinsic; and all other intrinsic properties supervene on them" (2009, 204). The o-terms are any *other* terms in the theory. Crucially, Lewis assumes that the t-terms are "implicitly defined by the theory" which means that for any t-term *T*, there is a true definition of the following form: 'x has *T* if and only if x has the property which meets the condition…*O*…', where the vocabulary used to spell out the condition in question is exclusively the o-vocabulary, i.e. a vocabulary consisting entirely of o-terms.[12]

RM2 Formulated

[10] I will assume in the text that the natural response to the empirical inadequacy of RM1 is to adopt one of RM2-4. But it is worth mentioning another idea which sometimes comes up at this point, viz. that one's knowledge of physical objects at the end of the day depends upon our causal interaction with them, and that in consequence one cannot know the intrinsic nature of these objects; for ideas along these lines Jackson 1998 and Langton 1998. I will not try to assess that idea here, beyond mentioning three points: first, these suggestions are subject to the prima facie difficulty that a causal explanation for our knowledge of physical objects does not by itself tell us much about the content or limits of that knowledge; second, it is plausible that suggestions along these lines will resolve themselves into a position akin to RM2-4—indeed, I read Lewis's comments at the beginning of 'Ramseyan Humility' as suggesting that that paper represents the best way to make sense of the causal argument; third, I am in any case sympathetic with the existing critical discussion of this idea in the literature; I have in mind in particular Van Cleve 2002 and Pereboom 2011.

[11] Lewis himself did not defend Russellian Monism, it is rather that the materials he describes can be used to formulate a version of that view.

[12] The definitions here are intended by Lewis as reductive definitions along the lines of 'x is a bachelor if and only if x is an adult unmarried male' (assuming that *is* a reductive definition of course).

With these ideas in place, our characterization of RM2, like that of RM1, proceeds in two stages. The first formulates a version of materialism by combining M1 with a view about physical properties that is suggested (but not entailed) by Lewis's discussion. According to this view, a physical property is a property expressed[13] by a predicate of the final theory (on the assumption, noted above, that the final theory is a physical theory)—for short, physical properties are final-theory properties. We may call the resulting version of materialism 'final theory materialism' or 'FT-materialism' for short. FT-materialism is true at a possible world *w* just in case for every property *G* instantiated at *w*, there is some final theory property (or complex of final theory properties) *F* instantiated at *w* such that *F* metaphysically necessitates *G*.

The second stage is to present RM2 as making three distinct claims about FT-materialism. The first is this:

RM2.a. FT-materialism is false, and false for reasons quite distinct from those involved in the conceivability argument.

The reason for RM2.a begins from some considerations that are prominent in Lewis's paper. Suppose we take the final theory, T, and write it as a single long conjunctive sentence; call this 'the postulate of T.' Suppose now we take the postulate of T and write it so that the t-terms are made explicit. The result would be sentence of the form 'T (t1, t2, t3…)'. Now suppose we replace all the t-terms in the postulate of T with variables; the result would be an open sentence of this form 'T(x_1, x_2, x_3…).' If we existentially quantify this open sentence we get what Lewis calls 'the Ramsey-sentence of T,' the theory that says that there exists x_1, x_2, x_3 … such that T(x_1, x_2, x_3…). If a *realization* of T is an *n*-tuple of properties that satisfy this open sentence, the Ramsey sentence of T says in effect that there is a realization for T.

Now, Lewis argues that there is more than one possible realization of T. Suppose that in fact the properties that realize T are the ordered triple <F, G, H>. There might, in an alternative possible world, be a distinct triple < F*, G*, H*> that realizes it. T will be true relative to this alternative possible world just as it is true at the actual world, it is simply that the realization of the theory are different. Why should we assume that <F*, G*, H*> are distinct from (i.e. non-identical to) < F, G, H>? After all, they realize the same total theory, and so there is no ordinary feature—no feature that may be captured by our total theory of the world—that distinguishes them. Lewis's

[13] Why are physical properties those *expressed* by a predicate of the final theory rather than those *named* by a term of that theory? The answer is that only then will knowledge of the theory tell us exactly what the properties are—indeed, this point is prominent in the reasoning from Lewis I am about to set out.

answer is that <F, G, H> are numerically different from <F*, G*,H*> even if they are duplicates with respect to qualitative features; he captures this by saying that properties constitutive of the first sequence have different *quiddities* from those constitutive of the second.[14]

Why do these considerations show that FT-materialism is false? Well suppose again that the actual realization of T is <F, G, H>. These properties are named by terms in T but are not expressed by any predicate of T. The reason is that T could be true at some possible world at which distinct *n*-tuple of properties realizes it, for example the world just described at which <F*, G*, H*> realize it. But now it follows that F, for example, is neither not expressed by T, and nor is it necessitated by any property expressed by T; *ditto* for G and H. So F, G and H are instantiated and yet are not physical properties, and FT-materialism is false.

The second claim of RM2 is this:

RM2.b While FT-materialism is false for the reason mentioned in RM2.a, there is a substitute thesis which is not false for that reason, and which preserves the spirit and structure of FT-materialism.

The substitute thesis—which we may call 'FT-materialism+'—is exactly like FT-materialism but with this difference: where FT-materialism says that the physical properties are those expressed by predicates of the final theory, FT-materialism+ say that the physical properties are *either* those expressed by such predicates or those denoted by expressions of the final theory. So, for example, while FT-materialism is false at the world at which <F*,G*,H*> realize T, FT-materialism+ is not false, since F, for example, is named there by T. Does FT-materialism+ preserve the spirit and structure of FT-materialism? For reasons analogous to those reviewed earlier, it certainly seems so.

The third claim is this:

RM2.c. While FT materialism is false for the reason mentioned in RM2.a, it is also false for another reason, viz., the conceivability argument; by contrast, the substitute thesis mentioned in RM2.b does not face this argument.

The argument for RM2.c parallels the argument given for RM1.c. The first premise of the conceivability argument says that it is conceivable that there be someone identical to me in all physical respects but different from me in respect of some physical property. Do we include the

[14] For discussion of Lewis on quidditism and related matters, see D. Locke 2009, forthcoming and the references therein.

properties that constitute the actual realization of T as physical properties or not? Suppose they are included. Then the conceivability claim loses its force—for how I am supposed to conceive of a situation I do not understand? Suppose they are not included. Then the argument loses its force against FT-materialism+ even if it remains successful against FT-materialism. Either way, therefore, the conceivability fails as an argument against FT-materialism+.

RM2 Assessed

Once again there are elements of RM2 that are not for everybody. Some philosophers will reject either Lewis's suggestion that t-terms are implicitly defined by o-terms or his quidditism. Others will reject our assumption (which was not Lewis's) that the final theory is a piece of physics.

But again I don't want to focus on these points, for there is a much simpler problem with RM2. This is that the sort of ignorance that a proponent of RM2 is committed to is (mere) *quiddistic* ignorance; that is, ignorance of the numerical identity of the properties in question, and about nothing else. It is not (what we might call) *qualitative* ignorance, i.e. ignorance of what goes beyond the mere numerical identity of the property in question. To put it vividly, with the exception of its numerical identity, any feature at all that F has at the actual world is also a feature that F* has in the alternative world we imagined. If F causes the heat death of the universe, so too does F*. If F is implicated in our suffering, so too is F*. If F is intrinsic or non-relational or higher order, so too is F*. Indeed, the *only* difference between F and F* is that they are numerically distinct; equivalently, the only property that F has which F* does not is the property of being F.

Why does its commitment to mere quiddistic ignorance undermine RM2? Well, consider RM2.c above. This says that while the conceivability argument defeats FT-materialism, it does not defeat FT-materialism+. Now this could only be true if the difference between them is relevant to the nature of consciousness. But it is hard to see how this could be so. The qualitative information contained in FT-materialism is identical to the qualitative information contained in FT-materialism+. The difference between them is that the former but not the latter leaves open the numerical identity of the realization of the final theory: FT-materialism+ names the actual realizers; FT-materialism does not. But the facts about consciousness are presumably qualitative if anything is. Hence the difference between FT-materialism and FT-materialism+, while real enough, cannot be relevant to the nature of consciousness.

One might point out that this objection is generated by the assumption that the quiddity of a property is its numerical identity; could one not operate with a thicker notion[15] of a quiddity

[15] For a discussion to the thicker notion of quiddity, see Chalmers 2012, 350, who introduced the phrase 'thick quiddity'.

according to which F is distinct from F* not simply numerically but in some further way as well? Perhaps one could, but this makes no difference to RM2. For suppose F has some thick quiddity that distinguishes it from F*. It is hard to see why this fact about F, whatever it is, would not be included in the final theory of the world the existence of which RM2 presupposes. After all, the only thing we are assuming about that theory is that it is, as we saw before "a true and complete inventory of those fundamental properties that play an active role in the actual workings of nature". Such an inventory would presumably have the resources to say what the thick quiddities of fundamental properties are, assuming they have them.[16]

So it would seem that the situation we are in at this point is this. RM1 makes an assumption about our epistemic situation that is implausible. RM2 avoids that problem by making assumptions that are defensible, even if philosophically controversial. But the problem is that we cannot use these assumptions in philosophy of mind; hence RM2.c is false.

Is there a version of RM that unlike RM1 is empirically adequate but that unlike RM2 leaves us with a useable form of ignorance? Well, a striking feature of the discussion so far is that, contrary to expectations, it has had nothing to do with either with Russell or with Monism! Since Russellian Monism must have *something* to do with Russell or Monism, weaving those elements into the formulation of the doctrine will (or so one might think) yield something that is an improvement on what we have considered so far.

Now, in the present context, referring to 'monism' is merely a way of marking the fact that RM is, as noted at the outset, not intended to be a traditional kind of dualism. So little of value may be wrung from the 'monistic' part of Russellian Monism. But it *is* possible to inject a bigger element of his lordship into proceedings than we have done so far. The result of doing so is the next version of RM I will consider.

RM3

The third version of Russellian Monism (RM3) starts off from a famous idea from Russell's *The Analysis of Matter*, viz., that "the aim of physics, consciously or unconsciously, has always been to discover what we may call the causal skeleton of the world" (1927, 391).[17] To say this is not to

[16] As Chalmers pointed out to me in discussion, one might argue here that, while the thick quiddity here is part of Lewis's final theory but is not part of physics. This is a possibility, but in the scheme of this paper it is to move away from RM2, and toward RM3. For one thing, as a matter of stipulation, RM2 requires, while RM3 does not, identification of the final theory and physics. Moreover, to implement the suggestion there must be some way to discriminate physics from the final theory; RM3 (as we will see) has such a way, while RM2 does not.

[17] As in the case of Locke, there are serious issues of interpretation here, but I will mostly ignore them. For discussion of Russell, see Demopoulos and Freidman 1985, and more recently Chalmers 2012.

deny that world contains things apart from its causal skeleton; the beast of nature may have flesh as well as bones. Rather it is to say that physics has no business talking about the flesh.

RM3 Formulated

With this in place, our characterization of RM3 proceeds as before in two stages. The first combines M1 with a theory about what it to be a physical property. The theory is that a physical property is a property expressed by a term of the physical theory in Russell's sense—for short, a physical property is a causal skeletal property. The resulting version of materialism—causal skeletal materialism or CS-materialism—is true at a possible world *w* if and only if for every property *G* instantiated at *w*, there is some causal skeletal property (or complex of causal skeletal properties) *F* instantiated at *w* such that *F* necessitates *G*.

The second stage is to present RM3 as making three claims. The first is:

RM3.a. CS-materialism is false, and false for reasons quite distinct from those involved in the conceivability argument.

The reason for RM3.a is closely analogous to the one given for RM1.a. Suppose the true physical theory, whatever it is, tells us about some sequence of events causally related to each other. As we noted, it certainly does not follow that the events do not have further properties. Moreover, for reasons closely related to the discussion of RM1.a above, one might think it very plausible that they do have such properties, i.e. one might think that if one event causes another there must be some property of that event in virtue of which it does. However, if we agree with this, it would seem that CS-materialism is false. For on the face of it these further properties neither are nor are necessitated by further causal structural properties

The second claim is this:

RM3.b While CS-materialism is false for the reason mentioned in RM3.a, there is a substitute thesis which is not false for that reason, and which preserves the spirit and structure of CS-materialism.

The substitute thesis—which we may call 'CS-materialism+'—is exactly like CS-materialism but with this difference: where CS-materialism says that the physical properties are causal skeletal properties, CS-materialism+ say that the physical properties are *either* causal skeletal properties or whatever properties events have in virtue of having causal skeletal properties. For reasons

analogous to those we have reviewed earlier, CS-materialism+ preserves the structure and spirit of CS-materialism.

The third claim is this:

RM3.c. While CS materialism is false for the reason mentioned in RM3.a, it is also false for another reason, viz., the conceivability argument; by contrast, CS-materialism+ does not face this argument.

The argument for RM3.c parallels the argument given for RM1.c. The first premise of the conceivability argument says that it is conceivable that there be someone identical to me in all physical respects but different from me in respect of some physical property. Do we include the properties in virtue of which causal skeletal properties obtain or not? Either way, therefore, the conceivability fails as an argument against CS-materialism+.

RM3 Assessed.

Once again, there are elements here that are not for everybody. Some philosophers for example will question the assumption that if one event causes another there must be some further property of the event in virtue of which it does.

But I don't want to focus on this, for there is a more serious problem with RM3. The problem concerns why physics is limited to describing the causal skeleton of the world in the first place. Russell's claim here is not empirical.[18] It is not that physics *as a matter of fact* only describes the causal skeleton, or that it only does so *currently*. Instead what is lying behind Russell's view is a theory about what empirical knowledge as such consists in.

Russell sets out this theory—sometimes called 'representative realism'—in the middle section of *The Analysis of Matter*. The first part of this theory says that in the first instance we have knowledge of the instantiation of particular properties and relations—qualities, as Russell called them—to which we bear a direct perceptual or quasi-perceptual relation. The qualities include: colors, shapes, experiential properties mentioned earlier, and logical and mathematical properties. The second part of the picture concerns how we might draw inferences about the external world from knowledge of the instantiation of these qualities. Russell's idea is that this sort of inference preserves mathematical structure but nothing else: "wherever we infer from perceptions, it is only structure that we can validly infer; and structure is what can be expressed by mathematical logic"

[18] That Russell's claim here is not empirical is consistent with what I take to be an obvious fact, viz., that *The Analysis of Matter* is motivated by developments within physics.

(1927, 254). So for Russell the reason that *physics* aims at describing structure is that *any* inquiry about the external world at all is based on a structural similarity between the qualities we know directly and items in the world.[19]

Now, to the extent that RM3 is connected to representative realism, it is fair to say it is unattractive, though for reasons of space I will simply mention the relevant points here rather than work them through. First, many contemporary philosophers reject representative realism outright; they hold against Russell that we can and do have non-inferential knowledge of the external world. Second, even if the first part of Russell's representative realism is true, it is not clear why the second part is: why are we limited to hypotheses about structure, as opposed to hypotheses of other sort? Third, even if both parts of it are right, Russell's picture is plausibly committed to dualism anyway, which drains it of interest as far as Russellian Monism is concerned, since, as we saw at the start, RM is supposed to be a view that gets away from dualism. If representative realism is true, if one instantiates an experiential property, one knows what it is in what Lewis much later called 'an uncommonly demanding sense'. But notoriously, if one knows what it is in that sense, it is hard to see how anything short of dualism could be true.[20]

So RM3 is implausible if it is given the background Russell gave it. Can one detach it from that background? I take David Chalmers to be suggesting this when he says that "physical descriptions of the world characterize the world in terms of structure and dynamics" (2010, 120), and goes on to explain that "in formal terms, a structural-dynamical description is one that is equivalent to a Ramsey sentence whose O-terms include at most spatiotemporal expressions, nomic expressions, and logical and mathematical expressions (2010, 120). Here, the reference to 'Ramsey-sentences and 'O-terms' should be understood in the way we saw Lewis does, and so Chalmers may be read as combining elements of Lewis and Russell, as follows. Like Russell, for Chalmers a causal-structural property as one expressed by a predicate of physics; unlike Russell, however, for Chalmers physics is understood, not in the light of representative realism, but rather as is a proper part of Lewis's final theory, i.e. the part in which the o-terms are restricted in the way indicated.

This version of RM3 certainly improves on the original. Moreover, while this version of RM3 uses ideas from Lewis, it avoids the problem we noted for RM2. The underlying reason is that

[19] One might wonder why for Russell the hypotheses of physics are causal, in addition to structural. The answer comes in passages like this, in which it is clear that for him to say that a hypothesis is causal says something about the way it is used rather than about its subject matter: "There is a causal relation whenever two events, or two groups of events…are related by a law which allows something to be inferred about the one from the other" (1927, 369).

[20] This sort of argument is in Lewis 1995; for discussion, see Stoljar 2009b

the class of o-terms is much larger for RM3 than it is for RM2. In the case of RM2, one arrives at the Ramsey sentence of the final theory by replacing with variables all the terms in it for fundamental properties; every other term is an o-term. But in the case of RM3, one arrives at the Ramsey-sentence by replacing with variables all the terms that are not spatiotemporal, logical or mathematical and nomic. Both Ramsey sentences permit more than one realization; in both cases, for example, it may be that <F, G, H> realizes the final theory in at this world, and <F*,G*, H*> realize it at some other world. But where for RM2, these realizations are qualitatively identical though numerically distinct, for RM3 these realizations might be qualitatively distinct too, so long as they are identical from the point of view spatiotemporal, nomic and mathematical properties.

However, while this version of RM3 is attractive, a serious problem for it emerges when we ask what 'spatiotemporal expressions' are. In another part of his work, Chalmers argues that spatial expressions "function to pick out that manifold of properties that serves as the normal causal basis of a corresponding manifold of properties in our spatial experience" (2012, 335). But, if we assume that this applies to spatiotemporal expressions generally, the result is that physics for Chalmers is somewhat open-ended. This by itself is no objection; open-endedness is appropriate for many purposes. But it is no good if one is out to defend RM3. For we now lose our reason to believe RM3.a, the claim that CS-materialism is false for a reason distinct from the conceivability argument. Suppose physics tells us about some sequence of causally related events, and that we agree (never mind why) that the events in the sequence have some further properties in virtue of which they stand in these causal relations. Why should these further properties not count as physical by Chalmers' lights? They too, after all, serve as the normal causal basis for the manifold of spatial, and spatiotemporal, experiences.

It might be replied that this objection neglects an important distinction, viz., between first- and second-order properties.[21] To illustrate the distinction, consider again the point mentioned in the course of discussing RM1, that when a chair is uncomfortable there must be some property of it in virtue of which it is uncomfortable. In this example, first-order property of the chair is its shape, i.e., the property that causes discomfort, while the second-order property is the property of having some property that causes discomfort. Now suppose again that physics tells us about some sequence of causally related events, and that these events must have some further properties in virtue of which is so. I argued that these further properties are physical on the ground that they are the normal causal basis for spatiotemporal experiences. But one might think (a) that the events in

[21] I am very much indebted to a discussion with David Chalmers here. It should be noted that my discussion in the text greatly simplifies the account of spatial expressions presented in Chalmers 2012; in particular, it does not discuss the distinction between concepts and properties. I will not try to deal with this issue here however.

question have two sets of properties, i.e. both first- and second-order properties; and (b) that CS-materialism as intended by Chalmers is committed only to the second-order properties being physical. If so, we retain our reason to believe RM3.a, for the first-order properties of the events will be instantiated and yet not be physical.

It is true that if this line of argument correct, we retain our reason to believe RM3.a. But by the same token, we lose our reason to believe RM3.c, the claim that CS-materialism+ does not face the conceivability argument. Unlike CS-materialism, CS-materialism+ will treat both the first- and second-order properties as physical, and it will evade the conceivability argument only if physical theory does not tell us about the first-order properties. But that does not seem right. It might be that physics is limited in various ways; there might be first-order (and indeed second-order) properties *of some kind* that physics does not tell us about. But to say that physics cannot tell us about the first-order properties *that cause our spatiotemporal experiences* is to say that physics cannot tell us about the causes of those experiences. But whatever else it does, surely physics can tell us about the causes of our experiences![22]

So it would seem that the situation we are in at this point is this. RM1 assumes a view about our access to the physical world that is empirically inadequate. RM2 makes no such assumption but has no application to philosophy of mind. RM3 in the form suggested by Russell is associated with an implausible epistemology and in the form suggested by Chalmers is either too open-ended or has the consequence that physics does not tell us about the causes of our experiences.

How to move forward? The only possibility I see starts from two general observations about the versions of RM we have examined so far. First, all of them respond to the conceivability argument in a particular sort of way, i.e. they try to make plausible the hypothesis that we are ignorant in a certain sort of way about the physical, and then exploit that hypothesis to undermine the conceivability argument. Second, all of them try to make that hypothesis plausible in what might be called an *a priori* sort of way, i.e. by a priori reflection on physical science, and related matters.

Can we develop RM by separating the first observation here from the second; in other words, instead of trying to argue that we are ignorant of the nature of the physical a priori, can we start by assuming this as an empirical (but not implausible) starting point? It is that sort of perspective that motivates the final version of Russellian Monism I will consider.

RM4

[22] Since Chalmers explains spatiotemporal expressions in terms of experiential expressions, there is the further worry that his version of RM3, like Russell's, is committed to dualism anyway. I will not try to press that concern here.

The final version (RM4) starts off from some ideas set out in Thomas Nagel's *The View from Nowhere*, particularly in passages such as this:

> The difference between the mental and the physical is far greater than the difference between the electrical and the mechanical. We need entirely new intellectual tools, and it is precisely by reflection on what appears impossible—like the generation of mind out of the recombination of matter—that we will be forced to create such tools. It may be that the eventual result of such exploration will be a new unity that is not reductionist. We and all other creatures with minds seem to be composed of the same materials as everything else in the universe. So any fundamental discoveries we make about how it is that we have mind and what they actually are, will reveal something fundamental about the constituents of the universe as a whole. (1986, pp. 52-3).

There are three ideas to take from this passage.[23] First, the sort of scientific theories that we have currently are incomplete in a fairly strong sense, particularly as regards consciousness. Second, while this is so, we might hold out the hope that in the limit of inquiry (a limit which we will perhaps never reach) a complete theory will be arrived at, a theory according to which the relation of consciousness to matter will be in some respects at least like the relation of the electrical to the mechanical. Third, we can take this sort of picture as implying that we are currently ignorant about theoretically important aspects of matter (or of what we call 'matter')—hence we find in Nagel a focus on the south pole (as we called it earlier) of the mind-body relation.

RM4 Formulated.

With these ideas in place, our characterization of RM4 proceeds in two stages. The first combines M1 with a view about physical properties that is suggested by (but not entailed by) Nagel's discussion. According to this account, a physical property is a property expressed by a predicate of our total current scientific theory—for short, physical properties are current-theory properties. We may call the resulting version of materialism 'current theory materialism' or 'CT-materialism' for short. CT-materialism is true at a possible world *w* just in case for every property *G* instantiated at

[23] This passage is quoted in Pereboom's recent discussion of Russellian Monism; see 2011, 116; indeed it was reading Pereboom that brought this passage to my attention. For a reason I mention in the last two paragraphs of the paper, is not quite clear to me that Nagel is endorsing Russellian Monism here, but there is no doubt he is endorsing something similar. (Nagel expresses some sympathy for a Russellian view in his 2000.)

w, there is some current theory property (or complex of current theory properties) *F* instantiated at *w* such that *F* metaphysically necessitates *G*.

The second stage is to present RM4 as making three distinctive claims about CT-materialism. The first claim of RM4 is this:

RM4.a. CT-materialism is false, and false for reasons quite distinct from those involved in the conceivability argument.

The reason for RM4.a is simply the sort of perspective taken up in Nagel's passage. If we think of ourselves as inquirers within a world, then it is difficult to deny that the picture of the world that we have developed to this point is incomplete in various ways—not for a priori reasons but for empirical though abstract reasons.

The second claim is this:

RM4.b While CT-materialism is false for the reason mentioned in RM4.a, there is a substitute thesis which is not false for that reason, and which preserves the spirit and structure of CT-materialism.

The substitute thesis—which we may call 'CT-materialism+'—is exactly like CT-materialism but with this difference: where CT-materialism says that the physical properties are those expressed by predicates of the current theory, CT-materialism+ say that the physical properties are *either* those expressed by such predicates or those expressed by the final theory, the theory that exists, as Lewis put it, in the way that yet-to-be-written poems do. Does CT-materialism+ preserve the spirit and structure of CT-materialism? As before, it certainly seems so.[24]

The third claim of RM.4 is this:

RM4.c. While CT materialism is false for the reason mentioned in RM4.a, it is also false for another reason, viz., the conceivability argument; by contrast, the substitute thesis mentioned in RM4.b does not face this argument.

[24] Is CT-materialism+ the same thesis as either FT-materialism or FT-materialism+? If the proponent of RM4 accepts the presuppositions of the proponent of RM2, the answer would appear to be 'yes', though I will not assume here that these presuppositions are shared.

The idea behind RM4.c is the same as that behind the counterpart theses in the other versions of Russellian Monism. The conceivability argument says that it is conceivable that there be someone identical to me in physical respects but for whom the coffee is less bitter. Are future physical properties included as physical? For reasons analogous to those we have seen already either answer to this question leaves the conceivability argument unpersuasive.

RM4 Assessed.

Once again there are elements here that are not for everyone. Lewis, for example, famously holds that current physics is very nearly complete, i.e. that every type of physical truth was known or nearly so, and that every contingent truth follows a priori from (and indeed can be defined in terms of) physical truths (see, e.g. Lewis 2009, fn.5).

I find optimism of this sort incredible but I don't want to focus on it here. For there is a more serious problem for RM4, sometimes called the 'structure and dynamics objection'. In Chalmers' canonical formulation it goes like this:

> First, physical descriptions of the world characterize the world in terms structure and dynamics. Second, from truths about structure and dynamics, one can deduce only further truths about structure and dynamics. Third, truths about consciousness are not truths about structure and dynamics (2010, 120).

If these premises are true, CT-materialism+ is no better off with respect to the conceivability argument than CT-materialism. The reason is that both theories will tell us that truths about consciousness are necessitated by truths about structure and dynamics; hence while we may well for other reasons find CT-materialism+ more plausible than CT-materialism, it leaves the issues surrounding the conceivability argument unaffected. Hence, RM4.c, and RM4 itself, should be rejected.

If the structure and dynamics objection is sound, the overall situation looks bleak. The Russellian Monist hoped for a theory that is (a) as elegant as materialism but (b) as comprehensive as dualism. But if what I have been saying so far is right, that hope is baseless. In particular, there are four versions of Russellian Monism, but all of them are implausible.

As against this, however, if we look harder we will see that the structure and dynamics objection ('SDO', henceforth) is more complicated, and much less plausible, than it appears. Explaining why this is so is the topic of my final section.

Structure and Dynamics Again[25]

Looking over the intellectual achievements of the twentieth century, references to structure are very common. We have (of course) the *Analysis of Matter*, the *Aufbau*, *Syntactic Structures*, *The Structure of Scientific Revolutions*, Lévi-Strauss and so on. In this century too, references to structure aren't hard to find: structuralism in philosophy of science, Sider on 'the fundamental structure of reality' etc.[26]

I mention this not to poo-poo structure; that would be silly. But the fact that references to it are ubiquitous should remind us that if a philosopher mentions 'structure' in the course of developing an argument or objection, the first thing you should do is ask what precisely is meant.

What then does 'structure'[27] mean in SDO? One suggestion is obviously in Russell: by 'structure' is meant (mere) *mathematical* structure. But this renders the first premise of SDO quite implausible. Physics does not characterize the world in terms of mere mathematical structure, because characterizations of that sort are mathematical rather than empirical.[28]

Another suggestion is that structure means *metaphysical* structure, i.e. a system (any system) of relations (i.e. *n*-place properties, where $n>1$). But on this interpretation, the SDO is subject to a number of objections, which I have set out in detail elsewhere (Stoljar 2006, 2009a). First, on this interpretation, the first premise says that physics tells us only about relational truths. But that is implausible at least if it is a thesis about physics as such—surely physics can and does tell about the one-place properties of the systems it studies? Second, it is false that from truths about relations you can only derive more truths about relations. For example, from the fact that a series of points in space are arranged in a particular way, you might derive the fact that that the region constituted by the points has a particular shape. Third, it is false that no truth about consciousness is a truth about relations, for some truths about consciousness are themselves relational.

Derk Pereboom (2011, 102, 112; see also Alter and Nagasawa 2011) has responded to these points by describing the notion of an absolutely intrinsic property, which is (roughly) a property which is (a) non-relational and (b) is not (a priori) derivable from any relational set of facts. In particular, Pereboom says, truths about absolutely intrinsic properties are not derivable from truths about relations and so the second premise of the SDO is true if it is read as saying that from truths

[25] 'Again' because I have discussed this objection twice before; see Stoljar 2006, 2009a.

[26] I will leave it as an exercise for the reader to chase down the references mentioned in this paragraph.

[27] The concept of dynamics is simply the concept of something changing its structure over time, so what we say will apply here *pari passu*

[28] This point is very closely related to a famous objection to Russell made in Newman 1928 according to which (very roughly) he confuses physics and mathematics. See Demopolous and Friedman 1985 and Chalmers 2012

about relations you may derive only truths that do not concern absolutely intrinsic or non-relational properties.

I agree that if you understand the premise that way it is true; indeed, on this interpretation it is analytic as Pereboom points out (2011, 113). But the key question is whether this helps the polemic purposes of a proponent of the SDO. And it is clear that it does not. For the only way in which this could help the argument is if the third premise were interpreted now as saying that the truths about consciousness are or include truths about absolutely intrinsic properties. But this is highly questionable. There is certainly some intuitive basis to the idea that truths about consciousness concern non-relational or intrinsic properties, but whether they concern *absolutely* non-relational or intrinsic seems to me something that is left open by anything we know or believe about consciousness, either as a matter of theory or as a matter of introspection.

David Chalmers has responded to these criticisms by accepting them so far as they go but suggesting they miss their target (2010, 210, fn.18). What he has in mind by 'structure' is neither mathematical nor metaphysical structure but rather spatiotemporal and nomic structure: "formally, a structural-dynamical description is one that is equivalent to a Ramsey sentence whose O-terms are limited to spatio-temporal expressions, nomic expression, and mathematical and logical expressions". Of course, we have looked at this idea already, during the examination of RM3. We criticized it in that context for being too open-ended or as entailing objectionably that physics does not tell us about the causes of our experiences—at least given Chalmers' suggestion that spatiotemporal expressions pick out whatever properties are the causal basis of the manifold of spatiotemporal experiences. Something similar is true when we consider it in the context of SDO. In particular, if we interpret this view as entailing that the properties that cause our experiences are physical, a proponent of RM4 will deny the third premise of the SDO: "we don't know currently what those properties are", he will say, "so we are in no position to assert that no truth about consciousness is a truth about them". On the other hand, if we interpret this view as supposing that physics tells us only about the properties of having some properties that cause experiences, the proponent of RM4 will deny the first premise of SDO: "surely physics can tell us about the causes of our experience", he will say, "so physics does not characterize the world merely in terms structure and dynamics in the relevant sense".

In sum, if by 'structure' one means either mathematical, metaphysical, or nomic and spatiotemporal structure, SDO is unpersuasive, and RM4 emerges as the most promising version of Russellian Monism. Still, one might wonder in closing whether there is a slightly different objection lying behind these animadversions about structure. This objection is that RM4 is not a genuine sort of Russellian Monism, not because of structure and dynamics, but because it does not

use key metaphysical ideas which each of RM1-3 uses. In particular, RM4 makes no theoretical use of the dispositional/non-dispositional distinction of RM1, the role/realizer distinction of RM2 and the structural/non-structural distinction of RM3. [29]

Actually, I have some sympathies with this. RM4 *is* different in this way from the other versions; and in other places I have treated Russellian Monism as a specific version of RM4, rather than the other way around (see Stoljar 2006).[30] Nevertheless there are two points to make. First, in view of the elusiveness of terms like 'structure' (etc.), it is unclear that one can use them to mark any natural division between RM1-3, on the one hand, and RM4 on the other. Second, even supposing there is a difference in this respect between RM4 and RM1-3, there remains an important similarity, viz. their common strategy for answering the conceivability argument. So perhaps our conclusion should be this: RM4 might not (*might not*) be a version of Russellian Monism, but it remains the closest thing that is plausible.

References

Alter, T. 2009 'Does the Ignorance Hypothesis Undermine the Conceivability and Knowledge Arguments?' *Philosophy and Phenomenological Research* 79 (2009), pp. 756-65

Alter, T and Nagasawa, Y 2012 'What is Russellian Monism?' *Journal of Consciousness Studies*, Volume 19, Numbers 9-10, 2012 pp. 67-95(29)

Alter, T, and Nagasawa, Y (eds.) forthcoming *Russellian Monism* Oxford University Press

Armstrong, D.M. 1961 *Perception and the Physical World.* London: Routledge

Armstrong D.M. 1968 *A Materialist Theory of Mind*. London: Routledge

Blackburn, S 1990 'Filling in Space' *Analysis* 50 (1990) pp. 62-5

Chalmers, D. 1996 *The Conscious Mind* New York: Oxford University Press

Chalmers, D. 1999 'Précis of *The Conscious Mind'*, *Philosophy and Phenomenological Research*, Vol. LIX, No.2, pp.435-438.

Chalmers, D. 2010 *The Character of Consciousness* Oxford University Press

[29] This point is reflected in an influential taxonomy of positions set out in Chalmers 2010. Chalmers would, I think, classify RM4 as a version of materialism—type-C materialism, in particular—and would set it aside from genuine Russellian Monism which would be classified as type-F monism.

[30] In his contribution to this volume, Derk Pereboom suggests a different way to respond to this objection, viz., to advance a version of RM4 that requires that the properties of which are ignorant are absolutely intrinsic. I have no objection to the idea that this *might* be true (the version of RM4 I discuss in the text permits it), but I am unconvinced by Pereboom's suggestion that it *must* be. For one thing, Pereboom's argument for this point apparently relies on the idea sometimes called 'revelation'—and this is an idea I have been critical of elsewhere (see Stoljar 2009a). But more generally, consider again Nagel's suggestion that the relation of consciousness to the physical might be like the relation of the mechanical to the electrical. It is very doubtful that the notion of an absolutely intrinsic property played a crucial role in the complicated empirical process that resulted in the unification of the mechanical and the electrical—why then suppose that it must play a role in an analogous process (as yet not undergone) that concerns consciousness?

Chalmers, D. 2012. *Constructing the World* Oxford: Oxford University Press. References are to the PDF version.

Demopoulos, W and Friedman, M. 1985 'Critical Notice: Bertrand Russell's The Analysis of Matter: Its Historical Context and Contemporary Interest, *Philosophy of Science*, 52 (1985) 621-639.

Feigl, H. 1967 *The "Mental" and the "Physical",* Minneapolis: University of Minnesota Press (Original publication 1956)

Foster, J. 1982. *The Case for Idealism* London: Routledge

Handfield, T. 2009 (ed) *Dispositions and Causes* Oxford: Clarendon Press

Holman, E 2008 'Panpsychism, Physicalism, Neutral Monism and the Russellian Theory of Mind' *Journal of Consciousness Studies*, Volume 15, Number 5, 2008 , pp. 48-67(20)

Langton, R. 1998 *Kantian Humility* Oxford: Oxford University Press

Lewis, D. 1995. "Should a Materialist Believe in Qualia", *Australasian Journal of Philosophy* 73, pp. 140-44

Lewis, D. 2009 "Ramseyan Humility", in David Braddon-Mitchell and Robert Nola (eds) *Conceptual Analysis and Philosophical Naturalism*, Cambridge, MA: MIT Press, pp, 203-22

Locke, D. 2009 "A Partial Defense of Ramseyan Humility" David Braddon-Mitchell and Robert Nola (eds) *Conceptual Analysis and Philosophical Naturalism*, Cambridge, MA: MIT Press, pp, 223-41

Locke, D. Forthcoming "Quidditism without quiddities", *Philosophical Studies* Available at: http://www.cmc.edu/pages/faculty/dlocke/Quidditism%20without%20Quiddities.pdf

Locke, J. 1975 *An Essay Concerning Human Understanding* ed. P.H.Nidditch Oxford: Oxford University Press

Lockwood, M., 1989. *Mind, Brain and the Quantum*, Oxford: Blackwell

Montero, B. 2010 'A Russellian Response to the Structural Argument Against Physicalism, *Journal of Consciousness Studies* Volume 17, Numbers 3-4, 2010 , pp. 70-83(14)

Nagel, T. 1986. *The View from Nowhere* New York: Oxford University Press

Nagel, T. 2000, "The Psychophysical Nexus", in Paul Boghossian and Christopher Peacocke (eds.) *New Essays on the A Priori*, Oxford: Clarendon Press, pp. 432-471

Newman, M.H.A. 1928 "Mr Russell's Causal Theory of Perception", *Mind* 37 pp.137-48

Pereboom, D. 2011 *Consciousness and the Prospects of Physicalism* New York: Oxford University Press

Russell, B. 1927 *The Analysis of Matter* London: Kegan Paul

Smart, JJC. 1959 'Sensations and Brain Processes' *Philosophical Review* 68: 141-156

Stoljar, D. 2001 'Two Conceptions of the Physical', *Philosophy and Phenomenological Research*, Vol. LXII, No.2, March 2001, 253-281.

Stoljar, D, 2006. *Ignorance and Imagination* Oxford University Press

Stoljar, D, 2009 'Reply to Alter and Bennett,' *Philosophy and Phenomenological Research* 2009 79 (3): 775-784

Stoljar, D, 2009a. 'The Argument from Revelation.' In Robert Nola and David Braddon Mitchell (eds) 2009 *Conceptual Analysis and Philosophical Naturalism* (MIT Press) pp. 113-138

Stoljar, D. 2010 *Physicalism* London and New York: Routledge

Stubenberg, L. (1997) 'Austria versus Australia: Two Versions of the Identity Theory' in K. Lehrer and J. Mare (eds) *Austrian Philosophy Past and Present: Essays in Honour of Rudolf Haller* Dordrecht: Kluwer, 125–46.

Van Cleve, J. 2002 'Receptivity and Our Knowledge of Intrinsic Properties', *Philosophy and Phenomenological Research* 65: pp. 218-37.