

MULTIPLE REALIZABILITY AND *NOVEL* CAUSAL POWERS

Ricardo Restrepo

Abstract

Framed within the dialectic of the causal exclusion argument (Kim 2005), this paper does two things. One, it clarifies some properties of multiple realizability based on its true origin (Turing 1950). And two, it challenges a form of argument Noordhof (1997), Clarke (1999), and Whittle (2007) employ to support the idea that the mental has causal powers not had by its physical realization base (*Novel*). The paper challenges *Novel* with ideas derived from multiple realizability, among others.

1. Introduction

Does the mental have novel causal powers its physical realization base does not have? Taking the lead from Noordhof (1997) and Clarke (1999), Whittle (2007) employs an argument which concludes that the mental does have such novel powers (*Novel*). The conclusion is exciting and has received relatively little critical attention. I will argue, however, that the grounds for *Novel* are unpersuasive. Building on Turing's remarks, an analysis of multiple realizability reveals a preferable alternative explanation for the cases Whittle, Noordhof, and Clarke use to support the novelty of the causal powers of the mental. *Novel* arises in the context of providing a solution to the problem of causal exclusion originally developed by Jaegwon Kim. It is a "problem" in that it aims to display an inconsistency in certain non-reductionist assumptions, and then argues by *reductio ad absurdum* that non-reductionism is not true. The challenge for the non-reductionist is to provide a solution that rescues the position's core commitments. The problem of causal exclusion says that since every effect has a sufficient physical cause, and there is no systematic overdetermination, mental properties must either be reduced to physical properties in order to be causal or be ontologically eliminated, on pain of inconsistency (Kim 2005). Whittle considers *Novel* to be an obstacle to her own proposed solution to the problem of causal exclusion. However, since *Novel* is ineffective, she may see this as an upside in the problem of causal exclusion dialectic, even though I am not a proponent of her solution, Noordhof's, or Clarke's.

2. Multiple Realizability

Multiple realizability is an important element in the contemporary conception of higher-level things (we will centrally deal with properties). Multiple realizability pertains to things which can be made real in more than one way, and its existence supports the view that properties in higher-level domains are distinct from the properties that realize them (Pereboom 2002). In this section I wish to develop some lessons from the origins of the conception of multiple realizability. In my view, these lessons invite us to take an appropriate stance toward *Novel* (to be articulated in the next section).

The conventional wisdom is that ‘[i]n a series of papers published throughout the 1960s, Hilary Putnam introduced multiple realizability into the philosophy of mind’ (Bickle 2006).¹ Claims similar to this one have been made by Block (1980; 1990), Kim (1992), Shapiro (2000: 635) and Funkhouser (2007b). However, while Putnam was the first to use the multiple realizability of the mental to argue for non-reductionism, he did not introduce the concept into the philosophy mind. Turing made distinctive use of the concept in the philosophy of mind no less than ten years earlier, when he was thinking about the possibility of building artificial intelligence in his article in *Mind*. Turing wrote:

The storage was to be purely mechanical, using wheels and cards.

The fact that Babbage’s Analytical Engine was to be entirely mechanical will help us rid ourselves of a superstition. Importance is often attached to the fact that modern digital computers are electrical, and the nervous system is also electrical. Since Babbage’s machine was not electrical, and since all digital computers are in a sense equivalent, we see that this use of electricity cannot be of theoretical importance... If we wish to find such similarities we should look rather for mathematical analogies of function (Turing 1950: 446).

¹ A reviewer for this journal noted that the multiple realizability of mental properties is ‘quite obvious and has been known all along’. The above quote, written by a prominent philosopher in the entry on *multiple realizability* in one of the top contemporary encyclopedias of philosophy, shows that this is not so, for it says it was only introduced in the 1960s (e.g. 1960; 1967). That there are many things which can share a property- that properties are multiply exemplified- is of course ancient information, about which Socrates inquired much. But this phenomenon, while conceptually close, cannot be assumed to be the phenomenon of interest to philosophers of mind when they speak of ‘multiple realizability’. It must be distinguished from the claim that there is multiple realization- which, according to a live approximation, is sameness of type through differences in the (lower-level) conditions that give rise to instances of that type (Funkhouser 2007a). This is precisely what Turing points to in his above remarks; for claiming that computing properties are realizable by wheels and cards as well as by electrical discharges, is of course to claim that two drastically different lower conditions give rise to instances of the same (computational) type.

As can be seen, this passage contains, in addition, key elements of functionalist and computational theories of mind which are commonly thought to have been introduced into the philosophy of mind by Hilary Putnam at least ten years later.

Theoretical importance, Turing argued, should not be attributed to just any property of considered machines. Two machines might weigh 10 kilograms but this could scarcely suffice for them to realize the same relevant type of computing machine. It is only some properties of physical machines that realize their being computing machines of the relevant types.

Similarly, to realize mentality, what is needed is the realization of certain, but not all, properties of things that have mentality. The hypothesis of computational cognitive science is that realizing certain computational properties of us is sufficient for having certain mental properties in question (a complementary formulation is found in Chalmers (1996: 309; Restrepo 2009)).

In this spirit, in a 1952 BBC discussion with Richard Braithwaite, Geoffrey Jefferson, and Max Newman, Turing said: ‘we are not interested in the fact that the brain has the consistency of cold porridge. We don’t want to say ‘This machine is quite hard, so it isn’t a brain, so it can’t think’’ (495). In order to create artificial intelligence or any other mental phenomenon, we must look at certain properties of the brain and not others. This opens the possibility that there are many ways in which mental properties can be realized.

The brain has certain causal powers in virtue of the properties it has. Normally, a living brain has the power to cause characteristic effects of having the consistency of cold porridge. However, this property and these powers are not important for psychology. Rather, examples of properties of the brain that are interesting for psychology are those that cause intelligent behaviour in persons who display it, those that realize the perception of redness in people who have those perceptions, feeling exhilarated, having pain, etc. The multiple realizability assumption has the fruitful effect of focusing our attention on those properties we wish to understand, and potentially replicate, without irrelevant properties confounding our attention.

Contrary to Turing’s claim, however, if electricity makes a positive causal contribution to our mental life, then it is of theoretical importance for mentality. An alternative multiple realization of mentality will have to duplicate those properties of

electrical transmission which contribute to mentality; never mind that having the consistency of cold porridge is not a property of importance. That positive causal contributions of multiply realizable properties are important features multiple realizations must duplicate is a recent development in the inquiry into the nature of multiple realizability. Shapiro (2000: 647) argues that philosophers who hold the multiple realizability of mental face the challenge of clarifying what multiple realizability amounts to. For if realizing kinds do not differ in the causally relevant properties then we do not have a legitimate case of multiple realizability, and if they do differ, then they are different kinds. "Steel and aluminum are not different realizations of the waiter's corkscrew because, relative to the properties that make them suitable for removing corks, they are identical" (Shapiro 2000: 644). Further, corkscrews can have different colors, but this is not sufficient for their being multiply realizable. Pereboom (2002: 524-525) counters that the difference is that colors do not make positive causal contributions to the nature of corkscrews- what they do- while being made of steel or aluminum does. Pereboom (2002: 525) concludes that "[a]ccordingly, *making a causal contribution to the nature of the thing that has it* might be the notion of causal relevance that is pertinent to a condition on multiple realization. This alternative conception would license steel and aluminum but not distinct colors as multiple realizations of a corkscrew, and for silicon and neural systems to count as multiple realizations of psychological features".

Just like the color of corkscrews, having the consistency of cold porridge does not realize mental properties; while electricity in machines that compute through electrical transmission does. In addition, while a special purpose adding machine that weighs 10 kilograms and a word processing machine that also weighs 10 kilograms have a similarity that may be of interest for some purposes, for other purposes, they are appropriately grouped differently. This interest-relativity is not ontologically degenerate. All things have real differences and similarities, which determine the sort of effects they have. Sometimes we are interested in some properties and some effects, and sometimes in others. It is to be expected that different grouping practices, *on the basis of real shared properties and causal powers*, are employed to serve our interests. This selection of properties in virtue of our interest and the fact that they fit into a causal

pattern is what Macdonald and Macdonald (2006) capture under their understanding of causal relevance.

Searle (1980; 1992) believes that such interest-relativity makes multiply realizable properties ontologically degenerate in comparison to physical properties. However, there are paradigmatic physical properties that are multiply realizable as well. Being a solid is a multiply realizable physical property. Being a solid can be realized by being made of Carbon and it can be realized by being made of Iron, for instance. For some purposes, these properties are grouped together and at other times they are grouped differently for different purposes.

Further, some possible realizations of a higher-level property do not always realize that property: at least in some cases, adequate background conditions need to be in place. Wheels and cards are one of the multiple realizations of digital computers; silicon-based electronic devices are another. Like silicon-based electronic devices, wheels and cards have the characteristic causal powers which enable them to realize the property of being a digital computer. However, in order for a collection of silicon-based electronic devices, and wheels and cards to activate the causal powers which enable them to compute a relevant function, they have to be appropriately put together into a computing machine in question. Just any organization will not do, of course; a specific one is required. Depending on where they are located within the machine, different causal powers of different wheels and cards, and electronic circuits will be activated.

In a particular computation, a few wheels and cards within the machine might move in such a way that we can narrowly attribute to them being the realization of the computation of the function in question. Shoemaker (1984) would call them “core realizations”. However, without the rest of the machine holding them appropriately positioned and timed relative to each other to generate certain effects, those parts would not realize the computation. These background conditions are composed of causal antecedents and part-whole relations of the object in question (in this case, a machine of a particular sort). Background conditions are quite important for core realizations, for without them core realizations would not be successful realizers of higher-level properties. The background conditions of the sort illustrated here together with the core realizations are what is sufficient for the complete realization of a higher-level property,

and channel the causal flow in such a way that the appropriate causal powers of the realization are activated. Shoemaker (1984) calls them ‘total realizations’.

We can see that core realizations by themselves will not be sufficient for higher-level properties they could realize; they need to be under appropriate conditions. When those conditions are in place, however, the realization is complete. It is total realizations that are sufficient for the things they could realize. When core realizations are not under appropriate background conditions, they cannot be reliably expected to realize the property in question and have its characteristic effect.

3. An Alternative Explanation to *Novel*

Whittle’s (2007) *Novel* runs as follows:

... suppose that C-fibres firing realizes pain in me, but sleepiness in my dog, whereas D-fibres firing realizes pain in my dog. One evening, I accidentally touch a hot hob and the pain/C-fibres firing causes me to cry out. Concurrently, my dog is also experiencing C-fibres firing, and yawning as a result. But if he had been in pain, he would have howled, so pain, unlike C-fibres firing, does cause him to cry out. This case is problematic not because C-fibres firing bestows causal powers not associated with pain, i.e. in this case, the causal power to make dogs sleepy; but because pain seems to have a causal power that C-fibres firing doesn’t, namely, the power to make my dog howl (Whittle 2007: 25-26).

We can see that Whittle is committed to the existence of multiple realizability. Whittle takes it that having pain could be realized by having C-fibres fire and by having D-fibres fire. Additionally, Whittle asks us to consider this as a case in which one instance of having C-fibres fire (the dog’s) does not realize having pain and does not have the characteristic effect of having pain (howling). Whittle concludes that having pain has a causal power that having C-fibres fire lacks. This, she thinks, is just a special case of mental properties having causal powers their physical realization bases lack. Consequently, the mental has novel causal powers its physical realization base does not have.

The problem with Whittle’s thought is that there exists a more obvious and simple alternative explanation for her scenario. All that is needed to explain the supposed existence of her scenario is the fact that in her chosen scenario the background conditions (including causal antecedents and part-whole relations) of having C-fibres fire in her and her dog are such that only when they fire in her are the causal powers

characteristic of having pain activated. The total realization was different in her and in her dog. If one changes these background conditions appropriately, having C-fibres fire in the dog would realize having pain and cause the dog to howl.

It might be true in a narrow stipulative sense that if the dog's C-fibres fire, pain would not be realized and the dog would not exhibit the characteristic behavior of having pain. But the alternative explanation is just that in the scenario she chose, the stipulated background conditions do not activate the causal powers relevant to having pain in the dog and that were those background conditions in the dog adjusted, having C-fibres fire in the dog would realize having pain and cause it to howl. For an argument with the conclusion that the mental has causal powers that its physical realization base does not have to work, we should not be able to straightforwardly explain the scenario by these means.

It is an uncontroversial truth that in different circumstances different causal powers of properties are activated, and that the activation of different causal powers of having C-fibres fire in different circumstances is no exception. One should expect quite different effects of having C-fibres fire between a case in which they are hooked up to realize pain and a case in which they are hooked up to realize sleepiness, as Whittle's scenario demands. But provided one adequately adjusts the peculiar way in which Whittle has stipulated her dog's C-fibres are hooked up, her dog's having pain would be realized by its having C-fibres fire and it would howl when they do.

In sum, the dog's having C-fibres fire does not realize his having pain and cause characteristic effects of pain. Whittle explains this fact by saying that having C-fibres fire does not have causal powers that having pain does have. A simpler explanation is just that Whittle has stipulated background conditions for the firing of C-fibers which prevent the causal powers characteristic of having pain from being activated. Were those background conditions to be adjusted, the firing of C-fibers would make pain be instantiated in the dog and cause it to howl.

4. The Empirical Grounds

Whittle thinks this kind of response does not work. She says:

Whilst it is no doubt true to say that, in the right circumstances, if the total realizer had been instantiated by my dog, then he would have howled, this does not establish that, in

the actual world, the total realizer has the causal power to make my dog howl. It may, for instance, be nomologically impossible for the properties of C-fibres firing to realize the pain role in dogs. So, granted such a scenario, claiming that this total realizer has the power to make my dog howl in this world, smacks of desperation (Whittle 2007: 26-27).

It is, however, puzzling why Whittle thinks her proposed nomologically impossible scenario should be taken for granted. Normally, humans and dogs have the same realization base for having pain (for humans see Rosenzweig *et al.* 2002; for dogs see Stafford 2007). So given that we can suppose that having C-fibres fire in Whittle is a possible realization of her having pain, we must be able to suppose that having C-fibres fire in her dog is a possible realization of having pain in her dog. The scientific understanding of the realization of having pain supports the idea that if having C-fibres fire is a nomologically possible realization of having pain in a human, then it is a nomologically possible realization of having pain in a dog. Consequently, given that having C-fibres fire in a human could realize having pain and cause pain behaviour, as Whittle supposes, having C-fibres fire in a dog could realize having pain in the dog and cause it to howl.

The only way in which one could grant Whittle what she asks us to grant is by applying in an *ad hoc* manner our understanding of multiple realizability. We would have to say that even though in the actual world the normal realization of having pain in humans and in dogs is the same, without further argument, in a human, but not in a dog, having C-fibres fire could realize having pain.

5. Methodological Dualism Galore

Whittle thinks the kind of alternative explanation proposed here for her chosen scenario has a fatal consequence:

This would be to abandon the idea that there is a mental property of pain. What we would have *instead* is the property of pain-in-humans and the property of pain-in-dogs...perhaps pain-in-humans has a very different quality than pain-in-dogs and thus warrants the ascription of different mental properties (Whittle 2007: footnote 12, emphasis added).

Compare Whittle's statement to someone who said that to suppose that having C-fibres fire happens only under certain circumstances has the unacceptable consequence that there is no such thing as having C-fibres firing itself, but C-fibres-fire-in-humans and C-

fibres-firing-in-dogs, and so on. The first thing to notice is that it is *true* that C-fibres fire only under certain conditions, just like everything else that has a cause and is a part of other things. So it should not be unacceptably surprising that pain likewise occurs under the conditions of human, of dogs, and so on.

It seems that by Whittle's standards we should be prepared to use hyphens for expressing just about any other idea about things in the world, since everything has causes and can be seen to be a part of other things (except, perhaps, the universe as a whole). However, I do not see what the point of that kind of use of hyphens would be, what confusion it would enlighten, what new fact would become clear, whether this would be a real replacement, or what value would be added.

Secondly, the fact that having C-fibres fire only happens under certain circumstances does not imply the elimination of having C-fibres fire *per se*, and its replacement by something else. Causes don't eliminate their effects, and wholes don't eliminate the existence of their parts. Similarly, having pain should not be eliminated and replaced just because it has causes and is a part of other things which form the total realization.

Thirdly, even if having pain has at least two varieties, in humans and in dogs, for example, this would not eliminate pain. There would still be a mental property, having pain, of at least two varieties. Consider the fact that there are atoms of tens of varieties, star cores of many varieties, and geometrical shapes of infinite varieties, for example. How would it entail from this multiplicity that there is no property of being an atom, or of being a star core, or of having a geometrical shape, or that no object has any of these properties (with their characteristic causal powers)? That this multiplicity with unity is commonplace is not only commonsense; it's also scientifically endorsed. For example, while different atomic numbers determine differences in atomic varieties, it is still true that atoms exist and that they are united in that they are composed of a nucleus with (a) proton(s) and (an) outer layer(s) of electrons. Further, there might well be a qualitative difference between dogs having pain and humans having pain, but it would still be a difference in their *having pain*, just as while there is a qualitative difference between Hydrogen and Helium, they are still atoms. If the answer to the above question is that there is no such entailment, and that this multiplicity with unity is *bona fide* and commonplace, then the unpersuasive and *ad hoc* move is in fact being made Whittle by

just assuming that somehow the mental-to- physical case is different from the rest of the physical world in this respect.

6. Noordhof's novel mental causal powers

Before Whittle, Noordhof (1997: 242) put forth the *Novel* argument like this:

Suppose that a causal role property C can be realized by three grounds: G1, G2 and G3. Let these grounds be causes of three other properties that variably realize a role-characterizing property S. Call these other realizing properties R1, R2, and R3. Then a common feature of variable realization is that although G1 - R1, G2 - R2 and G3 - R3 (where '-' means 'is causally related to') it is not the case that G1 - R2 nor is it the case that G2 - R3. The grounds of a causal role property are causally connected with some realizing properties of a role-characterizing property and not others. So although C is causally connected to S however S is realized, the grounds are not so connected. My thought is that it is in this sense that causal role properties introduce new causal relations. G1 may determine the instantiation of C because G1 is a member of the family of properties upon which C supervenes, but C introduces relations in which G1 couldn't possibly stand.

Perhaps a multiply realizable property has different causal powers than *one* possible realization base. Nevertheless the following theses seem to be unaffected: (i) that all the causal powers of the C coincide with (at least some of) the causal powers of the complete set of multiple realizers, and (ii) that in any particular case where C causes S, all the causal powers of C will be completely coincident with (at least some of) the causal powers C's ground, in virtue of which S is caused. To claim that there are causal relations between higher-level properties not coincident with causal relations between their realizers requires additional justification.

Suppose that there is some grounding property G1 causally related to R1, but not to R2, as Noordhof suggests. *On Noordhof's supposition, when C is realized by G1 it will not be able to cause R2.* The causal powers of C, when realized by G1, are completely limited by the causal powers of G1; not some other possible realization of C. The same is true of any other multiple realization of C: it will constrain the causal powers of C anytime it realizes C. It is hard to see a novelty worth asserting here.

7. Clarke's novel mental causal powers

Clarke puts his version of *Novel* like this:

... M must carry with it causal powers that no one of its base properties has. Suppose that M is realized in a human being by P1, in a lizard by P2, in an octopus by P3, and so on. Consider the causal powers of P1. Presumably, an exemplifying of P1 in a human will cause an exemplifying of P1*. What, however, about P1 in a lizard, in an octopus, etc.? Recall that P1 is a micro-structural property of a human being; it is, for example, a property that a creature has in virtue of the presence within that creature's brain of a certain pattern of neural activity in the structure that is characteristic of a human brain. It would seem to be impossible (at least nomologically) for P1 to be instantiated in a lizard or an octopus. Indeed, it appears that P1 could only be instantiated in a human being. If so, then P1 cannot be instantiated in a lizard, in an octopus, etc., and it cannot carry with it the power to cause in a lizard an exemplifying of P2*, the power to cause in an octopus an exemplifying of P3*, etc. P1, then, does not carry with it some of the causal powers that M carries with it. Or, to put it the other way around, M carries with it some causal powers that P1, one of the subvenient physical properties that realizes it, does not carry. And likewise for the relation of M's causal powers to those of any other subvenient physical property that realizes it (Clarke 1999: 304-305).

Suppose the realization bases of interest are total realizations and are necessarily different, as is most plausible when considering humans and octopi. Consider Babbage's machines and electronic computers. They provide different total realizations, but they can compute the same relevant functions, and in this important sense they are equivalent. Two adding machines with these variable realizations will have the property of processing the addition of 8 and 3. Call one realization in Babbage's wheels and cards machine *C1* and the other realization in the electronic machine *C2*. These two variably realizing properties, causally deliver the same effect: say an 11 paper print. Consequently, it is not necessarily true that two different total realizations do not have the same effect. Similarly, humans and octopi may totally realize being in pain differently, but that realization may have equivalent effects (pain behavior).

True, the pain behavior of humans and octopi have different total realizations and the total realization of pain in humans does not cause the total realization of pain in octopi, nor vice versa. But this case completely correlates with the fact that the pain of humans does not cause the pain behavior of octopi. So it does not support the idea that pain causes things its realizations don't.

Suppose a human has M in virtue of having P1, as Clarke supposes, and an octopus has M in virtue of having P3. Clarke's analysis is that this human's having M by having P1 will cause the human to have P1*, but will not cause P3*, whereas the octopus' having M by having P3 will cause it to have P3*.

However, just as the human's having M will not cause P3*, the octopus' having M will not cause P1*. The differences seen in the causal profile of M are perfectly

correlated with the differences in the causal profile of its multiple total realizations. When M is realized by P1, it will cause P1* and not P3*, and when M is realized by P3, it will cause P3* and not P1*.

We can account for the differential effects of M by considering that it is realized by P1 in one case and by P3 in another, as a consequence of the fact that Clarke supposes: that only one of them could be realized in the human and the other in the octopus and that whenever they are realized in their corresponding organism they have certain characteristic effects and not others. What causal powers does M have that cannot be fully accounted for by its realization base? When M is had by a human it causes P1* in the human and not in octopi, and when M is had by octopi it causes P3* in the octopi and not in the human. For *Novel* to work, it requires that M cause something that cannot be fully accounted for in this way.

As responded to Noordhof, on the supposition that M is multiply realizable, the following theses seem to be unaffected: (i) the causal powers of M will completely coincide with, and never outstrip, the causal powers of its realizers (or a subset thereof), and (ii) in any particular thing that has M, the causal powers of M will not be anything over and above the causal powers of its actual realization base. A case worth arguing for needs to provide grounds to deny these theses.

8. Reductionism

Whittle, Noordhof, and Clarke use the *Novel* argument on the way to providing a non-reductionist solution the problem of causal exclusion. Further, some of the remarks I have made have been in line with some of what Kim has argued (though Kim has not addressed *Novel* directly). So it may seem that my negativity towards *Novel* has the ultimate goal of supporting reductionism. This, however, is far from the case because I think Kim's argument is logically invalid for the reasons Restrepo (2012a; 2012b) expresses. While this is not the place to do full justice to the causal exclusion argument, since this paper is centrally about certain aspects of multiple realizability and the *Novel* argument, it is framed in the causal exclusion debate; so I must at least gesture at a key line of reasoning for not endorsing Kim's reductionist conclusion.

A key premise in Kim's argument is *Closure*: Each physical event has a sufficient physical cause (insofar as causes of events are sufficient) (Kim 2005). Of

course, given physicalism, every cause is physical. But being physical can be a fundamental physical property (a physical₁ property) or a property that is completely resultant from, but not reductively identical with, fundamental properties (a physical₂ property). Jackson illustrates this nicely:

We need... an extended sense (of the physical) because the patterns that economics, architecture, politics and very arguably psychology, pick out and theorise in terms of, include many that do not figure in the physical sciences. The reason is no mystery: it is that aggregation creates new properties ... because aggregations fall under patterns, kinds etc. that the items they are aggregations of do not fall under... Physicalists must allow that the world contains aggregations that have properties that are not physical₁ properties for the same reason, when all is said and done, that someone who holds, rightly, that a triangle is an aggregation of straight lines must allow that the triangle is not itself a straight line (Jackson 2006: 234, parentheses added).

The view I favor is that mental properties are physical₂ properties, which leaves the conclusion *that mental properties are physical causes* without its putative reductionist element. This is not inconsistent with Kim's (2005: 43) explicit conception of the physical as it figures in *Closure*, which he borrows from Papineau (2002). This conception is that the physical is to be understood as what is inorganically identifiable, or "identifiable non-mentally-and-non-biologically". That is, the physical is what can be referred to "independently of this specifically mental conceptual apparatus" (Papineau 2002: 41). This means that those concepts that do not use terms like *seeing* and *believing* are sufficient to refer to anything that exists in nature and participates causally in the world. Certainly, one can refer to triangles, like mental physical₂ properties, without explicitly using the term *triangle*, or *mental*. This could be done with terms like *lines put together in a certain way*, or *quarks and electrons put together in a certain way*. This, however, does not eliminate the causal relevance of physical₂ properties and a logically valid reductionist conclusion is blocked. The details of this "logical reply" are provided by Restrepo (2012a).

9. Conclusion

Whittle (2007) mentions *Novel* in developing a solution to the problem of causal exclusion. She, however, finds her own solution problematic because she thinks that the mental has powers not had by its physical realization base. The problem of causal

exclusion attempts to show that the mental does not have irreducible mental causal powers. To worry that one's response to this argument is problematic because of the supposed fact that the mental has *additional* causal powers is surprising in its own right. Nevertheless, it may be of theoretical comfort that *Novel* does not really pose an obstacle to any solution to the problem of causal exclusion. This is something that reflection on Turing's comments on building computing and thinking machines sheds some light on.

Ricardo Restrepo

Instituto de Altos Estudios Nacionales, Ecuador

ricardo.restrepo@iaen.edu.ec

References

- Bickle, J. (2006) 'Multiple Realizability' IN *Stanford Encyclopedia of Philosophy*, E. Zalta (Ed.), <http://plato.stanford.edu/entries/multiple-realizability/> . Accessed August 18, 2011.
- Block, N. (1980) 'Intro: What is Functionalism?' IN *Readings in Philosophy of Psychology* vol. 1, Cambridge: Harvard University Press.
- _____. (1990) 'Can the Mind Change the World?' IN *Meaning and Method*, G. Boolos (Ed.), Cambridge: Cambridge University Press.
- Chalmers, D. (1996) 'Does a Rock Implement Every Finite-State Automaton?', *Synthese*, 108, 309-333.
- Clarke, R. (1999) 'Nonreductive Physicalism and the Causal Powers of the Mental', *Erkenntnis*, 51, 259-322.
- Funkhouser, E. (2007a) 'A Liberal Conception of Multiple Realizability', *Philosophical Studies*, 132, 467-494.
- _____. (2007b) 'Multiple Realizability', *Philosophy Compass*, 2, 303-315.
- Jackson, F. (2006) 'On Ensuring that *Physicalism* is not a Dual Attribute Theory in *Sheep's Clothing*', *Philosophical Studies*, 131, 227-249.
- Kim, J. (1992) 'Multiple Realizability and the Metaphysics of Reduction', *Philosophy and Phenomenological Research*, 52, 1-26.
- _____. (2005) *Physicalism, or Something Near Enough*, Princeton: Princeton University Press.
- Macdonald, C. and Macdonald, G. (2006) 'The Metaphysics of Mental Causation', *The Journal of Philosophy*, 103, 539-576.
- Noordhof, P. (1997) 'Making the Change: the Functionalist's Way', *British Journal of the Philosophy of Science*, 48, 233-250.
- Papineau, D. (2002) *Thinking about Consciousness*, Oxford: Oxford University Press.

- Pereboom, D. (2002) 'Robust Non-Reductive Materialism', *The Journal of Philosophy*, 99, 499-531.
- Putnam, H. 1960 (1975) 'Minds and Machines' IN *Mind, Language, and Reality*, Cambridge: Cambridge University Press.
- _____. (1967) 'Psychological Predicates' IN W.H. Capitan and D.D. Merrill (eds.), *Art, Mind, and Religion*, Pittsburgh: Pittsburgh University Press.
- Restrepo, R. (2009) 'Russell's Structuralism and the Supposed Death of Computational Cognitive Science', *Minds and Machines*, 19, 181-197.
- _____. (2012a) 'Two Myths of Psychophysical Reductionism'. *Open Journal of Philosophy*, 2, in print. A version can be found at <http://philpapers.org/archive/RESTMO-3>. Accessed Feb. 4, 2012.
- _____. (2012b) 'Thinking about Physicalism'. *Open Journal of Philosophy*, 2, in print. A version can be found at <http://philpapers.org/archive/RESTAP>. Accessed Feb. 4, 2012.
- Rosenzweig, M., Breedlove, M. and Leiman, A. (2002) *Biological Psychology*, Sunderland: Sinauer Associates Inc.
- Searle, J. (1980) 'Minds, Brains and Programs', *Behavioral and Brain Sciences*, 3, 417-57.
- _____. (1992) *The Rediscovery of the Mind*, Cambridge: MIT Press.
- Shapiro, L. (2000) 'Multiple Realizations', *The Journal of Philosophy*, 97, 635-54.
- Shoemaker, S. (1984) *Identity, Cause and Mind*, Cambridge: Cambridge University Press.
- Stafford, K. (2007) 'Pain in Dogs: Its Recognition and Alleviation', IN *The Welfare of Dogs*, Dordrecht: Springer.
- Turing, A. 1950 (2004) 'Computing Machinery and Intelligence', IN J. Copeland (Ed.), *The Essential Turing*, Oxford: Oxford University Press.
- Turing, A., Braithwaite, R., Jefferson, G., and Newman, M. 1952 (2004) 'Can Automatic Calculating Machines Be Said to Think?' IN J. Copeland (Ed.), *The Essential Turing*, Oxford: Oxford University Press.
- Whittle, A. (2007) 'The Co-Instantiation Thesis', *Australasian Journal of Philosophy*, 85, 61-79, page numbers used from URL = http://personalpages.manchester.ac.uk/staff/ann.whittle/documents/coinstantiation_thesis.pdf.