

Multiple realization and the commensurability of taxonomies

John Zerilli¹

Received: 23 May 2017 / Accepted: 19 October 2017 / Published online: 30 October 2017
© Springer Science+Business Media B.V. 2017

Abstract The past two decades have witnessed a revival of interest in multiple realization and multiply realized kinds. Bechtel and Mundale’s (Philos Sci 66(2):175–207, 1999) illuminating discussion of the subject must no doubt be credited with having generated much of this renewed interest. Among other virtues, their paper expresses what seems to be an important insight about multiple realization: that unless we keep a consistent grain across realized and realizing kinds, claims alleging the multiple realization of psychological kinds are vulnerable to refutation. In this paper I argue that, intuitions notwithstanding, the terms of their recommendation make it impossible to follow, while also misleadingly insinuating that its application virtually guarantees mind-brain identity. Instead of a matching of grains, what multiple realization really requires is a principled method for adjudicating upon differences between tokens. Shapiro’s (J Philos 97(12):635–654, 2000) work on multiple realization can be understood as an attempt to adumbrate just such a method. While his “causal relevance” criterion can easily be mistaken for Bechtel and Mundale’s grain requirement, my analysis reveals exactly where and why these two tests diverge.

Keywords Multiple realization · Bechtel and Mundale · Neuroscience · Autonomy of psychology · Intertheoretic reduction

✉ John Zerilli
john.zerilli@gmail.com; john.zerilli@anu.edu.au

¹ School of Philosophy Coombs Building 9, Australian National University, Canberra, ACT 0200, Australia

1 Introduction

The multiple realization (“MR”) hypothesis asserts, at its baldest, that the same psychological state may be realized in neurologically distinct substrates (Polger 2009). Hilary Putnam’s (1967) ingenious suggestion that pain is likely to be a multiply realized kind (“MR kind”) rather neatly captures the thought—while presumably both mammals and molluscs experience pain, they’re likely to instantiate it in neurological systems of a very different sort. MR was played against a popular philosophical theory of mind in the 1960s which attempted to identify mental states with neural states. Since MR implies a many-to-one mapping from neural states to mental states, if it is in fact true that mental states are multiply realized, it follows that no clear identity relation can hold between them. As Bechtel and Mundale (1999, p. 176) frame the issue, “[o]ne corollary of this rejection of the identity thesis is the contention that information about the brain is of little or no relevance to understanding psychological processes”. When the MR hypothesis first came to prominence, its critics by and large accepted it as empirically correct, and merely denied its touted antireductionist implications. In recent years the debate has struck a new note, with many philosophers calling the empirical hypothesis itself into question. Bechtel and Mundale’s (1999) influential paper, followed quickly at the heels by Shapiro’s (2000) penetrating analysis of functions, perhaps did most to reignite the old controversy and drag MR back into the philosophical limelight. Bechtel and Mundale express what seems to be an important insight about multiple realization: that unless we keep a consistent grain across realized and realizing kinds, claims alleging the multiple realization of psychological kinds are vulnerable to refutation. In this paper I argue that, intuitions notwithstanding, the terms of their recommendation make it impossible to follow, while also misleadingly insinuating that its application virtually guarantees mind-brain identity. Instead of a matching of grains, what MR really requires is a principled method for adjudicating upon differences between tokens. Shapiro’s (2000) work on MR can be understood as an attempt to adumbrate just such a method. While his “causal relevance” criterion can easily be mistaken for Bechtel and Mundale’s grain requirement, my analysis reveals exactly where and why these two tests diverge.

2 Bechtel and Mundale’s grain requirement

Bechtel and Mundale appeal to “neurobiological and cognitive neuroscience practice” in the hope of showing how claims that psychological states are multiply realized are unjustified. Essentially, theirs is an argument from success: cognitive neuroscience’s method assumes MR is false, and the success of that method is evidence that MR *is* false. They argue that it is “precisely on the basis of working assumptions about commonalities in brains across individuals and species that neurobiologists and cognitive neuroscientists have discovered clues to the information processing being performed” (1999, p. 177).

Bechtel and Mundale examine both the “neuroanatomical and neurophysiological practice of carving up the brain”. What they believe this examination reveals is, firstly, that the principle of psychological function plays an essential role in both dis-

ciplines, and secondly, that “the cartographic project itself is frequently carried out comparatively—across species” (1999, p. 177), the opposite of what one would expect if MR were “a serious option”. It is the very similarity (or homology) of brain structures which permits generalization across species; and similarity in the functional characterization of homologous brain regions across species only makes sense if the claims of MR are either false or greatly exaggerated. For instance, “[e]ven with the advent of neuroimaging, permitting localization of processing areas in humans, research on brain visual areas remains fundamentally dependent on monkey research...” (1999, p. 195). “The clear assumption is that the neural organization in the macaque will provide a defeasible guide to the human brain” (1999, p. 183). Brodmann’s famous brain maps were based upon comparisons of altogether 55 species and 11 orders of mammals. If MR were true, “one would not expect results based on comparative neuroanatomical and neurophysiological studies to be particularly useful in developing functional accounts of human psychological processing” (1999, p. 178). They also argue that the ubiquity of brain mapping as a way of decomposing cognitive function points to the implausibility of the MR thesis. The understanding of psychological function is increasingly “being fostered by appeal to the brain and its organization” (1999, p. 191), again, the opposite of what one would expect “[i]f the taxonomies of brain states and psychological states were as independent of each other as the [MR] argument suggests” (1999, pp. 190–91).

In light of such considerations, Bechtel and Mundale (1999, pp. 178–79, pp. 201–204) resort to grains as a way of making sense of what they perceive to be the entrenched, almost unquestioning consensus prevailing around MR. They think that it can be traced to the practice of philosophers appealing to different grain sizes in the taxonomies of psychological and brain states, “using a coarse grain in lumping together psychological states and a fine grain in splitting brain states”. When Putnam went about collecting his various specimens of pain, he ignored the many likely nuances between them. At the same time, he had few compunctions about declaring them different at a neurological level. His contention that pain is likely to be an MR kind can only command our respect if we can be sure that when he was comparing his specimens from a neurological point of view he was careful to apply no less lenient a standard of differentiation than he applied when comparing his specimens from a psychological point of view. Bechtel and Mundale maintain that when “a common grain size is insisted on, as it is in scientific practice, the plausibility of multiple realizability evaporates”. As their examples of neuroanatomical and neurophysiological practice attest, scientists in these fields typically match a coarse-grained conception of psychological states with an equally coarse-grained conception of brain states. Despite the habit of philosophers individuating brain states in accordance with physical and chemical criteria, a habit no doubt originating with Putnam, this is not how neuroscientists characterize them. The notion of a brain state is “a philosopher’s fiction” (1999, p. 177) given that the notion neuroscientists actually employ is much less fine-grained, namely “activity in the same brain part or conglomerate of parts”.

A not unrelated factor is that the MR hypothesis often gets presented in a “contextual vacuum”. The choice of grain is always determined by context, with “different contexts for constructing taxonomies” resulting in “different grain sizes for both psychology and neuroscience”. The development of evolutionary perspectives, for instance, in

which the researcher necessarily adopts a coarse grain, contrasts with the much finer grain that will be appropriate when assessing differences among conspecifics:

One can adopt either a coarse or a fine grain, but as long as one uses a comparable grain on both the brain and mind side, the mapping between them will be correspondingly systematic. For example, one can adopt a relatively coarse grain, equating psychological states over different individuals or across species. If one employs the same grain, though, one will equate activity in brain areas across species, and one-to-one mapping is preserved (though perhaps further taxonomic refinement and/or delineation may be required). Conversely, one can adopt a very fine grain, and differentiate psychological states between individuals, or even in the same individual over time. If one similarly adopts a fine grain in analyzing the brain, then one is likely to map the psychological differences onto brain differences, and brain differences onto psychological differences (1999, p. 202).

There are various aspects of Bechtel and Mundale's overall analysis that could elicit a critical response. Various philosophers have supplied detailed criticisms of them. Ken Aizawa (2009, pp. 500–503), for instance, detects a tacit claim in their paper to the effect that unique cross-species localization of functions in the brain entails their unique realization. This is thought to be a non sequitur. It is true that, strictly speaking, what much of their paper succeeds in showing is the unlikelihood of "multiple *localization*", but two things can be said in response. Firstly, the criticism to some extent misses the point of their analysis. Bechtel and Mundale have *deliberately* opted for functional localization, i.e. brain activity in the same parts or conglomerate of parts across species, as the relevant standard by which to judge the sameness or difference of brain states, and they have done so in deference to cognitive neuroscience practice. Localization is for them the appropriate standard to adopt because it is at the right grain of description (see next section). Secondly, it is not actually easy to police the distinction between localization and realization, for neural localization is an important dimension of neural organization. Aizawa complains that Bechtel and Mundale do not provide independent evidence in support of any such proposition, but, all in all, it does not seem to be a particularly tendentious one.¹ It is true that, for all we know, functions could be localized in the same region in closely related species, and yet be organized very differently. But surely we would need some reason to think that this is likely.

Another criticism fastens upon the modality of the MR hypothesis. Bechtel and Mundale shape their discussion of MR around the evidence favoring it among existing biological systems, and hence construe the issue as an empirical one. Many proponents of MR, however, take it to be a claim about metaphysical possibility. Thus a particular version of the argument asserts that even if MR does not hold among existing biological systems, the *possibility* that cognitive states might be shared by built artifacts or alien life forms having very different physical structures *in itself* establishes the salience of the MR thesis. This version of the argument led some philosophers of AI to embrace the

¹ It only appears to be tendentious when a certain paradigm of realization and MR, the so-called "dimensioned" view, has one under its sway (see Gillett 2003).

further metaphysical claim that “mental processes are the operations themselves, and are not identified with whatever biological or other substances realize them” (Bechtel and Mundale 1999, p. 176). Following Bechtel and Mundale’s lead, and the contours of the recent debate surrounding MR, I shall not address this version of the argument or the metaphysical claim which it inspired here. The possibility that artifacts could have mental states is just the possibility that the identity theory is wrong, which is precisely the claim in dispute (see Polger 2009, p. 459).

A more serious criticism of their argument is that it speaks only to species that are homologous—it might demonstrate that the ubiquity of MR is questionable so long as we restrict our gaze to primates and rodents (i.e. mammals generally, related by common descent); but surely it does not succeed in showing that octopuses and humans realize their psychologies in anything like the same way (Kim 2002; Shapiro 2008). This is a powerful criticism, and one to which no satisfactory response has yet been given [although see Couch (2004), discussed below].

Still what no philosopher has so far called into question is Bechtel and Mundale’s use of “grains”. On the contrary, it is often implicitly—and sometimes even explicitly—endorsed by friend and foe alike (Couch 2004; Polger 2009; P. Godfrey-Smith, personal communication; C. Klein, personal communication; see also tacit approval in Aizawa and Gillett 2009, p. 573). Polger (2009) explains the motivation for the grain requirement in an illuminating way. Neuroplasticity has in recent times been thought to provide compelling evidence for the MR of mental states. He concludes that “contrary to philosophical consensus, the identity theory does not blatantly fly in the face of what is known about the correlations between psychological and neural processing” (2009, p. 470). The grains argument figures prominently in his reasoning. As he points out, it might be tempting to regard a phenomenon like cortical map plasticity—where different brain regions subserve the same function at different times in an individual’s history, say, after brain injury or trauma—as an existence proof of MR. But not if the point about grains is taken to heart. It all comes down to what we mean by “different brain regions” subserving “the same function”. Consider that recovered functions are frequently suboptimal. Genuine MR would indeed require the same psychological state to be underwritten by different neurological states; but suboptimality is evidence of difference underlying difference, not difference underlying sameness, as MR requires:

It’s true that this kind of representational plasticity involves the “same” function being mediated by “different” cortical areas. But here one faces the challenge leveled by Bechtel and Mundale’s charge that defenses of [MR] employ a mismatch in the granularity of psychological and neuroscientific kinds. If we individuate psychological processes quite coarsely—by gross function, say—then we can say that functions or psychological states are of the same kind through plastic change over time. And if we individuate neuroscientific kinds quite finely—by precise cortical location, or particular neurons—then we can say that cortical map plasticity involves different neuronal kinds. But this is clearly a mug’s game. What we want to know is not whether there is some way or other of counting mental states and brain states that can be used to distinguish them—no doubt there are many. The question is whether the sciences of psychology and neuro-

science give us any way of *registering the two taxonomic systems* (1999, p. 467, my emphasis).

3 Problems with the grain requirement: imprecise, impracticable, and misleading

But now the question is this: what, precisely, can it mean to use a “comparable” grain, or to keep a grain size “constant”, across both psychological and neurophysiological taxonomies? Polger’s motivation makes a lot of sense, to be sure, but talk of “registering” taxonomies (as of *aligning* classificatory regimes, or rendering distinct scientific descriptions *commensurable*, or however else one might care to put it) doesn’t shed any light on how the desideratum for consistent grains can actually be met. Since it is intended to serve in part as a methodological prescription, it’s important to know what to make of this requirement—metaphors won’t help us here. How, in *concrete* terms, is an investigator meant to satisfy such a condition as *this* on their research?

Let me begin by using pain and hunger as the MR kinds in question, both because they are states about which Bechtel and Mundale have something to say and because they have long been staples of the MR literature. In Sect. 2 I mentioned that when Putnam went about collecting his various specimens of pain, he ignored the many likely subtle differences between them—even as he had few scruples about declaring them different at a neurological level. Bechtel and Mundale’s complaint against Putnam therefore seems to be that when he says that pain (or hunger) is likely to be an MR kind, we can only go along with him if we can be sure that, when he was comparing his specimens from a neurological point of view, he was careful to apply no less lenient a standard of differentiation than he applied when comparing his specimens from a psychological point of view. But because in their view he didn’t do this, he opened himself to their challenge that he was unduly permissive in his classification of psychological states. This explains why Bechtel and Mundale chide that:

A human’s psychological state and that of an octopus might well be counted as the same insofar as they are associated with some general feature (such as food-seeking behavior, in the case of hunger). But with respect to other considerations, a human psychological state may be considered different from that of an octopus, even if we limit the scope to mere behavior. Food-seeking behavior for the octopus is different from food-seeking behavior in the human if one is concerned about such things as how one seeks the food, what foods are sought, under what conditions, etc. (1999, p. 203).

This seems intuitive and sensible. In the present context, therefore, the word “grain” seems to mean something like, “the basis for comparison”, or more precisely, “the respect under which we seek to compare a set of tokens”; and “same” grain means something like “equally permissive or stringent (bases of comparison)”. I think this takes us a step closer to what Bechtel and Mundale mean when they exhort us to employ the same grain across realized and realizing kinds. MR requires a comparison of a set of tokens from both a psychological and a neurological point of view; and we

must see to it (somehow) that these cross-disciplinary reference points are on talking terms. So far so good.

But now what does *this* mean? How can the bases for comparing a set of tokens (in the present context, two brains) be “on talking terms” or “commensurable” or “comparable” when the bases are delivered by two distinct disciplines (in the present context, psychology and neuroscience)? I’m going to illustrate what I take to be the most natural way in which such cross-disciplinary reference points can be brought into alignment by using an example where the two disciplines are more clearly distinct than psychology and neuroscience. In this way the nature of the problem will be brought into much sharper relief than if I were to stick with psychology and neuroscience. Suppose you have two tokens of fruit. The science of *botany* (say) could deliver descriptions under which the two are classified the same (e.g. from the point of view of *species*), but also descriptions under which they come out different (e.g. from the point of view of *varieties*). The first description could be said to apply a coarser grain than the second. Now imagine *economics* coming into the picture. The science of economics can likewise deliver descriptions under which both tokens are classified the same (e.g. both are forms of tradable fresh produce) or different (e.g. one, being typically the crunchier and sweeter variety, has a lower elasticity of demand than the other). Once again, the first description could be said to apply a coarser grain than the second. Perhaps, then, we could take it that botany and economics deliver descriptions at the same grain of analysis when their judgments of sameness or difference cohere in a given case. In the example, botanical descriptions via species classification would be furnished at the same grain as economic descriptions via commodity classification, so that species descriptions in botany are “at the same grain” as commodity descriptions in economics. By the same logic, *variety* descriptions in botany would be comparable to *elasticity* descriptions in economics.

This construal is fairly reasonable, I think, and it’s a serious, charitable attempt to make sense of Bechtel and Mundale’s recommendation. Still, there’s a big problem here. If this is all that “maintain a comparable grain” amounts to, it really does beg the question, for this is simply type-type identity by fiat. *Of course* such a recommendation will ensure that the mapping between psychology and neuroscience will be “systematic” (to use Bechtel and Mundale’s term), because on this account yielding concordant judgments of similarity or difference across taxonomies is simply what it *means* to apply the same grain. So *this* version of the grain requirement makes type-type identity a fait accompli, effectively obliterating all MR kinds from the natural order.

It’s just as well that I don’t think this is quite what Bechtel and Mundale had in mind when they made their move to grains. And yet they do say: “One can adopt either a coarse or a fine grain, but as long as one uses a comparable grain on both the brain and mind side, the mapping between them will be correspondingly systematic” (note that—it *will* be!). This sounds like someone with the utmost confidence in the grain requirement, which is of course what one *would* have if one thought grains could be legitimately matched in just this way. In the same passage they assert that, in the context of a researcher invoking a relatively coarse grain to equate psychological states across different individuals or species, “[i]f one employs the same grain...one will equate activity in brain areas across species, and one-to-one mapping is preserved...” (my

emphasis). “One will equate...”? How can they be so sure? My guess is that, while they do have something important to tell us about MR, a beguiling metaphor has led them to suppose that MR is easier to refute than it actually is (I’ll support this contention with a few examples in a moment.).

If not by this means, then, how else might two taxonomic descriptions be rendered commensurable? Perhaps when I suggested a moment ago that two sciences deliver descriptions at the same grain of analysis when their judgments of sameness or difference cohere in a given case I should have added “likely (to cohere)”. We would then have:

Any two sciences deliver descriptions at the same grain of analysis when their judgments of sameness or difference are *likely* to cohere in a given case.

This would at least solve the problem of inevitability, and makes sense of Bechtel and Mundale’s more cautious choice of words near the same passage I cited earlier, where they say, in the context of a researcher invoking a fine grain to differentiate psychological states, “[i]f one similarly adopts a fine grain in analyzing the brain, then one is *likely* to map the psychological differences onto brain differences, and brain differences onto psychological differences” (1999, p. 202, my emphasis). “Is likely to” is better than “will”. Still, this more cautious formulation doesn’t immunize them from the serious errors I point out below. It appears that the distinction between likelihood and inevitability must be handled very carefully in this debate, as is shown by Bechtel and Mundale’s following remarks on computer programs:

One often speaks of running the same program on different computer architectures, thereby generating a case of multiple realization. But in fact one has surreptitiously adopted a coarse grain in characterizing the program. *If one looks carefully enough* at performance measures, one will generally be able to identify differences in the way the program runs on different computers” (1999, pp. 202–203, my emphasis).

How carefully need we look before we can stop looking? Until the mapping is one-to-one between software and hardware? What if we don’t find a one-to-one mapping, having already looked more carefully at the software, but *could* have a one-to-one mapping if only we look yet a little more carefully? Do we stop, or keep going? How do we know? The road from a *likely* one-to-one mapping to an *inevitable* one-to-one mapping is here at its most precarious.

But let’s try to remedy this if we can. Bechtel and Mundale are obviously after some feature of taxonomic description that can, as it were, latch onto a corresponding feature of another taxonomic description. Perhaps this makes most sense if we replace the idea of a “grain” with the idea of a “level”, and more specifically a “level of description or abstraction”. A software program may be described at different levels (machine language, assembly language, compiler language, etc.), as indeed may its underlying hardware (microphysical events, transistors, chips/circuits, etc.). This at once provides a much clearer picture of how we might go about scoring off levels from one discipline against another (and thereby achieve that most elusive and sought-after match between grains). Higher levels of software will correspond to higher levels of hardware—and it need not be inevitable that two programs with the same compiler language will be the

same at the circuit level, even though the compiler language occupies (let's say) the corresponding rank within the software hierarchy that the circuit level does within the hardware one. Translating this thought into the present context, different *psychological* levels must be understood as tallying up with different *neurological* levels. When the levels of description correspond, the descriptions are at the “same grain”.

This does seem more promising, I agree, but I suspect that it is beset with more problems in turn. For example, can we really lay out these two sciences (psychology and neuroscience—indeed *any* two sciences!) on the Procrustean bed of “comparable levels”? And more worryingly, who gets to say whether a level within the hierarchy of one discipline really tallies with a level drawn from the hierarchy of another discipline? I fear that replacing “grains” talk with “levels” talk serves only to demystify temporarily. The problem has been merely postponed.

Of course matters aren't much helped by the reasonable suspicion that MR is the result of pairing *inconsistent* grains. For what is neuroscience if not a fine-grained description of psychology, and psychology if not a coarse-grained description of neuroscience? It is surely plausible that the neural and psychological sciences line up in something like this way, given that talk about the mind is really talk about the brain from a somewhat more abstract point of view.

What Bechtel and Mundale are ultimately trying to convey through their discussion of grains is the thought that claims of MR cannot be advanced willy-nilly—that there is an objective and standard way to go about verifying the existence of MR kinds and arbitrating disputes involving them. For the reasons just canvassed, however, it strikes me that talk of grains doesn't serve their purposes at all well. In fact they would have been nearer the mark had they said that what MR requires is some sort of principled *mismatching* of grains.

So far I've tried to indicate in what respects Bechtel and Mundale's grain requirement is imprecise and impracticable. Before I can show that the grains strategy is also misleading, and actually often gets things wrong, I need to set it against an account which demonstrably gets things right, indeed an account which even its detractors concede gets something very important right (e.g. see Gillett 2003, pp. 591, 592, 596, 597, 599, 600). Shapiro (2000) expresses with enviable lucidity what I think is the crucial insight towards which Bechtel and Mundale were uneasily groping.² Interestingly, some philosophers—e.g. Polger (2009)—write as if the grain requirement and Shapiro's own formula for MR were effectively interchangeable. This is a mistake: the two approaches deliver different judgments in nontrivial cases (as I'll illustrate in a moment).

² Shapiro ultimately uses his analysis to draw out a dilemma confronting the MR advocate. If a given functional kind is not multiply realized, the traditional argument for its autonomy and irreducibility falls away. If on the other hand the kind *is* multiply realized, the kind will not be a proper scientific kind, i.e. of the sort that can enter into laws. “Brittleness” might be a multiply realized property, but glass, steel and biscuits are each brittle in their own way: there can be no general science of brittle things. Ranging over such genuinely diverse physical realizations means the kind will not enter into laws (i.e. exhibit lawlike or projectable properties), except for those which are true analytically—such as *all mousetraps catch mice*, and *all eyes see*—and this in turn “undercut[s] the traditional motivation for admitting functional kinds into the ontologies of the special sciences” (Shapiro 2000, p. 637). In this paper I will not be addressing this issue, but see Couch (2009a, pp. 262–264) for some criticisms. In more recent work Shapiro himself seems to have backed away from this position (Polger and Shapiro 2016).

As Shapiro reminds us:

Before it is possible to evaluate the force of [the MR thesis] in arguments against reductionism, we must be in a position to say with assurance what the satisfaction conditions for [the MR thesis] actually are (2000, p. 636).

For him, “[t]he general lesson is this. Showing that a kind is multiply realizable, or that two realizations of a kind are in fact distinct, requires some work” (2000, p. 645). Furthermore, “[t]o establish [the MR thesis], one must show that the differences among purported realizations are causally relevant differences” (2000, p. 646). Shapiro’s concerns revolve around what motivates ascriptions of difference, and therefore sameness. The issue is important because the classic intuition pump that asks us to conceive a mind in which every neuron has been replaced by a silicon chip depends on our ascription of an interesting difference between neurons and silicon chips, apparently even where silicon chips can be made that contribute to psychological capacity by one and the same process of electrical transmission. His answer too, like Bechtel and Mundale’s, depends ultimately on context—in particular, the context set by the very inquiry into MR itself.

Shapiro (2000, pp. 643–644) argues that “the things for which [the MR thesis] has a chance of being true” are all “defined by reference to their purpose or capacity or contribution to some end”. This is the reason why carburetors, mousetraps, computers and minds are standard fare in the literature of MR. They are defined “in virtue of what they do”, unlike, say, water, which is typically defined by what it is, i.e. its constitution or molecular structure, and accordingly *not* an MR kind. Genuine MR requires that there be “*different ways to bring about the function that defines the kind*”. Truly distinct (indeed *multiple*) realizations are those that “differ in causally relevant properties—in properties that make a difference to how [the realizations] contribute to the capacity under investigation”. Two corkscrews differing only in color are not distinct realizations of a corkscrew, because color “makes no difference to their performance as a corkscrew”. Similarly, the difference between steel and aluminium is not enough to make two corkscrews that are alike in all other respects two different realizations of a corkscrew “because, relative to the properties that make them suitable for removing corks, they are identical”. In this instance, differences of composition can be “screened off”. Naturally there may be cases where differences of composition *will* be causally relevant (and it turns out that this will be important to the broader point I make below about where the grains strategy goes wrong). Perhaps rigidity is the allegedly MR kind in question. In that event, compositional differences will necessarily speak to how aluminium and steel achieve this disposition. The crucial thing to note here is that in a nontrivial sense MR *is* the context, because the very inquiry which seeks to determine whether a kind is multiply realizable *itself imposes an implicit contextual constraint*: MR makes *function* the primary comparative consideration (i.e. the specific point of view from which we will compare a set of tokens in the first instance)—not phenomenology, not behavioral ecology, or anything else for that matter. MR is after all a thesis about sameness and difference (Polger 2009); and any two particulars will both differ and resemble infinitely (Bechtel and Mundale 1999, p. 203). Whether two particulars are “the same” or “different” therefore depends on the aspect from which we choose to compare them. It just so happens that to ask whether a kind is an MR



Fig. 1 A waiter's corkscrew (a) and a winged corkscrew (b). Each contributes to the capacity of cork-removal in different ways

kind is already to have decided this preliminary question: what we want to know here is whether two tokens that serve a particular function do so in the same way. Explanatory considerations may of course fine-tune the *sort* of function that captures our attention (cork-removal, rigidity, vision, camera vision, etc.). But function here is our key preoccupation, and having settled on a specific function which a set of tokens can be said to perform, the all-important question on Shapiro's analysis is *how* the two tokens bring that function about. Each case must be judged on its own merits. Thus unlike the two corkscrews identical in all respects save color, which do not count as distinct realizations, waiter's corkscrews and winged corkscrews are enabled to perform the same task in virtue of *different* causally relevant properties, and therefore *do* count as genuinely distinct realizations of a corkscrew, one based on the principle of simple leverage, the other relying on a rack and pinions (Fig. 1).

Notice that to the extent Shapiro's causal relevance criterion envisages certain realizing properties being "screened off" from consideration in the course of inquiry, there is a sense in which the taxonomies of realized and realizing kinds may be said to be "commensurable" or "registrable" (no doubt explaining why some philosophers have simply confused commensurability with causal relevance). Thus when comparing the cork-removing properties of two waiter's corkscrews, compositional differences will not feature in the realizing taxonomy (if we accept Shapiro's characterization of the problem). So we have *cork-removal*, which features in what we may regard as a coarse-grained taxonomy, realized by two objects described by a "science" of cork-removal in which microstructural variations do not matter, hence which might also be regarded as a coarse-grained taxonomy. If on the other hand we were comparing the same corkscrews for rigidity, where one was made of steel and the other of aluminium, compositional differences *would* feature in the realizing taxonomy. Here we would have *rigidity*, which features in what we could well regard as a more fine-grained taxonomy than that encompassing cork-removal, realized by two objects described by a science in which microstructural variations really *do* matter (namely metallurgy), and which might also be regarded as a fine-grained taxonomy, at least more fine-grained than the fictitious science of cork-removal. But my point is this: commensurability nowhere appears as an independent criterion of validity in Shapiro's account of MR,

for it is an artifact of the causal relevance criterion, not a self-standing principle. Taxonomic commensurability is in fact an *implicit* requirement of the causal relevance criterion in the sense that it's taken care of once the proper question is posed. As an explicit constraint it is a will-o'-the-wisp.

Armed with this analysis, let's examine how Bechtel and Mundale attempt to refute the status of hunger as an MR kind. Putnam (1967) had compared hunger across species as diverse as humans and octopuses to illustrate the likelihood that some psychological predicates are multiply realizable. On the basis of their grains critique, however, Bechtel and Mundale suggest that hunger will not do the work Putnam had cut out for it; for "at anything less than a very abstract level", hunger is different in octopuses and humans (1999, p. 202). The thought is that a finer individuation of hunger refutes the existence of a *single* psychological kind, hunger, which can be said to cross-classify humans and octopuses. Thus they essay to challenge the cognitive uniformity which MR requires at the level of psychology.

Perhaps we might first note that when identifying a *single* psychological state to establish the necessary conditions for MR, nothing Bechtel and Mundale say actually *precludes* the choice to go abstract. If context is what fixes the choice of grain (as they are surely right to point out), who's to say that context couldn't fix the sort of grain that makes hunger relevant in an abstract sense? It may be tempting to think that a more detailed description of something is somehow more real. But there is of course nothing intrinsically more or less real about a chosen schema relative to others that might have been chosen (there is no reason to suspect, for instance, that a determinate has any more reality than a determinable). This consideration applies with no less force when we decide to let the sciences dictate what the relevant kinds will be, since different sciences can take both more and less abstract objects within the ken of their inquiries (Craver 2007). In fact one way of reading Bechtel and Mundale's paper (particularly Sects. 2–4) is precisely as suggesting that kind individuation *should* be fixed by the sciences. This point comes through pretty clearly when they discuss the individuation of neural states. In that specific context their paper provides excellent and detailed examples of how scientific taxonomies can be invoked in debates over MR. The philosophical notion of a brain state, they complain, is a "philosopher's fiction", because brain scientists by and large do not individuate brain states on the basis of physical and chemical criteria; it is enough for their purposes to employ the much coarser notion of "activity in the same brain part or conglomerate of parts". This is a clear demonstration of how regard for the sciences may frequently mean opting for less precise/more abstract descriptions of phenomena.³

Notice, incidentally, that when it comes to the individuation of psychological states (as distinct from neural states), Bechtel and Mundale's advice about heeding the sciences seems to get lost. The part of their paper where they do most to explicate a fine-graining procedure for psychological kinds (1999, pp. 201–204) nowhere exemplifies how the cognitive sciences can be invoked to guide individuation of psychological

³ Bechtel and Mundale's test at this point can be seen as a bowdlerized version of Shapiro's causal relevance criterion, which screens off causally irrelevant details when comparing realizing kinds (much as Bechtel and Mundale's "coarse-graining" procedure does).

kinds.⁴ Instead they rely on “kind splitting” (Polger and Shapiro 2016), where a given (higher level) kind (e.g. vision, as instanced in both humans and molluscs) is described in increasingly greater detail until the illusion that two tokens are the same is finally dispelled (allegedly demonstrating, for example, that human eyes and mollusc eyes are in fact different psychological kinds, not instances of the same overarching psychological kind “eye”). But kind splitting is gratuitous if it ignores the explanatory salience of the higher level kind (see below), and in none of their examples (see 1999, pp. 201–204) are the cognitive sciences actually consulted to check for this salience. This is precisely the point at which the cognitive sciences would be most instructive: we need to be sure whether a split is justified, and thus whether a split would really undermine the integrity of a given higher level kind. So at least in the context of individuating psychological kinds, it would appear that Bechtel and Mundale have failed to heed their own advice.

And yet there is a deeper problem with Bechtel and Mundale’s deployment of the grains strategy here. To repeat their complaint: “at anything less than a very abstract level”, hunger is different in octopuses and humans. But now why should *this* be relevant? Who would deny it? They themselves seem to be oblivious to the context which the very inquiry into MR makes paramount. They are not right to allege, as they do, that “the assertion that what we broadly call ‘hunger’ is the same psychological state when instanced in humans and octopi has apparently been widely and easily accepted without specifying the context for judging sameness” (1999, p. 203). The reason why hunger, pain, vision and so on were all taken for granted—assumed to be uniform at the cognitive level—is because MR made *function* the point of view from which tokens were to be compared. As Shapiro reminds us, “the things for which [the MR thesis] has a chance of being true” are all “defined by reference to their purpose or capacity or contribution to some end”. It was understood that, say in the case of pain, regardless of phenomenal, ecological or behavioral differences between human and octopus pain (I doubt any of which were lost on Putnam), all instances of pain in these creatures had something like *detection and avoidance* in common. This might be to cast pain at “a very abstract level”, but this just happens to be the context which the inquiry into MR itself sets. A similarly abstract feature is what unites all instances of hunger: let’s call it *nutrition-induction*. It is not that decades of philosophers had simply forgotten to specify the point of view from which these psychological predicates were being considered: it is rather that they simply didn’t need to, since all of them had read enough of Putnam and the early functionalists to know what they were about. Phenomenal and other differences that one might care to enumerate between these predicates come a dime a dozen. But the whole point of functionalism was to abjure the inquiry into essences and focus instead on the causal role of a mental state within the life of an organism. Yes, this is to compare tokens from an “abstract level”, but that’s what made functionalism intriguing to begin with. And if Shapiro’s analysis is any guide, it is really the *next* step in the endeavor to verify the existence of an MR kind that is the crucial one. Genuine MR requires that there be “*different* ways to

⁴ The nearest they come is this: “When comparing psychological states across different individuals, psychologists... tend to ignore differences and focus on commonalities” (1999, p. 202). If anything, however, one would have thought that this fact should actually *discourage* kind splitting (see text).

bring about the function that defines the kind”. So the follow-up question concerns *how* the relevant organisms achieve their detection and avoidance function, or nutrition-induction function, or whatever the case may be. It is in fact only by asking this next question that we can appreciate just how badly the grains strategy fares. The attempt to individuate hunger more finely would not refute the multiple realizability of hunger as between humans and octopuses if it could ultimately be shown that humans and octopuses achieve their nutrition-induction capacities in different ways (which seems extremely likely); for then the functional role of hunger would be played by two different realizers—the issue to which the MR inquiry is directed after all—*notwithstanding that the functional role can be more finely described within a more specific frame of reference*. Similarly, the attempt to individuate pain more finely would not refute the multiple realizability of pain as between humans and octopuses if it could ultimately be shown that humans and octopuses achieve their detection and avoidance capacities in different ways (which again seems extremely likely). The sort of type reduction that cuts its teeth by fine-graining psychological categories (i.e. kind splitting) is almost always *trivially* possible, since at the limit everything is a special case; but the move won’t succeed in refuting the existence of an MR kind at a higher level if the higher level functional kind can be brought about in causally distinct ways. Of course there are occasions when kind splitting may be mandated by the sciences, because the higher level kind has no taxonomic relevance. This makes sense if kind individuation should be fixed by the sciences (Couch 2009a, b). In such cases the split really would undermine MR at the higher level. But it’s important to emphasize that these cases require careful consideration and may turn out to be rare in any event (see Polger and Shapiro 2016, pp. 103, 104–105, 110–111; Couch 2009a, p. 267; 2009b, p. 514).

So we see that the grains strategy, to the extent that it involves fine-graining psychological states in order to undermine the cognitive uniformity required by MR, sets itself a very easy job indeed, and mischaracterizes the nature of MR by its neglect of function. Moreover Shapiro’s causal relevance criterion—which honors the core concerns motivating Bechtel and Mundale’s resort to grains—does *not* demonstrate that hunger (or pain) is type-reducible.

A good illustration of the grains strategy in action is provided by Couch’s (2004) attempt to refute the claim that the human eye and the octopus eye are distinct realizations of the kind eye. Conceding differences at a neurobiological level, the strategy again involves challenging the alleged uniformity at the cognitive level. As he explains, “[e]stablishing [MR] requires showing that...the physical state types in question are distinct [and] that the relevant functional properties are type identical. Claims about [MR] can be challenged at either step” (2004, p. 202). Reminding us that psychological states “are often only superficially similar”, and that “at a detailed level the neural differences make for functional differences” (2004, p. 203), he states:

Psychologists sometimes talk about humans and species like octopi sharing the same psychological states. However, they also recognize that there are important differences involved depending on how finely one identifies the relevant features...Establishing multiple realization requires showing that the same psychological state has diverse realizations. But we can always disagree with the

functional taxonomy, and claim there are psychological differences at another level of description (2004, p. 203).

Thus he relates that while the two types of eyes have similar structure in certain respects, both consisting of a spherical shell, lens and retina, they use different kinds of visual pigments in their photoreceptors, as well as having different numbers of them, the octopus having one in contrast to the human eye which has four. They also have different retinas. The human retina, with its rods and cones, allows the eye to focus light by bending the lens, thereby changing its shape. The octopus eye, with rhabdomeres instead of rods and cones, focuses light by moving the lens backwards and forwards within the shell. All these factors show up as differences in output, not just structure. The octopus, having only a single pigment, is colorblind, while its receptor's unique structure allows it to perceive the plane of polarized light. Retinal differences likewise make for functional differences, with very little information processing occurring on the octopus's retina, unlike the case of the human retina. This produces differences in stimuli and reaction times. So the two eyes might be similar, but when described with a suitably fine grain, he contends, they come out type distinct. In the result they are both physically *and* cognitively diverse, and so not genuine examples of MR.

Notice again that, contrary to what is claimed, it has not been demonstrated that type-type identity prevails here after all (on the understanding that the kind camera eye_{human} reduces to *its* distinct neural type, and the kind camera eye_{mollusc} in turn reduces to *its* distinct neural type). If anything what this foray into mollusc visual physiology succeeds in showing is that, relative to the kind camera eye, human camera eyes and octopus camera eyes count as distinct realizations(!), for, assuming Shapiro's causal relevance criterion applies, human camera eyes achieve the function of *camera vision* differently to the way octopus camera eyes achieve this function. Were we to attend to the original inquiry, which concerned whether human eyes and octopus eyes count as distinct realizations of the kind eye, Shapiro's own response, for what it's worth, is clear (2000, pp. 645–646): here we do seem to confront a genuine case of type-type identity, as Putnam himself assumed, because, relative to the function of *vision* (not *camera vision*), both humans and molluscs achieve the function the same way (namely, by camera vision!). Differences that would be relevant at the neural level between humans and molluscs when asking how camera vision is achieved can be conveniently screened off when the question is how vision, as distinct from camera vision, is achieved. Again if pain or hunger were the kind in question, it seems more likely than not that we *would* confront a case of MR (unlike with vision), as we conjectured earlier. Explanatory context dictates the function of interest, and the function is one that we have to assume is common to the tokens in question in order to get the inquiry into MR off the ground. Indeed if Shapiro's analysis is correct, with MR we're always asking how some common function is achieved by different tokens that *do that thing*. Where there is no common function the question of MR cannot so much as arise. The fact that the question *does* arise in all the cases we've considered is a powerful indication that we're dealing with functions which all the relevant tokens actually share. The grains strategy confuses matters by suggesting that in many cases involving putative MR kinds, psychological states can be individuated using a finer

grain of description. But if what I have been saying is right, this is not the proper way to refute a putative case of MR.

That mine is the correct assessment of the situation is not only attested to by Shapiro's analysis of MR, but also by the fact that it avoids the very mug's game Polger sought to eschew by embracing the grains strategy in the first place. If for any putative MR kind I am free to cavil with the choice of your size of grain ("oh, that's far too coarse for psychology", or "now that's really not coarse enough for neuroscience"), how is the resulting game any less of a mug's game than the one we were trapped in at the start? I myself have played a few of these games with philosophers. No one wins. Couch's remarks are telling: "we can always disagree with the functional taxonomy, and claim there are psychological differences at another level of description". So the game goes on.⁵ Yes, it is true that Bechtel and Mundale don't overtly admit to thinking their criterion eradicates MR kinds completely; and yet it would have been interesting to see them pointing out even a single instance of an MR kind they think rightly counts as such. But alas they don't.⁶ All their examples come out as type-type identities after being subjected to their test—even where it is highly likely that Shapiro's more precise test would lead us to conclude otherwise.

4 Conclusion

In sum, I think there's a real problem with the grain requirement. The central difficulty is that in the terms in which it's been put it is largely unworkable, and at best no more than a loose metaphor. For a recommendation intended to serve at least in part as a methodological reform, this is clearly unsatisfactory. I don't deny that Bechtel and Mundale were onto something. But whatever value their insight into MR might have has been obscured by their unfortunate formulation of the issue. Moreover, as I have tried to show, the formulation is unfortunate not *just* because it happens to be unworkable. More worryingly, the argument from grains distorts the truth about MR by encouraging the view that mind-brain identity comes for free once we invoke the "same grain" of description across both realized and realizing kinds. But when the insight to which this locution seems to point is expressed in terms that are intelligible and empirically tractable (namely, Shapiro's causal relevance criterion), mind-brain identity seems anything but a *fait accompli*. Grains talk makes it tempting to think MR is easier to refute than it in fact is. It is certainly true, as Bechtel and Mundale acknowledge, that context fixes the choice of grain (where by "grain" we mean the respect under which we seek to compare a set of tokens); but we are not ipso facto obliged to employ a consistent grain across realized and realizing kinds (since this is just about meaningless as far as a researcher into these matters would be concerned and raises a host of difficulties beside). Rather than matching grains, what MR really

⁵ In subsequent work, Couch (2009a, b) has been more careful in his remarks, and has brought out explicitly the importance of scientific taxonomy in the individuation of both neural and psychological kinds. Still I notice that in these papers he doesn't rely on granularity arguments, and indeed is even mildly critical of them (2009b, p. 267). When grains hold sway, MR stands little chance of receiving a fair hearing.

⁶ By contrast, Couch (2009b, p. 514) puts forward human eyes and pigeon eyes as plausible candidates for MR, and comments: "Accepting that this example is plausible is a revision from an earlier view of mine".

behooves us to do is to apply a principled method for adjudicating upon differences between tokens of a functional kind. Shapiro's work on MR shows us how to approach this important task.

Acknowledgements This paper was presented to the American Philosophy of Science Association on 3 November 2016, in Atlanta, Georgia. I am particularly indebted to Larry Shapiro and Tom Polger for helpful discussion, as well as three anonymous reviewers. Warm thanks also to Kim Sterelny and Paul E. Griffiths.

Funding This research was supported by an Australian Government Research Training Program (RTP) Scholarship.

Compliance with ethical standards

Conflict of interest No potential conflict of interest was reported by the author.

References

- Aizawa, K. (2009). Neuroscience and multiple realization: A reply to Bechtel and Mundale. *Synthese*, 167(3), 493–510.
- Aizawa, K., & Gillett, C. (2009). Levels, individual variation, and massive multiple realization in neurobiology. In J. Bickle (Ed.), *The Oxford handbook of philosophy and neuroscience* (pp. 539–81). New York: Oxford University Press.
- Bechtel, W., & Mundale, J. (1999). Multiple realizability revisited: Linking cognitive and neural states. *Philosophy of Science*, 66(2), 175–207.
- Couch, M. B. (2004). A defense of Bechtel and Mundale. *Philosophy of Science*, 71(2), 198–204.
- Couch, M. B. (2009a). Functional explanation in context. *Philosophy of Science*, 76(2), 253–269.
- Couch, M. B. (2009b). Multiple realization in comparative perspective. *Biology and Philosophy*, 24, 505–519.
- Craver, C. F. (2007). *Explaining the brain*. Oxford: Oxford University Press.
- Gillett, C. (2003). The metaphysics of realization, multiple realizability, and the special sciences. *Journal of Philosophy*, 100(11), 591–603.
- Kim, S. (2002). Testing multiple realizability: A discussion of Bechtel and Mundale. *Philosophy of Science*, 69(4), 606–610.
- Polger, T. W. (2009). Evaluating the evidence for multiple realization. *Synthese*, 167(3), 457–472.
- Polger, T. W., & Shapiro, L. A. (2016). *The multiple realization book*. Oxford: Oxford University Press.
- Putnam, H. (1967). Psychological predicates. In W. Capitan & D. Merrill (Eds.), *Art, mind, and religion* (pp. 37–48). Pittsburgh: University of Pittsburgh Press.
- Shapiro, L. A. (2000). Multiple realizations. *Journal of Philosophy*, 97(12), 635–654.
- Shapiro, L. A. (2008). How to test for multiple realization. *Philosophy of Science*, 75(5), 514–525.