

Against Conventionalism: Boundaries and Realism

TUOMAS E. TAHKO (www.ttahko.net)

ABSTRACT

This paper addresses three conventionalist criticisms against realism. They are all based on the idea that the manner in which we structure, divide or carve up the world is not grounded in any natural, genuine boundaries in the world. The first criticism that will be examined concerns the idea that there are no natural boundaries in the world, the second one focuses on the basis of our classificatory schemes, which the conventionalist claims to be psychological, and the third considers the significance of our particular features in carving up the world, such as physical size and perceptual capabilities. The purpose of the paper is to demonstrate that none of these objections succeed in undermining realism.

1. Introduction

Traditionally, realism attempts to uphold the natural boundaries that we observe in everyday life: apples, cats, mountains and stars are all objects with natural boundaries. It is, however, not easy to state the exact identity-conditions of these things. Mountains for instance do not seem to have a determinate boundary at the ground level, and a closer look will reveal that there is even vagueness concerning Tibbles the cat and her hair.¹ However, the problem that we will focus on is not the problem of vagueness, but rather a problem raised by an extreme conventionalist thesis concerning the mind-independence of the identity-conditions of objects and kinds that a realist would consider to ‘carve the nature at the joints’. The conventionalist thesis is that there are no such mind-independent identity-conditions and that all our efforts to determine natural boundaries are in fact conventional.

Extreme conventionalism is perhaps not a very popular view, but it has its roots in the work of Hilary Putnam and Michael Dummett. More recently, Achille Varzi (1997, Forthcoming) has put forward a novel, extreme conventionalist account, focusing especially on boundaries. More modest versions of conventionalism, i.e. views which take *some* boundaries to be mind-dependent, are fairly common. For instance, John Dupré’s (1993, 2002) species pluralism could be considered a form of modest conventionalism, as it takes the term ‘species’ to contain a variety of different ways to categorize biological organisms, many of which are entirely a matter of convention. It should be noted though that modest conventionalism is not necessarily incompatible with realism, which is why we will not discuss it in any detail. What *is* incompatible with realism is the idea that *all* objects and kinds lack mind-independent natural boundaries and it is entirely a matter of convention as to how we decide to categorize reality. Note that the problem at hand concerns both individual objects and kinds, but

¹ See for instance Tye (1990).

we will focus mainly on individual objects, as that is perhaps the more radical kind of conventionalism.

Accordingly, the central question is, are there any natural boundaries at all? Artificial boundaries are quite familiar to us and we seem to have no trouble in admitting that they are indeed artificial while at the same time we acknowledge their usefulness—borders of countries are the most obvious example. We may also express the artificial/natural boundary distinction in terms of *fiat* and *bona fide* boundaries, or *de dicto* and *de re* boundaries following Smith & Varzi (2000). It is important to recognize here that the *fiat/bona fide* distinction applies equally to the physical boundaries of objects and to the objects themselves: the physical boundary of an apple is, on the face of it, a *bona fide* boundary, and the individual apple is a *bona fide* entity (ibid., 402). However, this is exactly what the extreme conventionalist questions: when we look at the apple closely enough, it is clear that, far from the smooth boundary that it appears to have, we are in fact dealing with a very loose arrangement of molecules, and further, with a swarm of subatomic particles. Familiar topological problems highlight the problem at hand: when we cut the apple in half, which half of the apple is “open” and which one is “closed”?² Moreover, problems concerning composition are all grist in the conventionalist mill: When Tibbles the cat eats some fish, at what point does the fish become a part of Tibbles? Whichever criteria we apply, the extreme conventionalist will claim, is ultimately a matter of *fiat*; Tibbles may continue to exist, but its identity-conditions are not mind-independent (cf. Varzi Forthcoming, 12).

Here is how Varzi sees the situation:

If all boundaries were the product of some cognitive or social *fiat*, if the lines along which

2 See Smith & Varzi (2000: 406-8) for discussion on the open/closed issue.

we “splinter” the world depended entirely on our *cognitive* joints and on the categories that we employ in drawing up our maps, then our knowledge of the world would amount to neither more nor less than knowledge of those maps. The thesis according to which all boundaries—hence all entities—are of the *fiat* sort would take us straight to the brink of precipice, to that extreme form of conventionalism according to which “there are no facts, just interpretations”. On the other hand, to posit the existence of genuine, *bona fide* boundaries—to think that the world comes pre-organized into natural objects and properties—reflects a form of naïve realism that does not seem to stand close scrutiny. (Varzi Forthcoming, 15.)

The dilemma that Varzi puts forward here is the primary concern of this paper. Do we have to choose between extreme conventionalism and naïve realism? Well, perhaps there is a third way out: even if there are no exact natural boundaries, it does not mean that reality does not impose *any* constraints as to what there could be, and this is exactly how modest conventionalism could still be compatible with realism. Indeed, Varzi admits as much and suggests that even if it is implausible that the boundaries which we postulate would map the natural boundaries with perfect accuracy, it is just as implausible that they would be completely inaccurate (*ibid.*).

The extreme conventionalist, however, will not be happy with this compromise, and Varzi himself goes on to suggest that the approximate boundaries that we take to reflect reality’s joints and the manner in which we classify things are not due to some underlying regularity in reality, but are rather just psychological biases, completely decoupled from the structure of reality. Accordingly, the extreme conventionalist thesis is remarkably close to the Dummettian one of reality as an ‘amorphous lump’ (cf. Dummett 1981). However, Varzi at least wishes to distance himself from Goodman’s irrealism and Putnam’s relativism, he seems to be more sympathetic to Dummett’s and

Alan Sidelle's ontologies (Varzi Forthcoming, 18-21). For our purposes, the conventionalist thesis can be summarised as follows:

[CT] The world is a 'dough' and there is a number of ways to cut it. All of these ways to cut are neutral in terms of the structure of reality. How we choose to cut the dough depends on our of psychological biases.

In what follows I will offer a realist response to CT. There are three major points that are all crucial to the conventionalist stance. All of these are familiar from more modest versions of conventionalism, but Varzi combines them to produce the extreme conventionalist stance that we saw above. Firstly, we will consider whether there in fact are any natural, *bona fide* boundaries and suggest that fundamental particles are the best candidate. Secondly, we will attempt to settle why we classify things in the way we do and what our supposed psychological biases concerning our classificatory schemes are in fact grounded in. Thirdly, we will examine the possibility of alternative, alien classificatory schemes and consider whether the manner in which we classify things is unique to us, or at least unique to beings of roughly our size and with similar perceptual devices and rational capabilities. Finally, the results of the discussion will be evaluated.

2. Are there any natural boundaries?

The most obvious way to look for a simple refutation of CT in favour of realism is to look for natural, *bona fide* entities with mind-independent boundaries. Even a single example would do: it would give us a fixed point that would help in defining other boundaries. Where might we start looking for natural boundaries given the problems

that we encounter even with things like biological entities—perfect candidates for *bona fide* entities? It would maybe be best to look at smaller entities, atoms, perhaps. However, even in physics we might have reasons to think that our particular way to carve up the world has more to do with human psychology than with the boundaries in the physical world. For instance, different isotopes of the same element could arguably just as well be classified as different elements, different natural kinds. As Varzi (Forthcoming: 14) suggests, the problem is that there are too many differences in the world rather than too few, and to choose one over the others is to draw a *fiat* line.³

Thus, it seems that even elements are too arbitrary to be genuine *bona fide* entities, as it is not clear why some atoms count as distinct elements and others do not. We do not need to stop here though, for there are of course subatomic particles as well. Indeed, quarks and leptons, fundamental particles, might be the best candidates for *bona fide* entities with well-defined boundaries. Admittedly, the concepts that we use to define these particles are perhaps subject to human contingencies as well and there might be limits to the accuracy of our measurements concerning some of the crucial variables, such as charge, that we use to determine the natural boundaries of fundamental particles. Still, it would surely be too strict a requirement to insist that we must be able to *state* the exact natural boundaries of, say, electrons. It is certainly clear that we can determine these boundaries with an incredibly high accuracy. In other words, even though there may be epistemic constraints in effect here, it does not entail that the boundaries of electrons, for instance, are of a *fiat* sort. No doubt the conventionalist will insist on our ability to determine these boundaries nevertheless, but the problem may simply be that our concepts are fundamentally unable to reach a more accurate level. Consider *pi*: how could we hope to give its *exact* value when even ten thousand digits are not enough? It seems that the exact value of *pi* cannot be given like this at all.

³ See also Elgin (1995) on this.

Consequently, a different approach might be appropriate; perhaps there is indirect evidence for the claim that fundamental particles do have *bona fide* boundaries and are hence *bona fide* entities.⁴

In fact, I believe that there is an abundance of evidence that points towards the conclusion that there must be some *bona fide* entities in the world and that fundamental particles are indeed such entities. I contend that macroscopic objects would not be possible if this was not the case: the very existence of any kind of macroscopic objects speaks in favour of *bona fide* entities. Note that this has nothing to do with whether macroscopic objects themselves are *bona fide* entities; the argument concerns the physical possibility of macroscopic objects. Here is an outline of the argument:

1. There are macroscopic objects.
2. Certain things are physically necessary for the forming of macroscopic objects, e.g. the laws that govern molecular binding.
3. The relevant laws of physics require that fundamental particles have *exact* properties, such as electric charge.
4. Fundamental particles possess these properties by physical necessity.
5. *Fiat* entities could not have these necessary properties.
6. Since there are macroscopic objects, there must be *bona fide* entities.

(From 1-5.)

I take it that even the extreme conventionalist will accept the first premise; more precisely, we are here interested in the physical possibility of macroscopic objects. We

⁴ See also Sider (2009) and Cameron (2008) for discussion about indirect evidence against the conclusion that there would be no structure in the world.

will proceed to analyse the other premises in order.

2. Certain things are physically necessary for the forming of macroscopic objects, e.g. the laws that govern molecular binding.

In the light of what we know about the forming of macroscopic objects, certain things seem physically necessary: molecules must be able to form bonds, atoms must be able to form molecules, and subatomic particles must be able to form atoms. In virtue of what are these things possible? Well, the binding of molecules and atoms is dependent on the electron configuration of individual atoms, which in turn depends on the energy levels of specific electrons and is moderated by the Pauli Exclusion Principle.⁵ Similarly, the manner in which subatomic particles form atoms is dependent on the individual charges of subatomic particles—the negative charges of the electrons and the positive charges of the protons, where each proton consists of three quarks which make up the total charge of the proton.

3. The relevant laws of physics require that fundamental particles have *exact* properties, such as electric charge.

It follows from what was said above that the total charge of the atom has to be neutral. The picture gets somewhat more complicated when details about the underlying fundamental forces are introduced; for instance, the nucleus holds together in virtue of the *strong* force, which overpowers the repulsive forces between the positively charged

⁵ The Pauli Exclusion Principle states that no two identical fermions can have the same quantum number at the same time.

quarks. In any case, it is obvious that the forming of macroscopic objects is a delicate matter and would not be possible if subatomic particles were arranged in an arbitrary fashion. Even if we are unable to accurately state the charges of each subatomic particle which together form an atom, we do know that their sum has to be zero, otherwise the atom would not hold together. The upshot of this is that there has to be something, namely electrons and quarks, that possess an exact charge.

4. Fundamental particles possess these properties by physical necessity.

It is already apparent that the forming of any kind of macroscopic objects requires a considerable amount of orderliness on the microphysical level. The physical laws that govern the forming of atoms and molecules would not work if there were not *constants*, such as the charge of electrons. We identify elements by the number of protons in the nucleus and for the atom to be neutral in charge it must collect a certain number of electrons. The electrons are ordered into shells and the order of filling of electron energy states is governed by energy and the Pauli Exclusion Principle. Any study of the electron configuration will refer to the fact that it is *impossible* for two electrons to occupy the same quantum state, as stated by the Pauli Exclusion Principle. Everything that we see is dependent on this process of forming atoms, and the process would not be possible if electrons did not have physically necessary properties: electrons have an exact mass, exact charge, intrinsic properties such as spin and angular momentum, and they show no internal structure. Furthermore, the charge of an electron, $-1.6021892 \times 10^{-19}$ coulombs, is a fundamental physical constant: the charge of all other freely existing subatomic particles that have a charge is either equal to or an integer multiple of it. Accordingly, it is a feature of the physics of the actual world that electrons have

their charge by physical necessity. It may be that this is not *metaphysically* necessary—perhaps the laws of physics are not metaphysically necessary and there are possible worlds with alternative laws of physics—but all that is needed for the argument at hand is *physical* necessity. So, there must be entities that possess these necessary properties, but how does it follow from this that there must be *bona fide* entities? For this we need the final premise:

5. *Fiat* entities could not have these necessary properties.

Clearly, *fiat* entities as well can possess exact properties. Consider the city of London: it has a number of exact properties at any given time, such as the number of underground stations, annual budget and the number of citizens. There is, however, a fine difference between the exact properties that the city of London or other *fiat* entities may have and the exact properties that, for instance, electrons seem to have, for the latter are *necessary* whereas *fiat* entities can only possess contingent properties, or at least no substantial necessary properties. Consider the property ‘number of citizens’. We could circle any area of land and count the people living there, giving the property of the number of citizens in that area, and this may be the basis of all sorts of things, such as taxation. It is however completely arbitrary as to which area of land we decide to circle, there are no natural constraints for this. The property ‘number of citizens’ is, we might say, a *fiat* property: although it can be exact at a given time, it has no fixed requirements. It can change from time to time and we could even have decided not to introduce it at all. We could not do anything of the sort, say, in the case of the charge of electrons, because if the charge were different then the microphysical orderliness required for the existence of macroscopic objects would collapse. That is, the exact

actual charge of electrons is *necessary* for the emergence of macroscopic objects, whereas the population of London is thoroughly contingent. It could be objected here that, for instance, the *fiat* entity that consists of two electrons in my left hand surely has certain necessary properties, e.g. the sum of the charges of these electrons. However, this will hardly undermine the argument, for the necessity involved here is based on the necessary properties of the individual electrons, so the *fiat* entity will not have necessary properties in its own right. Accordingly, the modal content involved here might be the best indicator of *bona fide* entities that we can have.

6. Since there are macroscopic objects, there must be *bona fide* entities.

(From 1-5.)

We have now arrived at the conclusion of the argument: because there are macroscopic objects, there have to be *bona fide* entities that instantiate the exact, physically necessary properties that are needed for this macrophysical structure. Accordingly, the very existence of macroscopic objects speaks highly in favour of the existence of *bona fide* entities.

The argument has a number of interesting ramifications. Firstly, we can use the argument in two ways: as a general argument for the existence of *bona fide* entities of whatever type which are necessary for the macrophysical structure, or as a more specific argument towards the conclusion that electrons and other fundamental particles are in fact such *bona fide* entities. This could perhaps be extended into an argument for some sort of microstructural essentialism, but this is not the line we wish to pursue here. Here we only wish to establish the first point, although what was said above about electrons would be a promising place to start for establishing the second point as well. Regarding

the first point, the case seems fairly strong. We must have some entities which uphold the intricate structure required for the forming of macroscopic objects and to do this, these entities must be able to interact in a highly complex and stable manner. This interaction is possible only in virtue of a certain set of exact, necessary properties that these entities possess and these properties are also necessary for the macrophysical structure. *Fiat* entities do not seem to be capable of instantiating necessary properties of this type because they are, by definition, entities that could have been quite different; their boundaries and properties are *contingent*.

As to other ramifications of this argument, the most important ones for the current problem are of course the ones that undermine extreme conventionalism. I am not aware of any (extreme) conventionalist discussion of fundamental particles, which is rather disconcerting given that they are surely the best candidates for *bona fide* entities. Furthermore, even if our current physical theory about fundamental particles is mistaken, i.e. if the particles that we think are fundamental do after all have internal structure, the point that was made above would still hold. Whatever the fundamental particles are, they must be such that they are able to form atoms, and for them to be able to form atoms, they must have exact, necessary properties such as charges. More generally, subatomic particles are subject to a well-defined set of fundamental forces—their interaction is based on these forces—and where there is interaction that is based on certain necessary properties of the involved entities, then these entities must be distinct, *bona fide* entities with natural boundaries, regardless of whether or not we can state the identity-conditions of these entities.

The most important ramification of this account is perhaps that it effectively undermines the Dummettian view of reality as an amorphous lump. One cannot accept both the Dummettian picture and the idea that fundamental particles are *bona fide*

entities. Perhaps the extreme conventionalist could try to argue that the existence of electrons as *bona fide* entities is not required, but rather just that the amorphous lump contains local variations in such a way that the required necessary properties, i.e. negative and positive charges are present in certain regions of the amorphous lump.⁶ However, if the conventionalist concedes this much, then the amorphous lump picture is already undermined: surely a lump with fixed regional variations is anything but amorphous. If the extreme conventionalist thesis suggests that all our efforts to structure reality are based on *Gestalt* factors, then fixed local variations in the amorphous lump are also ruled out. Admittedly, this line would enable the conventionalist to deny the existence of electrons as *bona fide* entities, but it does entail the existence of something like *bona fide* properties.

There may in fact be *some* support for this type of an approach in contemporary physics, as the GRW interpretation of quantum mechanics suggests that what we have been calling particles may be nothing else than aspects of the behaviour of the wavefunction (Lewis 2006). However, as Peter J. Lewis puts it: ‘If the GRW theory is true, then particles and elephants are both instantiated by waves, but this provides no more reason to deny the existence of particles than to deny the existence of elephants’. The details are obviously more complicated than this, but what is clear is that whatever the fundamental structure of the world is, it contains features which enables the existence of macroscopic objects. Accordingly, the core thesis of extreme conventionalism is already refuted: there is structure in reality and it is according to the *de re* features of reality—be it particles or wavefunctions—that we carve it up. Indeed, Varzi claims that ‘The conventionalist stance simply entails that which of them [the individuals that we may postulate] come to play a role in our life is up to us.’

6 There are elements for a suggestion of this type in the work of O’Leary-Hawthorne & Cortens (1995) and Horgan & Potrč (2008).

(Forthcoming: 20), but it seems to me that it is not: we could try to ignore the structure present in the microphysical, but this would render our physical theories quite unable to do the job they were designed to do, namely, they would fail to be predictive.

3. Why do we classify things in the way we do?

We already have a fairly good case against extreme conventionalism, as it appears that fundamental particles are very likely candidates for *bona fide* entities. Now we turn to another serious claim, namely that our classificatory schemes are grounded in our psychological biases, *Gestalt* factors that do not represent the structure of reality in any way:

Consider the debate on unrestricted composition. There is no question that we feel more at ease with certain composites than with others. We feel at ease, for instance, with regard to such things as the fusion of Tibbles's parts (whatever they are), or even a platypus's parts; but when it comes to such unlovely and gerrymandered mixtures as Lewisean trout-turkeys, consisting of the front half of a trout and the back half of a turkey, we feel uncomfortable. Such feelings may exhibit surprising regularities across contexts and cultures. Yet, arguably they rest on psychological biases and *Gestalt* factors that needn't have any bearing on how the world is actually structured. (Varzi Forthcoming: 16.)

Is there any connection between how the world is structured and our evaluation of things such as trout-turkeys? Varzi argues that there might not be, as even though we initially feel uncomfortable about strange hybrids, we have nevertheless welcomed a variety of genetically manipulated plant-hybrids, such as orange-mandarins. Indeed, our intuitions and feelings of discomfort should not be relied on if we hope to determine the actual laws of reality; it is true that we are biased in our evaluations. However,

psychological biases like these have little to do with scientific practice—the very existence of genetically manipulated hybrids is proof enough. The actual structure of the world quite clearly *does* have a bearing on our scientific practices though. The case of the trout-turkey might not seem to be directly relevant here. After all, trout-turkeys are not supposed to be results of genetic manipulation, rather, they are just mereological sums consisting of two disconnected parts, the front half of some trout and the back half of some turkey. How do we evaluate the case of the trout-turkey in this context, then? The answer depends on our take on unrestricted mereological composition. Examples like trout-turkeys and the sum of one's nose and the Eiffel tower could certainly be seen as a *reductio* of unrestricted mereological composition, but we do not need to pursue that line of thought here. In any case, there may be good reasons for our initial, hostile reaction towards such entities: perhaps in these cases composition does *not* occur, there are no such entities. Indeed, I do not find unrestricted mereological composition plausible.⁷ To settle this, we would need a thorough discussion of unrestricted mereological composition, but because we cannot engage in this discussion here, it might be more interesting to modify the example so that it is neutral in regard to questions of mereology.

How should we modify the example? Well, instead of a mereological sum, we can consider trout-turkeys as a more traditional entity, a genetically manipulated hybrid. If we, for some twisted reason, wished to create a trout-turkey, the first question would obviously be whether it is *possible* to create such a creature. My guess is that it is *not* possible to do this, but even if it were possible, it is unlikely that the creature would be able to sustain itself, given the apparent differences between the habitats of trout and

7 See Tahko (2009a) for more discussion on unrestricted composition, there I argue that the vagueness argument against restricted composition fails and that we might be better off with restricted composition.

turkeys. Be that as it may, the fact that we feel uncomfortable about the idea of certain hybrids, whether they are just mereological sums or genuine hybrids, is entirely unrelated to our scientific classificatory schemes, which certainly do not concern mereological sums. Our question in this section is what are these schemes based on.

It was already mentioned above that an important tool of our classificatory framework is possibility: when we consider what sort of entities could exist, we do not decide this in terms of which entities we feel comfortable with, but rather in terms of which entities are possible. And how do we decide which entities *are* possible, leaving unrestricted mereological composition aside? Well, by examining the relevant sub-categories of possibility. In the case of trout-turkeys we would be interested in the *biological* possibility of these creatures, namely whether there could be a DNA sequence that produces trout-turkeys. This is of course (partly if not entirely) a matter for biological research. However, the space of possible organisms is also restricted by *physical* possibility, in other words we can rule out creatures that are not physically viable given the actual laws of physics. Certain insects above a certain size, for instance, would be ruled out, because their respiratory system would not be able to function; and there is a good chance that a trout-turkey would encounter similar problems.

We are now faced with a question: do we have any reason to believe that what we have just described is based on *Gestalt* factors rather than the actual structure of the world, i.e. the laws of nature? Admittedly, our theory about these laws of nature is fallible, and it might even turn out that something that we believed to be biologically or physically impossible is, after all, possible. Perhaps some of these mistakes are even due to *Gestalt* factors, but to claim that there is *no* structure behind them is like claiming that a monkey could have written *The Brothers Karamazov*—it is possible, if you put the monkey in front a keyboard and it randomly beats the keyboard to produce the book

in question, but it is not very likely. Similarly, it seems that our classificatory scheme must correspond to *something*, as otherwise it would be just gibberish. How could we possibly come up with such a sophisticated structure by coincidence? More importantly, our current scheme has enormous predictive power: we can predict a huge range of natural phenomena from chemical reactions to the movement of heavenly bodies. This predictive power must be based on something and the obvious explanation is that our classificatory system roughly corresponds with the world. Consider an example: Mendeleev's periodic table.

Mendeleev arranged elements into a table by their atomic mass and their chemical properties, which enabled him to predict the existence of a number of yet undiscovered elements as well as the chemical properties of these elements. To start with, Mendeleev had some established empirical information about certain elements, namely their atomic masses and chemical properties. It was a natural thing to do to examine the relationships between the elements, indeed, other similar attempts were being made around the same time as Mendeleev published his periodic table. What is interesting to us is how effective this system was in terms of making predictions about future empirical observations, namely undiscovered elements. The periodic table can be seen as a description of what is possible given certain building blocks (i.e. our knowledge about the atomic masses and chemical properties of certain elements).

The modal basis of Mendeleev's work were arguably the different possible states of affairs that could explain empirical observations.⁸ The likeliest explanation for the success of Mendeleev's classificatory scheme would appear to be that it is the *correct* scheme, correct in the sense that it corresponds with the structure of the world. This does not mean that we could not reach similar results with a very different classificatory

⁸ For a more extensive account on the modal basis of such tools, see Tahko (2008).

scheme, although there are certain pragmatic reasons to prefer Mendeleev's scheme, including the usual theoretical virtues such as simplicity. However, any scheme that differs from Mendeleev's so radically that it loses predictive power, never mind theoretical virtues, will simply be an incorrect way to characterize natural phenomena.

So, why do we classify things in the way we do? Certainly, psychological biases play a role here, but only a very modest one: we are quick to abandon them if they lack predictive power. If someone were to create a trout-turkey—let us assume that it is possible—we would soon acknowledge it, regardless of how uneasy we might feel about it. Consequently, our classificatory schemes, although always subject to revision, are fundamentally based on the actual structure of the world. It may be that we are unable to ever accurately state what that structure is, but it is nevertheless the basis of our classificatory efforts. It also seems that we are getting better at classifying things all the time, judging from the increasing predictive accuracy of our classificatory schemes.

Presumably, the conventionalist will deny the inference at hand: from the predictive power of our classificatory schemes to their approximate correspondence to reality as it is in itself. The conventionalist will no doubt acknowledge that it would be absurd to abandon the theories that offer best predictive power, but it could still be insisted that this does not mean that our best theories even roughly correspond with reality. This Humean position is difficult to refute; the conventionalist can always fall back to the point that there are no a priori reasons to think that our classificatory systems carve the world at the joints. Perhaps this is an unfair objection, as naturally we classify things in regard to empirical feedback rather than based solely on some a priori principles, but is there nothing else that we can say to convince the conventionalist? Well, perhaps there *are* some a priori principles that guide our classificatory schemes. These could be fundamental logical principles such as the law of non-contradiction. What is interesting

here is that the law of non-contradiction, for instance, does not appear to be strictly a principle of thought or language, as linguistic paradoxes such as the Liar have demonstrated. Or, at any rate, the principle seem to break down in these contexts. This vindicates the idea that the law of non-contradiction, if anything, is a principle concerning reality rather than our thoughts about reality, contrary to what someone like Dummett (1991: 1-2) would claim.⁹ The conventionalist will probably still maintain that we have no a priori reasons to think that *reality* rather than our thoughts about reality conforms to the law of non-contradiction, and perhaps we indeed do not have such reasons; however, our alternative here is utter scepticism and if this is the route the conventionalist wishes to take, then so be it. This plight is caused by the refusal to accept any empirical grounds for knowledge, but even those who question the validity of the law of non-contradiction act as if reality did conform to it. This is despite the fact that our thoughts and language sometimes do appear to be contradictory and hence it would certainly be within our powers to act against the apparent consistency of reality. But this line of thought will not get us much further, surely the conventionalist will have to provide a stronger case than the Humean one to convince the realist. Once empirical elements are accepted into our epistemology even in the most rudimentary form, we can quickly build on them with the help of the scientific method; perhaps the most important tool here is the predictive power of theories. We would hope that when faced with the choice between this realist view and utter scepticism, the decision will be easy.

4. Is our classificatory system unique?

Even if we can dismiss the utterly sceptical line as absurd, some doubts about just how accurately our systems of classification correspond with reality may remain. After all, the manner in which we classify things surely has something to do with our particular

⁹ For an extensive discussion of the interpretation of the law of non-contradiction, see Tahko (2009b).

psychological biases, which depend on our rational capabilities and physical characteristics. We are beings of a certain size, our senses are tuned in a certain way, and our brains have a certain capacity. Accordingly, it may be that beings who differ from us in regard to one or more of these features have a very different way of carving up the world.

Let us start from animals. We know that for instance dogs can hear frequencies that we are not able to hear, and on the other hand we can see things that they cannot see, namely certain colours. Presumably our rational capabilities also differ substantially from those of dogs. Naturally, there will be striking differences in how dogs perceive and classify the world and how we do. Just how striking will these differences be though? Dogs are unlikely to have any grasp of atoms or natural kinds, but they do clearly have some grasp of distinctions like edible and non-edible, friendly and hostile, leader of the pack and member of the pack. In other words, there is some overlap between how dogs carve up the world and how we do, even if there are also major differences. Many of these differences are based on physical factors which cannot be undermined, but some of them could perhaps be mapped to correspond with our framework of classification. For instance, a dog would determine what is edible primarily with his sense of smell, whereas we would often have to rely on taste, but the result will still be the same. In any case, the major problem in mapping the dog framework to ours is that dogs do not possess similar rational capabilities and thus their classificatory framework will necessarily be quite rudimentary compared to ours. Nevertheless, this does not mean that the dog framework would not rely on similar methods of classification, e.g. predictive power. For instance, a dog can learn to predict when there will be food available with the help of a number of markers: time of the day, a sound or a smell.

Perhaps a more interesting example than dogs would be beings of roughly similar intellect to us. As there are no obvious examples on Earth, we may take some sort of aliens as an example. Let us assume that they have roughly the same brain capacity and are similarly developed compared to humans. However, these aliens, although their physical constitution is roughly similar to ours otherwise, have no eyesight, instead they use a sophisticated sonar.¹⁰ Once again, there will necessarily be major differences between our and the alien framework of classification. For instance, colour concepts would be quite meaningless to the alien species, as their sonar would be to us. The key issue here are not everyday examples though, but rather scientific ones—examples that reflect our most sophisticated efforts of classifying. Consider elements: could the aliens have anything similar to the periodic table of elements which is so familiar and important to us? Their means to acquire information about the elements would certainly be different from ours, given that sight plays a major role for us, but then again things like colour concepts are completely unnecessary for stating something like ‘Hydrogen is the lightest element’. If we assume that the aliens have tools fit for their senses that enable them to examine atoms, then it would be strange if they did not have some grasp of what ‘the lightest element’ means. It does not matter in what way they come to know that there is an element which is the lightest element, namely hydrogen.

Provided that the aliens are at all scientifically minded, they will want to explain the same natural phenomena that we do. They might not have anything like our periodic table of elements though, perhaps their information about the elements would be arranged in a table according to the sound output that their sophisticated sonar microscope produces. In any case, whatever the format of their table of elements might be, it would still contain information about the relationships between different elements just as the periodic table does. We can assume this with confidence because the

¹⁰ A ‘Martian sonar’ like this has been discussed in McCulloch (1988).

relationships between different elements are crucial for an understanding of chemical reactions, which is what any intelligent being would surely hope to acquire. It would be possible for the aliens as it is for us to arrange elements by their electron configuration, and it is likely that the aliens could also predict still unobserved elements with the help of their table, as Mendeleev was able to do with his. The upshot of this discussion is that the superficial differences between classificatory systems are not important, as they could be mapped to correspond, translated to match each other. This is of course possible because the subject-matter of the systems is the same: the elements are the same for anyone who might wish to observe them. It is not relevant whether we would be able to perform the translation, all that matters is that there is a theoretical correspondence.

One further objection could be raised about both of our previous examples, as they deal with beings roughly the same size as us. One important limitation to the way we perceive the world is indeed our physical size. If there were beings, say, the size of atoms or the size of galaxies, they would surely have a thoroughly different point of view to the world. On the face of it, this is correct. Atoms and galaxies differ so radically from us in terms of size that we can barely comprehend just how small and large they are, respectively. Only some hundreds of years ago we did not even know about things like atoms and galaxies, and our system of classification did not, of course, include them. We only classified things that we could observe and understand; in this sense we indeed had a psychological bias dependent on our physical size. However, now we *do* know about atoms and galaxies, and even about entities much smaller and larger than these. We may not be able to fully grasp the scale of these things, but we can express it accurately in numerical form and in comparison to our own size. In fact, our most important classifications now arguably concern things that are either vastly smaller

or larger than us.

The question is, would beings much smaller or larger than us classify things in the same way as we do? Well, naturally they would face the same initial limitations due to their size as we did, but what if they did have some means to observe the same scale that we do? It seems to us that the answer is very clear: at first any kind of intelligent beings would be most interested in things of roughly similar size to themselves, but if they had the means to observe the same scale as we do, they would certainly wish to classify these things. It is obvious that there would be differences in all three systems of classification—ours, the atom-sized beings, and the galaxy-sized beings—but once again it would perhaps be possible to map these differences so that the systems would correspond. This is because each system, provided that we all could observe the same scale of things, would no doubt have roughly the same number of different *types* of entities. Each system would perhaps be more detailed in regard to its own scale, simply because it is easier to observe things roughly of our own size, but we would of course be more than happy to amend our own system of classification to accommodate any details that we might learn about from the systems of atom-sized beings and the galaxy-sized beings.

Yet another concern related to our physical size, and perhaps also to our perception of time, is that our conception of the persistence conditions of certain objects may be very different from creatures much smaller and larger than us, or creatures much more short- or long-lived than us. For instance, we would perhaps consider mountains to be unable to survive erosion which fully flattens them: in this case the mountain has been lost. However, for very short-lived, microscopic beings, something like bumps on a mattress would be rather similar as mountains are for us, yet we would not consider anything to have been lost were these bumps to be flattened. Similarly, for some

extremely long-lived and large creatures mountains might be quite like the bumps on a mattress are for us. Perhaps there are indeed some psychological biases here which we cannot overcome, as certain objects will simply not have a similar relevance for us and for beings of a very different size. However, the fact that we might not pay much attention to bumps on a mattress does not mean that they are not included in our ontology in some sense. We might not count them as *bona fide* entities and the argument at hand may be reason enough to dismiss mountains as *bona fide* entities as well—or perhaps we should count *both* the bumps on the mattress and the mountains as *bona fide* entities—but it is certainly not enough to motivate extreme conventionalism. In any case, I acknowledge that two object-candidates which have the same properties save for size would presumably have the same (or nearly the same) persistence conditions as well. The problem is to produce a plausible story about when composition occurs, but that is something that we cannot settle here.

Admittedly, the whole scenario is rather strange: it does not appear that intelligent beings the size of atoms or galaxies are possible, but then again this might just be a psychological bias due to our physical size! At any rate, we do know that there are organisms so small that we cannot observe them with our bare eyes. It may also be that, say, for beings the size of atoms it would be physically impossible to observe anything much larger than themselves and thus, if they were intelligent, their system of classification would be bound to atom-sized things, rendering it fundamentally alien to us. Nevertheless, although size and perceptual capabilities play a central role in our classificatory schemes, they can at least to a large extent be overcome with the help of tools, such as the microscope and telescope. Hence they are nothing more than a hindrance, the list of entities waiting to be classified is the same for everyone—size does not matter.

5. Realism refuted?

Where do we stand now that we have discussed three major conventionalist lines of criticism against realism? Not very far from where we started. The conventionalist will need some fairly strong arguments to undermine realism, for the alternative is not particularly attractive. As Varzi admits: ‘Surely the *intuitive* plausibility [of the conventionalist stance] is pretty low, and perhaps also its scientific tenability’ (Varzi Forthcoming: 21). Varzi continues by claiming that conventionalism nevertheless has some philosophical advantages, but these would need to be demonstrated thoroughly to give conventionalism any hope of overcoming its initial implausibility. Furthermore, the arguments that we have discussed appear to be inadequate to undermine the plausibility of realism.

We have seen that our system of classification is fundamentally grounded in reality; it *must* be grounded in reality, for otherwise it could hardly be so orderly—the predictive power of our classificatory schemes is a proof of this. It is an open question which entities are genuine, *bona fide* entities, we need philosophical inquiry as well as science to determine this, but it does seem that fundamental particles for the very least, whatever they are, must be *bona fide* entities. Otherwise macroscopic objects would not be possible in the first place, as we saw in the second section. Furthermore, although intelligent beings different from us might have very different systems of classification, their systems as well must share the same common ground. In conclusion, realism stands unrefuted, all major lines of criticism available for the conventionalist can be addressed.

References

Cameron, R. P. (2008). Turtles All the Way Down: Regress, Priority and

- Fundamentality. *The Philosophical Quarterly*, Vol. 58, No. 230, 1-14.
- Dummett, M. (1981). *Frege. Philosophy of Language*, 2nd edn. Cambridge, MA: Harvard University Press.
- Dummett, M. (1991). *The Logical Basis of Metaphysics*. Cambridge, MA: Harvard University Press.
- Dupré, J. (1993). *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*. Cambridge, MA: Harvard University Press.
- Dupré, J. (2002). *Humans and Other Animals*. Oxford: Oxford University Press.
- Elgin, C. Z. (1995). Unnatural Science. *Journal of Philosophy*, Vol. 92, No. 2, 289-302.
- Horgan, T. E. and Potrč, M. (2008). *Austere Realism: Contextual Semantics Meets Minimal Ontology*. Cambridge, MA: MIT Press.
- Lewis, P. J. (2006). GRW: A Case Study in Quantum Ontology. *Philosophy Compass*, 1/2, 224–244.
- McCulloch, G. (1988). What it is Like. *The Philosophical Quarterly*, Vol. 38, No. 150, 1-19.
- O’Leary-Hawthorne, J. & Cortens, A. (1995). Towards ontological nihilism. *Philosophical Studies*, Vol. 79, No. 2, 143-165.
- Sider, T. (2009). Ontological Realism. In Chalmers, D., Manley, D. and Wasserman, R. (Eds.). *Metametaphysics: New essays on the foundations of ontology* (pp. 384-423). Oxford: Clarendon Press.
- Smith, B. & Varzi, A. C. (2000). Fiat and Bona Fide Boundaries. *Philosophy and Phenomenological Research*, Vol. 60, No. 2, 401-420.
- Tahko, T. E. (2008). ‘A New Definition of A Priori Knowledge: In Search of a Modal Basis’, *Metaphysica*, Vol. 9, No. 2, pp. 57-68.
- Tahko, T. E. (2009a). ‘Against the Vagueness Argument’, *Philosophia*, Vol. 37, No. 2,

pp. 335-40.

Tahko, T. E. (2009b). 'The Law of Non-Contradiction as a Metaphysical, *The Australasian Journal of Logic*, Vol. 7, pp. 32-47.

Tye, M. (1990). Vague Objects. *Mind*, Vol. 99 (396), 535-557.

Varzi, A. C. (1997). Boundaries, Continuity, and Contact. *Noûs*, Vol. 31, No. 1, 26-58.

Varzi, A. C. (Forthcoming). Boundaries, Conventions, and Realism. To appear in Slater & Cambell (Eds.), *Carving Nature at Its Joints: Topics in Contemporary Philosophy*, Vol. 8. Cambridge, MA: MIT Press. Retrieved September 2009 from: <http://www.columbia.edu/~av72/papers/INPC_2008.pdf>.