

A Metaphysics for Semantic Internalism

Paul Tappenden paulpagetappenden@gmail.com

I think you will find a habit of hilarious Olympian controversy very helpful.

Bertrand Russell (1952)

Abstract The contemporary popularity of semantic externalism has arisen from so-called Twin Earth thought experiments which suggest that the representational content of a natural kind term cannot be wholly determined by processes within a speaker's body. Such arguments depend on the intuition that the extensions of natural kind terms cannot have changed as the result of the scientific investigation of natural kinds' constitutions. I demonstrate that this externalist intuition depends on an assumption about the mentality of isomorphic doppelgangers which has never been questioned but which is nonetheless arguably false. I develop an alternative view of the instantiation of mind which entails a revision of our understanding of the constitution of environmental objects. The picture seems to be fully coherent despite its oddity and I can find no good reason to reject it. The conclusion must be that the case for semantic externalism is thus less compelling than is often supposed.

Keywords Semantic Externalism - Twin Earth - Doppelganger - Natural Kind - Set Theory

1 Speaking of Experts

It's commonly accepted nowadays that the representational content of natural kind terms such as 'water' is wide. That is, the content cannot be determined solely by phenomena within a speaker's body such as cerebral processes. This thesis about the content of natural kind terms is the most well-established manifestation of what has come to be known as semantic externalism, a view which might be said to mark an Hegelian turn in analytic philosophy. That is how Tyler Burge characterised it, one of the pioneers and foremost proponents of the "Twin Earth" arguments which have

been pivotal in establishing contemporary semantic externalism (1979), p.73. This philosophical shift was to counter what Burge called ‘the elderly Cartesian tradition’ (ibid.) by which he referred not to Cartesian dualism but to the idea that mind is confined to the individual and so the environment is wholly external. Arguably Ludwig Wittgenstein was the harbinger of this twentieth century assault on Cartesian individualism but the Twin Earth arguments can seem to have provided the *coup de grace*.

Since semantic externalism is associated with such a profound dichotomy in modern philosophy it should not be surprising that an effective argument against it involves a novel and counterintuitive idea. That is what I shall present here and if it is to be rejected because it seems so strange then we should be clear about why the strangeness is not simply a matter of its conflicting with traditional habits of thought. I have not been able to find any sound reason to reject the idea; perhaps others can succeed where I have failed and thereby spare semantic externalism from this challenge.

The focus of attention will be the content of natural kind terms because they are usually thought to be most resistant to an internalist analysis and I shall take the example of ‘water’ because of its venerable status arising from Hilary Putnam’s seminal paper (1975), pp. 215-271. Quibbles about impurities, isotopes or the complex ionic structure of water can safely be set aside; it’s to be agreed for the sake of argument that water is H₂O. Putnam’s classic Twin Earth thought experiment, despite its recognised shortcomings and the plethora of variants to which it has given rise, will serve perfectly well to demonstrate the force of what follows and it should be clear that the argument will generalise to all other externalist claims supported in a similar manner.

Given that water is H₂O it is generally taken to follow that if, implausible as it may be, there existed a Twin Earth which were like Earth in every respect except that all the water-like substance had some other chemical constitution, XYZ, then utterances of the Twin-Earthly homophone of ‘water’ would not refer to water. This thought in itself does not establish semantic externalism. It could still be maintained that the content of ‘water’ and its Twin-Earthly homophone is determined by the discriminative abilities of expert chemists on each planet, such discriminative abilities being plausibly instantiated by the neural architecture of those chemists. Non-expert

utterers of 'water' and its homophone would lack those discriminative abilities but that would be no bar on the contents of the uttered terms being determined by the abilities of the relevant experts, as would be compatible with semantic *internalism*.

That's why Putnam set his famous thought experiment in 1750, before any such experts existed on Earth and on the corresponding version of Twin Earth. For some reason there's a forceful widespread intuition that if tokens of 'water' uttered now refer to all and only samples of H₂O then the same must have been true of tokens of 'water' uttered in 1750. That is, to put it another way, the content of 'water' cannot have changed as the result of the honing of the discriminative abilities of experts.

In his book defending semantic internalism Gabriel Segal sought to challenge that key intuition (2000), pp. 132-134. Attempting to make more acceptable the idea that the extension of tokens of 'water' in 1750 was less exclusive than nowadays Segal invites his readers to consider a very bizarre possibility. Allowing that a planet something like Twin Earth can exist, we are to imagine how things might have gone if it had actually existed, not necessarily populated, and technically advanced aliens had provided a 'helpful transport service' between Earth and Twin Earth prior to 1750 so that there was much to-ing and fro-ing before modern chemistry developed. Segal argues that Earthlings might well have come to refer to the liquid in the rivers and lakes on each planet as water and would have not seen fit to speak differently once the difference between the H₂O on Earth and the XYZ on Twin Earth had been discovered. He wrote:

The natural conclusion is that 'water' on Earth did not refer exclusively to H₂O in 1750. For it appears that it could have come to refer to both water and twater without a change in meaning (*ibid.*), p. 133.

Segal goes on to recognise that an externalist account of what might happen in the hypothetical case is possible but he suggests that the internalist story is more plausible. This is the best argument that I know of against the semantic externalist assumption that the extensions of utterances of natural kind terms do not change but on Segal's own admission the argument is at best inconclusive. Given that there was not an alien taxi service between Earth and some form of Twin Earth prior to 1750 the internalist must admit that there has been a change in the reference of tokens of

'water' since then. In that case it has to be said that when scientists discovered that passing an electric current through water caused hydrogen and oxygen to be given off separately at the two electrodes they were engaged in a process which would lead to a change in the reference of 'water' as well as the advancement of chemistry. That can seem an odd thing to say.

Further, suppose that in 1750 a philosophers' demon decided to change all the liquid in Earth's rainfall, oceans and so on from H₂O to XYZ for six months and then back again. According to semantic internalists nobody would have been deceived. All would have been correct in calling that liquid water at any time during the year. Again, that can sound odd. Of course, sounding odd is not in itself an insuperable problem but externalists may well wonder why internalism doesn't just lay down and die rather than flying in the face of intuitions which we are given no good reason to doubt.

If internalism were right then the extension of tokens of 'water' in 1750 was all and only samples of H₂O, XYZ *and anything else which would have been indistinguishable from H₂O in 1750 by those most expert on the constitutions of natural substances*. This seemingly counterintuitive idea could in principle be rendered more palatable if it could be shown that there exists an alternative interpretation of Putnam's Twin Earth scenario which makes it possible that samples which our ancestors called water in 1750, such as the contents of what they called Lake Michigan, *was not water*. I'll repeat the argument scheme to emphasise the point. The semantic internalist must argue that the extension of tokens of 'water' uttered in 1750 included samples of Putnamian XYZ, if it did indeed implausibly exist, and anything else indistinguishable from H₂O by the experts of those days. Segal has argued for exactly that but his argument is inconclusive and unlikely to be accepted by anyone who does not already believe, with Segal, that semantic internalism is right. But the claim that the extension of tokens of 'water' has changed since 1750 becomes more acceptable if it can be shown to be possible, given the implausible assumption of the existence of Twin Earth, that the very liquid in William Hogarth's gutters and puddles was not water but something else which would have appeared to be water to him and the most expert of his contemporaries. That sounds hard to believe but nonetheless the following argument is aimed at showing that there is a plausible way to support that idea.

Bear in mind that counterpart speakers on Putnam's Earth and Twin Earth in 1750 may be perfectly isomorphic doppelgangers. That is not a requirement of the argument but it must be a possibility if anything like Twin Earth is possible. Of course, counterpart doppelgangers must be different to the extent that Earthlings have H₂O in their bodies and Twin-Earthlings have XYZ in the corresponding places, but that could be avoided if a natural kind relevantly different from water were chosen so that a structural difference between doppelgangers was not entailed. Or it could simply be assumed that the somatic roles of H₂O and XYZ would have no differential effect on the doppelgangers' minds. Now for a leap of imagination.

2 Doppelgangers United

As a child I was a fan of Patrick Moore's popular astronomy television programme 'The Sky at Night'. I remember him once saying, with characteristic enthusiasm and eyebrow articulation, that he was sure that somewhere in the universe there was a planet just like Earth with someone just like him saying the very same words he was saying. The idea has a colourful history entertainingly recounted by Thomas Lepeltier (2010).

Let's begin by putting aside Twin Earth and imagine two perfectly isomorphic Earthlike planets. Water, water everywhere. Consider a pair of exact doppelgangers, one on each planet, and unconventionally suppose that they instantiate the single mental life of one subject, Tweedle. Charitably interpreting Tweedle's speech and behaviour imaginatively from afar, we can assign contents to his utterances, beliefs and desires. When Tweedle says 'I see a green apple before me' a single utterance is manifest as two isomorphic sonic emissions. We can understand Tweedle as speaking truly if we take him to be referring not to an apple but to a t-apple, a pair of twinned apples, and to be saying of that t-apple that it is t-green, in other words that both its component apples are green. Not just any pair of isomorphic apples would constitute a t-apple; the pair would have to be counterparts in matched environments so that they could, in principle, constitute a t-apple observed by a subject like Tweedle.

If Tweedle reaches out to grasp the t-apple then that intensional action is manifest as duplicate bodily movements and we can hypothesise that Tweedle reaches for the t-apple because he believes it's t-crisp and t-juicy and he desires to munch it,

where munching is manifest as matched mandibular motions. On a representational view of mind Tweedle's mental representations are most plausibly understood to be instantiated by t-cerebral processes. The contents of those representations are what interest us.¹

It's generally thought that the importance of mental content to psychology is in the explanation of behaviours. The idea is that the physical structures and processes which are the vehicles of content have a fundamental role in the causal determination of intensional action. It is important to see that this idea is preserved for what I'll call the unitary interpretation of doppelgangers. It's Tweedle's t-cerebral processes which cause the intensional behaviours manifest as twinned motions. It is not a matter of some whimsical t-causation. Here's the argument.

The idea of the causal role of content takes the everyday notion of cause for granted. It does not pretend to offer an account of natural necessitation, which David Hume taught us is so deeply problematic. The causal role of content-bearing representations is not supposed to be any more or less mysterious than the causal role of Humean billiard balls. Since our imagined pair of planets are hypothesised to have isomorphic histories all relevant Humean constant conjunctions are between t-events. If one t-ball t-strikes another t-ball then the second t-ball is caused to t-move by the first t-ball, in the everyday sense of 'caused'. Any intuition that somehow the "real" causation is between the events on each planet rather than between the t-events of a t-planet supposes that there can be a relation of natural necessitation between events but there is as yet no account of how such a relation between events could possibly exist.

As Tweedle describes in detail the t-garden in which he finds himself we see through our minds' eyes that his description corresponds exactly to the way we would describe each of the isomorphic gardens. Gesturing proudly to the t-spires beyond, Tweedle speaks of the t-town which is his home and which he calls Oxbridge. Suddenly he hears a voice behind him, 'Pardon me, but....'.

Turning around, Tweedle responds, 'Who are you, and what d'ya think you're doin' here?'. He addresses himself to the singular little girl whose duophonic

¹ There is currently discussion of an idea known as extended cognition, distinct from semantic externalism, which holds that some mental representations are not cerebral but include elements of an organism's environment. But that is controversial and anyway would still confine the locus of representations to the vicinity of an organism and so does not threaten the ensuing argument. For a selection of papers from both camps see (Menary 2010).

utterance has disturbed his reverie. And Tweedle thinks to himself, ‘What a scruffy girl; her shoes need polishing’. We assign to Tweedle a thought about one girl and her pair of t-shoes which are t-muddied by her adventures.

‘I’m Alice,’ she replies, ‘and I just wanted to ask if you know where and when the Mad Hatter’s Tea Party begins. I’m lost and I so very much want to be there on time’.

‘Then you’ll have to hurry,’ snorts Tweedle, ‘it’s one mile due north at four o’clock precisely’.

We notice that the eyes of Tweedle’s two bodies have swivelled toward twinned clock towers with weathercocks atop. The hands on the clocks show ten-to-four and mentally zooming out we see that twin tea gardens are indeed in the final stages of preparation a mile away in the direction of the north-pointer of each weathercock. We understand that Tweedle is referring to a t-event at a t-time and t-place in Oxbridge.

Things get curiouser. Tweedle and Alice spy a beguiling t-box which attracts their attention. Unbeknown to them, the interiors of the t-box’s component boxes are causally isolated from their surrounding environments which allows an anisomorphism to appear on the scene. One box contains a duck and the other a rabbit. So the t-box t-contains neither a t-duck nor a t-rabbit since the contents of the component boxes are neither both ducks nor both rabbits. Noticing a t-latch as they approach, Alice urges Tweedle, ‘Oh, do please open the box!’.

‘No, no,’ responds Tweedle nervously, ‘ladies first!’. Alice opens the t-box with a t-trembling t-hand but she’s a bit scared and so averts her gaze. Tweedle, however, watches eagerly from a distance as the t-lid t-opens. As the lids come off the boxes each component of Tweedle’s t-body registers different retinal patterns which have knock-on neural effects resulting in the instantiation of qualitatively different perceptual experiences and, hey presto! The single stream of experience which was Tweedle’s divides and Tweedle fissions into Tweedledum who sees a duck and Tweedledee who sees a rabbit.

Note that the different effects of the duck and the rabbit on the doppelgangers is not at odds with what I said earlier about causation. The different contents of the boxes are supposed to be causally isolated from the surrounding environments until their lids come off. So until that time the relevant constant conjunctions can be presumed to have always been between t-events with isomorphic components. Once

that symmetry is broken causal explanation has to be sought in terms of a causality based on the constant conjunctions of relevant events on each planet. Ducks are associated with retinal images of ducks and rabbits with images of rabbits.

At first brush the interpretation of doppelgangers charitable to the idea that they instantiate a single mental life seems to be coherent given a satisfactory analysis of personal fission, of which I shall say more later. And it is hard to see how Tweedle, prior to watching the t-box t-open, could have made any sort of observation or conducted any sort of experiment which would have indicated to him that he inhabited a t-town rather than a town. Nor is it easy to see how Tweedledum and Tweedledee could discover that they issued from a person who lived in a t-town, rather than the matched towns of Oxford and Cambridge which they themselves inhabit (though they of course both call their towns Oxbridge).

In the following sections I shall discuss whether there is any good reason to dismiss the idea that we ourselves are or could have been in a predicament similar to Tweedle's and I shall fail to find one. In the meantime, the implication of the Oxbridge duckrabbit story for semantic externalism is clear. If it's possible that Putnam's Twin Earth existed back in 1750 then it's possible that at that time Earthlings and Twin Earthlings were for-all-practical-purposes exact doppelgangers and the unitary interpretation can be applied. Back then, what we now call Earth would have been but one component of a t-planet. Our ancestors would have inhabited that t-planet and what they called Lake Michigan would have t-contained neither H₂O nor XYZ. What our ancestors called water would not have had the same constitution as what we call water. The predicament of the chemists first successfully probing that constitution would have been like that of Tweedle peering into the t-box. Those chemists would have fissioned and knock-on effects would have caused further widespread anisomorphism between Earth and Twin Earth with the result that the liquid in brooks and baths would have become H₂O for us on Earth and XYZ for our anisomorphic counterparts on Twin Earth. Notice that Putnam's famous thought experiment requires the causal isolation of the molecular mismatch between H₂O and XYZ, like the causal isolation of the boxed duck and rabbit.

Once the possibility of Twin Earth is invoked then the possibility of there being for-all-practical-purposes exact doppelgangers must be allowed. In which case, given the unitary interpretation of doppelgangers, it becomes perfectly intelligible how the

substance in what people called Lake Michigan in 1750 need not have had the same constitution as the substance in what people call Lake Michigan now. The oddity of the semantic internalist claim that the reference of tokens of ‘water’ has changed since 1750 as a result of scientific investigation is thus defused by an alternative interpretation of the strange but intelligible idea that Twin Earth and doppelgangers could exist, which has been the motor of the recent dominance of semantic externalism.

What remains to be done is to show that it’s credible that we could be in a situation similar to Tweedle’s and then consider whether there’s any good reason to reject that counterintuitive idea. This will require a novel proposal about the constitution of objects in observers’ environments and to prepare the way it will be useful to briefly review the place of the idea of isomorphic worlds in contemporary cosmology. That will serve to clarify a possible misconception about the nature of our local environment and show that the metaphysical scheme I shall introduce may have wider relevance than the philosophical debate which is the prime concern here.

3 A Cosmological Interlude

If there really were exact duplicate planets in the universe they would each have to have been subject to qualitatively identical causal influences for their entire history, setting aside very bizarre “quantum accidents” which we needn’t go into. That’s because natural processes are often chaotic in the sense of modern physical theory which means that a very tiny difference in causal influence would quickly lead to very different histories. Now, as it happens, according to contemporary cosmology there is a finite limit to the distance of events which could have had a causal influence on Earth because our local environment has apparently existed for a finite period. So we exist within a causally isolated cosmic bubble of finite radius, what cosmologists call a Hubble sphere. It is also considered plausible that our universe is spatially infinite in which case it could contain a countable infinity of isomorphic Hubble spheres of the sort we appear to inhabit and it is possible to calculate what their average separation would be.²

² It’s 10^{99} lightyears. For more details see (Tegmark 2007), p.104.

Even if we hold on to the conventional plural interpretation of doppelgangers we cannot now point towards the closest duplicate Earth. All lines of sight for us converge towards a very small region of the Big Bang. The geometry of what is visible to us, objects on the surface of our past lightcone, is like the surface of a globe. It's as if we were looking out from the north pole over the surface of the Earth with the hundreds of billions of visible galaxies scattered across that surface. The most distant galaxies would be in Antarctica and our line of sight would go on to the so-called last scattering surface, the source the cosmic microwave background. That region, to scale, would be situated along a line of latitude a few hundred metres from the south pole. The familiar idea that a real duplicate of Earth would be somewhere out there in a region towards which we could point is at odds with cosmology.

In what follows I shall assume that there might be any number of Hubble spheres isomorphic to the one we would conventionally think of ourselves as inhabiting. I'll call those isomorphic finite environments parallel worlds. Since parallel worlds are congruent each object in one will have corresponding objects at corresponding places in each of the other parallel worlds of that type and I'll call those corresponding objects parallel counterparts.

4 Concrete Sets

If your mind spanned a multiplicity of parallel worlds a single green apple you observe should somehow be a composite of a number of parallel counterparts, each in a distinct parallel world. A similar thought would apply to any other object you observe in your environment and so every environmental object would somehow be a composite of a number of parallel counterparts equal to the number of parallel worlds you spanned. Thus there would be an aspect of your observed environment which could be called environmental cardinality. If there were two parallel worlds the cardinality of your observed environment would be 2, and so on.

But what sort of composite object could such an observed apple be? In Tweedle's case the idea was that he was unaware of his twin-Earthly predicament and I introduced the 't' notation so that we could interpret him as referring to pairs of corresponding objects in twinned worlds. But in imagining ourselves to be in Tweedle's shoes we have to imagine a pair of apples somehow constituting a single

apple. So the metaphysical challenge posed by trying to make sense of the idea that we could be in a predicament like Tweedle's is that of understanding how an apple in our environment could be itself composed of a plurality of apples which are parallel counterparts. It would appear to be out of the question that an apple could be identified with an aggregate of apples. A proper part of an apple cannot be a whole apple and if an apple has a mass of four ounces an aggregate of two such must have a mass of half a pound. So the only option available would seem to be to propose that an apple could be a set.

The idea that environmental objects could be sets is not new. It was first proposed, so far as I know, by Willard Van Orman Quine (1969) pp. 31-32. He suggested that any particular could be construed as a self-membered singleton set. So could our composite apple be a multipleton set? If we envisage ourselves as spanning a number of parallel worlds we have to allow for the situation where all but one of those worlds is instantly annihilated, leaving us inhabiting a single world confronting an apple which does not have many parallel world counterparts as components. So if apples are to be sets we want any such remaining apple to be a set and Quine's proposal is the obvious candidate for being that set. So where we span many parallel worlds we want a single observed apple in our environment to be a set which takes singleton set apples as elements.

Further, if we allow that any permutation of parallel worlds which we span might be annihilated at any moment we want the remaining perceived apple to be observationally unchanged. There seems to be no reason to suppose that environmental cardinality is a property which is observable by a subject who spans many parallel worlds so cardinality need not be included as an observable property. That allows us to propose that where an apple is a multipleton set it has the same observable properties as all of its subsets. Every subset of the apple can then constitute an apple for an observer spanning just those parallel worlds which each include an element of that subset.

However, this does not allow for objects like the Oxbridge duckrabbit. The idea of contained anisomorphism between parallel worlds is required for the semantic externalist Twin Earth thought experiments and the aim here is not to challenge those thought experiments on the basis of rejecting the causal isolation required to envisage

natural kinds with different constitutions occupying corresponding roles in otherwise parallel worlds.

Suppose that we were in a situation like Tweedle's, spanning two parallel worlds, and we were faced with a box with a causally isolated interior. The idea so far is that that box would be a doubleton set with two singleton set elements which are boxes which are parallel world counterparts. But if one of the singleton boxes contained a duck and the other contained a rabbit it would not be true that all the subsets of our observed box had the same observable properties, because ducks and rabbits have different observable properties so the boxes would have different observable properties as regards their contents.

But there's no reason to assume that the doubleton box has *defined* observable properties for its contents where the contents of its singleton subset boxes differ. We have seen that it's impossible for a single observer to describe the contents of such a box. Any attempt at observation will give rise to personal fission into observers describing different contents. However, when Tweedledum describes his duck and Tweedledee describes his rabbit their descriptions may tally in some respects. The duck and the rabbit may both weigh two pounds, for instance, in which case the anisomorphic contents of the Oxbridge t-box would have both had the same mass. So this suggests that we can allow observable properties to be defined for the anisomorphic contents of objects which are parallel world counterparts so long as the those properties are the same for each of the anisomorphs.

These reflections lead to the following proposal for a metaphysics which is compatible with understanding our predicament as being that of single observers who span a multiplicity of parallel worlds:

Concrete Sets

Any environmental object *O* is a set of self-membered singletons which are parallel world counterparts and *O* possesses some observable property *F* if and only if all the subsets of *O* possess *F*.

It's necessary to note straight away that Concrete Sets is incompatible with David Lewis's Main Thesis in his *Parts of Classes* which states: 'The parts of a class are all and only its subclasses' (1991) p.7. An apple in our environment has parts such as

peel and pips which are not its subclasses according to Concrete Sets and so, sad to say, Lewis's idea is not for us, nice as it is.

According to Concrete Sets any object in our environment is a set of particulars which are parallel world counterparts. If we inhabit just a single world this idea reduces to Quine's proposal that particulars are self-membered singletons. If we are subjects who span a multiplicity of parallel worlds the particulars which are the elements of any given object in our environment will be isomorphic so long as no internal anisomorphism is sustained by causal isolation. But in that case it would be in virtue of what that those isomorphic objects would be numerically distinct? On our traditional understanding we imagine that it's possible for objects to be qualitatively identical and numerically distinct by being spatiotemporally separated.

If we're faced with a multipleton apple it would seem that all the singleton apples which are its elements are somehow in the same place. But if we allow that spatial regions can be subject to the Concrete Sets rule this problem is solved. If our environmental cardinality were greater than 1 the spatial region in our environment containing an apple would be a multipleton set of singleton spatial regions each of which contained an apple which would be an element of the observed apple. Since the singleton spatial regions would each be in a distinct Hubble sphere there would, as I mentioned earlier, be no problem with the idea of there being a spatial separation between them and so of the singleton apples they contained.

Given Concrete Sets it appears to be possible to coherently describe our everyday predicament as similar to Tweedle's, thus making available the challenge to semantic externalism which I've described. However, my experience is that when people are introduced to the unitary interpretation of doppelgangers they find it hilariously surreal. So it's necessary to unpack what might be at the root of this reaction and show that it is not founded on any sound reasons.

5 Qualitative Mind

The unitary interpretation of doppelgangers involves the idea that individual mental lives are not particulars, they are purely qualitative in the sense that they are instantiated universals rather than particulars. In which case, if we are to believe that material processes instantiate mental lives, isomorphic processes may instantiate one

and the same mental life just as they instantiate one and the same physical form. A particular mental life will be instantiated if there is an instance of it somewhere in the universe and multiple instances, multiple doppelgangers, do not multiply the number of mental lives; there is just a single mental life which is multiply instantiated. The existence of any one of the doppelgangers would suffice for the relevant mental life to be instantiated. Particular mental lives may be universals which are instantiated by cerebral processes or perhaps other sorts of physical phenomena yet to be discovered.

This is such a simple materialist thesis about the nature of mind that it's surprising that it seems to have been entirely neglected. What underpins the unquestioned belief that a mental life is a particular rather than a universal? There is much at stake here. Rejection of what Jacques Bouveresse has dubbed *le mythe de l'intériorité* (1976) has come to dominate both the analytic and continental traditions in Western philosophy. Such solid acceptance of a philosophical stance had better be based on more than a generally unquestioned intuition. So, why is it that the unitary interpretation of doppelgangers seems so surreal? I think several possible reasons can be identified all of which can now be challenged.

A first and quickly dismissed objection is this. How can two doppelgangers zillions of lightyears apart whose simultaneity we know, from Special Relativity, is entirely relative to an inertial frame, how can they share a single mind? Ought there not be some sort of causal link for that to be possible? But of course no causal connection is required for those doppelgangers to instantiate one and the same pattern of neural processes so this appeal to causality cannot be an objection to the doppelgangers instantiating one and the same mental life.

A second objection may seek to link the common assumption that a mental life is a particular with the idea that the environmental objects to which we indexically refer are particulars. I said that if we considered ourselves to be in a predicament similar to Tweedle's a concept of environmental cardinality is introduced, but that's not strictly true. It's rather that we have always assumed that the value of environmental cardinality is necessarily 1. That leads directly to the presumption that when we describe an environmental object such as an apple we are assigning properties to a particular. But Concrete Sets challenges that idea. If the environmental cardinality is greater than 1 objects in our environment become multipleton sets. Indexical reference to such objects then no longer picks out particulars to which we

assign properties; what is picked out by indexical reference is multipleton sets of particulars which are parallel counterparts and, according to Concrete Sets, those multipleton referents have the same observable properties as the particulars which are their elements. The *possibility* of objects in our environment not being particulars is part and parcel of the idea that mental lives *are* not particulars.

A third possible objection arises from our conception of space. When Patrick Moore talked of having a doppelganger on another planet he was no doubt thinking in terms of an endless continuous space with a duplicate Earth being far, far away. The unitary interpretation seems to imply that space can fold back on itself in a mindbending way and that could seem to contradict the very idea of continuously infinite space. But we have seen how modern cosmology confines our local environment to a causally isolated region and how the metaphysics of Concrete Sets can intelligibly construe our observed local environment as a set of isomorphic regions which appear to be “superposed”. There’s a further mindbending consequence of this which I haven’t yet mentioned. For any given pair of isomorphic Hubble spheres the environments beyond them will usually be anisomorphic. That means, according to Concrete Sets, that many physical properties will be undefined for the contents of regions of space which are outside our observed Hubble sphere if our environmental cardinality is greater than 1. Beyond the so-called event horizon which is the boundary of our observed Hubble sphere duckrabbits and other exotica would abound. A strange idea to be sure, but not unintelligible. We have seen that the observation of regions with such undefined properties will simply cause the fission of an observer into observers assigning definite properties to the regions which they observe.

But is the idea of personal fission manageable? There might seem to remain an important objection to the qualitative mind thesis and its associated unitary interpretation of doppelgangers because of the way fission is involved. Tweedle apparently underwent personal fission as a result of watching the Oxbridge t-box being t-opened by Alice and the conclusion about the stuff in William Hogarth’s puddles and gutters not being water depended on the idea that our ancestors underwent personal fission as a result of the investigations of chemists. Further, if our minds in fact span a large number or an infinitude of Hubble spheres we should expect personal fission to be continually happening in everyday life. That’s because

there would be continual divergence of the histories of Hubble spheres which were isomorphic up to a certain stage of their evolution. Divergence would arise from the expansion of the event horizons into anisomorphic environments and possibly from the occurrence of probabilistic processes within each Hubble sphere, if the standard stochastic interpretation of quantum mechanics were to be believed.

The idea of personal fission has long seemed deeply problematic. It might be argued that if we were to accept the possibility of ubiquitous personal fission then that would render rational future-directed action impossible. David Lewis has argued thus (1986, p. 209). However, there are now good reasons to believe that rational provision for the future is indeed possible for an informed subject about to undergo fission.³

In conclusion, I believe I have shown that it's plausible that sufficiently isomorphic doppelgangers instantiate a single mental life. In which case, if Putnam's Twin Earth did indeed exist it is possible that the stuff which our ancestors called water in 1750 was not H₂O, which makes it plausible that the extension of 'water' in 1750 was not all and only samples of H₂O even if Twin Earth did not exist. This thought threatens the widespread idea that some arguments for semantic externalism based on Twin World thought experiments are compelling. Those arguments have overlooked the unitary interpretation of doppelgangers and an immediate sense of implausibility does not give grounds for leaving it out of consideration. There seems to be no good reason to reject it.

³ The problem of making rational provision for the future in a world involving incessant personal fission has been a major challenge for the Everett interpretation of quantum mechanics. I have proposed a solution to that problem which involves a subject adopting an attitude to fission *as if* it were a probabilistic process (Tappenden 2011b). I have also applied a version of this analysis to the classic philosophical thought experiments on personal fission (Tappenden 2011a).

Acknowledgements

My thanks to David Papineau and Sarah Sawyer for comments on previous versions of this paper.

References

- Bouveresse, J. (1976) *Le Mythe de l'intériorité: Expérience, signification et langage privé chez Wittgenstein*. Paris, Minuit.
- Burge, T. (1979) Individualism and the Mental. *Midwest Studies in Philosophy* 4.1: 73-121.
- Lepeltier, T. (2010) *Univers parallèles*. Paris, Seuil.
- Lewis, D. (1991) *Parts of Classes*. Oxford, Blackwell.
- Lewis, D. (1986) *On the Plurality of Worlds*. Oxford, Blackwell.
- Menary, R. (ed.) (2010) *The Extended Mind*. Cambridge, Mass., MIT Press.
- Putnam, H. (1975) The meaning of 'meaning'. In K. Gunderson (ed.) *Language, Mind and Knowledge*. Minneapolis, University of Minnesota Press.
- Quine, W. V. O. (1969) *Set Theory and It's Logic*. Cambridge, Mass., Harvard.
- Russell, B. (1952) Advice to Those Who Want to Attain 80. *New York Times*, 18 May, §6: 13.
- Segal, G. (2000) *A Slim Book About Narrow Content*. Cambridge, Mass., MIT Press.
- Tappenden, P. (2011a) Expectancy and Rational Action Prior to Personal Fission. *Philosophical Studies* 153.2: 299-306.
- Tappenden, P. (2011b) Evidence and Uncertainty in Everett's Multiverse. *British Journal for the Philosophy of Science* 62.1: 99-123.
- Tegmark, M. (2007) The Multiverse Hierarchy. In B. Carr (ed.) *Universe or Multiverse?* Cambridge, CUP: 99 - 125.