

Personal Dualism and the Argument from Differential Vagueness

Paul Noordhof

Abstract: In *Causing Actions*, Pietroski defends a distinctive view of the relationship between mind and body which he calls Personal Dualism. Central to his defence is the Argument from Differential Vagueness. It moves from the claim that mental events have different vagueness of spatiotemporal boundaries from neural events to the claim that mental events are not identical to neural events. In response, I argue that this presupposes an ontological account of vagueness that there is no reason to believe in this context. I further argue that Pietroski's reasons for rejecting the possibility that mental events are vaguely constituted from neural events are inadequate. I go on to show how Pietroski's Personal Dualism is ill-equipped to deal with the problem of mental causation because of its apparently necessary appeal to *ceteris paribus* laws.

In a recent book, *Causing Actions* (2000), Paul Pietroski defends a distinctive view of the relationship between mind and body. He calls it *Personal Dualism*. This is a dualism of events deriving from considerations about the nature of persons. According to Pietroski, the concept of a person is a primitive although what it picks out is not ontologically basic [162, 164].¹ The justification for Personal Dualism seems to lie in the following theses.

- (a) Persons are the locus of freedom and they operate in the space of reasons [150].
- (b) Mental events cannot be described without bringing a thinking subject onto the scene i.e. they cannot be described impersonally. By contrast, neural events can be described impersonally [154].

¹ Pietroski (2000). Numbers in square brackets are to page numbers in this book.

- (c) Persons have a corporeal nature (are essentially embodied) [164-168].

The third element is required, according to Pietroski following P.F. Strawson, in order to provide the necessary preconditions for the application of our concept of persons, in particular, as the concept of something to which mental states should be ascribed ([160-168]; Strawson (1958); Strawson (1959), pp. 87-116). There are various standard objections to this line of thought but I won't pursue them here (e.g. Ayer (1963), pp. 116-118; Williams (1961), p. 125). For the purposes of argument, I shall assume that it can be bolstered in some way.

Instead, I want to focus on an important part of the justification for (b). It is that a certain view of the relationship between mental events and neural events cannot be correct, namely that mental events just are, or are constituted from, neural events. Pietroski calls his *Neuralism* [5-6]. He suggests that Neuralism is 'a capitulation to the idea that human actions are just one more species of impersonal occurrence' [178]. He writes:

... neuralism actually threatens the idea that reasons are causes by identifying mental events with events characterizable without reference to *persons*. Identifying reasons with impersonally characterised events threatens our view of ourselves as agents whose actions are *free* [12, italics in the original].

If Neuralism had been true, the thought appears to run, then there would have been nothing to stop us providing an impersonal description of mental events (in terms of the neural events from which they are constituted or with which they are identical).

By rejecting Neuralism, Pietroski implicitly rejects Davidson's Anomalous Monism. Davidson is quite happy to endorse Neuralism while at the same time arguing that mental states and events are reasons for a person's action. At the same time, Pietroski is unwilling to go the whole hog and adopt a Nonphysicalist Dualism. So we have the makings of an

intermediate position and, thereby, an interesting debate. The key to assessing Pietroski's attempt to steer his way between these two positions is his endorsement of a refined version of an argument in Jennifer Hornsby's work (Hornsby (1981)).² We may dub it the *Argument from Differential Vagueness*. It is his direct argument against Neuralism. In what follows, I shall be principally concerned with this argument and what follows from accepting its conclusion.

My discussion will proceed as follows. In the first section, I will outline two attempts to capture the characteristics of persons Pietroski emphasises, in particular, the idea that mental events cannot be characterised without bringing the thinking subject on the scene. I contrast this with Pietroski's own position. The role of the Argument from Differential Vagueness in its defence will become clearer. I will then turn to the argument itself and explain why it doesn't work. In the final section, I will outline why this gives rise to problems regarding Pietroski's treatment of mental causation.

1. Personal Dualism, Anomalous Monism and Emergent Property Dualism

To recall, Pietroski asserts that

- (b) Mental events cannot be described without bringing a thinking subject onto the scene i.e. they cannot be described impersonally. By contrast, neural events can be described impersonally [154].

At first glance, those convinced that mental events cannot be impersonally described do not seem forced to adopt Pietroski's position and reject Neuralism. Indeed, it might look as if (b) is only true if a rather limited understanding of neural events is at work. Some neural events might be impersonally described, but what of *the activation of S's*

² In a review of Pietroski's book (2002), Hornsby writes that her original argument rests on a different premise, but her remarks are rather brief. In any event, I think that Pietroski's argument is an understandable development of some of the reasoning in Hornsby (1981).

brain at location L or S suffering brain damage in the parietal lobe? It would seem that these are perfectly acceptable descriptions of neural events, and yet, given that the reference to S is a reference to a person, it is not obvious that they can be described in some other way. If we envisage that this event is identified to explain why a particular person is subject to epileptic fits or visual deficits, we are unlikely to suppose that the very same token events may occur in somebody else's brain. The kind of events we identify—and thereby, the essential properties they possess—is determined by the context of explanation. It seems implausible to assert that there are no contexts of explanation in which neural events might be personally identified, in particular, those which have significant ramifications for persons. I think we should be wary of adopting too sharp a distinction between what must be personally described and explained and what is part of the impersonal and objective world. It promises to fail to do justice to the way in which we are complex animals with a physical constitution and a consequent mental life. But, if this is right, then mental events may be neural events after all, just those which cannot be identified without reference to persons. Nothing rules this out.

I guess that Pietroski could insist that we may describe the neural events just mentioned in a way which does not refer to the person to whom they occur. But, to that extent, it seems that we could insist that a person's mental event could be specified without reference to a person, as in for instance, 'the yearning grew stronger'. Pietroski can claim that such a description does not explicitly pick out the yearning. It has an implicit context limiting the domain of quantification for the definite description so that there is a unique event which counts as the yearning. But even descriptions which make the individuation of mental events explicit don't seem to make essential reference to persons, as for instance, 'the foolish fear which occurred at 1 am at map reference x, y, z in the darkened room'. This description has slightly comic effect but comedy does not imply indescribability.

Other ways of capturing the distinctive characteristics of persons don't place the same emphasis on the impossibility of providing an impersonal

description through the rejection of Neuralism. One option is that of a Nonphysicalist Dualism, such as Emergent Property Dualism. According to the Emergent Property Dualist, there are personal properties over and above neural properties. So it is not just a matter of these events having an essential personal description by being events with a particular person as subject. Rather some of their constitutive properties are *sui generis* personal properties. These personal properties could be characterised in no other way but in terms of the perspective of S and subjects like S. They are related to neural properties by emergent psycho-physical laws.

Unfortunately, adopting this line is not immediately in tension with Neuralism. That would depend upon whether events were coarsely individuated so that they might have more than one distinct constitutive property or finely individuated in which case they could not. If events are coarsely individuated, Neuralism could still be true. Neural events would just, also, possess these personal properties. In which case, Emergent Property Dualism would not imply (b). Mental events could be described impersonally by mentioning their neural properties. It is just that they could not be so described if they are to be described in their mental aspects.

If events are finely individuated, then, according to the Emergent Property Dualist, mental events are distinct from neural events. In which case, Neuralism would be false. However, we could still pick out the neural events with which the personal events are correlated as a result of psychophysical law. The relevance of this point will be clearer in the next section. For now, we may note that, from the perspective of Emergent Property Dualism, the truth or falsity of Neuralism would seem far less interesting. It would all turn on whether events should be individuated thinly or thickly. All that matters is whether there are personal properties.

Davidson, like Pietroski, claims that the personal point of view cannot be abandoned without changing the subject and no longer talking about minds (see Davidson (1970), pp. 216, 223). Davidson does not quite put the emphasis on persons in the way that Pietroski does though. Instead

the key idea is that we adopt a scheme of interpretation by which we make sense of people's behaviour through ascribing beliefs and desires to them.³ Nevertheless, he need not reject Pietroski's stress on the importance of persons being taken as primitive. There is no radical disagreement here. As I have already mentioned, the disagreement lies in Davidson's espousal of Neuralism. Davidson would not view the identification of mental events with neural events as giving up on the personal perspective or only playing lip service to it. The perspective we adopt in interpreting people's behaviour is just another way of identifying neural events. Keeping the personal in view is refusing to give up on this perspective⁴.

There is one pretty straightforward reason why mental events are unlikely to be identical to neural events which I would like to get out of the way so we can get to the heart of the issue. Patterns recognised at the personal level need not be captured at the neural level. Neurophysiologists might be largely uninterested in the kind of things that are singled out at the personal level. If *neural* events are those which are categorised by the neurophysiologist, then mental events need not be neural events. Nevertheless it might be thought that they are constituted from neural events. That's still enough for Neuralism to be true. I presume that Davidson would be quite happy with the idea that the kind of neural states with which mental states are identical are not ones the neuroscientist would pick out. This would be part and parcel of the idea that mental categories find 'no echo' in physical theory. Nor, once we get this dimension into view, does the disagreement between Davidson and Pietroski disappear. Pietroski's rejection of Neuralism is also a rejection of the claim that mental events are constituted from neural events. The Argument from Differential Vagueness is supposed to work even against this claim.

Pietroski's reservation with Davidson's picture is that

3 For further discussion of, and reservations about this approach, see Noordhof (forthcoming).

4 For further discussion of how to understand this, see Noordhof (1999b), pp. 313-315.

it's not just the personal *perspective* I want. Davidson provides that. I'm interested in a picture according to which human actions *are* (by their nature) personal contributions—and not just causes that interpreters can 'describe as' the contributions of persons. I'm interested in a picture according to which human actions are exercises of freedom—and not just events that observers can 'think of' as exercises of freedom. I don't deny that one can build the personal perspective into an interpretationalist conception of the mental. But since this invites the worry that some events are just labelled as mental, either for certain pragmatic purposes or simply as a sop to our self-conception, I want to explore a position that clearly isn't just labelling certain physical events as mental—even at the cost of needing an alternative account of mental causation (Pietroski (2001), section 3).

In a nutshell, Pietroski is worried that Davidson's position is just a matter of labelling certain events as mental events because of the adoption of a certain scheme of interpretation (hard as this scheme is to give up). Pietroski wants reality to reflect the personal. He thinks that, on Davidson's view, it does not.

Pietroski's emphasis on reality reflecting the nature of persons might be thought to indicate that he finds Emergent Property Dualism more congenial than the differences I noted above would suggest. Nevertheless, it is clear that Pietroski rejects Emergent Property Dualism for an entirely familiar reason, namely, the threat of systematic overdetermination given that we have sufficient causes at the physical level (Pietroski (2002), footnote 6). His own version of Dualism is meant to steer between these two paths. It must somehow explain how the personal constitutes a genuine part of reality even though it cannot be seen as constituted by physical events and objects. Yet, at the same time, it must not transform itself into an Emergent Property Dualism. The Argument from Differential Vagueness is vital to this. Pietroski relies upon a supervenience claim to establish his Physicalist credentials [179]. The Argument from Differential Vagueness explains how this is not

sufficient for us to see the mental as merely constituted from the physical. It provides the basis for his rejection of Neuralism and support for (b). In the next section, we will scrutinise this argument.

2. The Argument from Differential Vagueness

In Pietroski's hands, the Argument from Differential Vagueness runs as follows.

- (1) If an event E is identical with an event F, then they have the same vagueness of spatiotemporal boundaries.
- (2) Mental events have different vagueness of spatiotemporal boundaries from any neural events.

Therefore

- (3) Mental events are not identical with (or even vaguely constituted from) neural events [173].⁵

I will not challenge premise (1) which seems to be an uncontentious application of Leibniz's law. Even if you don't believe that vague events (or indeed vague objects of any kind) exist, you can endorse premise (1). It's just that the same vagueness of spatiotemporal boundaries will be no vagueness in spatiotemporal boundaries. Instead, I shall focus on (2) and in the inference from (1) and (2) to (3).

Pietroski recognises that (2) is open to challenge. For one thing, we might just be currently ignorant of the precise boundaries of mental events. We could discover what they were by finding out whether we would still have a mental event of a certain type if certain kinds of neural events were absent. Slowly but surely we could zero in on what would be the most plausible candidate kinds of neural events with which particular mental events might be token identical. Our narrowing down to particular types of neural events may differ from person to person, creature to creature, species to species. Nevertheless, we might still end

⁵ See Hornsby (1981), pp. 66-73.

up with good grounds for putting forward a particular token identity from our study of an individual over time.

Pietroski's view of such a proposal is captured in the following passage. He writes

... it would be question-begging for neuralists to assume that each mental event has the spatiotemporal boundaries of some biochemical event, on the grounds that each mental event *is* some biochemical event. I see no independent reason for supposing that tryings [his candidate mental event] have much sharper boundaries than initial reasons suggest [175].

This seems a little quick. Our ignorance in this area appears quite plausible. Our understanding of mental events has grown up from two sources. First, our introspective appreciation of our own mental goings on. Second, our observation and attempt to make sense of the behaviour of others. Neither requires obtaining precise information about the spatiotemporal location of mental events. In the case of introspection, there is a fairly clear sense in which it does not provide a spatial presentation of mental events at all. So the Neuralist has an explanation for why it is not presently possible to say anything more precise. Neuralists are not just assuming in the face of counterevidence that the boundaries of mental events must be precise.

Although this flat-footed response has some plausibility, it relies upon reinterpreting Pietroski's intuition that mental events have vague boundaries in the way that other objects have vague boundaries. He writes:

... just as one can give a 'too specific' answer to the question of where a hurricane occurred, so one can give a too specific answer to the question of where a mental event occurred (Pietroski (2001), section 3).

If the situation were just a matter of ignorance in the way I specified, then a more specific answer as to where a mental event occurred would

eventually seem perfectly appropriate to us. It's just that we currently do not have good grounds for supposing that a particular mental event occurred precisely one inch in from the skull half an inch up from the ear. The mind might boggle that we would ever be in such a position.

For this reason, let us consider what would happen if we took Pietroski's claim more seriously and assimilated mental events to cases of apparently vague objects such as tornados, mountains and douds. There still appear to be problems with Pietroski's line of argument. It commits him to an ontological understanding of vagueness. That means that he should have reasons to reject the intuitive claim that the world itself is not vague. Vagueness is merely a consequence of our representation of the world.

The latter thought could be cashed out via an epistemic or semantic account of the vagueness. Now is not the time to try to alight upon the correct theory of vagueness. Nevertheless, I think it is instructive to consider for a moment in outline how these two types of theories might treat the case of mental events. Consider Timothy Williamson's famous epistemic theory. In a nutshell, Williamson's claim is that vague words are those for which we are unable to distinguish the concepts they actually express from other very similar possible concepts of things with different precise boundaries (Williamson (1994), pp. 226-237). Applied to the present case, the thought would be that we are unable to distinguish our concept of a particular mental event which includes spacetime point x, y, z, t from a possible concept of another mental event (or some very similar entity) which includes spacetime point x_1, y, z, t . Williamson suggests that we may explain our conviction that there isn't a precise boundary by noting that, if his view is correct, we cannot conceive of a way in which we would know, or even have reason to believe, that the boundary is drawn in such and such a place.⁶ Presumably that is precisely

⁶ It seems that some qualification is needed to this. Suppose there is an oracle who is capable of distinguishing between the various concepts of mental events and that we have reason to suppose that the oracle is reliable on these matters. Then we would have reason to believe that the boundary should be draw in such and such a place (see Sainsbury

what we would need to recognise that a particular boundary is not arbitrary.

Given how our mental concepts have developed, it seems all too plausible that the epistemic theory may be correct about them even if it does not carry conviction in every case of vagueness. There might be some features of our practice that make a concept of a particular mental event actually pick out an event with certain spatiotemporal boundaries rather than others. Yet we may be insensitive to what these features are.

The plausibility of Williamson's position partly relies upon the plausibility of his claim that we could know the precise concept expressed by a particular word or phrase even though we cannot distinguish between that concept and some other concepts that might be expressed by the word or phrase. If Williamson is wrong, then a semantic account of vagueness will be more plausible. Nevertheless, this need have no damaging ontological implications. Applied to our present case, the thought would be that the practice underpinning mental concepts leaves it unresolved exactly where in the brain a particular mental event is. There are a number of marginally different possibilities. Note this is not to preclude that a mental event may consist partly in relations to the environment. I'm just focussing on the brain side for ease of discussion. In any event, the remarks I made about the two venues in which mental concepts have developed—introspection and the interpretation of behaviour—makes it quite conceivable that our mental concepts should suffer this lack of determination.

The plausibility of a semantic theory of vagueness awaits a detailed account of the semantics of vague concepts. Significant problems have been raised for various attempts to provide this, for example, supervaluationism, degree theory and so forth (Sainsbury (1991), pp. 11-13; Sainsbury (1995a), pp. 33-46; Williamson (1994), pp. 96-164). If these prove to be insurmountable and Mark Sainsbury is right in his

(1995b), pp. 598-599). One answer would be to insist that in this case the experience of arbitrariness is dependent upon being unable to draw the distinction by ourselves.

conjecture that we will have to take the notion of a vague object as basic in the provision of such a semantics, then it may seem as if Pietroski's argument will be rather more forceful. Only the epistemic theory of vagueness would appear to stand between Pietroski and the conclusion he wants. Many find the epistemic theory simply incredible. However, I doubt that only the epistemic theory stands between Pietroski and his desired conclusion. This brings me to the validity of the Argument from Differential Vagueness.

The natural thought to have when faced with Pietroski's argument is that, even if mental events aren't neural events, they are vaguely constituted from them. This would sit well with the point made at the end of the previous section. Neuroscientists are unlikely to be interested in the categories delineated by our everyday concepts of mental events. So their demarcations will be different. But that doesn't stop neural events vaguely constituting mental events. If mental events are vaguely constituted from neural events, then we have vagueness in the world. We are not adopting a merely epistemic or semantic account of vagueness. We are recognising the kind of thing that Sainsbury suggests is necessary to provide the semantics of vague concepts. Pietroski thinks that, if he concedes that mental events are vaguely constituted from neural events, then the Neuralists have won.

Pietroski cites two reasons why this option can be closed off [177]. The first is that Gareth Evans and David Lewis, in his note on Evans' article, have shown that the suggestion is of dubious coherence (Evans (1978), pp. 176-177; Lewis (1988)). This a peculiar reason for Pietroski to give. Evans presents an argument purporting to show that there could not be vague objects. Vagueness, according to this view, would be an artifact of our ignorance or semantic indeterminacy. But if this is right, then Pietroski's Argument from Differential Vagueness does not work. For all that he has said, mental events are identical with or constituted from neural events. They cannot differ in vagueness of spatiotemporal boundaries because vagueness is not a real feature of the world. So

Pietroski cannot cite this as a reason on pain of losing the Argument from Differential Vagueness.

His second reason for rejecting the idea of vague constitution is that it could not answer our worries about mental causation. This point is obviously hard to assess outside of an extensive discussion of mental causation and the account we should provide of it.⁷ For the present purpose, we should just note the following. Taking mental events to be vaguely constituted from neural events puts them on a par with mountains and their rocklike constituents. It seems to me that there is considerable *prima facie* plausibility in holding that mountains are efficacious. If so, then the comparison with mountains would show that there was no greater problem with the efficacy of mental events than there was with the efficacy of mountains. This would be some step forward. As an *ad hominem* point, it is also worth noting that there is no reason to suppose that Pietroski's favoured account of mental causation would fail to apply to mental events so understood. According to Pietroski, mental properties are efficacious if they figure in *ceteris paribus* laws in the antecedent. Nothing in principle seems to rule this out in the case of vaguely constituted events. I shall discuss Pietroski's account in more detail in the next section.

There is one component to Pietroski's discussion which I have not yet mentioned. Pietroski, following Hornsby, suggests that, if there is considerable vagueness in the spatiotemporal boundaries of mental events, then there will be a number of different possible neural events each of which would be as good a candidate for being identical to the mental event. Here Pietroski and Hornsby envisage a Benacerraf-style argument. Paul Benacerraf argued that the number two cannot be identical to both $[\emptyset, [\emptyset]]$ and $[[\emptyset]]$ because $[\emptyset, [\emptyset]] \neq [[\emptyset]]$. Since there are no grounds for choosing between them, the number two is not identical with either of them (Benacerraf (1965), pp. 278-285). By the

⁷ See Noordhof (1999b) for my preferred approach, which would explain how vaguely constituted mental events were efficacious.

same token, Hornsby and Pietroski argue that token mental event m cannot be identical to token neural events n_1 and n_2 because events $n_1 \sqcap n_2$. Since, there are no grounds for choosing between them, m is not identical with either of them [176].

The idea that mental events may be vaguely constituted from neural events points to the disanalogy between the two cases. Suppose that Benacerraf was right in arguing that there were no considerations which favoured $2 = [\emptyset, [\emptyset]]$ over $2 = [[\emptyset]]$ or vice versa. Then it seems that it would be arbitrary to take up one of them. But that is not the situation with regard to our vaguely constituted mental event. It might be arbitrary to pick either n_1 or n_2 as identical with m . But that is not the proposal. The claim is not that m is identical with some neural event n_i . The claim is rather that m is vaguely constituted from neural events. There is no arbitrariness in saying that.

I should also point out that the suggestion I have made is not akin to the idea that mental events are identical to neural events which are, themselves, constituted from other neural events. I'm not claiming that any old collection of neural events can be taken to compose a larger scale neural event. The remarks I made about the different interests of neuroscientists and those ascribing beliefs, desires and other mental events should have made this clear. Nor am I claiming that mental events are just mereological sums of neural events. We are familiar with the thought that mountains are not just mereological sums of rocks. Mountains could persist if made up of slightly different rocks and would cease to exist if broken up into its constituent rocks. The mereological sum does not have these properties. I do not intend to challenge this picture here.⁸ Even if it is endorsed, there still is a sense in which rocks constitute mountains. I claim that things are no different in the case of

⁸ Hornsby is inclined to emphasise causal considerations in explaining why mental events cannot be understood as merely mereological sums of neural events (Hornsby (1985), pp. 52-57). I hold that these considerations are inconclusive (see Noordhof (1999a), Noordhof (1999b) for the basis for this view). In spite of this, I think her general point is correct.

mental events. There is still a sense in which neural events may vaguely constitute mental events.⁹

The possibility of vague constitution underlines how ill-suited reflections drawn from vagueness are to revealing the independence and importance of the personal. Suppose we deprecated the vagueness of the spatiotemporal boundaries of mental events of type M and decided to make them more precise. Suppose, for the kind of reasons mentioned earlier, we thought that mental events could not be identical with neural events. Instead, we stipulate that a particular mental event, m , is constituted from n_1 , o_1 , and p_1 rather than n_2 , o_2 and p_2 . The result gives a precise spatiotemporal location of m . Pietroski thinks that such a proposal would be mistaken on the grounds that would detract from the goal of personal ascriptions. He writes:

Action descriptions, however, purport to be descriptions of a person's contributions to history. And it cannot be a matter of mere decision which event is Nora's action of raising her hand. To say otherwise—that we can simply stipulate (from a third person perspective) which neural event is Nora's trying—is tantamount to saying that no event is Nora's action ... Our actions are tryings: and if these events are not biochemical, we should aim to understand this fact about ourselves [178].

But the stipulation would seem to lose us nothing. It is a move from taking mental events to be vaguely constituted from neural events to taking a mental event to be precisely constituted from neural events. Thinking of mental events as vaguely constituted from neural events captures no deep insight into the nature of action which precision would lose. So if there were something to be said for adopting a more precise

9 Hornsby claims that, if an event e_1 is located in the spatiotemporal region in which an event e_2 occurs, it does not follow that e_1 is part of e_2 (Hornsby (1985), pp. 59-60). Perhaps not, though she provides no clear cases. Nevertheless, if, in addition, the occurrence of e_2 noncausally depends upon the occurrence of e_1 , I think it follows that e_2 is partly constituted by e_1 . I take it that everybody can agree that the antecedent is met when e_2 is a mental event and e_1 a neural event.

terminology, Pietroski has not identified a consideration to counterbalance it.

Pietroski seems to feel that somehow, if we identify mental events with neural events, or take mental events to be precisely constituted from neural events, we will fail to capture the personal. But it is hard to see the basis for this. Either the personal is compatible with the impersonal or it is not. If it is, then an arbitrary precisification will still capture the generalisations and patterns we perceive at the personal level. There are just a number of ways in which this might have been done. If it is not compatible, then it doesn't matter whether things are precise or vague. There is still no way of fitting personal reasons in an impersonal world. Either way, it seems that there is nothing to lose. As we shall shortly see, there is something to gain.

3. Mental Causation and *Ceteris paribus* Laws

Physicalists face their own problem of mental causation. The standard way to pose it is to note that there is strong evidence that there are sufficient physical causes for any event. To avoid overdetermination, we must explain how this is compatible with there also being mental causes of some events. One way to do this is to identify the mental causes with physical events. This option is not open to Pietroski. As I have already noted, Pietroski's commitment to Physicalism leads him to claim that mental properties globally supervene upon physical properties. But because he denies Neuralism, he sticks at that. He thinks that we can't trace more specific relations between instantiations of mental properties and instantiations of physical properties. This is the kind of information that might help to explain how mental events can be efficacious in virtue of being intimately related to uncontroversially efficacious physical events. Since it is unavailable to him, he must provide some other basis for the claim that mental events are efficacious. The focus has to be on what holds at the mental level. In this context, appeal to *ceteris paribus* laws is bound to seem attractive. Pietroski takes up their cause. He argues

that mental events are efficacious because their interactions are governed by such laws.

Before I consider the details of Pietroski's position, it is worth pausing a moment on his relatively unreflective appeal to *ceteris paribus* laws. Davidson, Hornsby and others with similar concerns to Pietroski have denied that it is appropriate to look for laws, *ceteris paribus* or otherwise, to capture the explanatory force of attributions of mental events (e.g. Davidson (1970), pp. 207-209, 225; Hornsby (1981), pp. 73-74; Hornsby (1985), pp. 60-62). There are various dimensions to the rejection of an appeal to laws. I cannot hope to do justice to them here. Many, though, stem from the idea that explanations appealing to laws cannot capture the thought that agents are intelligible because they act upon reasons. If *ceteris paribus* laws did a good job of explaining how mental events were efficacious, then this concern would be worthy of serious examination. However, there are independent reasons for thinking that *ceteris paribus* laws do not explain how mental events are efficacious. So we don't need to go into the further matter.

To try to establish this point, let me turn to the details of Pietroski's position. Pietroski claims that causation is a relation between events, explanation a relation between facts [89]. Facts are true thoughts. They are abstract entities [90, 94-95]. Facts are individuated by the way subjects think of things [90]. Pietroski provides the following sufficient condition on explanation:

F1 explains F2 if F1 is an instance of the fact that an event of type T1 occurred; F2 is an instance of the fact that an event of type T2 occurred; and *ceteris paribus* if a T1 event occurs, then a T2 event occurs [10].

He characterises *ceteris paribus* laws as follows.

‘cp[$\forall x (Fx \rightarrow \exists y Gy)$]’ is non-vacuous if

- (i) ‘F’ and ‘G’ are otherwise nomological; and

- (ii) $\forall x (Fx \rightarrow \text{either } \exists y Gy \text{ or there is an interferer INT such that INT explains why } \neg \exists y Gy \text{ despite } Fx)$; and
- (iii) either some instance of F explains some instance of G or some instance of F in conjunction with an interferer INT explains some fact from which it follows that $\neg \exists y Gy$ [126].

The point of the first clause is to avoid counterexamples involving gerrymandered predicates. The heart of the proposal is the second clause. When things aren't equal there is an interferer that explains why there is no G. The interferer must be independently explanatory in other contexts. It cannot be tailor-made to explain why an otherwise faltering law does not face a counter instance [125]. The third clause just insists that the law does some explanatory work. There must be instances that obey it [127].

The standard worry about such an account of explanation is that merely appealing to laws of any kind plus matters of particular fact can't capture the full force of the explanatory relation. In addition, we need to identify something as a cause. Without this, things will be classified as explanations which intuitively are not. As Pietroski acknowledges, this would undermine his approach but thinks his account avoids the problem [142-143].

One case, which Pietroski discusses, is the pendulum law.

$$(PL) \quad P = 2\pi\sqrt{(L/g)}$$

where P is the period of the pendulum, L is the length of the string and g a gravitational constant. Intuitively, the length of the string explains the period of the pendulum but not vice versa. Yet the law backs inferences in either direction [141-142]. Pietroski, following Robert Cummins, draws a distinction between the law as stated and it being used as part of a transition theory: 'a theory that aims to explain certain changes of state in objects of the theory's domain' [142] (see Cummins (1983), pp. 1-14). Pietroski dubs laws which take this form: *transition laws*. In the present case, candidate transition laws relate changes in period to

changes in length and vice versa. Pietroski notes that these candidate transition laws are *ceteris paribus*. They would not govern changes in length, or in period, if, when the change is made, an additional source of drag is introduced [143]. Here, he argues, we find the required explanatory asymmetry. Introducing the source of drag explains why the period does not change in the predicted way as the length changes. However it does not explain why the length does not change when the period is changed [144]. So the candidate transition law which relates changes in period to changes in length fails clause (i) of his account of non-vacuous *ceteris paribus* laws. There is no interferer of the appropriate kind.

One preliminary concern I have about this suggestion is over how Pietroski cashes out ‘explanation’ in the claims he has just made about the interfering drag. What stops there being a *ceteris paribus* law relating introductions of drag with failure of the length to be different? We can imagine a scientist legitimately inferring as follows. The period has changed. So one would expect the length of the pendulum string to be different. It is not. So there must be some undetected source of drag which would explain why the pendulum string need not be different in length. Pietroski’s answer will be to identify another interferer which putatively shows how the candidate *ceteris paribus* law involving drag fails clause (ii) of his account of non-vacuous *ceteris paribus* laws. But then the question can arise all over again. To avoid a vicious regress, Pietroski must deny that we should understand the explanatory character of at least one relevant interferer in terms of *ceteris paribus* laws. He more or less acknowledges this feature of his account although not precisely concerning the issue I have raised [146]. He also appears to be committed to claiming that *ceteris paribus* laws cannot be developed into strict laws otherwise his method of characterising the asymmetry would have limited application. I shall accept these moves for the sake of argument.

The more significant problem Pietroski’s proposal faces can be illustrated by a case of Gabriel Segal and Elliott Sober’s (Segal and Sober

(1991), pp. 4-10). Segal and Sober claim that it is a *ceteris paribus* law that, if blue-eyed humans successfully mate, their off-spring will be blue-eyed. Yet being blue-eyed is not causally relevant to producing blue-eyed children. It is something to do with genetics. Pietroski's fix is to claim that parents with eyes cosmetically altered through surgery to the colour blue will provide counterinstances to this putative law which cannot be explained away by interfering factors. Hence it is not a *ceteris paribus* law at all [144]. Unfortunately, I don't think this works. First off, we could reformulate the *ceteris paribus* law to mention humans with congenitally blue eyes rather than merely blue eyes. I don't think there is any reason to suppose that having congenitally blue eyes is a gerrymandered predicate. If we reformulate the law in the way suggested, then we lose the counterinstances that allegedly can't be explained away by interferers. I guess Pietroski might insist that he could make a similar manoeuvre with regard to the reformulated putative *ceteris paribus* law. So let me provide a more general reason why pursuing the debate in this way is unlikely to be productive. Pietroski's response relies upon the fact that there may be more than one way in which the antecedent is realised. Yet, it is hard to see why this should settle whether or not the type of event mentioned in the antecedent is a cause. Suppose it's just impossible to surgically produce precisely the right tint. Still we would not say being blue eyed is a cause of blue-eyed children.

There is another problem. The reason why humans with surgically produced blue eyes fail to give birth to blue-eyed babies is that they have the wrong genetic material. Pietroski needs to explain why this wouldn't count as an interfering factor. If it could, then Pietroski's theory would yield the verdict that Segal and Sober's original candidate *ceteris paribus* law is, in fact, a *ceteris paribus* law. In his response to the paper in which I originally made this point, Pietroski appeals to the idea that identification of interfering factors come with defeasible assumptions about what counts as a normal case (Noordhof (2001), section. 3; Pietroski (2001), section 4, footnote 9). Interferers are there to explain the abnormal [124-131]. I don't see how this undermines the point. It is

plausible that having the wrong genetic material would qualify as an interferer which explains the departure from the normal case. A normal case of being blue eyed arises from a person's genetic material. A mother and father with surgically produced blue eyes engendering a brown-eyed child counts as an abnormal instance, explained by the parents having the wrong genetic material.

Of course it is possible that Pietroski may be able to add to his account of *ceteris paribus* laws so that he can deal with these problems. I question whether this is the right way to go. We seem happy to allow *ceteris paribus* laws without causality between antecedent and consequent. There is good reason for this. We are interested in noting certain patterns either for the purposes of rough prediction or because we want to investigate them further. Segal and Sober's initial claim that there is a *ceteris paribus* law holding between blue-eyed parents and blue-eyed children seems good in spite of Pietroski's point about surgical alteration. It is tempting to think that this is a circumstance in which all things aren't equal. If Pietroski wants to depart from our ordinary usage that's fine. But I think it would be better to be honest about this and state that one is interested in *patterns of causal relevance* rather than *ceteris paribus* laws. Appeal to *ceteris paribus* laws has the appearance of being respectful of scientific practice while showing how it appeases philosophical concerns. But the appearance is misleading. Philosophical concerns are forcing us to distort scientific practice by providing an account of *ceteris paribus* laws that has little plausibility outside of its capacity to assuage philosophical concerns.

One diagnosis of the source of Pietroski's problems with Segal and Sober's case is that Pietroski's approach ignores the vital question of the relationship between the events cited in the antecedent of a *ceteris paribus* law and the events mentioned in lower level laws governing the micro level. Is the relationship close enough for the parents being blue eyed to be causally relevant? The fact that being blue eyed is some causal distance from the mechanism whereby the genetic material of the parents has an influence on the genetic make up of the children suggests

not. In which case, any development of the notion of a pattern of causal relevance (as I have recommended we think about it) is going to need to appeal to the very factors that Pietroski has ruled not to be available in the case of mental events.

4. The Prospects for Personal Dualism

There are many ways Personal Dualism might be developed. I have questioned the way it has been developed in Pietroski's hands. It seems to me that the Argument from Differential Vagueness does not establish what Pietroski hopes, namely the falsity of Neuralism. Moreover, saddling Personal Dualism with the denial of Neuralism makes it ill-equipped to answer concerns about mental causation. If I am right, the challenge for Personal Dualists of Pietroski's stripe is to explain how the mental bears the right kind of relations to the neural for the mental to be efficacious given that the neural does not even vaguely constitute the mental. Perhaps it would be better not to give oneself that opportunity.¹⁰

University of Nottingham

References

- Ayer, A.J. (1963), *The Concept of Person* (London, Methuen).
- Benacerraf, Paul (1965), 'What numbers could not be', *The Philosophical Review*, 74, pp. 47-73, reprinted in P. Benacerraf and H. Putnam (1983, 2nd ed.), *Philosophy of Mathematics* (Cambridge, Cambridge University Press), pp. 272-294.
- Cummins, Robert (1983), *The Nature of Psychological Explanation* (Cambridge, Massachusetts, The MIT Press).

¹⁰ I would like to thank Paul Pietroski for his very helpful discussion of the issues I've raised and Colleen Noordhof for helping me to improve the final piece. I would also like to thank the AHRB Matching Research Leave Scheme and Mind Association for helping support the period of research leave during which this paper was written.

- Davidson, Donald (1970), 'Mental Events', in L. Foster and J. W. Swanson (eds), *Experience and Theory*, and in his (1980) *Essays on Actions and Events* (Oxford, Oxford University Press), pp. 207-225.
- Evans, Gareth (1978), 'Can There Be Vague Objects?', *Analysis*, 38, 4, p. 208, and in his (1985), *Collected Papers* (Oxford, Oxford University Press), pp. 176-177 [page references in text to latter].
- Hornsby, Jennifer (2002) Review of Paul Pietroski's *Causing Actions*, *Mind*, 11:441, 159-161.
- (1985), 'Physicalism, Events and Part-Whole Relations', in E. LePore and B. McLaughlin (eds), *Actions and Events* (Oxford, Basil Blackwell), pp. 444-458, reprinted in her (1997), *Simple Mindedness* (Cambridge, Massachusetts, Harvard University Press), pp. 46-62 [page references in text to latter].
- (1981), 'Which Physical Events are Mental Events', *Proceedings of the Aristotelian Society*, 81, pp. 73-92, reprinted with cuts in her (1997), *Simple Mindedness* (Cambridge, Massachusetts, Harvard University Press), pp. 63-77 [page references in text to latter].
- Lewis, David (1988), 'Vague Identity: Evans Misunderstood', *Analysis*, 48, 128-130.
- Noordhof, Paul (1999a), 'Probabilistic Causation, Preemption and Counterfactuals', *Mind*, 108, no. 429, pp. 95-125.
- (1999b), 'Causation by Content?', *Mind and Language*, 14, no. 3, pp. 291-320.
- (2001), 'Getting Personal', http://www.uniroma3.it/kant/field/pietroskisymp_noordhof.htm
- (forthcoming), 'Self-Deception, Interpretation and Consciousness', *Philosophy and Phenomenological Research*.
- Pietroski, Paul (2000), *Causing Actions* (Oxford, Oxford University Press).
- (2001), 'Reply to Noordhof', http://www.uniroma3.it/kant/field/pietroskisymp_replytonoordhof.htm
- Sainsbury, R.M. (1987, 1995a), *Paradoxes* (Cambridge, Cambridge University Press).

- (1991), *Concepts Without Boundaries*, Inaugural Lecture. (London, Published by the King's College Philosophy Department). The paper is reprinted in R. Keefe and P. Smith (1996, eds) *Vagueness: A Reader* (Cambridge, Massachusetts, The MIT Press), pp. 251-264.
- (1995b), 'Vagueness, Ignorance, and Margin for Error', *British Journal for the Philosophy of Science*, 46, pp. 589-601.
- Strawson, P.F. (1958), 'Persons' in H. Feigl & M. Scriven (eds), *Minnesota Studies in the Philosophy of Science, II* (Minneapolis, University of Minnesota Press), pp. 330-353.
- (1959), *Individuals* (London, Methuen).
- Bernard Williams (1961), 'Strawson on Individuals', in his (1973), *Problems of the Self*, (Cambridge, Cambridge University Press), reprinted from 'Mr Strawson on Individuals', *Philosophy*, 36.
- Timothy Williamson (1994), *Vagueness* (London and New York, Routledge).