

The Personal/Subpersonal Distinction Revisited: Towards an Explication

LARS DÄNZER

Abstract

The distinction between the personal and the subpersonal is often invoked in philosophy of psychology but remains surrounded by confusion. Building on recent work by Zoe Drayson, this paper aims to help further improve this situation by offering a satisfactory explication of the distinction that remains close to Dennett's original intentions. Reasons are offered for construing the distinction as applying to representational (as opposed to worldly) items, for not building contested theoretical assumptions into it, and for taking it to apply in the first instance to descriptive statements and only derivatively to explanations. An explication of the distinction that accords with these points is then developed, according to which the distinction should be drawn in terms of what personal and subpersonal-level statements are 'transparently about'. The theoretical role of this explication is discussed, and potential objections are addressed.

1. Introduction

An important distinction widely employed throughout the philosophy of mind and psychology is that between the personal and the subpersonal (for short: the 'P/SP distinction'). Originally introduced by Daniel Dennett (1969), this distinction has established itself in the field and is featured in some of its main introductory texts and anthologies (e.g., Bermúdez, 2005, 2006; Cain, 2016; Davies, 2005).

Unfortunately, the P/SP distinction is also surrounded by a serious lack of clarity, which manifests itself in a number of ways. First, the distinction is typically used without an accompanying definition; instead, we just get some paradigmatic examples and maybe

doi:10.1017/S0031819123000220 © The Author(s), 2023. Published by Cambridge University Press on behalf of The Royal Institute of Philosophy. This is an Open Access article, distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives licence (<https://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is unaltered and is properly cited. The written permission of Cambridge University Press must be obtained for commercial re-use or in order to create a derivative work.

some general characterizations whose exact status is left unclear.¹ Second, and relatedly, different potential criteria for drawing the distinction are often not properly distinguished. Third, it is plain that different theorists construe the distinction in importantly different ways, but typically without taking note of this fact.²

The unclarity besetting the P/SP distinction may sometimes be harmless, but other times they are not. Instead, they may frustratingly obfuscate the claims and arguments of those invoking the distinction, and they may even lead to serious confusion and error, as documented by Drayson (2014, p. 344). What's more, these unclaritys have also led some authors to question the usefulness of the distinction (e.g., Machery, 2009, pp. 24–5; Rey, 2001, p. 105).

Recently, some important steps have been taken to improve on this unfortunate situation, most notably by Zoe Drayson (2012, 2014).³ Drayson not only demonstrates the need for more clarity regarding the P/SP distinction, she also does much to clarify some of the different understandings of the distinction and to identify some of the confusions involved. What's more, by drawing attention to its original understanding in Dennett (1969, 1978), she presents important considerations for favouring some construals of the distinction over others.

Unfortunately, however, what Drayson does *not* offer is a satisfactory explication of the P/SP distinction that fits the bill. Her own proposal falls short, at least when measured against a list of reasonable criteria which she herself seems to accept – or so I'll argue below. Thus, we're left with a task: While Drayson has demonstrated the need for an explication of the P/SP distinction that does justice to its original understanding, such an explication is as yet lacking. The main goal of this paper is to take on this task. Building upon Drayson's insights, I aim to further improve our understanding of the P/SP distinction and to propose and defend a satisfactory explication.

¹ Are we given necessary and/or sufficient conditions or just some generalizations that may apply to typical instances of the categories? We typically aren't told.

² One indication of such divergences in understanding is that different theorists apply the distinction to very different things; a second indication is that when characterizations or definitions are being offered at all, they often differ substantively. See §§2–3 below.

³ Other helpful discussions include Hornsby (2000), Davies (2000a), Elton (2000), Skidelsky (2006) and Wilkinson (2015). Kriegel (2012) is also helpful but he explicitly focuses on what I below call the ontic construal of the distinction.

The Personal/Subpersonal Distinction Revisited

The plan is as follows. Sections 2 and 3 argue, following Drayson's lead, in favour of construing the P/SP distinction as applying to representational (as opposed to worldly) items and as thin, i.e., not embodying contested views about the mind. Section 4 argues that, within the representational domain, the classification of explanations as personal or subpersonal should be determined entirely by the status (personal or subpersonal) of its component descriptive statements (contra Drayson, 2012). This leaves descriptive statements as the most basic domain of the distinction and the rest of the paper is thus concerned with explicating the P/SP distinction for such statements. After considering some unsatisfactory proposals in section 5, my favoured explication is presented in section 6. Finally, section 7 discusses the theoretical role of the proposed explication and section 8 addresses possible objections.

Before we get going, however, I should say more about how I think of the project of this paper and about the methodological assumptions I adopt. My project is one of explication or conceptual re-engineering along broadly Carnapian lines (Carnap, 1962; Brun, 2016). That is, my question is how we *should* construe the P/SP distinction in light of a number of criteria of adequacy. In particular, inspired by Carnap's list, I adopt the following three criteria: (i) clarity and simplicity; (ii) theoretical utility; (iii) agreement with the original understanding of the distinction as displayed in the writings of Dennett and those who stay reasonably close to his understanding.⁴

To accept the last criterion is not, I emphasize, to construe the task as mainly exegetical, i.e., to capture as well as possible what Dennett (or some other author) had in mind. After all, the criterion is just one among three. Still, to accept it is to insist that we want an account that is *reasonably close* to what Dennett and those picking up on his distinction had in mind. Also, while I think there are good reasons for adopting this criterion – continuity with its original meaning is surely *one* relevant consideration in deciding how we should define a term – I acknowledge that it is not mandatory. After all, it might be said that the distinction has acquired a use, especially in its application to states and processes (cf. Drayson, 2014, 341–4), that has rather little to do with Dennett's original one but which is no less theoretically important for that. This is likely true and although it is perhaps debatable whether we should really use the terms 'personal'

⁴ Importantly, I am not suggesting that there is a single, definite understanding of the distinction shared by these theorists. However, there is often significant overlap in aspects of their application of the distinction and these overlaps should be respected, I suggest, on pain of changing the subject.

and ‘subpersonal’ in such a different way, that debate is probably not worth having. Instead, I take the lesson to be that any attempt at explicating ‘the’ P/SP distinction will have to be more specific in identifying its target (i.e., it does not make sense to evaluate an explication in terms of how well it accords with *all* uses of the distinction). This is what the third criterion does.

In fact, *even if* all my three criteria are accepted, it is still far from clear that they entail a determinate decision as to which explication of the P/SP distinction is best. The considerations involved in applying the criteria and weighing them against each other are likely not decisive. Accordingly, my claim in this paper is a modest one: while I think that the account I provide in this paper deserves the title ‘P/SP distinction’ at least as much as any other, I also take it that there’s little point in insisting that my proposal captures *the right*, or *the only* way to explicate the distinction. As I see it, what ultimately matters is that those invoking the distinction make it clear what meaning they are attaching to it so that the distinction can further discussion rather than confound it. My goal is to present one strong candidate for such a meaning that so far has not been clearly articulated. Moreover, whether or not my proposal is ultimately accepted, the discussion should help clarify the various choices to be made in the course of assigning a determinate meaning to the distinction as well as some of their pros and cons.

2. The Distinction’s Domain: Representational vs Ontic Construals

One useful distinction for bringing some structure into the various construals of the P/SP distinction, emphasized by Drayson, concerns the *objects* to which the distinction applies. With respect to this issue, we can broadly distinguish two versions of the distinction. On the one hand, the distinction can be understood as applying to *the linguistic means* that we use to talk about the mind or brain, i.e., to such representational (and semantically evaluable) items as descriptions, explanations (*qua* representations),⁵ or theories. We can call this the

⁵ Throughout this paper I invoke a *representational* conception of explanation as opposed to an *ontic* conception, according to which explanations are worldly items such as mechanisms (e.g., Craver, 2014). My reason for doing so is that, as noted below, there are good reasons for construing the P/SP distinction representationally and therefore as applying to explanations *qua* representational and not to explanations *qua* ontic. However,

The Personal/Subpersonal Distinction Revisited

representational conception of the P/SP distinction. This representational conception is operative, for instance, when it is said that folk-psychological description/explanation in terms of an agent's mental states – e.g., her beliefs, desires, emotions – is located at the personal level, whereas the descriptions/explanations we find in (classical) computational-representational psychology are pitched at the subpersonal level.

On the other hand, the P/SP distinction can also be construed as applying to 'worldly items' like states, events, and processes. We may call this the *ontic conception* of the P/SP distinction. This ontic conception is operative, for instance, when it is said that beliefs and desires (among others) are personal-level states, whereas the states postulated by cognitive scientists (e.g., Marr, 1982) to be involved in early visual processing are subpersonal-level states.

The importance of distinguishing between the ontic and the representational conception of the P/SP distinction may not be immediately obvious, which is probably why it is so often overlooked. Doesn't the representational version of the distinction automatically yield an ontic version of the distinction, when we construe personal- and subpersonal-level states (for instance) as *states posited by personal-level and subpersonal-level explanations*, respectively? No. The problem, as Drayson (2012, pp. 8–12) points out, is that it is perfectly coherent to hold – as proponents of some varieties of functionalism and the identity theory do – that the states posited by personal-level explanations (e.g., beliefs) are *identical* to the states posited by subpersonal-level explanations (e.g., certain computational-representational states). A proponent of such a view could thus grant that there is a distinction to be drawn between personal-level and subpersonal-level *claims*, while also maintaining that there is no corresponding distinction between personal- and subpersonal-level *states* (processes, *etc.*), because the two kinds of claims are really just superficially different ways of talking about the *same* states (processes, *etc.*).⁶ (We shall return to this point in §4 below.)

note that I'm not thereby taking any stance on the question which lies at the heart of the relevant debate over the nature of explanation, *viz.*, whether the representational or the ontic conception of explanation is more fundamental. Even proponents of the ontic view grant that in one of its meanings the word 'explanation' applies to representational items and that is all I am assuming here.

⁶ More fully stated, the problem is that on the kind of view under consideration, all personal-level states would also be subpersonal-level states. But a reasonable P/SP distinction for states should presumably be exclusive

As Drayson (2012, 2014) and others have pointed out, Dennett's (1969) original understanding of the P/SP distinction is clearly representational in character, not ontic. He introduces the labels 'personal' and 'subpersonal' to mark the distinction between the 'explanatory level of people and their sensations and activities' on the one hand and the 'sub-personal level of brains and events in the nervous system' on the other (1969, p. 93). Consideration of the context makes it clear, moreover, that the distinction is meant to apply more broadly to 'modes of description or explanation' (1969, p. 89), 'stories' (1969, p. 189) or 'modes of discourse' (1969, pp. xiv, 189).⁷ This leaves little doubt about the representational nature of the distinction. There is also an important thread in the literature which follows Dennett in construing the distinction representationally, exemplified by such authors as Jennifer Hornsby (1997, 2000), Martin Davies (2000a, 2000b, 2005), Matthew Elton (2000, 2003), and José Bermúdez (2000, 2005), all of whom have done much to lend the distinction its current influence.⁸ Therefore, if we want an explication of the P/SP distinction that accords with its original

at least for paradigmatic examples, i.e., it should entail that beliefs, desires, *etc.* are personal *and not also subpersonal.*

⁷ The P/SP distinction is introduced in section 11 of the book, which is entitled 'Personal and Subpersonal Levels of Explanation: Pain'. This section is the last of Part I of the book, entitled 'The Language of Mind', which focuses on developing 'the notion of a distinct mode of discourse, the language of the mind, which we ordinarily use to describe and explain our mental experiences, and which can be related only indirectly to the mode of discourse in which science is formulated' (p. xiv). In accordance with this, Dennett reminds us in the very first sentence of section 11, that '[t]he aim of Part I has been to describe the relationship between the language of the mind and the language of the physical sciences' (p. 90), and he goes on to say that the objective of section 11 is 'to consolidate the gains of Part I by illustrating them in application to a particular mental phenomenon: pain' (p. 90). When, in discussing this example, he then distinguishes between the 'explanatory level of people and their sensations and activities' and the 'sub-personal level of brains and events in the nervous system', it is virtually impossible to interpret this other than as an application to the realm of explanation of the distinction between the two modes of discourse which lies at the heart of Part I of the book. The interpretation is confirmed when, later in the book, Dennett talks explicitly about 'the personal, mental language level' (p. 113).

⁸ This is not to say that these authors are clear on the distinction between the representational and the ontic construal and explicitly adopt the former, but only that the representational construal figures prominently

The Personal/Subpersonal Distinction Revisited

understanding, we clearly have good reason to adopt a representational construal.

3. The Distinction's Status: Thin vs Thick Construals

When we look at the authors who work with a representational conception of the distinction, however, we find that there is a further important distinction to be drawn, one concerning the intended role or status of the distinction. Whereas a *thick* version of the P/SP distinction builds substantive theoretical assumptions into the distinction, a *thin* version does not. We can illustrate the contrast by considering the following familiar claim:

- (1) Whereas folk-psychology operates at the personal level, representational-computational psychology ('RCP') operates at the subpersonal level.

Depending on the author, this claim can have either of two very different statuses. On the one hand, it is sometimes offered as a largely *uncontroversial observation* regarding folk-psychology and RCP, i.e., an observation that does not imply any substantive view regarding the *relation* of folk-psychology and RCP. This is the case, for instance, when authors like Bermúdez (2005) and Davies (2000a) appeal to the P/SP distinction to formulate what they consider as a central question, or challenge, to be addressed by the philosophy of psychology, *viz.*:

How do the personal-level descriptions and explanations of folk-psychology relate to the subpersonal-level descriptions and explanations to be found in various areas of cognitive science?

Bermúdez (2005) calls this the *integration challenge* and places it at the centre of his influential survey of the field of philosophy of psychology. Since the challenge itself is supposed to be theoretically neutral – i.e., acceptable to all parties – the P/SP distinction it invokes must be a thin one.

On the other hand, (1) can also be put forward as a *substantive thesis* that carries controversial implications regarding the relation between folk-psychology and RCP. This is the case, for instance, with McDowell (1994) and Hornsby (1997, 2000), for whom (1)

in their writings. Bermúdez (2005) in particular shifts back and forth between the representational and the ontic construal.

effectively means that folk-psychology is strongly autonomous from, and not reducible to, RCP.

The difference between thin and thick versions of the P/SP distinction is clearly important, but often overlooked. Drayson points towards it by claiming that Dennett's original version of the P/SP distinction, in contrast to some of its later interpretations, is neutral on the relation between the personal and the subpersonal level (2012, pp. 8–12, 2014, p. 341). I agree with her on that point, but I also think that Dennett (1969) is less than fully clear on the status of his distinction and that this explains why it has been understood in both thin and thick ways. As noted above, Dennett's (1969) P/SP distinction is apparently meant to coincide with his distinction between two 'modes of discourse' that plays a prominent role in the first part of his book (cf. Dennett, 1969, pp. xiv, 189–90). He introduces this distinction to consider the question of how the two modes of discourse are related – which strongly suggests that no particular answer to this question is already definitionally built into the distinction and that it is therefore thin in this respect. However, there is also reason to interpret his distinction as thick in another respect, *viz.*, as embodying the assumption that ascriptions of mental states and processes (e.g., belief, inference, *etc.*) are coherent only at the personal level, but not the subpersonal one. Thus, in his (1969) he sometimes characterizes the two modes of discourse simply as 'the language of mind and the language of science' (Dennett, 1969, p. 189), respectively, and he emphatically rejects the 'homuncular' attribution of mental states or processes (e.g., belief, inference, *etc.*) to functional subsystems as misguided.⁹ That said, evidence against this interpretation comes from his later writings, especially the fact that Dennett (1978), in a radical change of mind, comes to fully endorse homuncular subpersonal explanations as coherent and theoretically fruitful. This favours interpreting even Dennett (1969) as not building the incoherence of mentalistic, homuncular talk at the subpersonal level into his P/SP distinction – otherwise, we would have to assume that he substantially changed his understanding of the distinction between 1969 and 1978 without giving any note of this (as Hornsby, 2000, effectively suggests).

In sum, there is good reason to interpret Dennett's P/SP distinction as thin throughout. There are two additional reasons in favour of a thin construal of the distinction. First, as noted above, a thin

⁹ However, Dennett (1969) endorses as legitimate subpersonal explanations that ascribe *content* to states and events (possibly functionally characterized) in the nervous system. See also §5 below.

The Personal/Subpersonal Distinction Revisited

construal seems to be operative in some of the most influential literature on the P/SP distinction, especially Bermúdez (2005) and Davies (2000a). Moreover, a thin construal has an important strategic advantage. Whereas any thick version of the distinction is bound (by definition) to be rejected by some theorists as incorporating false assumptions, a thin version should be acceptable (almost) universally. At the same time, the thin version of the P/SP distinction still allows us to formulate any substantive view regarding the personal and the subpersonal levels we happen to favour (a point to which I return in §7). It is just that such views are not already ‘built into’ the distinction. Thus, by adopting a thin version of the distinction we arguably have much to gain – *viz.*, widespread acceptance of the distinction – but little to lose.

To sum up, there are good reasons for adopting a thin, representational construal of the P/SP distinction. My question in what follows is how we should explicate the distinction in light of this.

4. The P/SP Distinction for Descriptions and Explanations

Saying that we want a thin, representational version of the P/SP distinction leaves some important questions open. One derives from the fact that, within the representational domain, the P/SP distinction is commonly applied to both *descriptions* and *explanations*. For instance, as noted in §2, Dennett (1969) takes the distinction to apply not only to explanations but to descriptions as well, as reflected in his talk of ‘modes of description or explanation’ (1969, p. 89). And, following his lead, many refer to the P/SP distinction explicitly as a distinction between ‘levels of descriptions or explanations’ (Skidelsky, 2006, p. 116) or shift back and forth between talking of ‘levels of description’ and ‘levels of explanation’ (e.g., Bermúdez, 2005). This raises a question: how should we think of the relation between the P/SP distinction as applied to descriptions on the one hand and as applied to explanations on the other?

In thinking about this question we need an account of how descriptions and explanations (*qua* representations) relate to each other. For our purposes it is useful to start from the plausible idea that descriptions and explanations both come in the form of *statements*, *viz.* *descriptive statements* and *explanatory statements*, respectively.¹⁰ Moreover, an explanatory statement can plausibly be construed as a

¹⁰ More generally, we could say that they come in the form of *intentional contents*, which may be conveyed by bits of language but also by such things

complex statement that is made up of descriptive statements. Thus, a paradigmatic explanatory statement takes the form ‘P because Q1, ... Qn’, where P describes the explanandum (what is to be explained) and Q1, ... Q2 describe the explanantia (the explaining factors).

Given this framework as background, the first and most straightforward idea is that the P/SP distinction for descriptive statements is basic and that explanations inherit their status as personal or sub-personal from the status of their component descriptions. Call this the *Simple View*. I think this is the view we should adopt. To see why, it is helpful to consider a proposal that conflicts with it, *viz.*, that of Drayson (2012).

According to Drayson (2012), the P/SP distinction should be understood in the first place as ‘a distinction between two kinds of psychological explanation, one horizontal and the other vertical’ (2012, p. 14).¹¹ Personal-level explanations are horizontal, where ‘horizontal explanations attempt to account for an event’s occurrence by citing a sequence of preceding events’; subpersonal-level explanations are vertical, where ‘vertical explanations attempt to account for a thing’s features by citing its components’ (2012, p. 14). More specifically, personal-level explanations are horizontal explanations that proceed by ‘ascribing psychological predicates to whole persons’ (2012, p. 7), whereas subpersonal-level explanations are vertical explanations that aim to account for a person’s capacities by breaking the person up into functionally specified components that are themselves characterized in psychological terms, i.e., ascribed psychological predicates. Thus, subpersonal-level explanations treat certain components of a person as ‘subpersons’ or ‘homunculi’ that may have goals and knowledge and that can calculate, make decisions, and communicate with each other. However, Drayson also notes that the P/SP distinction ‘not only allows us to distinguish between vertical and horizontal psychological explanations in general, but also enables us to clarify individual instances of psychological predicate ascription’ (2012, p. 8). More specifically, she says, it allows us to distinguish ‘ascription of a

as diagrams or graphs, for instance. However, for the sake of simplicity, I will continue to talk about statements.

¹¹ It is worth noting that Drayson has not reiterated this account in her (2014) paper on the P/SP distinction. That said, that later paper is mainly concerned with describing how other theorists have understood the distinction, whereas her 2012 paper is more explicitly about how the distinction *should* be understood. Moreover, Drayson has never retracted her 2012 account, so it is certainly worth discussing.

The Personal/Subpersonal Distinction Revisited

particular psychological predicate to a part of person from ascription of the same psychological predicate to a whole person' (2012, p. 8).

Thus, we find in Drayson (2012) both an account of the P/SP distinction as it applies to explanations and as it applies to descriptive statements. But note that Drayson's account conflicts with the Simple View regarding the relation between the P/SP distinction for descriptions and for explanations: by construing the P/SP distinction for explanations as an instance of the horizontal/vertical distinction, she builds conditions into the distinction that go beyond the status of its component descriptive statements.

However, this feature makes the account vulnerable to objections. Take first the idea that personal-level explanations must be horizontal and consider this explanation:

- (2) Paul is able to reliably identify flying mammals (as such) because he knows that bats are the only flying mammals and he knows what bats look like.

Surely, this should count as a personal-level explanation, but it is far from obvious that it is horizontal. Horizontal explanations, remember, 'attempt to account for an event's occurrence by citing a sequence of preceding events' (Drayson, 2012, p. 2). But, on the face of it, (2) does not account for Paul's ability by citing *preceding* mental states of Paul, but rather by citing states that obtain *simultaneously* with Paul's ability. To be sure, this construal of (2) is debatable, but the point is independent of whether it is ultimately correct. Instead, the point is simply that (2) should surely count as a personal-level explanation *whether or not* it is horizontal.

The problem is even more obvious with the claim that subpersonal explanations must be vertical. Cognitive science is full of horizontal explanations in which both explanandum and explanans are pitched at the subpersonal level. For instance, relevance theory (e.g., Sperber and Wilson, 2002) seeks to explain how accurate representations of what a speaker said are generated in the hearer's mind through processes in multiple more-or-less modular systems dedicated to the analysis of incoming speech and mind-reading. Since such explanations are not vertical but horizontal, Drayson's account oddly implies that they are *not* subpersonal (nor personal, of course). But that's surely an unacceptable implication.

So while personal explanations are certainly often horizontal, and subpersonal explanations often vertical, building these characteristics into the personal/subpersonal distinction for explanations has little to recommend itself. In fact, we can generalise the point to give a direct argument for the Simple View. What underlies the above

objections to Drayson's account is the following intuitive principle: an explanation in which both the explanandum and the explanans are located at the personal level should clearly count as a personal-level explanation, and *mutatis mutandis* for the subpersonal level. Any account that conflicts with the Simple View is bound to run into similar problems vis-à-vis this principle as Drayson's account.¹² Therefore, we have good reason to adopt the Simple View as our account of how the P/SP distinction for explanations relates to the P/SP distinction for descriptions.

In light of this conclusion, the rest of the paper focuses on how we can explicate the P/SP distinction for the domain of descriptive statements. With such an account in hand, the extension to explanations will be straightforward.¹³

5. Towards a P/SP Distinction for Statements: Unsatisfactory Proposals

We are looking for an account of the P/SP distinction that is thin and applies to descriptive statements. (For the sake of brevity, I omit the qualifier 'descriptive' in what follows and simply talk of statements.) We may start, in this section, by getting clear on why some *prima facie* promising proposals won't do.

As already noted, Dennett (1969) in places suggests that his P/SP distinction coincides with the distinction between 'the language of mind and the language of science' (Dennett, 1969, p. 189). Taking

¹² This is true, for instance, of the Hornsby/McDowell view according to which it is partly constitutive of the concept of personal-level explanations that they operate in a distinctive way, *viz.* by invoking rational norms. Since this account of folk-psychological explanations is contested (see, e.g., Fodor 1987 for a different view) this view also has the undesirable consequence that it makes it contentious whether paradigmatic folk-psychological explanations actually qualify as personal-level. This is another manifestation of the fact that the Hornsby/McDowell account of the P/SP distinction is a thick one.

¹³ Rather than a simple two-fold classification into personal- and subpersonal-level explanations, we will get a more fine-grained fourfold classification that takes account of the fact that both the explanans and the explanandum of an explanation can be either at the personal or the subpersonal level.

The Personal/Subpersonal Distinction Revisited

our clue from these remarks, we might try to draw the P/SP distinction for statements along the following lines:

First Shot

A statement is located at the personal level iff it is cast in 'mental language'.

A statement is located at the subpersonal level iff it is cast in 'the language of science'.

However, we've already seen why any account along these lines won't work. According to First Shot, a statement that uses 'mental language' to ascribe a mental state or process *to a subsystem or component* of a person (i.e., a statement of the sort associated with 'homuncular' explanations) comes out as personal-level – which is clearly the wrong result. Indeed, as we saw, Dennett himself came to endorse homuncular explanations as an important variety of subpersonal explanation in his (1978).

Let us next consider Drayson's suggestion for drawing the P/SP distinction in the realm of statements (2012, p. 8), which we have encountered in the previous section:

Second Shot

A statement is located at the personal level iff it involves attribution of a psychological predicate to a whole person.

A statement is located at the subpersonal level iff it involves attribution of a psychological predicate to a part of a person.

In thinking about this proposal let us start by focusing on Drayson's characterisation of subpersonal-level statements. A first point to note is that Drayson uses the notion of a 'psychological predicate' in a broad sense, in which it covers not only predicates that ascribe psychological states, processes, *etc.* to something. As she rightly points out (Drayson, 2012, p. 10), subpersonal explanations are not limited to explanations in terms of homuncular subsystems that are said to 'judge', 'decide', 'interpret', *etc.* Instead, they also include explanations in terms of interacting internal (often functionally or computationally specified) states that are ascribed certain *contents*. Since Drayson intends her account of subpersonal-level explanations to cover the latter explanations as well, her notion of a 'psychological predicate' must be understood to cover predicates that ascribe *content* as well.

At a purely terminological level, one might worry that this broad use of the term 'psychological predicate' in an account of the subpersonal level is unfortunate because it encourages conflation of these two importantly different kinds of subpersonal-level

explanation.¹⁴ But there is also a substantive worry concerning Drayson's characterisation of subpersonal-level statements. By requiring that such claims must involve 'psychological predicates', she excludes statements that are couched in purely non-intentional, e.g. neuroscientific, terms from the category of the subpersonal. But this is in clear conflict with how the category of the subpersonal has commonly been understood. Tellingly, Dennett's original characterization (1969, p. 103) of the 'subpersonal level' as the 'level of brains and events in the nervous system' does not involve any such restriction, and the context makes it clear that he intends the 'subpersonal level' to cover, *inter alia*, explanations that are cast fully in the non-intentional vocabulary of neuroscience. Similarly, Hornsby (1997, 2000), Davies (2000a), Bermúdez (2005), and Skidelsky (2006), among many others, all consider non-intentional neuroscientific claims and explanations as clear instances of the subpersonal.

Why, then, does Drayson limit subpersonal statements to those involving 'psychological predicates'? The reason, it seems, is that she construes the P/SP distinction as a distinction *within the domain of the psychological*. Her concern seems to be with the distinction between personal and subpersonal *psychology*. However, while it is certainly true that the P/SP distinction as originally understood by Dennett and others can be *applied* to the psychological, it is not *restricted* to it. We would need a good reason to abandon this established construal of the distinction, but Drayson fails to provide one.¹⁵ Moreover, as we'll see later (§7, fn. 20), reflection on the

¹⁴ In fact, this kind of conflation afflicts Drayson's own account of the early history of the P/SP distinction. Drayson claims in multiple places (2012, pp. 7–8; 2014, p. 339) that Dennett introduced the P/SP to make room for subpersonal homuncular explanations. However, as already noted, Dennett (1969) in fact emphatically *rejected* such explanations as misguided. What he meant to endorse were only subpersonal explanations that ascribe *content* to states and events (possibly functionally characterized) in the nervous system. While it is true that he soon came to revise his view on homuncular explanations (Dennett, 1978), Drayson's historical claim is nonetheless mistaken.

¹⁵ Drayson suggests that the point of the P/SP distinction is 'to emphasise that there is a type of psychological explanation which is not folk-psychological' (2012, p. 7) and 'to clarify individual instances of psychological predicate ascription' (2012, p. 8). However, the distinction can play these roles equally well without a definitional restriction of the subpersonal-level to the psychological. We can simply talk about 'subpersonal *psychological* descriptions/explanations' (as opposed to simply 'subpersonal descriptions/explanations') for these purposes. In fact, it is a common

The Personal/Subpersonal Distinction Revisited

theoretical point of the thin P/SP distinction provides additional reason against Drayson's restricted understanding of the subpersonal level.

There is an obvious way to modify Second Shot in response to the criticism that a subpersonal-level statement *may* characterize components of a person in psychological terms, but *need not* do so:

Third Shot

A statement is located at the personal level iff it involves attribution of a psychological predicate to a whole person.

A statement is located at the subpersonal level iff it involves attribution of a predicate – psychological or otherwise – to a part of a person.

However, this account still faces two problems.¹⁶ For one thing, consider claims of the following sorts:

(3) Paul is in such-and-such (type of) neural state.

(4) Paul has such-and-such an internal functional organization.

Intuitively, such claims are clear examples of subpersonal-level claims. However, Third Shot doesn't yield this verdict, since the subject of these predications is a person, not an internal part of a person. The moral of these examples is that the P/SP distinction doesn't neatly line up with a difference in subjects of predications in the way Third Shot has it. There are subpersonal-level claims whose subject of predication is a whole person, *viz.*, claims ascribing a *structural property* to a person.

One might try to fix this problem by going disjunctive and defining a subpersonal-level claim as one that involves *either* the attribution of a property to a part of a person, *or* the attribution of a structural property to a person. But not only is such a move suspicious – isn't there something that the cases covered by the disjuncts have in common, which should be used in the definition instead? – but it also still runs into the second problem.

The second problem with Third Shot is that there are many claims that we would want to classify as personal or subpersonal in which the

suggestion that the subpersonal level can be further subdivided, into the level of subpersonal psychology and the level of neurobiology, for example (e.g., Bermúdez, 2005, pp. 28–9).

¹⁶ They also apply to Second Shot, Drayson's unmodified proposal.

subject of predication is *neither* a whole person *nor* a component of a person. Here are some examples:

Paul's belief that there was no more beer in the fridge triggered his desire for revenge.

The result of the computation carried out by the visual module is a 2½-D sketch.

The upcoming exam was at the forefront of Paul's mind when he walked home.

Such examples suggest, I think, that the general idea for defining the P/SP distinction that is exemplified in Third Shot – *viz.*, that the distinction is to be drawn in terms of a statement's subject and predicate – is hopeless. The form of both personal- and subpersonal-level statements is simply too varied. We need a different approach.

6. A Better Proposal

Fortunately, we need not look far. Dennett, remember, introduces the P/SP distinction by distinguishing the 'explanatory level of people and their sensations and activities' from the 'sub-personal level of brains and events in the nervous system' (1969, p. 93). Note that Dennett here characterizes the distinction in terms of what the respective explanatory levels are *of* or *about* – that is, in terms of their *subject matter*. Moreover, although this is never emphasized, this is how the P/SP distinction is very often introduced (e.g., Bermúdez, 2005, p. 28; Davies, 2000a, p. 88).

The idea, then, is to draw the P/SP distinction for statements in terms of their subject matter. What are these subject matters? Taking into account the lessons from our previous discussion, the following characterisation suggests itself:

Fourth Shot

A claim is located at the personal level iff it is about the mental events or intentional actions of a person.

A claim is located at the subpersonal level iff it is about the parts or internal organization of a person.

A number of clarifications are in order. Let's start with the notion of a 'mental event'. Here, the term 'event' is used in a broad and somewhat technical sense in which it also covers *processes* and *states*.¹⁷ As

¹⁷ Such a broad notion of 'event' is familiar from event semantics (e.g., Maienborn, 2011).

The Personal/Subpersonal Distinction Revisited

for what counts as a *mental* event (process, state), the account invokes the familiar, well-established understanding of this notion conveyed by introductory textbooks in the philosophy of mind. This understanding is typically conveyed by an overview or list of paradigmatic states, events, *etc.* that count as mental in the relevant sense, including such familiar items as sensations, perceptions, propositional attitudes, emotions, volitions, thoughts, decisions, *etc.* (e.g., Kim, 2011; Mandik, 2013; Kind, 2020).

Also, in accordance with the common understanding of the P/SP distinction, *intentional actions* are included in the subject matter of personal-level statements. Such actions often form the explanandum of personal-level explanations in terms of mental events. From a systematic perspective, their inclusion is warranted by the widely accepted view that intentional actions conceptually require the presence of appropriate mental events (Davidson, 1963).

The notion of aboutness or subject matter also requires further comment. While there are technical explications of the notion in the literature (e.g., Lewis, 1988; Yablo, 2014), I take the intuitive understanding of a statement's subject matter to be sufficient for the purposes of this paper. Still, two points are worth noting. First, it is a familiar point that a claim can be about something that does not exist. Thus, just as a claim can be *about Pegasus* although Pegasus does not exist, so a claim can be *about a person and her mental states* even if – as some philosophers have claimed – there actually should happen to be no persons or no mental states. In other words, on the present account, the claim that folk-psychology is located at the personal-level of description could be accepted even by an eliminativist about the mental (or some aspect of it). This is good news given our goal of providing a thin explication of the P/SP distinction.

Second, we may distinguish between a claim being *wholly* (or *entirely*) about a given subject matter, and a claim being (*at least*) *partly* about a subject matter. Intuitively, for a claim to be wholly about a subject matter **m** is for **m** to exhaust the claim's subject matter. That is, there is nothing besides, or in addition to, **m** that is also part of the claim's subject matter. Conversely, a claim is (at least) partly about a subject matter **m** iff **m** is (at least) part of, but does not necessarily exhaust, the claim's subject matter. That is, there may be other things besides, or in addition to, **m** that are also part of the claim's subject matter. By applying this distinction to Fourth Shot we can generate a corresponding distinction between statements that are *wholly* personal/subpersonal and statements that are *partly* so. This is also a welcome result in view of claims such as:

Paul's belief that snow is white is realized by an internal representation with the content <snow is white> that plays the belief-role.

Intuitively, such a claim is partly personal and partly subpersonal, which is exactly the result we are getting. Thus, while the proposed account is 'exclusive' in the sense of entailing that a statement cannot be at once *wholly* personal and at least partly subpersonal (or *vice versa*), it allows for a statement to be at once *partly* personal and *partly* subpersonal, which seems to be just the right result.¹⁸

Fourth Shot also avoids all the problems faced by the proposals considered in the last section and is thus the most plausible account considered so far. It allows the use of psychological predicates at both the personal and subpersonal level but doesn't require them at the subpersonal level. Claims like (3) and (4) involving structural properties are correctly classified as subpersonal, and there is also no restriction on the subject-predicate structure of personal- and subpersonal-level statements.

Unfortunately, the proposal nonetheless faces a pernicious problem, which is a variation on a theme noted earlier, *viz.*, the reasons why a representational P/SP distinction cannot straightforwardly be extended to the ontic realm (§2). In connection with Fourth Shot the problem arises as follows. Clearly, an adequate account of the P/SP distinction should get the established extension of the distinction right, at least for clear cases. Thus, it should entail that claims like (5) below are located at the personal level, but *not* at the subpersonal level, and *vice versa* for (6):

- (5) Paul believes that snow is white.
- (6) There is some internal state of Paul with the content <snow is white> that plays the belief-role.

The problem with Fourth Shot is that it doesn't give us this result in a *suitably uncontroversial way*, that is, unless we tie it to substantive and controversial assumptions about the semantics of mental terms. For instance, according to some functionalists, claims (5) and (6) are analytically equivalent and *really about the same things*. That is, on such a view, (5) and (6) are *both* about Paul's internal organization and also *both* about Paul's mental states. Thus, on such a view, Fourth

¹⁸ The distinction is also not exhaustive. This accords with the common understanding that the P/SP distinction is a distinction within a particular restricted domain – roughly, the domain of statements about people and/or their interiors – and that the distinction doesn't apply to statements outside this domain.

The Personal/Subpersonal Distinction Revisited

Shot entails that (5) and (6) are both personal- and subpersonal-level at the same time.

Note that this problem arises because we're treating Fourth Shot as a candidate explication for the *thin* construal of the P/SP distinction, not a thick one. There are, of course, also views about the semantics of mental terms on which Fourth Shot delivers the right result for (5) and (6), since, according to these views, statements about persons and their mental states are *not* statements about the persons' internal organization, nor *vice versa*. (And, indeed, authors like McDowell and Hornsby hold exactly such views.) Thus, proponents of such views might want to make Fourth Shot the basis of a *thick* version of the P/SP distinction by tying it directly to these views. However, since we want a *thin* version of the distinction, the problem with Fourth Shot stands.

It is also worth noting that the problem applies equally to other *prima facie* attractive ways of drawing the P/SP distinction for statements, such as that of Drayson considered earlier. While Drayson does not offer an explicit definition, her notion of a 'psychological predicate' is most naturally understood as a *predicate that ascribes a mental state (event, etc.) or mental content* to something. But, according to the functionalist view under consideration, (6) ascribes a mental state – and thereby a 'psychological predicate' in the relevant sense – to Paul and thus comes out as personal-level on Drayson's account.

Fortunately, however, the problem with Fourth Shot can be solved without abandoning the idea of drawing the P/SP distinction in terms of a statement's subject matter. For, to stay with the example, note that even if functionalists are right about the relation between (5) and (6), it is also indisputable that they are not *obviously* right. *Maybe* attributions of mental states to persons turn out, on analysis, to be about internal organization – but they are clearly not *overtly* or *transparently* so. Conversely, *maybe* it turns out, on analysis, that some claims about a person's internal organization are claims about his mental states – but, again, they are clearly not *overtly* or *transparently* so.

Thus, I suggest that the way to overcome the problem with Fourth Shot is to distinguish what a claim is *transparently* (or *overtly*) about, from what a claim may be about *non-transparently*, where the target notion of *transparent* aboutness is to be understood along the following lines: a statement S is transparently about a subject matter **m** iff the fact that S is about **m** is obvious to speakers who

fully understand S and grasp the concept **m**.¹⁹ Thus, we get the following account:

Final Shot

A claim is located at the personal level iff it is *transparently* about the mental events or intentional actions of a person.

A claim is located at the subpersonal level iff it is *transparently* about the parts or internal organization of a person.

To illustrate how this account works, let's consider how it applies to some paradigmatic examples of personal- and subpersonal-level claims. First, take (5) above. It certainly looks as if anyone who fully understands (5) will agree that it is about a mental state of Paul. After all, the statement is obviously about what Paul believes and it is hard to see how anyone could understand (5) without seeing this. Moreover, belief is a paradigmatic mental state, so it will be equally plain to those who understand (5) and grasp the concept of a mental state that (5) is about a mental state of Paul. On the other hand, it is certainly possible to fully understand (5) while denying that it is about Paul's parts or internal structure, as witnessed by the fact that many philosophers actually do so (e.g. Hornsby and McDowell). Therefore, (5) qualifies as personal-level but not as subpersonal-level according to Final Shot.

Next, consider (7):

(7) Paul's visual module computes a $2\frac{1}{2}$ -D sketch of his surroundings.

Again, it is hard to see how anyone might fully understand (7) without also agreeing that it is about a part of Paul. After all, it will surely be obvious to anyone who understands (7) that it is a statement about Paul's visual module. Moreover, since a module is *by definition* a functionally specified sub-system/part of a cognizer, it will also be plain to anyone who fully understands (7) (and thus knows what is meant by 'module') that it is about a sub-system/part of Paul. On the other hand, it is certainly *not* plain that (7) is about a mental state of Paul, as witnessed by the fact that many philosophers (such as Hornsby and McDowell) deny that it is. Therefore, (7) qualifies as subpersonal-level but not personal-level according to Final Shot.

Similarly, Final Shot plausibly classifies all the test case statements we have considered so far in the right way. Since my claims about what counts as the 'right' way were based on considerations of how

¹⁹ Thus, note that in the construction 'S is transparently about **m**' the term designating the subject matter, **m**, occurs in a *hyperintensional* context.

The Personal/Subpersonal Distinction Revisited

the P/SP distinction was originally understood by Dennett and those remaining close to his understanding, this is another way of saying that the account does well by our third criterion laid out in §1. That Final Shot plausibly provides a thin, representational account of the distinction is, of course, also crucial in this regard. Moreover, the account seems reasonably clear and simple and thus also does well by the first criterion. What is perhaps less clear at this point is whether the account also satisfies the second criterion, that is, whether it gives the P/SP distinction a useful theoretical role to play. So let's turn to this issue before considering some potential objections to the account.

7. The Role and Relevance of the Proposed Distinction

What, if anything, is the *point* of the P/SP distinction according to the explication offered in the previous section? To answer this question, remember what was said in §2 about how the thin version of the P/SP distinction figures in the literature, *viz.*, that it is invoked in formulating a central question, or challenge, to be addressed by the philosophy of psychology, *viz.*, *the question how the personal-level descriptions/explanations of folk-psychology relate to the various kinds of subpersonal-level descriptions/explanations to be found in cognitive science* (the 'integration challenge'). It is exactly because this question is supposed to be theoretically neutral – i.e., acceptable to all parties to the debate – that it must make use of the thin version of the P/SP distinction.

We can now see how the explication proposed here makes sense of this question. To illustrate, consider again claim (1) above:

- (1) Whereas folk-psychology operates at the personal level, representational-computational psychology ('RCP') operates at the subpersonal level.

On the proposed account, this claim *serves to draw attention to the fact that the claims of folk-psychology and RCP differ in their respective overt subject matter*. In other words, it stresses that folk-psychology is overtly about one thing (the mental states, *etc.* of a person), whereas RCP is overtly about another (the internal parts or organization of a person). It is this observation – the difference in overt subject matter – that immediately gives rise to the specific version of the integration challenge concerning the relation between the claims of folk-

psychology and those of RCP. And similarly for other branches of subpersonal-level cognitive science, including neuroscience.²⁰

Thus, the explication of the P/SP proposed here has a clear-cut, albeit modest, theoretical role to play: it serves to highlight important differences in the overt subject matter of various claims pertaining to the mind. But is it also important in ways that go beyond this modest role? While this is certainly an important question, it is not one to be addressed here. To see why, consider what it would mean for the distinction to be important in such a way. Presumably, it would mean that the distinction figures in some substantive philosophical truth about the mind (or perhaps some sound argument for such a truth). There are many potential candidates for such claims (or arguments) in the literature, such as the claim that psychological predicates cannot coherently (and literally) be used in subpersonal-level claims, but only in personal-level claims (Dennett, 1969), or the claim that personal-level claims cannot be analysed in terms of, or reduced to, subpersonal-level claims (Hornsby, 1997, 2000). The crucial point for our purposes, however, is this: plainly, the question whether the P/SP distinction is important in ways that go beyond its modest role cannot be separated from the evaluation of substantive claims and arguments in the philosophy of mind and psychology. It therefore lies squarely beyond the scope of this paper.

An important corollary of this point is also worth noting. Just as there is no quick way to establish that the P/SP distinction is important in ways that go beyond its minimal role, so there is no quick argument showing that the distinction is *not* important in any such way. After all, such an argument would have to show that no substantive claim or argument featuring the P/SP distinction is correct – and no one, I trust, thinks that a quick argument of this sort is possible. Therefore, we should not only be suspicious of supposedly quick arguments *for* the importance of the distinction, but equally of supposedly quick arguments *against* it (e.g., Fodor, 1975, pp. 52–3; Machery, 2009, pp. 24–5).

²⁰ We can now see why Drayson's restriction of the subpersonal level to *psychological* (i.e., intentional) descriptions (see §5) receives no support from consideration of the distinction's theoretical role. On the contrary, the distinction is as useful for highlighting the difference in overt subject matter between folk-psychology and neuroscience as it is for highlighting the corresponding difference between folk-psychology and psychological branches of (subpersonal) cognitive science.

8. Objections and Replies

Let me now address some possible worries and objections. To start with, one might object that the proposal is open to counterexamples. In particular, note that the proposal classifies a statement such as (8) as subpersonal:

(8) Paul has a tear in his right calf muscle.

But this looks like the wrong result on intuitive grounds. The problem, it might be said, is that (8) is not about some part or aspect of a person's internal structure *that is in any sense constitutive of, or bears on, Paul's personhood or mental events* and therefore should not count as subpersonal (nor as personal, of course).

One reaction to this objection would be to modify Final Shot in the way suggested by the objection so as to exclude (8) and the like from the category of the subpersonal. Thus, one might define subpersonal-level statements as being overtly about some part or aspect of a person's internal structure *that is constitutive of, or bears on, her personhood or mental events*. However, I think there are reasons for not going this way. How exactly, and where, the distinction between aspects of internal structure that are, and are not, constitutive of personhood and/or mental events is to be drawn is bound to create problems and controversy. For one thing, it is difficult to see how we might exclude claims about calf muscles from the domain of the subpersonal while not also excluding some claims about neural tissue – after all, individual neurons and their connections are arguably just as dispensable for personhood and mental conditions as are calf muscles. For another, many theorists hold that *embodiment* is a central constitutive aspect of cognition (see, e.g., Shapiro, 2011) and they will likely deny that any clear distinction can be drawn between aspects of embodiment that are, and others that are not, relevant in this regard. Accordingly, they might resist the claim that (8) is about a bodily part that is non-constitutive for cognition. As these examples indicate, the modification envisaged would not only make the account less clear and simple, it would also conflict with our goal of making applications of the P/SP distinction as theoretically neutral as possible.

In light of this we should ask whether classifying (8) as subpersonal is really such a great cost. I think not. While it is certainly true that (8) is not a *clear* instance of a subpersonal-level claim, it is arguably also not a *clear non-instance*, i.e., one on which the literature provides a clear negative verdict. Instead, claims like (8) are usually *just not*

considered in discussions that invoke the P/SP distinction. Also, classifying (8) as subpersonal does not subtract from the ability of the distinction to play its assigned role and also does not seem to involve any other serious cost. Conversely, as noted, there are clear costs associated with modifying Final Shot to exclude claims like (8) from the category of the subpersonal. Therefore, counting (8) as subpersonal seems like a price that we should be willing to pay for the sake of greater clarity, simplicity, and theoretical neutrality.

Next, someone might worry that Final Shot involves dubious assumptions pertaining to the notion of ‘transparent aboutness’. The account presupposes that it is ‘obvious’ to everyone who ‘fully understands’ a given statement that it is about certain things, and *not* obvious to them that it is about certain others. However, one might worry that this presupposition is dubious. Haven’t we learned from experimental philosophy, for instance, that there is much more variability in the intuitive judgments of competent speakers with respect to various questions than philosophers have commonly assumed?

This line of objection alerts us to the fact that more needs to be said about who counts as ‘fully understanding/grasping’ statements and concepts for the purposes of the account and why certain things are supposed to be obvious to them. However, with the required clarifications in place, there are good reasons to expect the required agreement in the relevant judgements about subject matter – or so I want to argue.

As indicated by my earlier illustration of the workings of the account, there are two steps involved in the judgements of subject matter that are relevant for the account. The first is a mainly ‘disquotational’ step, as in judging that (5) is about what Paul believes, and (7) is about Paul’s visual module. I don’t see any reason to expect disagreement in these disquotational judgments. In fact, even someone who has no idea what the word ‘module’ means in the context of psychology will agree that statement (7) is about Paul’s visual modules.

The next, less mechanical step involves classifying the disquotational subject matter as belonging to the category ‘mental event or intentional action of a person’ or the category ‘parts or internal structure of a person’. It is here, I take it, that disagreement might be thought to arise and where more must be said about the conditions on ‘full understanding’. There are two main issues to consider. The first concerns the kinds of concepts that typically figure in subpersonal-level claims, such as *module*, *functional subsystem*, *internal representation*, *cerebral cortex*, *neural pathway*, etc. For the account to deliver the desired verdict, it must be the case that someone who

The Personal/Subpersonal Distinction Revisited

fully understands the concept of a module, say, must thereby also know that a module is (by definition) an internal subsystem of a cog-nizer. This strikes me as eminently plausible. If someone thinks a visual module is a skin condition, say, she simply fails to understand the concept. Thus, it certainly seems plausible that clear cases of sub-personal-level claims will come out as such according to Final Shot.

The second issue concerns the classification of certain events, states, *etc.* as *mental*. For the account to deliver the intended results, it must be plain to those with a grasp of the relevant concepts that belief, for instance, qualifies as a mental state. This also seems plausible. As I said earlier, the target notion of the mental is anchored in a list of paradigmatic instances typically conveyed by introductory textbooks. To fully grasp the notion, we can stipulate, is to be familiar with that list. Therefore, plausibly, it should be obvious to those who grasp the notion that belief qualifies as a mental state, and similarly for the other paradigmatic mental events (states, *etc.*) on the list.

However, it might be objected that there is actually *disagreement* among those familiar with the relevant list whether all the listed items actually belong there. For instance, some theorists have argued for the view that all mental states are necessarily conscious and that, therefore, 'standing' or 'non-occurrent' beliefs actually are not mental states (Strawson, 2010). Accordingly, they would deny that statement (5) is about a mental state of Paul when read as attributing a standing belief. More generally, the worry might be that the notion of the mental is too contested to underwrite the agreement in judgements of subject matter required by Final Shot.

There are different possible responses to this line of objection. First, one might point out that those who deny that standing beliefs are mental states presumably also would deny that ascriptions of standing beliefs belong to the personal level. If so, the objection actually supports the account in one respect, *viz.*, by showing that judgements about whether a statement counts as personal really depend, as the account has it, on judgements about whether a statement is overtly about the mental domain. From this perspective, what the objection shows is only that the account does not make the classification of statements as personal-level fully theory-neutral (contrary to what one might have hoped), but not that it should therefore be rejected. In addition, it could be noted that the view underlying the objection is an extreme minority position and so the vast majority of philosophers could still agree that (5) should be classified as personal-level.

Another possible response is to say that those who deny that standing beliefs are mental states are best understood as proposing a *revision* of the relevant concept of the mental. On this view, they do

not really disagree that standing beliefs are mental *as that concept is commonly understood* but propose that the concept *should* be understood differently. While this interpretation may be controversial, it has plausibility in light of the fact that standing beliefs are typically regarded as *paradigmatic* instances of the mental.

A third possible response is to modify Final Shot to make the classification of ascriptions of standing beliefs (and similarly other states whose status as 'mental' might be disputed) as personal-level independent of debates about the concept of the mental. This response could be motivated by what was said in the last section regarding the intended role of the thin P/SP distinction. Assume we are interested in how ascriptions of standing belief relate to subpersonal-level descriptions of a person's cognitive architecture, a familiar question that has attracted much discussion in philosophy of mind (e.g., Dennett, 1978; Fodor, 1987) and that is of interest independently of whether standing beliefs are counted as 'mental'. Invoking the P/SP distinction seems to be just as useful and pertinent in addressing this issue as it is in addressing a corresponding issue regarding ascriptions of conscious states, *viz.*, to highlight that there is a *prima facie* difference in subject matter between the two kinds of statement: one is overtly about a property/state of a person (whether or not it is called 'mental') and the other overtly about the person's internal structure. Accordingly, we might want to make sure that ascriptions of standing beliefs come out as uncontroversial instances of personal-level statements.

One radical implementation of this third response would be to eliminate any reference to the mental (and to intentional action) in the specification of the personal level. To support this idea, it might be said that the role of the P/SP distinction on the present account really comes down to highlighting the contrast between *claims that are overtly about a person's internals and claims that are not*. However, adopting this idea would require us to count statements about such things as the height, weight, time of birth, *etc.* of a person as personal-level, which presumably takes us too far from the common understanding. Moreover, spelling out the idea in a satisfactory way actually results in an account that is cumbersome and non-intuitive.²¹

²¹ A purely negative characterization of the personal level ('statements not overtly about a person's internal structure') would allow in claims about inanimate objects (e.g., 'My desk is white'). Also, we must allow for statements that are *partly* personal-level even though they are also partly about internal parts or structure (e.g., 'Paul's feelings of anxiety are

The Personal/Subpersonal Distinction Revisited

Accordingly, I think a better implementation of the third response is to simply *stipulate* more precisely what counts as ‘mental’ for the purposes of Final Shot. In fact, to avoid fruitless verbal disputes, we can rephrase the account by using the term ‘mental*’ instead of ‘mental’ and stipulate that an item counts as *mental** *iff* it is included in the usual list of paradigmatic examples mentioned earlier or is sufficiently like these paradigmatic examples. This stipulation ensures that ascriptions of all the items on the standard list come out as clear-cut cases of personal-level statements, which is the desired result.

To sum up, I think none of the objections considered in this section provides a compelling reason to reject or substantially modify the proposed explication of the P/SP distinction. I conclude, therefore, that the explication has much to recommend itself all things considered.

9. Conclusion

Despite its wide influence, the P/SP distinction has long been surrounded by an unfortunate lack of clarity. In this paper I have tried to improve on this situation in three main ways, focusing on what I’ve called the *representational* conception of the distinction. First, I have stressed distinguishing *thick* from *thin* versions of this distinction, depending on whether or not they incorporate substantive or controversial assumptions, and I have offered reasons for giving preference to the thin construal. Second, I have proposed that we should take the distinction to apply in the first instance to descriptive statements, and only derivatively to explanations. Finally, I have pointed out problems with *prima facie* plausible proposals for drawing the thin P/SP distinction for descriptive statements and I have proposed an explication that avoids these problems. The distinction, I have argued, should be drawn in terms of what the statements are *transparently about*: personal-level statements are transparently about a person’s mental events (processes, states) or intentional actions, whereas subpersonal-level statements are transparently about a

caused by over-activity of his amygdala’). To accommodate these points, we would need something like the following complex proposal: a statement S is (partly) personal-level iff it is (partly) overtly about some event E involving a person P such that S is not thereby (i.e., in being about E) overtly about P’s internal condition.

person's parts or internal organization. The point of the distinction, thus understood, is to highlight the difference in overt subject matter that separates the typical claims of folk-psychology from those of various branches of cognitive science, which in turn puts into focus the question of their relation.

Acknowledgements

The research for this paper was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation), Grant/Award Number DA 1983/1-1. I'd like to thank Sascha Fink, Fabian Hundertmark, Christian Nimtz, Raphael van Riel, Thorsten Sander, Peter Schulte, Lukas Schwengerer, Thomas Spitzley and audiences in Bielefeld, Essen, and Magdeburg for helpful feedback on earlier versions. I'm also grateful to two anonymous reviewers for this journal, whose comments led to substantive improvements to the paper.

References

- José Luis Bermúdez, 'Personal and Sub-Personal: A Difference without a Distinction', *Philosophical Explorations*, 3 (2000), 63–82.
- José Luis Bermúdez, *Philosophy of Psychology: A Contemporary Introduction* (New York: Routledge, 2005).
- José Luis Bermúdez (ed.), *Philosophy of Psychology: Contemporary Readings* (New York: Routledge, 2006).
- Georg Brun, 'Explication as a Method of Conceptual Re-Engineering', *Erkenntnis*, 81 (2016), 1211–41.
- Mark J. Cain, *The Philosophy of Cognitive Science* (Cambridge, UK: Polity Press, 2016).
- Rudolf Carnap, *Logical Foundations of Probability* (2nd edition), (Chicago: University of Chicago Press, 1962).
- Carl F. Craver, 'The Ontic Account of Scientific Explanation', in Marie I. Kaiser, Oliver R. Scholz, Daniel Plenge and Andreas Hüttemann (eds.), *Explanation in the Special Sciences: The Case of Biology and History* (Dordrecht: Springer, 2014), 27–52.
- Donald Davidson, 'Actions, Reasons, and Causes', *The Journal of Philosophy*, 60 (1963), 685–700.
- Martin Davies, 'Interaction without Reduction: The Relationship between Personal and Sub-Personal Levels of Description', *Mind & Society*, 1:2 (2000a), 87–105.
- Martin Davies, 'Persons and their Underpinnings', *Philosophical Explorations*, 3:1 (2000b), 43–62.

The Personal/Subpersonal Distinction Revisited

- Martin Davies, 'Cognitive Science', in Frank Jackson & Michael Smith (eds.), *The Oxford Handbook of Contemporary Philosophy* (Oxford: Oxford University Press, 2005), 358–94.
- Daniel C. Dennett, *Content and Consciousness* (London: Routledge & Kegan Paul, 1969).
- Daniel C. Dennett, *Brainstorms: Philosophical Essays on Mind and Psychology* (Cambridge, MA: MIT Press, 1978).
- Zoe Drayson, 'The Uses and Abuses of the Personal/Subpersonal Distinction', *Philosophical Perspectives*, 26:1 (2012), 1–18.
- Zoe Drayson, 'The Personal/Subpersonal Distinction', *Philosophy Compass*, 9:5 (2014), 338–46.
- Matthew Elton, 'The Personal/Sub-personal Distinction: An Introduction', *Philosophical Explorations*, 3:1 (2000), 2–5.
- Matthew Elton, *Daniel Dennett: Reconciling Science and Our Self-Conception* (Cambridge, UK: Polity Press, 2003).
- Jerry A. Fodor, *The Language of Thought* (New York: Thomas Y. Crowell, 1975).
- Jerry A. Fodor, *Psychosemantics: The Problem of Meaning in the Philosophy of Mind* (Cambridge, MA: MIT Press, 1987).
- Jennifer Hornsby, *Simple Mindedness: In Defense of Naive Naturalism in the Philosophy of Mind* (Cambridge, MA: Harvard University Press, 1997).
- Jennifer Hornsby, 'Personal and Sub-Personal; A Defence of Dennett's Early Distinction', *Philosophical Explorations*, 3:1 (2000), 6–24.
- Jaegwon Kim, *Philosophy of Mind* (3rd edition) (Boulder, CO: Westview Press, 2011).
- Amy Kind, *Philosophy of Mind: The Basics* (New York: Routledge, 2020).
- Uriah Kriegel, 'Personal-Level Representation', *Protosociology*, 28 (2012), 77–114.
- David K. Lewis, 'Statements Partly about Observation', *Philosophical Papers*, 17:1 (1988), 1–31.
- Edouard Machery, *Doing without Concepts* (Oxford: Oxford University Press, 2009).
- Claudia Maienborn, 'Event Semantics', in Claudia Maienborn, Klaus von Heusinger, and Paul Portner (eds.), *Semantics: An International Handbook of Natural Language Meaning, Vol. 1* (Berlin: De Gruyter, 2011), 802–29.
- Pete Mandik, *This is Philosophy of Mind: An Introduction* (Hoboken: Wiley-Blackwell, 2013).
- David Marr, *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information* (San Francisco: W.H. Freeman, 1982).
- John McDowell, 'The Content of Perceptual Experience', *The Philosophical Quarterly*, 44 (1994), 190–205.
- Georges Rey, 'Physicalism and Psychology: A Plea for a Substantive Philosophy of Mind', in Carl Gillett and Barry Loewer (eds.),

Lars Dänzer

- Physicalism and Its Discontents* (Cambridge: Cambridge University Press, 2001), 99–128.
- Lawrence Shapiro, *Embodied Cognition* (Oxford: Routledge, 2011).
- Liza Skidelsky, 'Personal-Subpersonal: The Problems of Inter-Level Relations', *Protosociology*, 22 (2006), 120–39.
- Dan Sperber and Deirdre Wilson, 'Pragmatics, Modularity and Mind-Reading', *Mind & Language*, 17 (2002), 3–23.
- Galen Strawson, *Mental Reality* (2nd edition), (Oxford: Oxford University Press, 2010).
- Sam Wilkinson, 'Dennett's Personal/Subpersonal Distinction in the Light of Cognitive Neuropsychiatry', in Carlos Muñoz-Suárez and Felipe De Brigard (eds.), *Content and Consciousness Revisited* (Cham: Springer, 2015), 111–27.
- Stephen Yablo, *Aboutness* (Princeton: Princeton University Press, 2014).

LARS DÄNZER (lars.daenzer@uni-due.de) is a postdoctoral researcher in philosophy at the University of Duisburg-Essen, Germany. His publications include *The Explanatory Project of Gricean Pragmatics (Mind & Language, 2021)*, *Utterance Understanding, Knowledge, and Belief (Ergo, 2017)*, and *Sentence Understanding: Knowledge of Meaning and the Rational-Intentional Explanation of Linguistic Communication (Mentis, 2016)*.