

Constructing Persons: On the Personal–Subpersonal Distinction

Mason Westfall, University of Toronto

This is a pre-print of a paper to be published in *Philosophical Psychology*. Please cite the published version when possible.

What's the difference between those psychological posits that are 'me' and those that are not? Distinguishing between these psychological kinds is important in many domains, but an account of what the distinction consists in is challenging. I argue for Psychological Constructionism: those psychological posits that correspond to the kinds within folk psychology are personal, and those that don't, aren't. I suggest that only constructionism can answer a fundamental challenge in characterizing the personal level—the plurality problem. The things that plausibly qualify as personal are motley. Other attempts at accounting for the personal level either cannot accommodate this plurality, or cannot explain what unifies the personal. Given arguments others have given for a pluralistic conception of folk psychology, constructionism explains and predicts this plurality in a systematic and unified way, thereby solving the plurality problem.

Philosophers and cognitive scientists routinely distinguish between the personal and the subpersonal. Some of what we posit as occurring in my mind is 'me' and some of it isn't. What's the difference between the mental stuff that's 'me' and the mental stuff that's 'below' that? I'd like to answer that question.

Though the personal–subpersonal distinction¹ is widely deployed—in, for example, metaethics (Ferrero 2009), economics (Ross 2007) and psychiatry (Hughes 2011)²—it's less often taken up as the object of theorizing itself. For some theorists, that may be fine. If the distinction is merely an intuitive or heuristic one, then we may feel scant pressure to ensure that it is well-understood. However, those of us who believe the distinction is a substantive one, a distinction that does philosophical work, cannot rest so easily. Perplexities arise in trying to explain what the distinction consists in. Should we be unable to resolve the problems, we may be forced to give up the distinction, casting theorizing that relies on it into disrepute. I take that to be an unhappy result, but one that can be avoided.

I propose that we should understand the personal–subpersonal distinction as constructed by our folk psychological capacities: those psychological kinds that correspond to the posits of folk psychology are personal, and those that don't, aren't. I call this view 'constructionism'. I'll argue that only constructionism has the requisite structure to answer a fundamental challenge in char-

¹ I'll sometimes call it 'the distinction'.

² These citations were found in (Drayson 2012).

acterizing the personal level. The things that plausibly qualify as personal are motley. Other attempts at accounting for the distinction cannot accommodate the plurality essential to the personal level. I call this the ‘plurality problem’. By noticing that folk psychology is itself pluralistic, the constructionist can unify and explain the pluralism of the personal level.

1 Clarifying the Distinction

People deploy the ‘personal–subpersonal distinction’ in different contexts, and bring to bear different assumptions about it. I am interested in a class of uses—primarily in empirically informed philosophy of mind and cognitive science—that are meant to capture the same thing. How this class of uses intersects with other uses is an interesting question, but posterior to the project of this paper. My goal is to give an account of this class. We can then query the relationship to other uses. In order to fix ideas, I’ll present a paradigmatic example. The class of uses at issue can be identified as those that are relevantly similar to this paradigm.

Consider the early visual system and conscious visual experience. David Marr (1982), in his seminal work, offers a computational account of the visual system. On Marr’s account, the visual system performs computations to recover three dimensional percepts from two dimensional retinal activation. Some particular two dimensional retinal activation is consistent with myriad three dimensional scenes. In order to recover percepts, the visual system encodes ‘assumptions’ about the sorts of environments we’re likely to inhabit. These ‘assumptions’ enable the visual system to make fallible but ecologically valid ‘inferences’ about the causes of the pattern of retinal activation. For example, the visual system ‘expects’ things to be lit from overhead.³ This ‘expectation’ enables the visual system to resolve otherwise ambiguous scenes, e.g. a convex shape lit from above, induces the same retinal activation as a concave shape lit from below.⁴

Conscious visual perception is the result of these fallible computations. When I experience a scene, I don’t need to resolve ambiguity due to things like light source. My experience is not characterized by a two dimensional array from which *I* need to reconstruct percepts. I just get the percepts. Of course I may have an ambiguous experience, as perhaps occurs when I am unsure whether I recognize a friend walking towards me down the street.⁵ But this depends on prior resolution of ambiguity my visual system does ‘behind the scenes’. To see a convex shape lit from above, I just open my eyes and take a look.

³ See, for example (Ramachandran 1988a, 1988b; Kleffner and Ramachandran 1992).

⁴ See (Thomas, Nardini, and Mareschal 2010) for developmental work on this topic.

⁵ See (Morrison 2016) for a discussion of these cases.

Marr himself does not explicitly invoke the personal–subpersonal distinction, but I read him as tacitly appealing to it:

[L]et us look at the range of perspectives that must be satisfied before one can be said, from a human and scientific point of view, to have understood visual perception. First, and I think foremost, there is the perspective of the plain man. He knows what it is like to see, and unless the bones of one’s arguments and theories roughly correspond to what this person knows to be true at first hand, one will probably be wrong. . . . Second, there is the perspective of the brain scientists, the physiologists and anatomists who know a great deal about how the nervous system is built and how parts of it behave. The issues that concern them. . . . must be resolved and addressed in any full account of perception. And the same argument applies to the perspective of the experimental psychologists. (Marr 1982, 4)

My primary aim is not Marr exegesis, but rather using a Marrian story about vision as a paradigm of the personal–subpersonal distinction. This understanding is relatively widespread in the philosophy of perception. For example, Peacocke (1994, 311) describes the ‘subpersonal computations in the visual system’ responsible for depth perception in the context of a Marrian framework, and Bermúdez (1995, 344) discusses ‘subpersonal computational states of the visual processing system, as understood by Marr’.

The point of the personal–subpersonal distinction is to mark the difference in kind between psychological operations that are *done by the person*, and those that are not. We can distinguish both between psychological states and psychological processes in this way. The assumptions encoded in my visual system are psychological states that are not properly ascribed to me. I *also* believe that things are usually lit from overhead, but that isn’t necessary to resolve low-level scene ambiguity. Indeed, in known illusions, the visual system resolves scene ambiguity in a way that conflicts with the subject’s beliefs. Additionally, the computations performed by my visual system are subpersonal psychological processes. *I don’t perform those computations, my visual system does.* In what follows, I will primarily discuss psychological states, though I intend my discussion to apply to processes *mutatis mutandis*.⁶

⁶ Though I leave open the relation between the distinction as applied to states and as applied to processes. Perhaps personal states are at times involved in subpersonal processes. I take my discussion to be consistent with this possibility.

2 The Plurality Problem

A fundamental challenge in giving an account of the distinction is that it applies variously. Persons don't just have beliefs and desires. They also experience pains and perceptions, moods and emotions. They have skills and personality traits. They have preferences and make inferences. They intend and decide. Persons are variegated.

The 'plurality problem' is the problem of accounting for this diversity in an account of the distinction. Proposals may appear promising when we restrict our attention to only some psychological kinds, but often founder when extended to the whole motley.

Take the proposal that the personal states are just the conscious states. This proposal may be extensionally adequate for perceptual states and pains. Clearly, though, character traits need not be conscious. So although the phenomenal proposal may be a necessary condition on *perceptual* states being a personal states, it does not explain what makes a state personal in general.

Another proposal explains the personal–subpersonal distinction in functional terms: personal states are subpersonal states meeting some further functional specification.⁷ The appeal is obvious. Functional and intentional resources are arguably naturalistically kosher. Moreover, the diversity of topics that have struck philosophers as apt for a functional analysis might give us hope that such an account could have the requisite generality to address the plurality problem.

Indeed, this appeal has been noticed in a number of domains. For example, Jerry Fodor (1975) proposes that to believe that P is to have an internal representation of P tokened in one's belief box. An internal representation, for Fodor, is a string in the language of thought, and 'being in one's belief box' is a metaphorical shorthand for a specification of the functional role of beliefs. David Velleman (1992) proposes that an action is a movement that is appropriately caused by something playing the role of the agent. We identify the functional role of the agent in action, and then identify an aspect of psychology, in non-agential terms, that plays that role. Peter Carruthers (2015) proposes that the conscious states are the globally broadcast states, i.e. the states that are made available to consumer systems. 'Availability to consumer systems' is a functional notion.

⁷ Traditionally, functionalists have been either analytic functionalists—who take the functional roles to be determined by our folk mental concepts (e.g. (Lewis 1966))—or psychofunctionalists—who take the functional roles to be discoveries of empirical psychology (e.g. (Fodor 1968)). I'll argue that we should expect that no unified functional profile captures the personal level, so the discussion in the text applies indifferently to analytic and psychofunctionalism.

In each case we identify an aspect of the personal level, and then account for it by proposing a functional role in other terms, often the terms of (subpersonal) cognitive psychology. Many philosophers have hoped that we can give such a specification of every aspect of the mind. Supposing we did, we might think that we have a reductive account of the personal level in functional terms.

This last thought would be a mistake. However optimistic one is about the prospects of the functionalist project in each particular domain, that optimism should not carry over to identifying a common functional role occupied by all and only personal states. In order to explain the distinction, one needs to explain why *these* functional profiles are the personal ones. Were the functionalist to complete their project, they would be left with a massive disjunction of functional profiles corresponding to the myriad personal kinds. Though this disjunction may be extensionally adequate, it is not explanatory.

Put another way, perhaps what it is to believe that P is to have an internal representation of P tokened in one's belief box (Fodor 1975); perhaps what it is for something to be an action is for it to be appropriately caused by something playing the role of the agent (Velleman 1992); perhaps the conscious states just are the globally broadcast states (Carruthers 2015). But that does not explain why these functional profiles, and not others, should characterize kinds that are attributed to *me*. What distinguishes the personal from the subpersonal in a particular domain is a different question from what the personal–subpersonal distinction consists in. This is the plurality problem.⁸

How seriously should we take the plurality problem? One attitude is: not seriously at all. Sure, we attribute all sorts of states to persons, but only as a kind of loose talk. Analogously, we might attribute digestion to Noam, but surely nothing metaphysically interesting follows. That's just loose talk. Perhaps the same is true of the kinds of attributions to persons that I've suggested generate the plurality problem. I don't find this attitude attractive. The difference between attributing digestion to Noam, and the psychological attributions reviewed here is that while the former may be loose talk, the latter is not. When I say that Noam perceives, or draws an infer-

⁸ A different proposal would be to understand the distinction in terms of two different kinds of explanation, as Drayson (2012; 2014) suggests. I consider this suggestion in detail in Section 4.

ence, or is extroverted, I *mean it*. When I say that Noam digests, perhaps I don't.⁹ Those who, like me, mean it when they attribute diverse psychological kinds to the person must take the plurality problem seriously. The alternative attitude unpalatably condemns work that appeals to the distinction to the category of loose talk. For example, when Boghossian (2014, 2–3) characterizes inference as a personal-level process, when Curry (2021, 7896–7904) considers whether beliefs are identical to subpersonal or personal states, when Lyons (2016) argues that personal but not subpersonal beliefs can be part of a subject's evidence, when theorists debate whether hypotheses about social cognition should be understood as personal or subpersonal,¹⁰ I take it they are not engaging in loose talk. If not, though, then the plurality problem must be faced.

3 A Constructionist Account of the Distinction

I believe constructionism is the way to solve the plurality problem. Before explaining what I mean in detail, we need some terminology from social metaphysics.

Consider the fact that the bit of matter in my pocket is a dime. This fact is 'social', in the sense that it obtains in virtue of a variety of facts about social organization. Among them are facts about this bit of matter's history—how it was produced—as well as facts about the broader norms and rules that we collectively accept—e.g. which modes of production count as producing dimes. Brian Epstein (2015) argues that the fact that my bit of matter is a dime depends on these two kinds of facts in different ways. Facts about my dime are partially *grounded* by facts about how the bit of matter was produced. The production process is what makes something a dime. In order for the mint to be able to produce dimes, though, we must collectively accept various rules and conventions that govern money. Epstein calls this second kind of dependence *anchoring*.¹¹

Social facts are anchored and grounded. The fact that we collectively accept that being over 65 counts as being a senior citizen anchors a general rule, or 'frame principle', that anyone over 65 is a senior citizen, and the fact that my grandpa is over 65 grounds the fact that he is a senior citi-

⁹ An anonymous reviewer suggests that in fact we *do* mean it when we say that Noam digests, and that this can be treated analogously to the treatment I offer of psychological attributions below. Readers who share this assessment are invited to substitute an example they *do* think is mere loose talk—perhaps talk of people 'metabolizing'. The important point, for my purposes, is that we have good reason to resist treating personal psychological attributions as loose talk, whether or not the same holds for digestion.

¹⁰ For contributions to this debate, see (Herschbach 2008; Spaulding 2010, 2015; Musholt 2018).

¹¹ See Epstein (2015, Chapter 6) for a discussion of grounds and anchors for money.

zen.¹² This framework enables us to clarify the view I am arguing for, and to consider its implications.

Recall the plurality problem. The states that count as personal are various, so it is difficult to see how to unify them—to explain what the personal level consists in. I suggest that the solution to the plurality problem is psychological constructionism. The personal–subpersonal distinction is anchored by folk psychology. Kinds that correspond to those posited by folk psychology are personal; those that don’t, aren’t. Derivatively, a psychological explanation is personal just if it adverts to kinds found in folk psychology, and subpersonal otherwise.

I am adapting Epstein’s framework, which is developed for *social* construction, and applying it to *psychological* construction. The primary difference is in how anchoring works. Rather than facts about *collective acceptance* anchoring the constructed domain, facts about individual psychology anchor the constructed domain—for the personal–subpersonal distinction, facts about folk psychology. Facts about the personal–subpersonal distinction are not anchored by facts about collective acceptance, but by facts about our psychology.¹³ A social constructionist view would enjoy some of the explanatory benefits I’ll suggest my view has, but I prefer the psychological version because the facts about social organization that a social version would advert to

¹² Perhaps anchors are just a special case of grounds. See (Epstein 2015, Chapter 9) for an argument against this view. If one thinks that anchors are species of ground, they can simply read ‘anchors’ and ‘grounds’ as referring to two different kinds of grounding. Epstein’s framework itself is neutral on this question, as am I, but it eases exposition to maintain different terminology.

¹³ An anonymous reviewer raises the question ‘whose psychology?’. This question becomes pressing to the extent one thinks that substantial interpersonal or cultural variation exists with respect to folk psychology, as argued in e.g. (Wu and Keysar 2007; Heyes and Frith 2014; Heyes 2018; Curry 2020). To the extent that this turns out to be correct, the constructionist should be prepared to accept a degree of indeterminacy or relativism about the personal, in a way structurally similar to (Curry 2020). However, the kind of variation that would license such relativism about the personal is more fundamental than the kind of variation at issue for Curry’s relativistic interpretivism about belief. Cultural variation in the propensities to offer e.g. dispositional or situational explanations of behavior (see Krull et al 1999) would *not* suffice to license relativism about the personal. Rather, we would need cultural variation in the *ontology* of folk psychology, for example, a culture whose folk psychology did not explain behavior by appeal to beliefs *at all*. Though I think in general we find predominantly commonality, see Section 6 for a potential example of this kind of variation. See (Lavelle 2021) for discussion.

arguably already need a personal–subpersonal distinction, and so circularity worries are more serious.¹⁴¹⁵

This proposal solves the plurality problem, because our folk psychology is itself pluralistic. Others have argued for a pluralistic conception of folk psychology independently. In brief, pluralistic approaches to folk psychology hold that the classical discussions of folk psychology exhibited an unduly narrow focus on predicting and explaining behavior on the basis of propositional attitude ascriptions. According to pluralists, folk psychology serves a variety of goals. Our distinctively social forms of intelligence involve predicting and explaining behavior, but also regulating behavior, inducing beliefs and desires in others, and cooperating or competing with them. We accomplish these various goals by deploying many cognitive strategies, among them attributing propositional attitudes, but also attributing personality traits, deploying stereotypes, relying on features of the situation, constructing models of individuals, and so on.¹⁶ As Andrews puts it:

Folk psychology is a social competence, which includes the ability to identify, predict, explain, justify, normalize, and coordinate behavior... The social competencies of folk psychology are subsumed by a number of different cognitive mechanisms, and one's degree of success as a folk psychologist is a function of the number of compe-

¹⁴ I address the question of circularity for my own view in section 3.2. I'm pessimistic a social constructionist view could reply in the same way.

¹⁵ For this reason, I prefer to construe 'mindshaping' proposals like (McGeer 2007; Zawidzki 2013), which I will invoke below, as pointing out that normative and regulative dimensions of social intelligence constitute important and often overlooked components of folk psychology, rather than taking folk psychology to primarily be a social practice, not a component of cognition. Of course, there are many such social practices that are themselves interesting things to theorize about. The important point for my purposes is just that the pluralist about folk psychology can take on board important insights from discussions of mindshaping without losing the distinction between folk psychology as a cognitive system and the kinds of social practices that it supports. Readers who judge it best to give up this distinction should likewise give up the distinction between psychological and social construction, though they could still endorse a version of constructionism. I think circularity worries would be pressing for such a view, but perhaps nevertheless superable. Thanks to an anonymous reviewer for encouraging me to clarify this point.

¹⁶ See (Andrews, Spaulding and Westra 2020) for a helpful overview of pluralistic approaches to folk psychology. See also (McGeer 2007; Maibom 2009; Andrews 2012; Zawidzki 2013; Spaulding 2018a; Westra 2018).

tencies mastered and the degree of facility with the various competencies. (Andrews 2012, 11)

My goal is not to argue for pluralism about folk psychology, but to assume it in service of my argument for constructionism. Note though, that only a pluralistic conception of folk psychology has the requisite structure to answer the plurality problem. So we find additional motivation for a pluralistic conception of folk psychology, since it's apparently necessary to account for the personal–subpersonal distinction.

To see how constructionism works in detail, consider a specific pluralist proposal about folk psychology: model theory (Godfrey-Smith 2005; Maibom 2009; Spaulding 2018a, b). According to model theory, folk psychology is constituted by constructing a variety of cognitive models of other kinds of people and individuals, and relying on them for a variety of purposes. There is a 'core folk-psychological model' that can be used as is, or elaborated and varied depending on the circumstances (Godfrey-Smith 2005, 10). As Godfrey-Smith puts it:

The core model includes such features as a distinction between beliefs and desires or preferences, the idea of sensory input and behavioral output, and the characteristic dependence of action on perceptions, memories, goals, and temptations. But it is easy to introduce more detail, and this can be done in a range of different ways.

(Godfrey-Smith 2005, 10)

Assuming model theory, the constructionist offers the following answer to the plurality problem: experience, behavior, memories, goals, temptations, and other kinds that may get added to our models, depending on circumstances, are personal in virtue of figuring in the cognitive models that constitute our folk psychological competence. Given the plurality of kinds that figure in our cognitive models, constructionism predicts and explains the plurality that characterizes the personal level.

Constructionism avoids the shortcomings of functionalism. The functionalist needed, quixotically, to identify a functional profile associated with all and only the personal kinds, or, failing that, to stipulate that the personal level is identical to a highly disjunctive functional kind. Because folk psychology is pluralistic, we can account for both the unity and plurality of the personal level. The personal is unified, insofar as all of the personal kinds are quantified over by folk psychology. The kinds posited by folk psychology are diverse, so we should expect a domain consti-

tuted by folk psychology to be diverse. Indeed, that is what we find, as the plurality problem makes vivid.¹⁷

The support for constructionism provided by the plurality problem is tentative. Perhaps other proposals I have not considered could offer satisfying explanations of plurality. If so, further work would be needed to adjudicate between these competing explanations. I am pessimistic that such alternatives will be forthcoming, since the most obvious and plausible competing proposals don't work. But I haven't considered every possible proposal, and readers who find constructionism implausible may hope that a heretofore unarticulated proposal can answer the plurality problem. They are warmly invited to articulate it, and in so doing undercut much of my argument for constructionism. In the absence of such a proposal, though, constructionism is preferable to no solution to the plurality problem at all. So, tentatively, we should accept constructionism.

3.1 Clarifying Constructionism: Triviality

It's easy to hear this proposal as expressing a triviality. *Of course* those states that we tend to attribute to persons are also the ones that we think are attributable to persons. You haven't given an *account* of the distinction, you've merely described how we think. The proposal is not trivial though. What it is for a state to be personal rather than subpersonal concerns the metaphysics of mind. The proposal that the right answer adverts to how we think about one another is not trivial. To see this, it's useful to consider constructionist views in other domains.

Imagine being a constructionist about water. That is, consider the view that *what it is* to be water is to be the sort of thing we think of as 'water'. If my proposal were trivial as applied to the personal, the analogous view about water should be trivial as well. Not only is it not trivial, it's false. Water is not constituted by our folk theory of water, but by its chemical structure.

Now imagine being a constructionist about money. This involves holding—plausibly—that a proper account of the kind will make ineliminable reference to our thought and practice involving money. It is also—persistent popular misunderstandings notwithstanding—*not* to say that money is 'unreal', or 'imaginary', or 'all in our head'. More philosophically, saying that money

¹⁷ The constructionist proposal could be construed as an unusual functionalist proposal. The relevant functional role would be specified in terms of occasioning folk psychological representations. I myself prefer to use 'functionalism' to refer more specifically to views that identify a kind with its overall functional profile. Doing so eases exposition of my view. I take this to be a verbal matter, though. If one were to insist that I am arguing for a specific kind of functionalism as against other kinds of functionalism, I would not object.

is socially constructed does not commit one to *antirealism* about money. *Of course* money is real —‘as real as anything’ (Barnes 2017, 2423)—but, the constructionist suggests, the features of reality that constitute money’s reality are social, rather than chemical.

Constructionists about the personal–subpersonal distinction hold that the personal is more like money, and less like water. That’s not trivial. I say ‘more like’, because I endorse psychological construction, rather than social construction for the personal–subpersonal distinction. In my view, the distinction is anchored by folk psychology, rather than broader facts about our social practices.¹⁸

Water and money are relatively uncontroversial examples. Other topics exhibit active controversies between constructionist and non-constructionist views.

Take the metaphysics of race. Quayshawn Spencer (2019a) argues that races are biologically real. On Spencer’s view, races are human continental populations. Moreover, human continental populations are biologically real, in the sense that they are ‘epistemically useful and justified entit[ies] in a well-ordered research program in biology’. By contrast, Sally Haslanger (2000, 2012, 2019) defends a social constructionist account of race. On Haslanger’s view, race is a social category, not a biological one.¹⁹

Haslanger and Spencer agree that our folk concepts of races track races, and do so imperfectly. After all, our folk conceptions are messy, and somewhat contradictory. If our folk concepts *perfectly* tracked the underlying reality, then philosophy and science would be much easier than they actually are. A proposal about what race is that was *completely* divorced from our folk usage, though, could fairly be accused of changing the subject. The disagreement is about what kinds of

¹⁸ One reason *not* to include extra-psychological facts among the anchors is that the personal–subpersonal distinction arguably contributes to delimiting the social sphere. If that turns out to be right, then we would have reason to resist including social facts in our specification of the distinction, to avoid circularity.

¹⁹ One argument Haslanger presents for her constructionist view is that it can *explain* and *unify* the diversity in racial categories across time and space (Haslanger 2000). Though the analogy is imperfect, and the arguments are logically independent, my argument for constructionism about the personal–subpersonal distinction is similar in spirit to Haslanger’s.

facts determine the race facts. Spencer says biological facts determine the race facts; Haslanger says social facts.²⁰

These sorts of disagreements are not unfamiliar in the philosophy of cognitive science. Take a debate about emotions. One tradition in psychology is the ‘basic emotion paradigm’.²¹ Basic emotions are psychological states, characterized by their functional role, that have been selected by evolution to register ‘core environmental themes’ like danger, onto adaptive behavioral repertoires like avoidance. These emotions are psychologically real, and figure in cognitive psychological explanations.²²

By contrast, some psychologists are constructionists. On this view, psychological theorizing should not quantify over emotions when offering scientific explanations. Rather, our experience of emotion and thought about emotion should be explained in terms of other psychological kinds.²³ So, emotional experiences are to be understood as conceptualizations or construals of core affect or bodily changes, in light of the situation subjects find themselves in. As Feldman Barrett puts it, emotions are not natural kinds; they are social kinds (Barrett 2006).

The constructionist and basic emotion theorist agree that emotions are real. Moreover, they agree that our emotion concepts (imperfectly) track emotions—we sometimes misunderstand one another, and often succeed in understanding one another. They disagree about which facts constitute the emotional facts. Basic emotion theorists take some emotions to be psychological kinds. The constructionist, instead, holds that emotions are constituted by our conceptualizations and construals.

These examples illustrate that constructionism about the personal–subpersonal distinction is not trivial. Perhaps it’s trivial that our folk concepts of personal states (imperfectly) track personal states—if personal states exist. After all, a view on which this were not the case could fairly be accused of changing the subject. Admitting this much is to leave open the metaphysical question:

²⁰ More precisely, Spencer offers a ‘radically pluralist’ theory of race, on which *one* kind of race talk—‘OMB race talk’—tracks biological facts, and other kinds of race talk don’t (Spencer 2019b, 213). Thanks to an anonymous reviewer for encouraging me to clarify this point.

²¹ Historically, this tradition is associated with Charles Darwin (1872/1998). Among contemporary psychologists, Paul Ekman (1999) is a prominent exponent.

²² For an overview, see (Tracy and Randles 2011).

²³ Historically, this tradition is associated with William James (1884). Among contemporary psychologists, Lisa Feldman Barrett (2017) is a prominent exponent.

what sorts of facts determine the personal–subpersonal distinction? Psychological constructionism answers that the distinction is anchored by facts about folk psychology.

3.2 Clarifying Constructionism: Circularity

Another reaction is to hear my proposal as circular. In talking about folk psychology, surely I am covertly appealing to a personal–subpersonal distinction, and as such illicitly assuming what was supposed to be explained. The circularity complaint can be elaborated in two different ways.

First, the complainant might observe that folk psychology itself adverts to persons. Second, they might think that *in order to specify* folk psychology, we must make use of the distinction.

The first elaboration is mistaken. References to the personal within our folk theory do not constitute a circularity. Such references occur within an intentional context, and so do not presuppose an independently specifiable distinction. In the same way, we might explain truths about Sherlock Holmes in terms of the books Sir Arthur Conan Doyle wrote. The books purport to refer to Holmes, but that isn't circular.

The second elaboration requires more delicacy. I take the thought to be that 'folk psychology' refers to a personal capacity. So to pick it out, we must rely on the personal–subpersonal distinction. I am not persuaded for two reasons. First, because even if folk psychology is a personal capacity, we do not need to rely on the personal–subpersonal distinction to grasp folk psychology. What we need to rely on is *folk psychology*—our capacity to understand others, and their psychological properties. But, it seems to me, the personal–subpersonal distinction is not ordinarily an aspect of folk psychology. Rather, the distinction is a theoretical posit. Many people have not studied philosophy or cognitive science, and so may not know about the personal–subpersonal distinction. Nevertheless, they are in a position to grasp others' capacity for interpersonal understanding—i.e. to grasp others' folk psychology. So our capacity to identify folk psychology does not rely on a prior understanding of the personal–subpersonal distinction.

Second, we could specify folk psychology scientifically. On this approach, folk psychology is a cognitive system, with a particular internal structure, posited by psychologists to explain a suite of observed capacities. Folk psychology is an empirically motivated posit in psychology, akin to the early visual system. I've just told you what folk psychology is without adverting to the distinction.²⁴

²⁴ These two thoughts are complimentary. We can think about others' folk psychological capacities using our own folk psychological capacities, and in doing so we identify a phenomenon that psychologists then go on to investigate scientifically.

The complainant may point out that the capacities this empirical posit explains are personal capacities, and so re-raise their initial complaint. But this objection must be dismissed. The complainant is not objecting to constructionism in particular. Rather, they are objecting to any account of the personal–subpersonal distinction in terms of empirical psychology. Any such account will involve identifying and explaining persons’ capacities in terms of cognitive subsystems with particular features. That is, it will involve offering psychological explanations. This objection would prove too much. It’s not circular to offer an explanation of the personal–subpersonal distinction in terms of empirical psychology, even though psychological explanation involves identifying persons’ capacities in order to explain them.

Folk psychology may correspond to a kind posited in folk psychology. Indeed, I expect that it does. That is, folk psychology, as specified by the cognitive scientists, may itself contain a representation of folk psychology.²⁵ If so, then the constructionist holds that folk psychology is a personal system, since a system is personal if it is represented in folk psychology. But the dialectically relevant point is that we did not have to rely on the personal–subpersonal distinction to say what folk psychology is. We just asked the cognitive scientists.

4 Drayson’s Objection: Explanations and States

Zoe Drayson (2012) argues against thinking about the personal–subpersonal distinction as I recommend. In fact, we should not think about the distinction as in the first instance between psychological *states*. Rather, the personal–subpersonal distinction is primarily a distinction between two styles of *explanation*. We can define a ‘derivative’ distinction between kinds of states, but only by rejecting ‘one of the most common’ metaphysical interpretations of the explanatory distinction (Drayson 2012, 1). I don’t find Drayson’s argument persuasive, and worry that Drayson’s own account struggles with the plurality problem. I’d like to explain why.

4.1 The Distinction and Styles of Explanation

On Drayson’s view, the personal–subpersonal distinction is a special case of a more general distinction between horizontal and vertical explanation (Drayson 2012, 2). Horizontal explanations are causal explanations: citing earlier events to explain the occurrence of a later event. Vertical explanations, instead, involve explaining a thing’s capacities or dispositions in terms of its parts

²⁵ This representation would presumably not be an exhaustive characterization of the internal structure of folk psychology. That would not be computationally tractable, nor would it be psychologically plausible. We have to *do psychology* to figure out the internal structure of folk psychology.

and components (Drayson 2012, 2). Drayson adapts the horizontal–vertical distinction from Jaegwon Kim (Drayson 2012, 2).²⁶

As applied to psychological explanation, Drayson avers that folk psychology corresponds to horizontal explanation and cognitive psychology corresponds to vertical explanation. She argues that the personal–subpersonal distinction is a special case of the horizontal–vertical distinction. Are we explaining why a person behaved as they did? If so, we are engaged in horizontal—i.e. personal—explanation. Or, are we explaining a person’s cognitive capacities? If so we are engaged in vertical—i.e. subpersonal—explanation. This understanding of the distinction is primarily motivated by fidelity to Daniel Dennett’s (1969) seminal discussion of the distinction.²⁷ There, Dennett is explicit he means the distinction to be between two ‘levels of explanation’ (Dennett 1969, 93).

If we assume that the distinction is between two styles of explanation, it does not follow that distinct states and processes correspond to those styles of explanation (Drayson 2012, 8). Indeed, some views entail that there will *not* be such a distinction. For example, if one thinks explanations are instrumental, and so their terms might fail to refer, one can endorse the personal–subpersonal distinction between kinds of explanation without committing to the existence of personal or subpersonal states at all (Drayson 2012, 9). Or one could endorse personal explanation, but think that things like belief are not best thought of as mental *states*. Helen Steward (1997), for example, thinks the ontological commitments of folk psychology are to abstract entities, rather than ‘unobservable concrete ones’ (242). As Drayson puts it ‘[o]n this view, to be in a mental state is to instantiate a certain relation to a proposition, rather than to possess an internal state’ (Drayson 2012, 9).

These considerations counsel caution in making assumptions about how exactly the parts of our theory fit together. I do not think, though, that they undermine a distinction between personal and subpersonal states. I’ll argue that Drayson is too pessimistic about the state distinction on the assumptions she takes to make trouble for it, and that there’s reason to take the state distinction, not the explanatory distinction, to be primary.

²⁶ Kim’s distinction is originally between two kinds of determination—‘horizontal causation’ and ‘vertical determination’ (Kim 2007, 36). I take it that styles of explanations can be derivatively characterized as explanations that purport to advert to the relevant kinds of determination.

²⁷ Though, of course, this distinction was already present in some form in Wittgenstein (1953/2010) and Ryle (1949/2009), among others.

4.2 Fodor and Dennett Against the State Distinction?

Drayson finds purported reasons for not applying the distinction to states in the work of Jerry Fodor and Daniel Dennett, though their reasons are opposed. For one, terms for personal states refer to subpersonal states (Fodor 1975, 198); for the other, they don't refer at all (Dennett 1969, 96). These thoughts require some unpacking.

For Fodor, 'having a particular propositional attitude is being in some computational relation to an internal representation' (Fodor 1975, 198). We can give a reductive explanation of personal states—propositional attitudes—by identifying them with subpersonal states.²⁸ As Drayson puts it, 'when we identify the posits of personal and subpersonal explanations...we lose any notion of a distinction between personal and subpersonal *states*' (Drayson 2012, 11, emphasis original). If personal states like propositional attitudes *just are* subpersonal states, then there's no distinction between personal and subpersonal states. Better, thinks Drayson, to reserve the distinction for the well-understood distinction between kinds of explanation.

Drayson's line of thought is questionable. It's a substantive, exciting claim that personal states *just are* subpersonal states. Supposing it is true, though, we don't lose a distinction between personal and subpersonal states. Personal states may just be subpersonal states, but not all subpersonal states are personal states. On this view, personal states are subpersonal states *that meet some further conditions*. There's a perfectly good distinction between the subpersonal states that meet those further conditions, and those that don't. So we can endorse Fodor's proposal that propositional attitudes are subpersonal states occupying a particular functional role, while maintaining a distinction between those states and subpersonal states that don't occupy the relevant functional roles.

Drayson points out an ambiguity in our use of 'personal' and 'subpersonal' as applied to states. Given the possibility of reducing personal states to subpersonal states meeting some further conditions, we should be clear about whether we are using 'subpersonal states' to refer to all subpersonal states, or only subpersonal states that are not themselves personal states. This is an important ambiguity, but however it is resolved, it is not a threat to the distinction between personal and subpersonal states.²⁹

²⁸ Recall my earlier discussion of functionalism.

²⁹ On an inclusive understanding of 'subpersonal states', the set of personal states is a subset of the set of subpersonal states. On an exclusive understanding of 'subpersonal states', the set of personal states and the set of subpersonal states are disjoint.

As I read her, Drayson ultimately agrees with the preceding, but she thinks we should mark the distinction between states in other terms, not ‘personal’ and ‘subpersonal’. Rather, she suggests we use Stephen Stich’s (1978) terminology: ‘doxastic’ and ‘subdoxastic’ (Drayson 2012, 12). Drayson thinks it is more intellectually hygienic to reserve ‘personal’ and ‘subpersonal’ for styles of explanation, and use ‘doxastic’ and ‘subdoxastic’ to pick out different kinds of states.

At this point, the matter is primarily terminological, but it’s worth lingering on why Stich’s terminology is inapt. In short, Drayson’s proposal runs afoul of the plurality problem. ‘Personal’ and ‘subpersonal’ naturally extend beyond propositional attitudes in a way that ‘doxastic’ and ‘subdoxastic’ do not. Among a person’s attributes are their doxastic states, but persons also have personality traits, skills, desires, urges, experiences and so on. The personal level involves all of these. We could stipulate that all of the personal attributes are ‘doxastic’ if we wanted, but this would be unmotivated. Skills don’t qualify as personal because they’re deep down ‘really like beliefs’.

To the extent we are concerned with fidelity to a distinction’s origins, we have affirmative reason to reject Stich’s distinction as a replacement. For Stich, two features distinguish beliefs from subdoxastic states: access to consciousness and inferential integration (Stich 1978). Neither of these features can overcome the plurality problem. Character traits, preferences and skills clearly need not be accessible to consciousness, nor does inferential integration even seem to make sense when applied to character traits, skills or inference itself. None of this is an objection to Stich. Stich may be right that beliefs and subdoxastic states differ in these respects, and that this constitutes good reason to retain the concept of belief in our psychological theorizing. But I do think this gives us reason for hesitance in using Stich’s distinction in lieu of a personal–subpersonal distinction between states.

The situation with Dennett is more delicate. Dennett’s distinction is, primarily, one between different kinds of explanation. It’s less clear that the kinds of explanation he has in mind are horizontal and vertical, though. I’ll suggest this is good news for Dennett, because understanding the personal–subpersonal distinction in vertical and horizontal terms is fraught. First, let’s consider the exegetical question.

Dennett is apparently comfortable with horizontal subpersonal explanations. For example:

There should be possible some scientific story about synapses, electrical potentials and so forth that would explain, describe and predict all that goes on in the nervous system. If we had such a story we would have in one sense an extensional theory of behaviour, for all the *motions* (extensionally characterized) of the animal caused by the activity of the nervous system would be explicable and predictable in these extensional terms, but

one thing such a story would say nothing about was *what the animal was doing*. (Dennett 1969, 78)

Granted, the ‘scientific story’ Dennett imagines is non-intentional, and so arguably not psychological. But on Dennett’s view, the content-involving explanation is a *reinterpretation* of the non-intentional explanation.³⁰ Assigning content provides ‘an invaluable heuristic advantage’, but there is no need for the content assignment to be a personal matter (Dennett 1969, 79). If assigning content is primarily of heuristic value, then we should expect the resultant explanation to be a horizontal subpersonal one.³¹

Interpreting Dennett as thinking personal explanations are horizontal is also fraught. Pain is his most thoroughly considered example (Dennett 1969, 90–6). He is diametrically opposed to Fodor’s view:³²

An analysis of our ordinary way of speaking about pains shows that no events or processes could be discovered in the brain that would exhibit the characteristics of ‘mental phenomena’ of pain, because talk of pains is essentially non-mechanical, and the events and processes of the brain are essentially mechanical. (Dennett 1969, 91)

Dennett’s insistence that ‘talk of pains is essentially non-mechanical’ is striking. However plausible, it is difficult to square with construing personal explanation as horizontal, i.e. causal. This is not an absentminded remark. Dennett is expressing the core of his attitude toward personal and subpersonal explanation:

Abandoning the personal level of explanation is just that: *abandoning* the pains and not bringing them along to identify with some physical event. The only sort of explanation in which ‘pain’ belongs is non-mechanistic; hence no identification of pains or painful sensations with brain processes makes sense, and the physical, mechanistic explanation can proceed with no worries about the absence in the explanation of any talk about the

³⁰ This appears to be a very early articulation of ideas about intentional explanation he went on to develop in greater depth (Dennett 1989, 1991). For a more detailed discussion of Dennett’s attitude towards content assignments, see (Elton 2000).

³¹ I consider this verdict in more detail in [4.3](#).

³² Though he is writing before Fodor, so the disagreement is tacit. One could endorse both Fodor’s view about propositional attitudes and Dennett’s view about pains. My concern is with their attitudes toward the personal–subpersonal distinction, not towards pains or propositional attitudes as such.

discrimination of unanalysable qualities. What is the physical explanation to be? Something like this. When a person or animal is said to experience pain there is afferent input which produces efferent output resulting in certain characteristic modes of behaviour centering on avoidance or withdrawal. (Dennett 1969, 94, emphasis original)

It's difficult to read this as expressing a view on which personal level explanation is horizontal causal explanation. Indeed, Dennett concludes his discussion by going even further: '[g]iven this interpretation it is in one sense true that there is no relation between pains and neural impulses, because there are no pains; "pain" does not refer' (Dennett 1969, 96).

I think this is good news for Dennett, because there is reason to doubt that the personal–subpersonal distinction is an instance of vertical and horizontal explanation. Some personal explanations are vertical, and some subpersonal explanations are horizontal. These constitute counterexamples to Drayson's proposal.

4.3 The Subpersonal is not (Only) Vertical

The existence of subpersonal horizontal explanations is clear, on reflection. Horizontal explanations are just causal explanations. Causal explanations often advert to causes other than personal psychological states.

Recall Marr's account of the early visual system. One thing Marr does is offer the resources to give a vertical explanation of some visual capacities. But those same resources make available many horizontal explanations. My current visual experience involves the experience of many particular edges, for example. I suggest a causal explanation exists adverting to particular computations performed by my visual system that were causally responsible for the edges that show up in my visual experience. I'm not in a position to give that causal explanation in any detail, but it nevertheless exists.

Consider the situation more generally. Dennett is surely right that for any bodily movement there is a sufficient neurological explanation. In principle, causal explanation of my bodily movements can be given adverting just to facts about neural activation, together with the facts about my physical environment. Of course, we aren't currently in a position to give many explanations like this, but in principle we could.³³

³³ Note that the existence of such subpersonal causal explanations is a presupposition of the causal exclusion argument—originally raised by Norman Malcolm (1968) and popularized by Jaegwon Kim (1989, 1992, 2000)—and the vast literature on mental causation it precipitated. See Bennett (2007) for a useful review.

More epistemically tractable examples are available. Consider Brent, in a foul mood. Brent explains his mood by adverting to his depleted serotonin reserves, caused by the neurochemical interactions MDMA produced, after he took a pill his friend gave him at the concert. Depending on which neuroscience courses Brent has taken, he may be able to give a somewhat detailed and accurate explanation of the causal mechanisms that produced his foul mood.³⁴

These explanations are diachronic causal explanations that advert to subpersonal states and processes. Are they personal explanations or subpersonal explanations? Drayson's view apparently classifies them as personal, in virtue of being diachronic causal explanations. It's substantially more plausible, though, to treat these as subpersonal explanations. If so, then they constitute counterexamples to Drayson's proposal.

Moreover, we have a good explanation for this result. I suggest—*contra* Drayson—that personal and subpersonal explanations are to be understood, derivatively, in terms of the kinds posited in the explanation—personal explanations advert to personal states, subpersonal explanations, subpersonal states. This proposal delivers the plausible result for these cases, and motivates giving an account of the distinction as applied to states that is prior to the distinction as applied to explanation.

4.4 The Personal is not (Only) Horizontal

The proposal that personal explanations are horizontal sounds plausible when we restrict our attention to explaining behavior by appeal to propositional attitudes, but we offer many different kinds of personal explanations. Once the plurality of personal explanations is appreciated, Drayson's proposal looks much less appealing. More directly: we often offer *vertical* personal explanations. Rather than narrowly focusing on predicting and explaining behavior, folk psychology is a more holistic means of understanding people.

³⁴ An anonymous reviewer points out that, depending on the details of the case, it is possible to imagine that adverting to serotonin levels has become enmeshed in Brent's folk psychology, such that this explanation would, by the lights of the theory advanced here, qualify as personal. Perhaps so. The dialectically relevant question is whether there are also cases in which the explanation adverted to is not enmeshed in a subject's folk psychology. If so, then we have identified counterexamples to Drayson's proposal. I myself see no obstacle to someone offering this kind of explanation without incorporating serotonin into their folk psychology. The objection to Drayson maintains its force, though, even if such explanations were only ever given by professional cognitive scientists in the lab. So one can instead imagine much more complex and intricate causal explanations that could be offered by experts but that resist being assimilated to our folk psychological explanatory frameworks.

Suppose I want to explain the fact that John is a good philosopher. Perhaps I could offer a horizontal explanation—mention his studying under an advisor renowned for their pedagogy, or his time spent reading literature which partly explains his prose, or what have you. Just as naturally, though, I could offer a vertical explanation. I could point to the subcapacities and traits he has that together constitute his being a good philosopher. Perhaps he’s a good philosopher because he’s a careful listener, comes up with clever objections, and never loses sight of the bigger picture. As Drayson characterizes vertical psychological explanations, they explain ‘complex psychological capacities, such as depth perception or language acquisition, by breaking them down into simpler subcapacities that combine to produce the complex phenomena’ (Drayson 2012, 5). Clearly, the explanation just offered fits this characterization.

Again, we should ask whether the proposed explanation is personal or subpersonal, and again, I suggest the plausible result runs contrary to Drayson’s proposal. Drayson is apparently committed to categorizing this as a subpersonal—because vertical—explanation. My own inclination is to treat this as a personal explanation. I suggest the explanation for this inclination is that the subcapacities adverted to are capacities of the person, not of subsystems. So again, we have reason to think that the distinction as applied to states and processes is prior to the distinction as applied to explanations.³⁵

So I think we have strong reasons not to assimilate the personal–subpersonal distinction to the vertical–horizontal distinction.

5 The Subpersonal and the Nonpersonal

³⁵ One might object that the subcapacities adverted to in this explanation are themselves psychological, whereas the subcapacities Drayson is interested in are nonpsychological. I do not think that this is correct, though. Drayson herself stresses that the vertical explanations characteristic of subpersonal explanation involve subcapacities that are themselves psychological (though subpersonal). Here is a representative quotation:

Functional analysis in psychology involves a particular kind of decomposition: the decomposition of the person into subpersons to whom we ascribe the sorts of psychological predicates that can explain the personal-level capacities. (Drayson 2012, 5)

So I don’t think that Drayson could sidestep the counterexample by pointing to the psychological nature of the subcapacities at issue. Thanks to an anonymous reviewer for encouraging me to consider this objection.

So far, I have advertised an account of the personal–subpersonal distinction, but I have only offered an account of what makes a state personal. What, if anything, distinguishes subpersonal from nonpersonal states?³⁶ One could accept constructionism about the personal level, consistent with a number of attitudes towards the subpersonal, and the motivations for constructionism are consistent with any of these attitudes. As such, I'd like to remain ecumenical. Readers can conjoin a number of attitudes towards the subpersonal with constructionism about the personal. That said, let me briefly sketch three attitudes about the subpersonal.

A first attitude is eliminativism. Perhaps nothing in particular unifies the subpersonal as distinct from the nonpersonal. If not, then we might give up on the category of 'subpersonal' states. Eliminativism about the subpersonal is consistent with constructionism, since we might nevertheless hold that the personal is importantly different from the nonpersonal, and that this distinction is constructed by folk psychology. No corresponding distinction exists between the 'subpersonal' and nonpersonal states. Eliminativism is perhaps the most conservative proposal. Given that I've only argued for a distinction between the personal and nonpersonal, why multiply distinctions?

That said, I think there are reasons to maintain a distinction between the subpersonal and the nonpersonal. The nonpersonal is a capacious category, and one might reasonably suspect that something marks out the subpersonal from the rest. To put it rhetorically, the assumptions encoded in the visual system are interestingly different from the rest of the nonpersonal world. For those who feel this pull, let me suggest two ways we might mark this distinction.

First, one might hold that there is a psychological domain. I don't know how exactly to delimit the psychological domain, but perhaps it is characterized by representational states.³⁷ If so, then we can specify the subpersonal as the nonpersonal that is nevertheless psychological. That is, a state is subpersonal just if it is within the psychological domain, and not a personal state. Insofar as we have plausibly delimited the psychological domain—something I have not done—this proposal will deliver reasonable verdicts concerning subpersonal states and processes. For example, insofar as we offer a plausible account of the psychological domain, Marr's story of visual processing will fall within it. As such, the assumptions encoded in the visual system will qualify as subpersonal states.

³⁶ Thanks to an anonymous reviewer for encouraging me to address this question.

³⁷ This proposal is intuitively too broad. Many artifacts represent, for example. As such, we can decompose this proposal into two views, one on which the subpersonal is broader than we might have pretheoretically thought, and another on which some further restriction is introduced on the kinds of representations that are psychological. Perhaps their intentionality is 'non-derived' (Searle 1983).

Second, one might be inspired by Dennett and Drayson's account of the particular kind of vertical explanation essential to cognitive psychology. Suppose that cognitive psychology involves identifying distinctively intelligent capacities of agents, and explaining how they are realized by relatively less intelligent subcapacities. Our accounts of these subcapacities will involve positing states and processes internal to them. Proposal: the states and processes that constitute the subcapacities that realize our personal capacities are subpersonal, provided these subcapacities are not themselves personal.³⁸ This proposal emphasizes the 'sub' in 'subpersonal'. Subpersonal states are those 'below' the personal ones. If one accepts this kind of vertical explanation of personal capacities, then the place to look for the subpersonal states is within the subcapacities that realize the personal capacities. This proposal is, in one sense, a conciliation with Drayson. In adopting it, we would be agreeing with Drayson that this kind of vertical explanation plays an important role in our understanding of the subpersonal. However, the conciliation is only partial. This account still primarily attaches to states and processes, not styles of explanation. The specific kind of vertical explanation is a means of identifying the subpersonal states. I take this to be a good feature for reasons discussed in the previous section: both personal and subpersonal explanations can be both vertical and horizontal. This proposal vindicates that thought while Drayson's doesn't.

None of these is a fully worked out account of the subpersonal. More would have to be said to motivate one and adjudicate between them. My goal has been to gesture towards some ways in which a constructionist could extend their account of the personal to the subpersonal.

6 Objection: Implicit Attitudes

If there are personal psychological states that do not correspond to the representations in folk psychology, then constructionism is false. Implicit attitudes are perhaps the most worrisome example for the constructionist. As the story goes, psychologists have discovered an important class of attitudes that explain broad swaths of behavior. Presumably, the fact that implicit attitudes were recently discovered tells against their featuring in folk psychology. Nonetheless, they are plausibly personal states. So, the objector would have it, constructionism is false.

Assessing this objection requires delicacy. Very little about implicit attitudes is uncontroversial, and, as I'll indicate, the significance of the objection depends on how ongoing controversies are resolved. I'll consider two representative proposals about implicit attitudes that have different implications for constructionism. In neither case, I'll argue, do implicit attitudes pose a for-

³⁸ The caveat that the subcapacities are not themselves personal is required to avoid the objection in 4.4.

midable challenge to constructionism. My goal is not to settle outstanding questions about implicit attitudes. Rather, I intend to argue that we should be optimistic about constructionism however those questions are resolved.

The primary question, in assessing the significance of implicit attitudes for constructionism, is whether implicit attitudes fall under preexisting categories in folk psychology. Eduard Machery (2016) has argued that implicit attitudes are traits.³⁹ According to Machery, we should understand implicit attitudes as suites of dispositions to behave, attend and cognize in particular ways (Machery 2016, 111), *not* as a newly discovered category of mental states.⁴⁰ Traits are determined by a heterogeneous collection of mental states and processes—beliefs, emotions, self-control processes, and so on (Machery 2016, 112). My extroversion is not another mental state along with my beliefs, but rather a trait determined by a complex of mental states and processes that collectively ground the particular suite of dispositions that is identical to my extroversion.

If implicit attitudes are traits, then they are unproblematic for constructionism. Character traits are a familiar component of folk psychology. We regularly understand others in terms of their character traits.⁴¹ On this understanding, implicit attitudes are a ‘rediscovery’ of a category already found in folk psychology. If so, then implicit attitudes are personal, and they also figure in folk psychology. So implicit attitudes are wholly consistent with constructionism.

Now suppose that implicit attitudes are not covered by our preexisting folk psychology. For example, Neil Levy (2015) argues that implicit attitudes are ‘patchy endorsements’.⁴² Patchy en-

³⁹ For other arguments that implicit attitudes are not a novel psychological kind, see (Mandelbaum 2013, 2016), who argues that implicit attitudes are beliefs. My discussion of implicit attitudes and constructionism assuming Machery’s view applies *mutatis mutandis* to Mandelbaum’s view.

⁴⁰ Machery argues that the trait model offers a better explanation of the empirical evidence concerning indirect measures. My goal is not to assess Machery’s argument, but to consider the implications of Machery’s view for constructionism.

⁴¹ For discussion of the role character traits play in folk psychology, see especially (Westra 2018).

⁴² Levy (2017) argues that implicit attitudes are *not* attributable to agents for the purposes of assigning moral responsibility. It’s not obvious how Levy is conceiving of the relationship between the personal and the attributable, but we can assume for the purposes of considering the objection that the objector either takes the personal to outstrip the attributable, or is not persuaded by Levy’s argument.

dorsements constitute a *sui generis* mental kind, neither association nor belief. While they exhibit some content-driven transitions—and so are not associations—Levy holds that they are not inferentially promiscuous, an essential characteristic of belief (Levy 2015, 813). So patchy endorsements cannot be assimilated to preexisting folk psychological categories. As such, they are apt to raise a problem for constructionism.

Even supposing that implicit attitudes are a *sui generis* kind, I don't think they are troubling for constructionism. The constructionist should say that as an understanding of implicit attitudes proliferates, they *have become*, or *are becoming* personal states. Increasingly, people are understanding one another in terms of implicit attitudes. As that happens, representations of implicit attitudes become enmeshed in our folk psychology, and—according to the constructionist—implicit attitudes become personal states. Diffusion of the concept IMPLICIT ATTITUDE is apt to be uneven. As a result, some indeterminacy and context sensitivity as to whether implicit attitudes are personal can be expected. The idea that states might become personal as folk psychology changes over time is not obviously objectionable. Indeed, I find it to be congenial.⁴³

The objector may continue to press the objection, claiming that even in counterfactual circumstances in which we did not include implicit attitudes in our folk psychology—or in the past—implicit attitudes were personal states. I'll focus on the counterfactual version for simplicity:

The Counterfactual

If we did not include implicit attitudes in our folk psychology, implicit attitudes would still be personal.

The objector suggests that this counterfactual is intuitively true, and so constructionism is false. I suggest that this counterfactual is ambiguous between two readings. On one reading, the con-

⁴³ An anonymous reviewer raises a more general version of this issue. In addition to implicit attitudes, various sorts of neuroscientific explanations have recently been proliferating in society—e.g. people explaining romantic feelings by appeal to oxytocin, or various features of their moods in terms of hormones. I think the kind of story I offer for implicit attitudes is also available for these examples. To the extent that these kinds of explanations become enmeshed in our folk psychology more generally, the constructionist should say that such posits become personal. One might wonder whether this plasticity undermines the distinction between the personal and the subpersonal in general. While I think the constructionist should admit the possibility of indeterminacy and context-sensitivity on the margins, nevertheless I think an important psychological distinction remains between folk psychology and cognition more generally, though I cannot mount a full argument for that claim here. As such, I think a corresponding distinction between the personal and the subpersonal remains, even if we admit a substantial degree of plasticity.

structionist agrees that this counterfactual is true. On the other, it is false, but it's not obvious that this result is problematic. We can either assess the counterfactual holding our folk psychology fixed, or adopting the folk psychology of the counterfactual situation.⁴⁴ On the more natural reading, we assess the counterfactual in light of *our* folk psychology.⁴⁵ On this reading, the constructionist agrees that the counterfactual is true. Regardless of how people in the imagined possibility think, *we* think in a way that determines the personal level for *us*. So the constructionist can happily agree with the objector that, on a natural reading, the counterfactual is true.

The constructionist also admits that there is a reading on which the counterfactual is false. If we adopt the folk psychology of the people in the imagined possible world, then the counterfactual is false.⁴⁶ At this point, though, the objection has lost its force. The objection was supposed to be that the constructionist cannot explain the intuitive truth of the counterfactual. As we've seen, the constructionist has a straightforward explanation, and it's hard to see how to motivate the claim that no false reading exists without simply begging the question against constructionism.

7 The Distinction and Normativity

Some readers have surely been scanning these pages in vain for a discussion of normativity. Many philosophers think a special connection holds between the personal and the normative.⁴⁷ Given such a connection, isn't it foolish for the constructionist to offer an account of the distinction without mentioning normativity? I don't think so, but constructionism should be of interest even to readers who do.

Distinguish between the claim that there is a special connection between the personal and the normative, and the claim that the personal–subpersonal distinction is best-explained in normative terms. A constructionist can happily accept that a special connection exists between the personal and the normative. For example, the view that reason giving explanations only make sense at the personal level is consistent with constructionism. For a more normative approach to be a *com-*

⁴⁴ See (Epstein 2015, 118–20) for an analogous discussion. I find Epstein's perspective to be similar in spirit to two-dimensional approaches to counterfactuals. See (Kaplan 1989; Schroeter 2019).

⁴⁵ This reading corresponds to assessing possibilities 'within a frame' on Epstein's model.

⁴⁶ This corresponds to a 'frame-shifting' reading on Epstein's model.

⁴⁷ See, for example, (Hornsby 2000) who thematizes the personal as connected to reason giving explanations.

petitor to constructionism, we would need an explanation of the personal level in normative terms. In order for such an explanation to be non-circular, the normative approach would need a way of getting a grip on e.g. reason giving explanations that did not involve persons. I'm pessimistic about such an account.

Some might also worry that normative approaches struggle with the plurality problem. While normative approaches are well-suited to capture our explaining actions by citing an agent's reasons, what about the personal explanations we give that are not obviously normative?⁴⁸ For example, introversion and extroversion are character traits that do not have any obvious normative valence. The same goes for my preference for vanilla over chocolate ice cream. What reason is there to prefer vanilla over chocolate ice cream? Sometimes the personal involves predicting behavior on the basis of statistical generalizations—citing what people usually do, not what they ought to do. In my view, examples like these are more naturally understood as non-normative. These considerations are far from decisive, though. Arguments could be made that each of these examples has, deep down, a normative character.⁴⁹

Rather than pursuing that dispute here, though, let me close on a more conciliatory note: constructionism makes room for a normative understanding of the distinction. A normative development of constructionism would involve arguing that folk psychology itself must be understood in normative terms. Though I would say this is a minority position among contemporary theorists of social cognition, it also has a storied history,⁵⁰ as well as a recent revival in interest.⁵¹ This normative constructionism offers an explanation of why the personal is distinctively normative: because it is constructed by folk psychology, and folk psychology is fundamentally normative.

⁴⁸ For related reasons, Drayson (2012) suggests that the distinction at issue for people like Hornsby—who Drayson calls 'The Pittsburgh School' is just a different distinction than the personal–subpersonal distinction at issue for people like Dennett and Drayson (Drayson 2014, 340–1).

⁴⁹ Thanks to an anonymous reviewer for encouraging me to think more carefully about these issues.

⁵⁰ Donald Davidson (1973) and Daniel Dennett (1989) are perhaps the most eminent exponents of this approach.

⁵¹ See (McGeer 2007) and (Zawidzki 2013), though I tend to read these authors as emphasizing that there is a normative *dimension* to folk psychology, rather than urging a wholesale understanding of folk psychology in normative terms, although a reasonable stronger reading is also available.

So even those more attracted to a normative understanding of the distinction can accept a version of constructionism, and avail themselves of constructionism's solution to the plurality problem.

8 Conclusion

I've argued for constructionism about the personal–subpersonal distinction. The primary motivation for constructionism is the plurality problem—how can we explain the diversity inherent in the personal level? Constructionism offers an elegant explanation. The personal level is unified by our folk psychology, which is itself pluralistic. As such, the constructionist predicts and unifies the plurality we find. Moreover, insofar as the personal–subpersonal distinction plays an ineliminable role in theorizing across many domains, so too does folk psychology.

Acknowledgements

I would like to thank Jennifer Nagel, David Barnett, Diana Raffman, Kristin Andrews, Evan Westra, Brendan de Kenessey, Ian Phillips, Jeremy Pober, Chaz Firestone, Mohan Matthen, audience members at the 2021 Pacific APA, and two anonymous reviewers for this journal for their thoughtful feedback, which greatly improved this paper.

References

- Andrews, K. (2012). *Do apes read minds?: Toward a new folk psychology*. MIT Press.
- Andrews, K., Spaulding S., & Westra E. (2021). Introduction to folk psychology: Pluralistic approaches. *Synthese*, 199(1), 1685-1700.
- Barnes, E. (2017). Realism and social structure. *Philosophical Studies*, 174(10), 2417–2433.
- Barrett, L. F. (2006). Are emotions natural kinds? *Perspectives on psychological science*, 1(1), 28–58.
- Barrett, L. F. (2017). The theory of constructed emotion: an active inference account of interoception and categorization. *Social cognitive and affective neuroscience*, 12(1), 1–23.

- Bennett, K. (2007). Mental causation. *Philosophy compass*, 2(2), 316–337.
- Bermúdez, J. L. (1995). Nonconceptual content: From perceptual experience to sub-personal computational states. *Mind & Language*, 10(4), 333–369.
- Boghossian, P. (2014). What is inference?. *Philosophical studies*, 169(1), 1–18.
- Carruthers, P. (2015). *The centered mind: What the science of working memory shows us about the nature of human thought*. Oxford University Press.
- Curry, D. S. (2020). Interpretivism and norms. *Philosophical studies*, 177(4), 905-930.
- Curry, D. S. (2021). How beliefs are like colors. *Synthese*, 199(3), 7889-7918.
- Darwin, C. (1872/1998). *The Expression of Emotion in Man and Animals*. Harper Collins.
- Davidson, D. (1973). Radical interpretation. *Dialectica*, 313–328.
- Dennett, D. C. (1969). *Content and consciousness*. Routledge.
- Dennett, D. C. (1989). *The intentional stance*. MIT press.
- Dennett, D. C. (1991). Real patterns. *The Journal of Philosophy*, 88(1), 27–51.
- Drayson, Z. (2012). The uses and abuses of the personal/subpersonal distinction. *Philosophical Perspectives*, 26(1), 1–18.
- Drayson, Z. (2014). The personal/subpersonal distinction. *Philosophy Compass*, 9(5), 338–346.
- Ekman, P. (1999). Basic emotions. In T. Dalgleish & M. Power (Eds.), *Handbook of Cognition and Emotion*. Wiley.
- Elton, M. (2000). Consciousness: Only at the personal level. *Philosophical Explorations*, 3(1), 25–42.
- Epstein, B. (2015). *The ant trap: Rebuilding the foundations of the social sciences*. Oxford University Press.
- Ferrero, L. (2009). Constitutivism and the Inescapability of Agency. *Oxford studies in metaethics*, 4, 303–333.
- Fodor, J. (1968). *Psychological Explanations*. Random House.
- Fodor, J. A. (1975). *The language of thought*. Harvard University Press.

- Godfrey-Smith, P. (2005). Folk psychology as a model. *Philosophers' Imprint*, 5(6), 1-16.
- Haslanger, S. (2000). Gender and race: (What) are they? (What) do we want them to be? *Noûs*, 34(1), 31–55.
- Haslanger, S. (2012). *Resisting reality: Social construction and social critique*. Oxford University Press.
- Haslanger, S. (2019). Tracing the Sociopolitical Reality of Race. In J. Glasgow, S. Haslanger, C. Jeffers, & Q. Spencer, *What is Race? Four Philosophical Views*. Oxford University Press
- Herschbach, M. (2008). Folk psychological and phenomenological accounts of social perception. *Philosophical Explorations*, 11(3), 223–235.
- Heyes, C. (2018). *Cognitive Gadgets*. Harvard University Press.
- Heyes, C. M., & Frith, C. D. The cultural evolution of mind reading. *Science* 344(6190), 1243091.
- Hornsby, J. (2000). Personal and sub-personal; A defence of Dennett's early distinction. *Philosophical Explorations*, 3(1), 6–24.
- Hughes, J. C. (2011). *Thinking through dementia*. Oxford University Press.
- James, W. (1884). What is an Emotion?. *Mind*, 9, 188-204.
- Kaplan, D. (1989). Demonstratives. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan*. Oxford University Press.
- Kim, J. (1989). Mechanism, purpose, and explanatory exclusion. *Philosophical perspectives* 3, 77–108.
- Kim, J. (1992). The nonreductivist's trouble with mental causation. In J. Heil and A. Mele (Eds.), *Mental Causation*. Oxford University Press.
- Kim, J. (2000). *Mind in a physical world: An essay on the mind-body problem and mental causation*. MIT press.
- Kim, J. (2007). *Physicalism, or something near enough*. Princeton University Press.

- Kleffner, D. A., and Ramachandran, V. S. (1992). On the perception of shape from shading. *Perception & Psychophysics*, 52(1), 18–36.
- Krull, D. S., Loy, M. H., Lin, J., Wang, C., Chen, S., & Zhao, X. (1999). The fundamental attribution error: Correspondence bias in individualist and collectivist cultures. *Personality and Social Psychology Bulletin*, 25(10), 1208-1219.
- Lavelle, J. S. (2021). The impact of culture on mindreading. *Synthese*, 198(7), 6351-6374.
- Levy, N. (2015). Neither fish nor fowl: Implicit attitudes as patchy endorsements. *Noûs*, 49(4), 800–823.
- Levy, N. (2017). Implicit bias and moral responsibility: Probing the data. *Philosophy and Phenomenological Research*, 94(1).
- Lewis, D. K. (1966). An argument for the identity theory. *The Journal of Philosophy*, 63(1), 17–25.
- Lyons, J. (2016). Unconscious evidence. *Philosophical Issues*, 26(1), 243–262.
- Machery, E. (2016). De-Freuding implicit attitudes. In M. Brownstein & J. Saul (Eds.), *Implicit bias and philosophy*. Oxford University Press.
- Maibom, H. (2009). In defence of (model) theory theory. *Journal of Consciousness Studies*, 16(6-7), 360-378.
- Malcolm, N. (1968). The conceivability of mechanism. *The Philosophical Review*, 45– 72.
- Mandelbaum, E. (2013). Against alief. *Philosophical studies*, 165(1), 197–211.
- Mandelbaum, E. (2016). Attitude, inference, association: On the propositional structure of implicit bias. *Noûs*, 50(3), 629–658.

- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. MIT Press.
- McGeer, V. (2007). The regulative dimension of folk psychology. In D. D. Hutto & M. Ratcliffe (Eds.), *Folk psychology re-assessed*. Springer.
- Morrison, J. (2016). Perceptual confidence. *Analytic Philosophy*, 57(1), 15–48.
- Musholt, K. (2018). The personal and the subpersonal in the theory of mind debate. *Phenomenology and the Cognitive Sciences*, 17(2), 305–324.
- Peacocke, C. (1994). Content, computation and externalism. *Mind & Language* 9(3), 303–335.
- Ramachandran, V. S. (1988a). Perceiving shape from shading. *Scientific American* 259(2), 76–83.
- Ramachandran, V. S. (1988b). Perception of shape from shading. *Nature*, 331(6152), 163–166.
- Ross, D. (2007). The economics of the sub-personal: two research programs. In B. Montero & M. D. White (Eds.), *Economics and the Mind*. Routledge.
- Ryle, G. (1949/2009). *The Concept of Mind*. Routledge.
- Schroeter, L. (2019). Two-Dimensional Semantics. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*, Winter 2019, ed. by Edward N. Zalta. Metaphysics Research Lab, Stanford University.
- Searle, J. (1983). *Intentionality: An Essay in the Philosophy of Mind*. Cambridge University Press.
- Spaulding, S. (2010). Embodied cognition and mindreading. *Mind & Language*, 25(1), 119–140.
- Spaulding, S. (2015). Phenomenology of social cognition. *Erkenntnis*, 80(5), 1069–1089.
- Spaulding, S. (2018a). *How We Understand Others: Philosophy and Social Cognition*. Routledge.
- Spaulding, S. (2018b). Mindreading beyond belief: A more comprehensive conception of how we understand others. *Philosophy compass* 13(11), e12526.

- Spencer, Q. (2019a). How to be a Biological Racial Realist. In J. Glasgow, S. Haslanger, C. Jeffers, & Q. Spencer, *What is Race? Four Philosophical Views*. Oxford University Press
- Spencer, Q. (2019b). Spencer's Reply to Glasgow, Haslanger, and Jeffers. In J. Glasgow, S. Haslanger, C. Jeffers, & Q. Spencer, *What is Race? Four Philosophical Views*. Oxford University Press
- Steward, H. (1997). *The ontology of mind: events, processes, and states*. Oxford University Press
- Stich, S. P. (1978). Beliefs and subdoxastic states. *Philosophy of Science*, 45(4), 499–518.
- Thomas, R., Nardini M., & Mareschal D. (2010). Interactions between 'light-from-above' and convexity priors in visual development. *Journal of Vision*, 10(8), 1–7.
- Tracy, J. L., & Randles, D. (2011). Four models of basic emotions: a review of Ekman and Cordaro, Izard, Levenson, and Panksepp and Watt. *Emotion Review*, 3(4), 397–405.
- Velleman, J. D. (1992). What happens when someone acts?. *Mind* 101(403), 461–481.
- Westra, E. (2018). Character and theory of mind: An integrative approach. *Philosophical Studies* 175(5), 1217–1241.
- Wittgenstein, L. (1953/2010). *Philosophical Investigations*. Wiley.
- Wu, S. & Keysar, B. (2007). The effect of culture on perspective taking. *Psychological science*, 18(7), 600-606.
- Zawidzki, T. W. (2013). *Mindshaping: A new framework for understanding human social cognition*. MIT Press.