

Transformative Experiences and the Equivocation Objection

*Inquiry**

Yuri Cath

La Trobe University

Abstract: Paul (2014, 2015a) argues that one cannot rationally decide whether to have a transformative experience by trying to form judgments, in advance, about (i) what it would feel like to have that experience, and (ii) the subjective value of having such an experience. The problem is if you haven't had the experience then you cannot know what it is like, and you need to know what it is like to assess its value. However, in earlier work I argued that 'what it is like'-knowledge comes in degrees, and I briefly suggested that, consequently, some instances of Paul's argument schema might commit a fallacy of equivocation. The aim of this paper is to further explore and strengthen this objection by, first, offering a new argument—the *modelling argument*—in support of it, and then by evaluating a range of replies that might be given to this objection on Paul's behalf. I conclude that each reply either fails or, at best, only partially succeeds in defending some but not all instances of Paul's argument schema. In closing, I consider how we might revise Paul's concepts of transformative experiences and choices in response to this conclusion.

Keywords: Transformative experiences, knowing what it is like, first-personal perspectives, imagination, cognitive modelling,

Introduction

Suppose you are faced with a choice between two life paths. If you choose the first path your life will continue in much the same way as it does now. However, if you choose the second path, your life will come to include a new experience that, compared to all your previous experiences, is distinctive with respect to *what it is like* to have this experience. You might be considering, for example, whether to remain childless or, instead, to become a parent for the first time. Or perhaps you are simply considering whether to have a new kind of experience which will not change who you are in any significant way, but which will still be different from any of your past experiences, like tasting vegemite or smelling a skunk for the first time.

* This is a preprint of an article whose final and definitive form is published in *Inquiry* (2022 online first: <https://www.tandfonline.com/doi/full/10.1080/0020174X.2022.2107063>).

How should you approach decisions of this kind, where one of the actions you are choosing between might lead to a new experience with new ‘what it is like’ properties? Paul (2014, 2015a) develops an important critique of the following decision-making strategy one might naturally reach for: First, one makes judgments about what it would be like to have the new experience, and the subjective value of the outcome of having an experience of that kind. Second, one makes judgments about what it is like to continue to have the kinds of experiences that make up your current life, and the subjective value of the outcome of continuing to have those kinds of experiences. Third, with these pairs of subjective value judgments to hand, one then chooses between the two relevant actions (i.e., the actions that are each meant to lead to one of these respective outcomes) in the way recommended by normative decision theory, that is, one chooses the action that best maximizes expected utility.

Paul acknowledges our inclination to try to make decisions in this way, however, she argues that this strategy is irrational given that these are *epistemically transformative* experiences, that is, experiences where one cannot know what it is like to have the experience without having had that experience oneself.¹ Here is a simple way of reconstructing the key part of Paul’s general argument schema:

The Knowledge Premise (KP) You can know what it is like to Φ only if you have Φ -ed yourself.

The Value Premise (VP) You can approximately determine the subjective value of what it is like to Φ only if you know what it is like to Φ .

Therefore,

¹ Paul also has a notion of *personally transformative* experiences, which are a species of epistemically transformative experiences distinguished by the further property of being experiences that would change who you are, or your core preferences, in some significant way. Compare, for example, eating durian for the first time versus renouncing one’s worldly life to become a nun or a monk. Paul’s primary interest is in personally transformative experiences, but her argument is meant to apply to all choices involving the broader class of epistemically transformative experiences.

The Intermediate Conclusion (IC) You can approximately determine the subjective value of what it is like to Φ only if you have Φ -ed yourself.

So, as applied to Paul's most famous example (Paul 2015a), we would get the conclusion that you can approximately determine the subjective value of what it is like to have a child only if you have had a child yourself. And, with this intermediate conclusion in hand, Paul can quickly establish her final conclusion that trying to decide whether to have a child for the first time on the basis of comparing its subjective value with the subjective value of remaining childless is not rational. For if the above reasoning is sound then one lacks any epistemic access to one of the subjective values in that comparison.

Both (KP) and (VP) are prima facie plausible. As Paul discusses, (KP) receives support from the same kinds of intuitions that support the idea that Jackson's (1986) Mary cannot know what it is like to see something red until she sees something red. And (VP) is plausible given Paul's characterisation of 'subjective values' as "the values of what it is like to have the experiences, or of what it is like to be in these experiential states" (Paul 2014: 11).

However, in previous work (Cath 2019) I argued that at least some instances of Paul's argument schema are in danger of committing a fallacy of equivocation. After briefly reintroducing this equivocation objection (§1), my aim here is to strengthen it by, first, offering a new argument—*the modelling argument*—in support of one its key assumptions (§2), and then by critiquing a range of different replies that one might try to give to this objection on Paul's behalf (§§3-6), each of which is related to importantly different ideas in Paul's works. In closing, I make some suggestions about how we might revise Paul's concepts of transformative experiences in light of the preceding discussion.

1. The Equivocation Objection

In Cath (2019) I grant that there are available interpretations of (KP) and (VP) which are each plausibly true. However, I suggest that there is no one univocal interpretation of the premises on which they are both plausibly true. The motivation for this claim comes from the idea that WIL-knowledge comes in different degrees.

Most notably, I distinguish three different grades—gold, silver, and bronze—of *knowledge of experience*, that is, knowledge of propositions of the form ‘*w* is way it feels to Φ ’. And I argue that while knowing-WIL ascriptions are, paradigmatically, used to attribute gold knowledge of experience, in many contexts we will ascribe WIL-knowledge to someone when they only have silver or bronze knowledge of experience.

Consider the experience of going to war. On my view, Mary has *gold knowledge of experience* with respect to this experience type when she knows, for some way *w*, that *w* is a way it feels to go to war, *and* her knowledge is based on her own direct acquaintance with the phenomenal properties of her own experiences of going to war. And, at the other end of the spectrum, Mary has *bronze knowledge of experience* when her knowledge of this form is based only in sources like testimony and reading books like, say, Marlantes’ (2011) *What it is Like to Go to War*.

However, between the gold and the bronze cases, Mary has *silver knowledge of experience* if she knows, for some way *w*, that *w* is a way it feels to go to war, and her knowledge is based not only in sources like testimony and books, but also in her acquaintance with the phenomenal properties of her own experiences distinct from, but relevantly similar to, the target experience of going to war. Suppose that Mary is an ambulance officer. In reading Marlantes’ book, Mary might learn of interesting similarities between her own experiences (e.g., in relation to dealing with people in shock with traumatic injuries), and experiences of going to war. In which case, Mary can not only know that *that* way (as described in, say, Marlantes’ book) is a way it feels to go to war, but she can also know that *that* way feels relevantly similar to *this*, where ‘this’

picks out the phenomenal properties of one of her own experiences (either a current experience that Mary is attending to, or a past experience that she is recalling in her imagination).

With the above ideas in place, we can note that there will be three different available interpretations of (KP) and (VP), respectively. So, for (KP) as applied to the parent example we get the following interpretations:

(^GKP). You can have *gold-standard* knowledge, for some way w , that w is a way it feels to have a child only if you have had a child yourself.

(^SKP). You can have *silver-standard* knowledge, for some way w , that w is a way it feels to have a child only if you have had a child yourself.

(^BKP). You can have *bronze-standard* knowledge, for some way w , that w is a way it feels to have a child only if you have had a child yourself.

And for (VP):

(^GVP). You can approximately determine the subjective value of what it is like to have a child only if you have *gold-standard* knowledge, for some way w , that w is a way it feels to have a child.

(^SVP). You can approximately determine the subjective value of what it is like to have a child only if you have *silver-standard* knowledge, for some way w , that w is a way it feels to have a child.

(^BVP). You can approximately determine the subjective value of what it is like to have a child only if you have *bronze-standard* knowledge, for some way w , that w is a way it feels to have a child.

Starting with (KP), note that (^GKP) is trivially true and (^SKP) and (^BKP) are both trivially false, given how the notions of gold, silver, and bronze standard knowledge are defined. If Paul's

argument is to be sound then, (VP) must be interpreted as (^GVP). For the mere truth of (^SVP) or (^BVP) will not do because even if these claims are true inserting them into the argument for (IC) will result in an invalid argument when (KP) is interpreted as (^GKP).

But the problem is that, in contrast to (^GKP), (^GVP) is far from being a trivial truth. Suppose one possessed lots of silver and bronze, but no gold, knowledge of propositions of the form ‘*w* is a way it feels to Φ ’. Why couldn’t this knowledge put one in a position to make at least approximately accurate judgments about the subjective value of experiences of Φ -ing? One may have never been a parent but why couldn’t activities like reading books, talking to parents, having distinct but relevantly similar experiences, and letting all these resources inform one’s imaginings of what it is like to be a parent,² put one in a position to have various forms of silver and bronze knowledge of what it is like to be a parent³ and, on the basis, to make at least approximately accurate judgments about the subjective value of being a parent?

The crucial issue then for Paul is whether there is a good case to be made for (^GVP), which will involve arguing that silver and bronze WIL-knowledge are never sufficient for one to be able to even approximately determine the subjective value of the target experience. And, of course, our interest is not just in the specific instance of (^GVP) that we get when consider the experience of being a parent, but rather the much wider set of experiences that Paul identifies as potentially transformative experiences, which include such examples as becoming a parent, gaining or losing access to a sensory modality, getting married, getting divorced, going to war, being

² For support for the claim that imagination can help us to know what it is like to have experiences we have not had ourselves, see Kind (2020, 2021).

³ See Krishnamurthy (2015) and Harman (2015), for related claims that by such means one could come to know *something* about what it is like to be a parent. Harman, for example, discusses her own experiences in looking after her younger sister and a child’s friend. In response Paul (2015b, 518) suggests that Harman’s experiences sound like cases of *alloparenting* which, being a type of parenting, is consistent with Paul’s view that one can know what it is like to be a parent only if one has been a parent. The weakness, as I see it, of Paul’s reply is that even if her diagnosis of Harman’s specific examples is right there will still be various kinds of nearby experiences which are not rightly described as being cases of any form of parenting, but which could still give one some partial knowledge of what it is like to be a parent. See §4 below for related discussion.

betrayed or betraying someone else, leaving home for College, going to graduate school, changing careers, and losing a loved one. However, before considering replies to the equivocation objection, I want to first set out what I take to be an important argument one can give in support of it.

2. The Modelling Argument

To help introduce this argument, it will be useful to underline some simple but significant points concerning the interpretation of Paul's argument and her background views on WIL-knowledge.

Starting with (KP), one should not interpret this premise as the claim that one can know what it is like to Φ at time t_1 only if one is Φ -ing at t_1 . For that claim would commit one to the implausibly sceptical view that no one ever retains knowledge of what it is like to Φ after the experience of Φ -ing finishes. Relatedly, (VP) should not be interpreted as the claim that one can determine the subjective value of Φ -ing at t_1 only if one knows what it is like to Φ at t_1 on the basis of one's Φ -ing at t_1 . For that idea would lead to an implausibly sceptical position according to which one can only ever accurately assess the subjective value of Φ -ing when one is currently Φ -ing.

Paul follows Lewis (1998) in thinking that when you have an experience of Φ -ing for the first time, and you thereby come to know what it is like to Φ , you will also come to possess new abilities to "imagine, recognize, and cognitively model possible future experiences of this kind" (Paul 2015b: 476). And it is specifically one's abilities to imagine and cognitively model possible future experiences of Φ -ing, which allow one to accurately make judgments about the subjective value of possible future experiences of Φ -ing (Paul 2014: 26).

Imagine that Stephanie is considering whether to go on holiday to the Isle of Skye again, having holidayed there many times before, or whether to just stay home. Paul allows that Stephanie

could rationally make this decision by comparing the respective subjective values of these two different outcomes because both outcomes involve types of experiences that she had before, which means that Stephanie can cognitively model possible future instances of these experiences and make accurate judgments about their subjective value on that basis.

One might try to account for the accuracy of Stephanie's subjective value judgments by thinking of them as drawing only on mere semantic memories about how it feels to have these experiences. However, such a view would make Paul's position vulnerable to the objection that other sources of such descriptive information, like testimony, could help us to make accurate judgments about the subjective value of experiences we have not had ourselves. This is why it is important to Paul that, if we are not currently having a given type of experience that we have had before, then exercising our abilities to imaginatively simulate possible experiences of this kind is required for us to grasp and assess their subjective value.⁴

But now think about the quality of the cognitive modelling one can do using these abilities. The modelling argument starts with the simple observation that such cognitive models will almost always be highly imperfect models of the target experiences they aim to simulate and represent.

Suppose we think of this cognitive modelling as, at least partly, involving mental acts of imagining what it is like to Φ where these imaginings are occurrent experiences which count as being simulations of Φ -ing in virtue of sharing phenomenal properties with the very experiences of Φ -ing they are modelling.⁵ If we conceive of cognitive modelling in this way,

⁴ One way of challenging Paul's views would be to argue that one can approximately grasp the subjective value of possible future experiences without engaging in any cognitive modelling at all. See e.g., Campbell and Mosquera (2020) for a critique of this kind that appeals to the fact that Paul allows for some exceptions to the idea that one can only assess the subjective value of a possible future experience by cognitively modelling it in one's imagination. I think part of the significance of my argument in this section is that it could still go through even if there were no such exceptions.

⁵ On Goldman's (2006) influential account for one process P to simulate another process P^* is for P to duplicate, replicate, or resemble P^* in some significant respect, and also for P to thereby fulfill a function or purpose (in our

as I think Paul does, then when one is acquainted with the phenomenal properties of one's mental act of imagining one can thereby be acquainted with phenomenal properties that are also phenomenal properties of the target experiencing of Φ -ing. But, even so, there will undoubtedly be many phenomenal properties involved in the experience of Φ -ing that are not found in one's experience of imagining that experience.

Compare what it is like to see the face of someone you know well versus what it is like to imagine seeing their face, or what it is like to see a room you know well versus what it is like to imagine that room when you are not in it. The literature on aphantasia (see e.g., Zeman et al. 2015) suggests that people might differ greatly in the quality of the visual images they can generate using their imagination, with some people perhaps even lacking this ability altogether. But even if we restrict our attention to those who have good abilities to generate visual imagery there will still be significant differences for them between an actual visual experience of seeing someone's face versus an experience of imagining such an experience, with the later experience containing less information, and different phenomenal properties, etc.

Parallel points apply even more strongly when we consider much more complex experiences—like having a child, or going to war, or getting divorced—which are stretched out over long periods of time, involve multiple perceptual modalities, and phenomenology connected to non-perceptual sources like emotions and cognitive processes, etc. Imagination can be a useful tool for modelling possible experiences of a type one has had before. But, like most non-ideal models, the cognitive model is no replacement for its target, and if it simulates its target, it does so in spite of the many differences that obtain between model and target. In which case, if we are to avoid the sceptical conclusion that we can only grasp the subjective value of an experience at the time as which it is occurring, then the imperfections of such cognitive models

case the purpose would be set by the subject's intention to imagine what the target experience is like). And the sharing of relevant phenomenal properties is clearly one way for one experience to resemble another. Following Fischer's (2006) distinction this is a 'concrete' rather than a merely 'abstract' simulation.

must be consistent not only with their being able to grant us epistemic access to the experiences they model, but also the subjective values of those experiences.

These points support, I submit, a *parity premise*, namely, that if imperfect cognitive models of experiences of Φ -ing of the kind that Paul favors (i.e., models of Φ -ing built using knowledge and imaginative abilities gained from past experiences of Φ -ing) can help one to approximately determine the subjective value of possible future experiences of Φ -ing then, in principle, other imperfect cognitive models of experiences of Φ -ing (i.e., models of Φ -ing built using knowledge and imaginative abilities from distinct but similar experiences) could do the same.

The *modelling argument* takes this parity premise and combines it with Paul's own *anti-sceptical premise* that imperfect cognitive models of experiences of Φ -ing of the kind she favors (built using knowledge and imaginative abilities gained from past experiences of Φ -ing) can help one to approximately determine the subjective value of possible future experience of Φ -ing. Together these premises establish the conclusion that, in principle, other imperfect cognitive models of experiences of Φ -ing (built using knowledge and imaginative abilities gained from distinct but similar experiences) could help one to determine the subjective value of possible future experience of Φ -ing.⁶

This argument supports the equivocation objection by supporting its assumption that possessing the right kinds of silver and bronze WIL-knowledge, together with one's imaginative abilities, could, in principle, put one in a position to approximately determine the

⁶ See Mathony and Messerli (2022) for a recent critique of Paul's views which also focuses on cognitive modelling. Mathony and Messerli's arguments rely on two key ingredients: (1) the assumption that there is a tight link between (i) our abilities to cognitively model possible future experiences, and (ii) our abilities to predict how happy possible future experiences would make us feel; and (2) a meta-analysis of empirical research on 'affective forecasting' which supports the claim that whether an experience is transformative or not makes no difference to our abilities to predict how happy it would make us feel (either way they are inaccurate). I cannot discuss Mathony and Messerli's interesting paper at length here, but it is worth noting that a key difference between their arguments and mine is that my critique of Paul does not rely on (1). This is significant because one important way Paul could reply to their critique would be to challenge their account of the links between (i) and (ii).

subjective value of a given experience. And this argument also points us to an important constraint on replies to the equivocation objection, namely, that any considerations they offer in support of (^GVP) should not be so strong that they threaten to undermine Paul's own anti-sceptical assumptions that after having Φ -ed someone can still know what it is like to Φ , and grasp the subjective value of Φ -ing.

3. The Distance Reply

Turning now to replies to the equivocation argument, I want to first consider a way of defending (^GVP) that appeals simply to the idea that any shared properties between one's past experiences and the relevant experience one has not had, will be too thin and inconsequential for those past experiences to serve as a good basis for cognitively modelling the missing experience. In which case, no amount of silver or bronze WIL-knowledge, or imagination, will suffice to bridge that gap.

This reply works well when the missing experience is tied to a particular sensory modality that one either lacks access to altogether, or which one only has very partial or restricted access to (e.g., consider Mary's visual experiences when she is still in her black and white room). Paul (2014: 106–7), for example, gives the case of a blind saxophonist who is given a “one-time-only chance to have retinal surgery to become sighted”. Paul (2014: 107) writes that the saxophonist “cannot assign an experiential value [to any given visual experience] because he lacks the capacity to imaginatively represent the nature of this lived experience”. And this kind of case can be used to support the following specific instance of (^GVP):

(^GVP) You can approximately determine the subjective value of what it is like to have visual experiences in general only if, for some way w , you have *gold-standard* knowledge that w is a way it feels to have visual experiences.

The idea that the saxophonist would need to have had at least some gold-standard knowledge of visual experience in order to determine the subjective value of visual experiences does seem plausible. For any similarities between visual experiences and other types of experiences that the saxophonist has had will plausibly be too thin and uninformative for him to use those similarities as a basis for building a cognitive model of the way it feels to have visual experiences that would be informative enough to help him to even approximately determine the subjective value of, say, seeing a sunset.

The problem is that such experiences are only a small subset of the transformative experiences to which Paul's argument is meant to apply. For most of the experiences that Paul identifies as "transformative experiences" would not involve the gain or loss of a sensory modality, or other radically new sensory experiences. Paul's case of having children is a case in point, for that is a case where there could be substantive overlaps between one's past experiences and the experience one hasn't had, and the new experience does not involve any events of gaining access to a new sensory modality.

So, if Paul's argument schema is to have the kind of broader significance it is clearly meant to have, it is crucial that (^GVP) can be defended with respect to a much wider range of transformative experiences than just cases involving radically new sensory experiences. And the distance reply cannot do that.

4. The Keystone Reply

Unlike the distance reply, the related keystone reply can allow that the missing experience may share many phenomenal properties with one's past experiences. The claim that this reply rests on is not that there are very few similarities in total between one's past experiences and the transformative experience, but that the differences that are there make all the difference. More precisely, this reply will defend any relevant instance of (^GVP) by appealing to something like the following three claims:

1. The target experience has certain distinctive phenomenal properties that are not shared by any of one's past experiences (although they may share many other phenomenal properties).
2. These distinctive phenomenal properties play a very significant role in determining the subjective value of the target experience.
3. Any shared phenomenal properties between the target experience and one's past experiences will not be a good basis for forming a cognitive model of these distinctive phenomenal properties that would enable one to even approximately grasp the subjective value of the target experience.

The parenting experience is, again, a good case for illustrating this reply because one might think that there are distinctive feelings involved in the experience of having children and that one (1)-(3) all hold true when applied to these feelings. Paul herself also offers a number of suggestions as to what these special feelings might be including: (i) feelings of love and awe upon seeing and touching one's newborn child for the first time, (ii) feelings of shock and numbness upon seeing and touching one's newborn child for the first time, (iii) feelings caused by hormonal changes in parents, and (iv) "special, intense, and unique feelings of parental love, care, personal engagement, and responsibility" (Paul 2015b: 485) that develop and change over time.

One point to note here, and one that Paul is sensitive to, is that no one of these feelings is plausibly an essential part of the larger experience type of being a parent. As, obviously, many of these feelings will not be shared by all parents, and even for those parents who have the relevant feeling it may only be a small part of their overall parenting experience and for a short time (e.g., the feelings associated with short-term hormonal changes), and so on. And these points might put some pressure on (2), that is, the idea that these special feelings play a very significant role in determining the subjective value of being a parent.

But while no one of these distinctive sub-experiences is an essential part of the experience of being a parent, having at least some significant amount of these sub-experiences is plausibly part of the paradigmatic experience of being a parent. In which case, for many people when they try to imagine what it would be like to be a parent it seems plausible that what they are partly interested in is the subjective value of having a parenting experience which includes these distinctive feelings.

The real concern, I think, for the keystone reply is how often (3) will be true. It is worth remembering that properties can, of course, share properties, and this point applies to phenomenal properties as much as any other kinds of properties. A person who has not been a parent may not have had the specific feelings of love and concern that parents (paradigmatically) have for their children. But they will likely have had experiences of love and concern for other people (including perhaps people they have cared for like older parents or younger siblings) that share many properties with those parental experiences of love and concern. And those distinct but related experiences—together with one's imagination and sources of further information like testimony and literature—might form a good enough basis for them to build a useful cognitive model of the way it feels to have those distinctive parental feelings of love and concern.

This model will be imperfect. But, recalling the lessons of §2, we can note that the same would be true of the cognitive models that, say, a parent of one 8-year-old child would construct when trying to recall what it is like to have the special feelings involving in having a newborn child. For the parent's experience of imaginatively simulating those feelings will very likely share only some, and not all, of the phenomenal properties involved in those original distinctive feelings. And if we want to avoid the sceptical conclusion that at this point in time such a parent could not even approximately grasp the subjective value of such feelings then, I submit, we

should also grant that the imperfect cognitive models of those who have never been parents could, in principle, suffice for them to grasp the subjective value of such feelings.

5. Specificity Replies

In this section I want to consider two replies which are similar in spirit as they each rely on different ways of interpreting the embedded question in the knowing-WIL ascription in (VP). The common strategy being that (^GVP) can be defended when it is interpreted in the right way.

5.1 *The specific experience reply*

The *specific experience reply* relies on different ways of precisifying the ‘Φ’ in a ‘S knows what it is like to Φ’ ascription and argues that (^GVP) is true when we focus on a sufficiently specific characterisation of the target experience of Φ-ing. For example, one might claim that someone can only ascertain the subjective value of being the parent of *the particular child* they would have if they were to become a parent only if they become a parent. Paul (2014: 79) herself briefly appeals to this kind of specificity point in relation to (KP), writing:

So if you choose to become a parent, when you choose, not only do you not know, in a general sense, what it will be like to be a parent, you don’t know what it will be like to have the particular child that you’ll end up having.

In Paul’s discussion of this specificity point her claim is just that it exacerbates the original problem of not knowing what it will be like to be a parent in general. But the question now is whether this specificity point by itself can be used to answer the equivocation objection. So, focusing on the parent example, a proponent of this specific experience reply might allow that (^GVP) is false when it interpreted as (^GVP₁) but they will claim that it is true when interpreted as (^GVP₂):

(^GVP₁) You can approximately determine the subjective value of what it is like to be a parent only if, for some way *w*, you have gold-standard knowledge that *w* is a way it feels to be a parent.

(^GVP₂) You can approximately determine the subjective value of what it is like to be a parent *of the particular child you would have if you were to become a parent* only if, for some way *w*, you have gold-standard knowledge that *w* is a way it feels to be a parent *of that particular child*.

So interpreted, (^GVP₂) is hard to deny. But, of course, for the resulting argument to be valid we would also need to interpret (^GKP) in the same way, and together these two premises would only establish:

Intermediate Conclusion (IC₂). You can approximately determine the subjective value of what it is like to have *the particular child you would have if you were to become a parent* only if you have *been a parent to that particular child* yourself.

The problem is, so interpreted, the truth of this conclusion does not undermine the original decision strategy that Paul opposes. For one might accept this conclusion and still reasonably hold that one can rationally make a decision about whether to have a child on the basis of comparing the subjective value of the experience of having a child *in general* with the subjective value of the experience of remaining childless *in general*.

Relatedly, this specificity reply threatens to generate sceptical consequences that conflict with Paul's anti-sceptical commitments identified in §2. Imagine someone saying to Stephanie: "But no two holidays on Skye are exactly alike so it is irrational to let your decision be informed by your attempts to imagine a new holiday there". This would be an implausibly sceptical position, and it is one that Paul herself would reject.

If we required a perfect matching between our past experiences and our possible future experiences in order for us to be in a position to make reasonable judgments about the subjective value of those possible future experiences then we would rarely, if ever, be in such a position. But we can be in such a position, which suggests that the specificity point is of

limited interest when it comes to making decisions on the basis of our abilities to imagine possible future experiences of Φ -ing when we have Φ -ed before. And, if the parity premise from §2 holds, then the parallel point extends to abilities to imagine possible future experiences of Φ -ing which one possesses in virtue of having had distinct but relevantly similar experiences, and consulting testimony, etc.

5.2 *The specific subject reply*

The *specific subject reply* relies on different ways of interpreting the covert pronoun, or ‘PRO’, in the embedded question of the relevant ‘S knows what it is like to Φ ’ ascription. It is standard to distinguish two different readings of embedded question constructions with an infinitive verb phrase: a generic ‘one’ reading (‘S knows what it is like for *one* to Φ ’) and a reading that is anaphoric on the main subject (‘S knows what it is like for *S* to Φ ’). The specific subject reply appeals to this ambiguity and claims that while instances of (GVP) might be false when interpreted along the lines of (GVP_3), they will be true when interpreted along the lines of (GVP_4):

(GVP_3) Stephanie can approximately determine the subjective value of what it is like *for one* to Φ only if, for some way w , Stephanie has gold-standard knowledge that w is a way it feels for *one* to Φ .

(GVP_4) Stephanie can approximately determine the subjective value of what it is like *for Stephanie* to Φ only if, for some way w , Stephanie has gold-standard knowledge that w is a way it feels *for Stephanie* to Φ .

Importantly, this reply does not rely on any notion of the *de se* or special modes of presentation (we will consider replies of that kind in §6.2), as indicated by the fact that we can state (GVP_4) above using just ‘Stephanie’ rather than the indexical ‘her’ (note that I will use the indexical at times below for ease of expression). All that the specific subject reply relies on is the idea that

who we are can sometimes change how it feels for us to Φ , such that the WIL-properties of the token experiences of Φ -ing of two different people (or the same person at different times) may differ given differences in their past experiences, cultural backgrounds, mental states, etc. To take a very simple example, two people who go caving will have very different experiences if one of them suffers from claustrophobia, and the other does not, and these differences will likely mean that caving has a very different subjective value for each of them.

So, a proponent of the specific subject reply might appeal to this idea and argue that reading books, talking to people about their experiences, and relating one's own experiences to the target experience etc., can at best tell Stephanie what it is like for *one* to, say, be a parent or go to war, or maybe it what it was like for *that* person, but it can't tell Stephanie what it would be like for *Stephanie* to do these things because those experiences might differ for her given her individual psychology and history, etc. And, if that is the case, then Stephanie won't be able to grasp the subjective value for her of being a parent, or going to war, by learning about what it is like for other people to have these experiences.

Obviously, who we are can sometimes change our experiences. But this fact doesn't support a good strategy for defending (GVP_4). One reason for thinking this is that knowledge of who we are, and our past experiences, can usefully inform our attempts to imagine what an experience of Φ -ing would be like for ourselves. So, Stephanie might have never have been caving but if she knows that she is claustrophobic, and if she has had experiences of feeling claustrophobic in other small spaces, then it seems likely that she will be able to build a cognitive model of caving that is sufficiently informative for her to grasp the subjective value for herself of caving.

Another reason is that we shouldn't make too much of the fact that our experiences can vary depending on who we are. In particular, that fact is consistent with acknowledging that there can still be significant *patterns* in people's individual experiences of, say, being a parent, or going to war, etc. After all, one of the reasons that people form things like parent support

groups, and returned services leagues, is that they find comfort in spending time with those who have had similar experiences. And, given the similarities that can obtain between people's individual experiences of Φ -ing, knowing what it was for other people to Φ can be a useful guide for someone when they try to construct a cognitive model of what Φ -ing would be like for them.

6. First-Person Replies

My aim now is to consider whether there is a promising strategy for defending (G VP) which appeals to notions like “first-personal perspectives”, “subjective modes of presentation”, or “experiential points of view”, which are very prominent notions in Paul's work (see e.g., Paul 2016, 2017a, 2017b, 2017c). So, one might naturally think that such notions are a good resource to look to in trying to offer a defence of (G VP) on Paul's behalf. But to get clear on whether this is the case, we need to first make a distinction between two significantly different ideas that philosophers can have in mind when they talk about first-personal perspectives in relation to WIL-knowledge.⁷

On the one hand, we might think of the first-personal perspectives involved in knowing what it is like to Φ as certain distinctive phenomenal properties that partly constitute the overall way w that it feels to Φ . I will call this the *experiential conception* of first-personal perspectives because, on this conception, a first-personal perspective is something that is a part of the phenomenal character of the target experience itself.

⁷ The distinction I make below between (i) *experiential* versus (ii) *attitudinal* conceptions of first-personal perspectives in relation to WIL-knowledge is related to the one Cappelen and Dever (2017: 330) make between what they call (i*) notions of ‘subjective experience’ in the Nagel tradition versus (ii*) notions of the ‘De se’, ‘self-locating attitudes’, and ‘essential indexicality’ in the Perry/Lewis Tradition. However, as I construe (i) it covers a distinctive species of what Cappelen and Dever have in mind when they talk about (i*). Also, while I am sympathetic to Cappelen and Dever's claims about the unimportance of first-personal perspectives to understanding empathy, my arguments do not rely on any of their criticisms of (ii*). Roughly, what I argue in this section is that, even if plausible, neither (i) nor (ii) can help Paul to address the equivocation objection.

One example of phenomenal properties that would be relevant to this experiential conception is what philosophers have in mind when they talk about the *subjective* aspects of the phenomenal character of our experiences, as opposed to their *qualitative* aspects (Levine 2001, Kriegel 2005). For example, Kriegel (2005: 24) makes this distinction in relation to an experience of looking at the blue sky:

When I have a conscious experience of the sky, there is a bluish way it is like for me to have or undergo my experience. I suggest that we distinguish two aspects in this “bluish way it is like for me”: (i) the *bluish* aspect, which we may call the experience’s *qualitative character*, and (ii) the *for-me* aspect, which we may call its *subjective character*.

For Kriegel the subjective character of a phenomenally conscious experience is what makes it a phenomenally conscious experience *at all*, whereas its qualitative aspects are what make it the specific phenomenally conscious experience that it is.

Another example of phenomenal properties that would be relevant to an experiential conception of first-personal perspectives is the kinds of perspectival features involved in what it feels like to view or hear something from one’s given spatial location, body, or bodily position. This notion of a first-person perspective is invoked in passages in Paul’s work like the following (2016: 108):

What is a consciously centered, experiential point of view? An example, couched in terms of camera angles, can help to bring the idea out. Think of the sort of view that you get with a Go Pro, a type of digital camera designed for filming action while being immersed in it. Or take an immersive computer game. An “immersed first personal viewing angle” is a distinctive and important camera angle that you get from, metaphorically, occupying the boots of your character in a computer game. The view is as though you were looking out the eyes of the character, seeing the world as it sees the world.

This kind of ‘experiential point of view’ is a constituent part of the overall phenomenal character of one’s experience of seeing a blue sky that is, arguably, distinct from not only the

qualitative aspects (the bluishness of the sky), but also its supposedly subjective aspects (the for-me-ness of the experience). So, one might think that there are subjective and perspectival features of the phenomenal character of our experiences, and either of these features could be something that someone has in mind when they talk about ‘first-personal perspectives’ in relation to WIL-knowledge.

On the other hand, we might think of a first-person perspective as something involved in having a special kind of first-personal thought *about* one’s experience. On this conception, a first-personal perspective is something that helps to explain the difference between merely knowing that *w* is a way it feels for *one* or *that person* to taste Durian, versus knowing that *w* is a way it feels *for me* to taste Durian (where, in contrast to the specific subject reply discussed in §6, the indexical ‘me’ is now meant to be essential to the kinds of states of knowledge and thought that we are interested in here). I will call this the *attitudinal conception* of first-personal perspectives because, on this conception, a first-personal perspective is characterised as something involved in having a *de se* propositional attitude *about* an experience, rather than a part of the experience itself. These two different notions of first-personal perspectives—the experiential and the attitudinal—might well be intimately connected in various ways, but it will be useful to consider each of them separately when assessing whether either notion can play a role in defending (^GVP).

6.1 The experiential conception of first-personal perspectives

If one were to appeal to this notion in defending (^GVP) the suggestion would be that these kinds of ‘first-personal’ phenomenal properties involved in the missing experience will be too different from the phenomenal properties of any of one’s past experiences for one to use one’s past experiences as a basis for modelling the relevant missing experience. However, this is not a promising strategy because the ‘first-personal’ properties of our experiences vary little across different experiences.

Take the kinds of perspectival properties that Paul points us to with her Go-Pro example and consider Jackson's Mary before she has left her room. Let us imagine that Mary has a certain white vase in her room and, hence, she had many black-and-white visual experiences of seeing this vase from different 'immersed first personal viewing angles'. One day, this vase is secretly painted bright blue by someone else, and then left in its original location for Mary to discover. When Mary sees the vase, her resulting experience will be revelatory and 'she will not say "ho, hum"' (Jackson 1986: 291), but all the 'wow' factor here stems from the new qualitative property (the bluishness) that Mary is now acquainted with for the first time. The perspectival properties involved in this visual experience (e.g., those involved in seeing the blue vase from a certain spatial location and angle, and with one's body in a certain position etc.) are, if not identical, at least extremely similar to ones she was already acquainted with.

The general point is that the perspectival properties of experiences will often vary very little across different experiences. In which case, these properties look to be good candidates for being 'projectible' from one kind of experience type to another using imagination and cognitive modelling, and that makes them bad candidates for defending (^GVP) because it is easy to grasp the perspectival properties of experiences one has not had before.

And the same points apply even more strongly to the supposed for-me-ness properties of our conscious experiences that Kriegel identifies. This is because, as noted earlier, the subjective character of an experience is meant to be what makes it a phenomenally conscious experience at all. In which case, the subjective or for-me-ness character of our experiences is something that is *invariant* across *all* phenomenally conscious experiences.

6.2 The attitudinal conception of first-personal perspectives

If one were to defend (^GVP) by appealing to the attitudinal notion of a first-person perspective the suggestion would be that (^GVP) is true once we interpret it along the lines of (^GVP₅) where

now (unlike the specific subject reply in §6.2) the indexical ‘me’ is meant to be playing an ineliminable role in characterising the relevant state of knowledge:

(^GVP₅) I can approximately determine the subjective value of what it is *for me* to Φ only if, for some way w , I have gold-standard knowledge that w is a way it feels *for me* to Φ .

The problem here is that the same kinds of considerations already identified for rejecting (^GVP) in §§1-2, and (^GVP₄) in §5.2, can be redeployed against (^GVP₅). For if the way that it would feel for me to Φ is relevantly similar to other different experiences I have had, and to the experiences of other people, then even if I have not Φ -ed I may be able to have lots of silver and bronze WIL-knowledge not only of propositions of the form is ‘ w is a way it feels to Φ ’ but also of the form ‘ w is a way it feels for *me* to Φ ’. And that knowledge might suffice to put me in a position to grasp the subjective value of Φ -ing not only for people in general but also for myself.

There are some kinds of indexical knowledge about what it is like to Φ that I cannot have until I actually Φ . For example, I cannot know propositions of the form ‘this is what it feels like for me to Φ *now*’ until such time as I am Φ -ing. But that knowledge will also be lost as soon as I stop Φ -ing. In which case, one cannot appeal to such knowledge in trying to defend (^GVP) without coming into conflict with the anti-sceptical points identified in §2, namely, that people can both retain their knowledge of what it is like to Φ , and assess the subjective value of Φ -ing, after their experience of Φ -ing has finished (cf. Lewis’ 1998 critique of what he calls “the third way” of missing the point of Jackson’s knowledge argument).

So, despite their prominent role in Paul’s work, first-personal perspectives—in either the experiential or the attitudinal sense—are not going to help to answer the equivocation objection.

Conclusions

I submit that the equivocation objection is now strengthened and remains unanswered with respect to many instances of Paul's argument, including her famous example of deciding whether to become a parent for the first time. The modelling argument strengthened the key assumption that silver and bronze knowledge of experience could suffice to allow someone to grasp the subjective value of experiences they have not had. And the modelling argument also helped us to show how each of the replies we considered in §§4-6 succumbed to a kind of shared dilemma, namely, they relied on assumptions that, if sound, would either make knowledge of the subjective values of our experiences too hard to come by (in violation of the anti-sceptical premise identified in §2), or too easy to come by (violating Paul's assumptions about the inaccessibility of such knowledge in the absence of experience).

Where the equivocation objection fails (§3) are cases involving radically new sensory experiences. However, such experiences are a small subset of the kinds of cases that Paul wants to class as transformative experiences. It may be that there is some other more successful reply that would extend to all of Paul's examples of transformative experiences. But assuming that is not the case, what wider morals should we draw from the partial success of the equivocation objection?

One moral I think we should draw here is that some rethinking of the concepts of transformative experiences is in order. In Paul's work transformative experiences are presented as if they are experiences of Φ -ing such that, in any normal circumstances, one cannot know what it is like to Φ without having Φ -ed oneself and, hence, one cannot even approximately grasp the subjective value of Φ -ing without having Φ -ed oneself. However, our discussion suggests that, for many supposedly transformative experiences, people can, in principle, build cognitive models of what it would be like to have those experiences

that would be good enough for them to at least make approximately accurate judgments about their subjective value.

That said, we shouldn't underestimate the significant difficulties that can be involved in trying to build such cognitive models. There are all sorts of problems that can come up when trying to imagine an experience one hasn't had. For example, people will vary in their skills at imagining experiences they haven't had (Kind 2020), putative information (from sources like testimony, art, and theories) about what the target experience is like can be mistaken or misleading, you may not have had any relevantly similar experiences to compare the target experience with, or you may simply not have access to such information (perhaps due to biases and limitations connected to one's social location, see Kind 2021 and Ngo 2017), and so on. In many situations then a person will not be in a position to form a good cognitive model of the relevant experience they have not had, and it may be very difficult for them to improve their abilities to form cognitive models of that experience. But those points are all consistent with the claim that it is *possible*, in normal circumstances, for a person to be in a position to form a good cognitive model of an experience they have not had.

With these issues in mind, I suggest we need to think of transformative experiences less as a fixed class of experiences such that no person could ever know what one of those experiences is like without having had that experience themselves. Rather, we should think of whether an experience would be a transformative experience for a particular person given their *current* abilities to build informative cognitive models of that experience. If, in their current context, a person only has an ability to form a very poor cognitive model of what it is like to Φ , then Φ -ing will be a transformative experience for them in that context. Similarly, if, in that context, they tried to use a cognitive model in

deciding whether to Φ they would be criticisable for essentially the same kinds of reasons that Paul sets out so forcefully.

However, it may not be the case—and here is where we part ways with Paul—that it would be hopeless for such a person to try and improve their cognitive modelling of Φ -ing by seeking out new sources of information about what it like to Φ like stories, testimony, and theories, or distinct but partially similar experiences, etc. It is not easy to build an informative cognitive model of an experience one has not had, but for many kinds of experiences—including many of the experiences Paul identifies as ‘transformative experiences’—it may be possible and valuable to do so.⁸

References

- Campbell, T. & Mosquera, J. 2020. “Transformative experience and the shark problem.” *Philosophical Studies* 177 (11):3549–3565.
- Cappelen, H. & Dever, J. 2017. “Empathy and transformative experience without the first person point of view.” *Inquiry: An Interdisciplinary Journal of Philosophy* 60 (3): 315–336.
- Cath, Y. 2019. “Knowing What It is Like and Testimony.” *Australasian Journal of Philosophy* 97 (1):105–120.
- Fisher, J. C. 2006. “Does simulation theory really involve simulation?” *Philosophical Psychology*, 19(4): 417–432.
- Goldman, A. 2006. *Simulating minds: The philosophy, psychology and neuroscience of mindreading*. Oxford: Oxford University Press.
- Jackson, F. 1986. “What Mary Didn't Know.” *Journal of Philosophy* 83: 291–95.
- Harman, E. 2015. “Transformative Experiences and Reliance on Moral Testimony.” *Res Philosophica* 92 (2): 323–339.
- Kind, A. 2020. “What Imagination Teaches.” In John Schwenkler & Enoch Lambert (eds.), *Becoming Someone New: Essays on Transformative Experience, Choice, and Change*. Oxford: Oxford University Press: 133–146.
- 2021. “Bridging the Divide: Imagining Across Experiential Perspectives.” In Christopher Badura & Amy Kind (eds.), *Epistemic Uses of Imagination*. New York: Routledge: 237–259.
- Kriegel, U. 2005. “Naturalizing Subjective Character.” *Philosophy and Phenomenological Research* 71 (1): 23–57.
- Krishnamurthy, M. 2015. “We Can Make Rational Decisions to Have a Child: On the Grounds for Rejecting L.A. Paul’s Arguments.” In R. Vernon, S. Hannan & S. Brennan (eds.), *Permissible Progeny*. Oxford University Press.

⁸ Thanks to an anonymous reviewer for this journal for a number of helpful suggestions. Thanks also to Laurie Paul, David Braddon-Mitchell, Brian Weatherson, and others in audiences at Deakin University, the University of Queensland at the “Epistemic & Personal Transformation: Dealing with the Unknowable & Unimaginable” workshop, the University of St Andrews, and the “Phenomenal Knowledge and Testimony” workshop run by the *How Does it Feel?* research group at the University of Duisburg-Essen and the University of Liverpool.

- Levine, J. 2001. *Purple haze: The puzzle of consciousness*. Oxford: Oxford University Press.
- Lewis, D. 1998. "What Experience Teaches." In *Mind and Cognition: An Anthology*, 2nd edn, ed. W.G. Lycan, Oxford: Blackwell: 449–519.
- Marlantes, K. 2011. *What It is Like to Go to War*, New York: Atlantic Monthly Press.
- Mathony, M. & Messerli, M. 2022. "Transformative Experiences, Cognitive Modelling and Affective Forecasting." *Erkenntnis*:1–23.
- Ngo, H., 2017. "Simulating the Lived Experience of Racism and Islamophobia: On 'Embodied Empathy' and Political Tourism." *Australian Feminist Law Journal*, 43(1), pp.107–123.
- Paul, L. A. 2014. *Transformative Experience*. Oxford University Press.
- 2015a. "What You Can't Expect When You're Expecting." *Res Philosophica* 92 (2):1–23.
- 2015b. "Transformative Choice: Discussion and Replies." *Res Philosophica* 92 (2): 473–545.
- 2016. "Who Am I? The Immersed First Personal View." In *Power and Limits of Artificial Intelligence*, edited by Antonio M. Battro and Stanislas Dehaene, 106–132. Vatican City: Pontificiae Academiae Scientiarum Scripta Varia.
- 2017a. "First personal modes of presentation and the structure of empathy." *Inquiry: An Interdisciplinary Journal of Philosophy* 60 (3): 189–207.
- 2017b. "De se preferences and empathy for future selves." *Philosophical Perspectives* 31 (1):7–39.
- 2017c. "Phenomenal Feel as Process." *Philosophical Issues* 27 (1): 204–222.
- Zeman, A. Z. J., Dewar, M., & Della Sala, S. 2015. "Lives without imagery: Congenital aphantasia." *Cortex* 73: 378–380.