

Self-Locating Content in Visual Experience and the ‘Here- Replacement’ Account

Draft – Jonathan Mitchell (jonathan.mitchell@manchester.ac.uk) - forthcoming in *Journal of Philosophy* (please do not quote without permission)

Introduction

Are there forms of self-consciousness to be found in experiential states? According to one affirmative answer, there is a kind of self-consciousness present in perceptual experiences in virtue of their perspectival character. Focusing on the visual case, some philosophers claim that certain types of visual experiences have self-locating and so first-person (or *de se*), spatial contents. For example, standing in front of an object (x) located in egocentric space, visual experience represents not merely the object and its spatial location (e.g. x as in front), but also the perceiver’s spatial location – their perspective – relative to the object. The content of such a visual experience is approximated in relational egocentric terms as $\langle x \text{ as } in \text{ front of } me \rangle$, and so involves a self-locating component. Following John Schwenkler’s framing of this view, we can label this the *Self-Location Thesis* (SLT hereafter).¹ We will have occasion to clarify the SLT further in due course.

We can follow Schwenkler in contrasting the SLT with the so-called *Minimal View*. According to this view, the relevant visual experiences admit of characterizations without self-locating contents, where the principal strategy is to *monadize* the relevant contents.² In

¹ See Schwenkler 2014: 139. For adherents of the SLT see Cassam 1994: 45; Husserl 1989: 166; Peacocke 2000: 264-7; 2014; Hurley 1998: 207-47; Schwenkler 2014: 137-155; Bermudez 2002: 89 (see also Bermudez 1998: Ch.5 which draws on Gibsonian ecological perception; see section 4 for discussion).

² See Perry 1986; Campbell 1994: 4.1; 2002: 9.3; Evans 1982: 232-3; Musholt 2015: Ch.3 (see also Wittgenstein 1921: 5.6-5.641). Beyond the spatial case, Jean-Paul Sartre (2004) also argues that ‘pre-reflective’

a visual experience in which one is standing in front of an object (x) located in egocentric space, the experience represents the object and its monadic spatial location (e.g. x as in front); it has ‘nothing to say’ on the location of the perceiver. Note, the relevant contents are still specified by egocentric spatial predicates, such as *in front, to the left, to the right, behind*. It is just that such egocentric spatial properties as figuring in experience are, on this view, monadic rather than relational, and so are not given subject-reflexive qualifications by way of first-person indexicals (i.e. me, my, I). This caveat is essential since the Minimal View does not deny that visual experience is perspectival or (spatially) egocentric, but rather contests whether the relevant visual experiences are self-locating in virtue of that perspectival character.

This paper critically examines the SLT with respect to *dynamic-reflexive* visual experiences, which involve the movement of an object toward the location of the perceiving subject. These cases are noteworthy because they create problems for the monadizing strategy of the Minimal View. The main aim of this paper is to offer an interpretation of these cases which resists attributing them self-locating content, arguing for a replacement with a non-conceptual equivalent of the indexical ‘here’ (the *b*-replacement account).³

The structure of the paper is as follows. Section 1 clarifies the SLT, arguing that the view should be construed as positing nonconceptual *de se* content, which meets a

experience is impersonal, in that it does not contain a *de se* component in its content; he writes ‘there is no place for *me* at this level; the *me* appears only with the reflective act...’ (2004: 13). See Peacocke (2014: 34-5; 220-23) for discussion of Sartre’s views within this context.

³ NB: ‘impersonal’ accounts of the contents of visual experience are not committed to no-self or no-subject views of the metaphysics of experiences, expressed by the claim that experiences either do not require subjects or are not ‘owned’ by subjects. The metaphysics of the domain of subjects and experiences is a different issue from the theory of the intentional content of the relevant experiences.

Perceptual-Demonstrative constraint. Section 2 then considers dynamic-reflexive visual experiences, arguing that a self-locating interpretation of their content is plausible, not least because the monadizing strategy of the Minimal View (along with some other attempts) fails. Section 3 presents an alternative *b*-replacement account of their content. Finally, section 4 considers an extension of the *b*-replacement account to cases of visual kinesthesia.

Before proceeding, let me highlight the import of the project. What follows contributes to the ongoing project of determining the admissible contents of experience, specifically visual experiences. Dynamic-reflexive visual experiences are not usually discussed at length in such contexts, so an in-depth study of their content is welcome. More specifically, however, such visual experiences naturally admit of a self-locating interpretation (as will become apparent in what follows). As such, it is incumbent on those who are sceptical about the presence of first-person content in visual experience to consider such cases and suggest alternative interpretations. The *b*-replacement account offered here is one such alternative.

1. The SLT and the nonconceptual *de se*

To begin with, note that the kind of content that both the SLT and the Minimal View are concerned to dispute is the personal level content of visual experiences, rather than subpersonal content (or ‘information’) processed by the visual system. Whether subpersonal perceptual states have self-locating content (or carry self-related information) is not what is at issue. Connected to this, the supposed differences in personal level content – between say ‘x as in front’ and ‘x as in *front of me*’ – should be reflected in differences in phenomenal character, that is in what-it-is-like for the subject to enjoy the relevant visual experiences (this will be important later in considering specific arguments).

Schwenkler's framing of the SLT is instructive: 'let the *Self-Location thesis* be the claim that, simply in virtue of its perspectival character, visual experience can include the location of the perceiver among its face value contents'.⁴ The restriction that visual experiences have self-locating content *solely in virtue of their perspectival character* highlights an important caveat. Namely, that the correctness of the SLT does not turn on the presentation of any 'self' object, specifically one's body (or parts thereof), in the relevant visual experiences.

Christopher Peacocke, another advocate of the SLT, expresses this point as follows:

A subject reflexive state or event can have a content that refers to the subject of that state or event without the subject also being given at the same time in some other, further way, be it perceptually, or in some other demonstrative fashion made available by some conscious state. (Peacocke 2014: 12)

Also, in a later passage:

A first person notion can also enter the content of experience even when the subject is not perceiving any part of his own body. You can experience a ball as coming through the air towards you even if you have no sensation in any part of your body, and do not perceive any part of your body. (Peacocke 2014: 121)⁵

Finally, Schwenkler says the following:

⁴ See Schwenkler 2014: 139. Note the scope qualifier 'can'. A stronger version of the SLT would be that *all* visual experience includes self-locating content. This paper only considers the weaker claim.

⁵ See also Peacocke 2014: 36; 2000: 264 (see also Bermudez 1999: 108).

In saying that it [visual experience] does this ‘simply in virtue of its perspectival character’ we are alluding to an idea...that visual experience can be self-locating even when ‘the self’ is entirely out of view; that is, even when the perceiver’s body is nowhere in the field of vision. (Schwenkler 2014: 139)

These passages highlight an important caveat: any supposed *de se* component, as figuring in the content of experience, should be identification-free (or non-criterial). That is to say, *de se* components of contents should not ‘in their subject use’ (to borrow a phrase from Wittgenstein),⁶ leave open the possibility of misidentification, as would be the case if such visual experiences ‘of oneself’ were based on perceptual-cum-demonstrative identifications (in the latter case, we have to provide for the possibility of error).

Let me expand on the above points with examples. Standing in a room, I may be unaware that mirrors are reflecting various objects therein, or at least unaware which mirrors are reflecting which objects (as in a house of mirrors attraction at a funfair). In front of me, I see the back of a tall man wearing a duffel jacket. As it turns out, this is a reflection of my back. My visual experience has a content approximating to ‘that man with the duffel jacket’ in which the demonstrative ‘that man’ refers to myself, but which I fail to identify as such; I identify the figure as someone else. However, once appraised of the situation (*vis-à-vis* the mirrors), I recognize that the man in the duffel jacket is, in fact, the back of myself, and so (re)identify the figure *as myself*. In the latter case, the perceiving subject is given to themselves by way of a perceptual (re)identification of *something*, namely a perceptually presented object with specific visible properties, *as oneself* (‘that’s me’ I might say on the basis of my visual experience). So, in one, not entirely unnatural way of putting

⁶ Wittgenstein 1958: 66-7 (see also Shoemaker 1968: 555-67).

it, I spatially locate something as myself, namely ‘that man over there with the duffel jacket’, although I do so by means of a perceptual-cum-demonstrative identification.

Alternatively, so the line of thought of the SLT goes, the same visual experiences – both when I fail to recognize that the man is myself, and when I do – have a self-locating egocentric spatial content approximated by ‘that man with the duffel jacket *in front of me*’, where the relevant *de se* component is supposedly identification-free. Regarding this *de se* aspect of the content, there is no perceptual-cum-demonstrative identification of something as myself which grounds it. Instead, the sense in which the experience represents oneself is, so the SLT claims, not in virtue of taking something I can (and in the second case do) see as *myself* – identifying something *as* myself. Instead, the sense in which the experience represents oneself is (supposedly) *solely in virtue of its perspectival character*, that is in virtue of the experience representing its objects egocentric spatial position relative to the location of the perceiver. It is in this sense that the visual experience is supposed to have self-locating content.⁷

The above reflections generate a constraint on cases of visual experiences with putatively self-locating content, which the SLT should meet:

Perceptual-Demonstrative Constraint (PDC hereafter): the sense in which the relevant class of visual experiences have self-locating content is *not* in virtue of a perceptual, or otherwise demonstrative, (re)identification of *something* as *oneself*.

This constraint also reflects the thought, reflected in visual phenomenology, that one does not (or at least need not) *see oneself* in visual experiences. As such, the contention of the SLT, reflected in the PDC, might be expressed by saying that it isn’t *necessary* to have any

⁷ See Peacocke 1992: 70-2.

part of one's body in view in order for the phenomena of self-location in visual experience to get off the ground. Importantly though, the constraint clarifies that even in those cases where one does see what one (re)identifies as oneself, as in the second mirror case, or in more common-garden variety cases of directly perceiving one's body (or parts thereof), the supposed self-locating content of visual experience doesn't (or at least doesn't necessarily) turn on this feature.⁸ Put otherwise, the presence of a self-locating component in the content of visual experience does not turn on some perceptual-cum-demonstrative (re)identification of a perceivable object *as oneself*. Instead, the relevant self-locating contents, which visual experiences putatively have solely in virtue of their perspectival character, is identification-free; indeed, it is in virtue of this that self-locating contents are taken to be genuine *de se* contents.⁹

Having made the above clarification of the SLT, let me now detail how we should think of the *de se* component, specifically its status as a *nonconceptual de se*. First, note that defenders of the SLT are not claiming such experiences are constituted by 'I-thoughts', where that would involve deploying the first-person concept – say in the form of judgements to the effect that 'that object is to the right of *me*'. We can capture this by saying that the relevant visual experiences, on the SLT, are supposed to be *nonconceptual*

⁸ Note this way of formulating the SLT is compatible with the thought that when one's body, or a part of it (e.g. a nose, torso, pair of arms) are in view, and are perceived from a characteristically first-person angle, it is possible to have a properly first-personal awareness of them, which in some important sense doesn't depend on a (re)identification of them *as one's own* (e.g. identification free forms of putative bodily awareness).

⁹ In a related context, Cassam (1994: 55) asks whether spatial perception is compatible with 'body blindness'. Peacocke and Schwenkler maintain that not only that this is possible, but in line with the PDC should claim that the *de se* content of spatial perception does not turn on any perceptual-cum-demonstrative identification of a body *as one's body* (cf. fn.8). Self-locating visual perception is, therefore, taken to be compatible with body blindness. I will accept this claim here, although more needs to be said on this issue.

relative to the *de se* component. This is important because it allows us to bypass the broader question of whether the contents of visual experiences are best construed as (wholly) conceptual or nonconceptual with respect to their non *de se* aspects (e.g. their representation of fine-grained colour properties). As such, even if one were to think that a visual experience of a willow tree *as a willow tree*, necessarily involves the deployment of specific conceptual recognitional capacities, one might still, in line with the SLT, think that the visual experience involves a *non-conceptual* self-locating component, such that its content approximates as <willow tree in front of *me*>.

However, regardless of whether such visual experiences include co-present I-thoughts – they typically do not – might not the relevant *de se* component be conceptual in some other sense? Arguably not, since this would lead to the SLT violating the PDC. Given we are considering the level of visual *experience*, then the relevant concepts in play are likely best thought of as conceptual capacities, along the lines of short-lived recognitional capacities. For example, perhaps a visual experience as of a *red apple*, involves the passive drawing into operation, in experience, of recognitional capacities for which the subject possesses the relevant sortal properties.¹⁰ Nevertheless, even if such a model of the contents of visual experiences were plausible for the relevant sense-perceptible properties, this cannot be the model for *de se* contents since such conceptual capacities are identification-dependent. Remember, the *de se* as it figures in visual experience is supposed to turn, according to the SLT, *solely on its perspectival character*, not on some putative perceptual-cum-demonstrative (re)identification of something *as oneself*, say in virtue of its perceptible features – it is supposed to be identification-free.

¹⁰ See McDowell 1994: 59; 2009: Ch.14. Such recognitional capacities arguably allow the content of experience to be taken up in propositional attitudes by the use of context-sensitive demonstrative phrases like ‘that shade’, or their mental equivalents (e.g. what is singled out by acts of attention). For further discussion, see Kelly 2001a: 397-420, and Peacocke 2001: 239-64.

Perhaps there is an alternative way of construing the SLT along conceptualist lines. However, let's accept that the sense in which the SLT intends to theorize the structure of the *de se* component in the content of the relevant visual experiences is as nonconceptual at least in the sense that it does not involve co-present I-thoughts or conceptual recognitional capacities. This position has the benefit that the SLT can claim that non-human animals and human infants can enjoy visual experiences with the relevant self-locating contents, despite lacking mastery of the first-person concept. In that sense, the SLT concerns itself with a primitive level of self-representation, as a level of self-representation below the level of application and mastery of first-person indexical concepts.¹¹

2. The Campbell-Perry strategy and dynamic-reflexive visual experiences

This section first expands on the Campbell-Perry strategy of monadizing spatial contents. It then introduces a specific class of visual experiences which admit of a *de se* interpretation and explains why the Campbell-Perry strategy does not work. It also considers alternative strategies which likewise fail.

Here is a passage articulating the strategy in question:

Although vision provides a great deal of information about oneself, information whose most direct articulation uses the first person, this is not because the visual information already employs the first person. Rather, the egocentric frame [of reference] used in vision employs monadic spatial notions, such as 'to the right', 'to the left', 'above', 'in front', and so on, rather than relational notions, such as 'to my right', 'above me', 'in front of me', and so on.' (Campbell 1994: 119)

¹¹ There is more to be said on this issue (see Peacocke 2014; Bermudez 1998 and Hurley 1998 for discussion).

Also, in a later work, Campbell argues for the view as follows:

Now in stating the spatial content of vision, we do not seem to need these relational notions. We do not need the general conception of something's being to the right or left of an arbitrary subject. Rather, we need the more primitive monadic egocentric terms. These are notions such as 'x is to the right', 'x is below', and so on. (Campbell 2002: 184).¹²

The strategy of monadizing any putative self-locating visual contents, which include relational egocentric notions, is primarily motivated by considerations of parsimony. More specifically, we should see the relevant considerations of parsimony as related to what is (minimally) required to support perception-action explanations. Unpacking this, the Campbell-Perry strategy, as it will be understood here, turns on the following claim: we can get the same perception-action intentional explanations with visual-spatial contents that are *monadic* rather than *relational*. Given this, construing the spatial contents of vision in terms of self-locating relational egocentric notions is surplus to requirements. It involves positing an additional *de se* component that is not necessary to explain the relevant

¹² See also Perry 1986. Note for Perry the monadic egocentric terms still *concern* the subject, forming part of the *context* which determines the truth conditions of propositions which includes subject's spatial location. But that self-location is not explicitly represented (and so is not part of the content), it is instead an 'unarticulated constituent'. As such, the 'information' provided by visual experience is only *implicitly* self-relational.

perception-action links; these can be explained just as well with monadic egocentric notions.¹³

Consider the following example, drawn from haptic perception. When moving house, I attempt to lift a piece of furniture. I pick up the piece of furniture at one end and enjoy a haptic experience of its being heavy; it resists being brought to full elevation and requires effort to sustain at a substantial angle of incline. In exacerbation, I exclaim “it is too heavy” and drop the furniture back down. There are (in this context) two options for the content of the experience; either the content involves only the monadic notion ‘heavy’, or it involves a relational notion ‘heavy-for-me’. How are we to decide which is preferable?

To aid us in answering, consider the following extension of the example. Much to my dismay, my brother easily lifts the furniture, retorting “it is not heavy at all”. Being of a philosophical bent, I come to realize that heaviness is a relational property. Nothing is heavy *per se*, but rather heavy (or not) relative to a particular subject and their ability to lift, or otherwise manipulate, the relevant object (physical particulars have non-relational objective weights, but that is something different). Do my realizations, therefore, count in favour of the *de se* interpretation, that my haptic experience has the content ‘x is *heavy-for-me*? Not obviously. The fact that an explication of the property of heaviness reveals that it is a relational property does not mean that haptic experience represented it as such (that the content was self-involving in this way).¹⁴ Instead, it is as plausible that my haptic experience had a monadic content ‘x as *heavy*’. Indeed, what I realize is that a property of an object which seemed monadic is, in fact, relational (or dyadic). We might put this by saying that while the metaphysical nature of heaviness is relational, it does not wear that

¹³ There is a connected worry that insofar as self-locating contents necessitate *self*-representational capacities then this is cognitively overdemanding with respect to more ‘primitive’ subjects who nonetheless enjoy visual experiences (I return to this point in section 3).

¹⁴ See Shoemaker 1996: 254.

relational nature on its phenomenological sleeve. Further to this, we can maintain the relevant perception-action links on the monadic interpretation. My dropping the couch back to the floor is arguably explained as well by a haptic perception involving the representation of the monadic property (its seeming *heavy*), as one in which there is an explicit *de se* component (its seeming *heavy-for-me*).¹⁵ So, we have a case in which the Campbell-Perry strategy works. What we can draw out of this discussion is the way the strategy works in reference to monadizing contents and considerations of parsimony, turning principally on maintaining the relevant perception-action links.

With the above in mind, let me now introduce the cases that will be the focus for the rest of this section (and the following one) namely *dynamic-reflexive* visual experiences. We can start with a distinction between static and dynamic visual experiences. A paradigmatic example of static visual experience would be a visual experience as of a stationary object located in front of the perceiver (these are the cases that are usually discussed).¹⁶ Contrastingly, a case of dynamic visual experience is one where the object is perceived in locomotion, say as moving from left to right across the perceiver's visual field. A simple example could be a visual experience as of an aeroplane moving horizontally across the sky (from left to right).

Dynamic-reflexive visual experiences involve a visual experience of an object moving toward *where the perceiving subject is located*. A paradigmatic example is a visual experience of a golf ball as moving through the air towards the perceiver, as they stand on the fairway. Let us first say that dynamic-reflexive visual experiences have a content which is of the

¹⁵ Note, one might prefer to analyse this case of haptic perception along the lines of an affordance, say *liftable*, rather than the attribution of a property to an object. However, the point stands: in perceiving the couch as heavy I don't experience it as *not-liftable-for-me*, but just as *not-liftable*.

¹⁶ See Musholt 2015: Ch.3 for a development of the Campbell-Perry strategy in arguing against the SLT in cases of static visual experience.

general structure $\langle(A) \text{ moving towards } (B)\rangle$.¹⁷ As such, there are two spatially located objects, and one of them, (A), is moving across space towards the other, (B). Further, it is essential to emphasize that in the cases under consideration, the perceiving subject (B), is stationary (I consider putative ‘self-movement’ in section 4).

Let us consider another example. A cat stalking a mouse may enjoy a visual experience of the mouse as located in egocentric space, as ahead of where it is, and will be careful to remain stationary to not frighten its prey. Suddenly the mouse makes an ‘error of judgement’ and moves toward the cat. When this happens, it is plausible the cat has a dynamic-reflexive visual experience with a content of the general form $\langle(A) \text{ moving toward } (B)\rangle$, where (B) is where the cat is located. Clearly, in an example of this kind (as above), there is no suggestion that the first-person concept figures in the relevant contents; any putative self-locating component will be nonconceptual.

Now, the SLT will offer a *de se* reading of the spatially located object (B). The content of dynamic-reflexive visual experience as including a self-locating component would be of the form $\langle x \text{ moving towards } me\rangle$. What considerations support this? First, we can consider the unavailability of the Campbell-Perry strategy. Remember, the Campbell-Perry strategy is principally that of monadizing the relevant contents. How would this work? Remember, the theory-neutral specification of the structure of the content is $\langle(A) \text{ moving towards } (B)\rangle$. Perhaps the strategy would be to get rid of one of the *relata*, either (A) or (B). Given (B) is defined as where the perceiving subject is located, perhaps we get rid of (B). However, this clearly won’t work. ‘(A) moving towards (B)’ cannot be re-interpreted as ‘(A) moving toward...’. Moving towards, at least in its spatial sense, is necessarily a

¹⁷ Matters are, in fact, a little more complicated. Dynamic-reflexive visual experience is a form of event perception, insofar as the perception of movement is necessarily that of a temporally extended event as unfolding. For our purposes, we need not worry about this complication, since both the SLT interpretation and alternative accounts can accept that dynamic-reflexive visual experiences are forms of event perception.

moving towards *something*. Likewise, getting rid of (A), and having merely ‘moving towards (B)’ will be a non-starter. Insofar as the relation of moving towards is a spatial relation, it only makes sense as *something* in space moving towards *something* else in space. On this basis, ‘moving towards’ is necessarily a 2-place spatial relation. Insofar as we delete one of the terms (or objects) of the spatial relation we fail to have an instance of the relation in question.

Supporting the above points, the necessity of ‘moving towards’ being a 2-place spatial relation is reflected in the phenomenology. It is not as if ‘moving towards’ is merely explicated, in reflection, as a 2-place relation, whereas visual experience somehow presents ‘moving towards’ as only a 1-place relation (or as not relational at all). Insofar as there are cases of visual experience in which an object (A) is perceived (in locomotion) as moving towards where the perceiving subject is located (B), then to preserve the distinctiveness of such experiences we arguably need both terms in the content to reflect the phenomenology. If this is correct, then the strategy of appealing to non-relational monadic spatial predicates to characterize the relevant contents is not an option.

The defender of the Campbell-Perry strategy might, however respond, that we are relying too much on the particular form of the words ‘moving towards’, in trying to capture the content of such experiences. Perhaps a monadic notion like *approaching* would be as plausible a candidate (and one that does not lend itself so easily to a *de se* interpretation). However, at least one problem with the claim that we could capture the content of dynamic reflexive visual experiences in terms of <x approaching> is that such a content is ambiguous between dynamic visual experience and specifically dynamic-*reflexive* visual experience. I might enjoy a visual experience in which I see a ball approaching the fairway, and so which would have the approximate content <ball *approaching*>, but this would be a dynamic visual experience, rather than a reflexive one.

Alternatively, the defender of the Campbell-Perry strategy might claim that ‘moving closer’ is available, where the content would be $\langle x \text{ moving closer} \rangle$. Again, though we face the same problem. I can certainly have a visual experience of the golf ball *moving closer* to the green, as contrasted with a visual experience of a golf ball *moving closer* to where I am located. A content of the form $\langle \text{ball moving closer} \rangle$, with only one object-term, is ambiguous between these two cases, that is ambiguous between dynamic and dynamic-reflexive visual experiences. Of course, once we recognize this, and formulate the content as $\langle (A) \text{ moving closer to } (B) \rangle$, the SLT is off the ground since it can argue that what most plausibility fills the (B) position, as distinguishing dynamic-reflexive visual experience, is *me*, such that the content would be $\langle (A) \text{ moving closer to } me \rangle$.¹⁸ The lesson to be drawn from the above is that dynamic-reflexive visual experiences are resistant to the Campbell-Perry strategy of using, as Perry puts it, ‘an n-place predicate or concept to deal with an n+1 ary relation.’¹⁹ It is unclear what n-place predicate could replace the relational predicate ‘moving towards’, whilst distinguishing the type of case in question.

However, even if the above is correct, the SLT still has more work to do because the critic of the SLT might appeal to non *de se* specifications of what occupies the (B) position in the general content form $\langle (A) \text{ moving towards } (B) \rangle$. So, the defender of the SLT also needs to show that these readings are problematic. In what follows, I argue that two such strategies for replacing the *de se* fail.

Consider a non *de se* reading of the content, specified as the ‘ball moving towards *somewhere*’. The first problem is that we no longer clearly have a case of dynamic-reflexive

¹⁸ Note, if someone prefers *moving closer* to *moving towards* in terms of specifying the relevant relation for dynamic reflexive visual experience that is fine, as long as it is kept in mind that *moving closer* will, for the contents in these contexts, have be specified in the general form $\langle (A) \text{ moving closer to } (B) \rangle$ for the reasons given.

¹⁹ Perry 1986: 221

visual experience, defined as involving a visual experience of an object moving toward where the perceiving subject is located. To be clear, this definition does not surreptitiously build in a *de se* component by stipulation; the condition ‘where the perceiving subject is located’ does not presume that such a self-location is represented as a *self*-location, as necessarily part of the content of the visual experience. However, it does place a constraint on what counts as an instance of the type. The idea that in a case of dynamic-reflexive visual experience, the location of object (B) can merely be specified as ‘somewhere’ does not meet that constraint. This is demonstrated if we consider that the content <ball moving towards *somewhere*> – where somewhere is an arbitrary spatial location (akin to anywhere) – is ambiguous between a putative case of dynamic-reflexive visual experience, where the somewhere happens to be where the perceiver is located, and dynamic visual experience *per se*, say the case of seeing a golf ball moving towards *the green* (which is undoubtedly *somewhere*). The proposed non *de se* content, therefore, fails to individuate the relevant class of dynamic-reflexive visual experiences.

Further to the above, the non *de se* interpretation does not support the relevant perception-action links. Consider that the action of raising one’s arms as a response to *what one sees* makes sense, so the defender of the SLT will argue, if the content is <ball coming towards *me*>, in a way that it wouldn’t if the content were <ball coming towards somewhere> or <ball coming towards the green>. This is so even in the case where the *somewhere* turns out to be where the perceiver is located. Since a visual experience which leaves it open or does not sufficiently determine for its subject where that ‘somewhere’ is, does not support intentional action-explanations, in the same way as one which does.

Similar problems affect a replacement of the *de se* with a notion like ‘someone’. Consider the following example. I may be stood on the fairway with numerous other golfers. So, there is a range of possible individuals whose spatial locations could be the referent of the general subject ‘someone’, including my own (there is reference

indeterminacy). Yet there is a phenomenological difference between enjoying a visual experience of the ball moving towards where I am located, and enjoying a visual experience of the ball moving towards where Bill, or any other subject, is located. Moreover, this is no surprise since enjoying a visual experience of the ball as moving towards Bill is a case of dynamic visual experience rather than dynamic-*reflexive* visual experience. So, the ‘someone’ (analogously to the somewhere) replacement fails to respect the distinction between the cases in question. It fails to respect the way that dynamic-reflexive visual experience, as a type, is defined as a visual experience as of an object moving toward where the perceiving subject is located.

Further to this, it will not help to reply that as it turns out the ‘someone’ in fact refers to the perceiver themselves, even though their location is not part of the content *as their location*. As in the ‘somewhere’ case, a visual experience which leaves it open or does not (sufficiently) determine for its subject which ‘someone’ is in question will not support relevant intentional action explanations in the way that contents with *de se* components do. For example, say my visual experience has the content <ball moving towards someone>. On this basis, and taking that content at face value, it would be rational for me to judge that *someone* (whomever the term refers to) should take evasive action to avoid being hit. But insofar as that visual experience leaves it open precisely who that someone is, then even if it turns out to be me, taking those contents at face value wouldn’t provide me with a non-inferential reason for ducking, or moving in some appropriate way so as to avoid being hit, in a way that the *de se* interpretation <ball moving towards *me*> does.

In summary, the defender of SLT can argue for their position as follows. The *de se* interpretation of dynamic-reflexive visual experiences not only preserves the distinctness of the type of experiences in question but provides us with a plausible account – which neither of the impersonal varieties of ‘somewhere’ or ‘someone’ do – of why perceiving subjects have non-inferential reasons for action on the basis of such experiences. Consider

again the case of a visual experience of the golf ball moving towards where I am located. In such a case, a typical response would be to duck or to take measures to avoid being hit (e.g. raising one's arm, or jumping out of the way). The precise course of action I take in any particular instance is context-dependent (and subject to background beliefs). Yet, so the defender of a *de se* interpretation might argue, it is only because my visual experience represents the ball as moving towards *me*, as moving towards my location represented *as my location*, as having self-locating content, that my actions are considered rational responses to *what I can see*. Such self-locating contents give me non-inferential reasons for acting in the way I do, and so preserve the link between perception and action in these cases.²⁰

Given the discussion of this section, the SLT is in good shape concerning its interpretation of the content of dynamic-reflexive visual experiences. The goal of the next section is to provide an alternative replacement, which fairs significantly better than those considered above.

3. The *h*-replacement for dynamic-reflexive visual experience

3.1 *The h-replacement account*

²⁰ For the purposes of this paper I am going to assume that it is not open to the Campbell-Perry strategy to claim that the relevant *non de se* 'impersonal' component of the content satisfying the (B) position could be 'this body' or 'this person'. Such formulations would be, as the demonstrative term makes clear, *identification-dependant* in a way that arguably makes the content of dynamic-reflexive experience overly complex. It would also mean that a subject who for whatever reason could not *demonstratively identify* their body or themselves (so to speak), either in general or on specific occasions, could not enjoy such experiences, which seems implausible. Although more needs saying on this issue.

The central claim of this section is that it is possible to replace the supposed self-locating component in the content of dynamic-reflexive visual experience with a non-conceptual impersonal locational predicate, modelled after the indexical concept 'here'.

To get started, let me introduce some points about 'Here'-thoughts from Gareth Evans:

To understand how 'here'-thoughts work we must realize that they belong to a system of thoughts about places that also includes such thoughts as 'Its *F* over there', 'It's *F* up there to the left', 'It's *F* a bit behind me'. 'Here'-thoughts are merely the least specific of this series. We may regard this as an *egocentric* mode of thought.' (Evans 1982: 153)

The subject conceives himself to be in the centre of a space (at its point of origin), with its co-ordinates given by the concepts 'up' and 'down', 'left' and 'right', and 'in front', and 'behind'. We may call this 'egocentric space', and we may call thinking about spatial positions in the framework centring on the subject's body 'thinking egocentrically about space'. A subject's 'here'-thoughts belong to this system: 'here' will denote a more or less extensive area which centres on the subject' (Evans: 1982: 154).

Evans emphasizes that 'here'-thoughts' are embedded within a (holistic) system of egocentric spatial notions. Applying these ideas to the level of experience, a visual experience which involved some non-conceptual analogue of 'here', would embed that non-conceptual analogue within visual experiences that represented a range of egocentric spatial locations, such as 'over there', 'to the left', 'to the right', 'behind', 'in front'. Put otherwise, it is only in the context of visual experiences that represented such egocentric spatial properties that some non-conceptual equivalent of 'here' could get a grip. Building

on this, we should add that the non-conceptual analogue of ‘here’, as putatively figuring in the content of a visual experience, would not only be embedded within a holistic system of egocentric spatial notions, but would play a unique role within that system, namely referring to what Evans calls *a more or less extensive area which centres on the subject*.

With this framework, we can sketch a non *de se* reading of the content of dynamic-reflexive visual experiences. First, it bears emphasising that the claim is not that the indexical concept ‘here’ figures in the content of such visual experiences, but rather that a non-conceptual analogue does. Let’s keep this straight, by introducing a piece of notation for this feature, namely *b*. So, the claim is that the content of the visual experience of a ball as moving through the air towards where the perceiver is located can be given in the following way: <x moving towards *b*>.

To respond to an objection that we are unjustifiably introducing components into the contents of visual experiences, we can specify what individuates *b* (as it figures in those contents) by the following reference rule:

Reference rule for *b*: *b*, as it figures in the relevant visual contents, refers to where the perceiver is spatially located, as that more or less extensive area which centres on the subject.

This reference rule allows us to respect the principle in theory of intentional content that the introduction of a component into a content is legitimate *iff* the relevant experience is assessable for correctness or veridicality relative to the component (i.e. as determining, in part, the correctness conditions). So, in a case where the perceiver is subject to a visual hallucination of a dynamic-reflexive type, it might *seem to the subject as if* x is moving towards that more or less extensive area which centres on the perceiving subject, but as it turns out they are spatially located elsewhere. Contrastingly, a correct content for a dynamic-

reflexive visual experience will be one in which *x really is moving towards* the more or less extensive area centring on the perceiving subject.

The next feature of this account worth emphasizing is that it should not be confused with the claim that the content of the experience, relative to *b*, is itself structured by way of a mixed demonstrative-descriptive component (as given in terms of the reference rule). In the visual experience, whose content is <*x moving towards b*>, *b* is individuated by its reference rule. This reference rule, as stated, is a structured demonstrative-descriptive content. However, *b*, as it figures in the content of visual experience, is not itself structured. Put otherwise, *b*, as it figures in experience, is not as that structured descriptive-demonstrative *where the perceiver is spatially located, as that more or less extensive area which centres on the subject*. Instead, *b*, as it figures in the content of those visual experiences we are considering, is *unstructured*.²¹

Given we understand this point, we should also recognize that *b* is not akin to any perceptual-cum-demonstrative identification, of the form *this body (or part thereof)* occupying the relevant location. There is no suggestion that in the visual experience with the content <*x moving towards b*> the subject perceptually, or otherwise, demonstratively identifies *this body*, as being that towards which *x* is moving. It is not as if the subject has to first identify a referent for *b*, say as the location at which *this body* is, to then be able to enjoy an experience with the content <*x moving towards b*>. To see this point, consider that it is possible to imagine cases of dynamic-reflexive visual experience in which the perceiving subject has no sensation in any part of their body, and does not perceive any part of their

²¹ This analysis follows Peacocke (2014: 10), who makes a similar observation concerning how to think about the way 'now' figures in the content of present-tense utterances. Note it is *b*, as it figures in the content of those visual experience that is claimed to be *unstructured*, which is not to say that the *overall content* is unstructured (after all, if the content of experience contains different components then this suggests it is structured).

body. The subject has a visual field and is located in egocentric space, but the subject's 'body' is entirely out of view. Still, it would make sense – so the *b*-replacement account claims – to characterize the content of a dynamic-reflexive visual experience as enjoyed by them as $\langle x \text{ coming toward } b \rangle$.²² Put otherwise, a dynamic reflexive visual experience having a content which includes the non-conceptual component *b* does not do so in virtue of a perceptual presentation or otherwise demonstrative (re)identification of something as occupying the perceiver's location.

Finally, the *b*-replacement account can respect the type of experiences that dynamic-reflexive visual experiences are in a way that those previous attempts to depersonalize the relevant contents did not. A key issue there was preserving the fact that we are dealing with a two-place spatial relation, in which there are two terms (or 'objects') figuring in the content, which stand in a spatial relation. The *b*-replacement account respects this by preserving that general form $\langle (A) \text{ moving towards } (B) \rangle$, but provides a non *de se* interpretation in terms of $\langle x \text{ coming towards } b \rangle$. Further to this, the account has the benefit that it doesn't fall prey to the reference indeterminacy worries that undercut replacements in terms of 'somewhere' or 'someone'. This is a consequence of the fact that *b*, like 'here', is a genuine indexical in a way that 'somewhere' or 'someone' is not; there could be no confusion with some other spatial location – *b* fixes the relevant location.²³

3.2 *The Action Challenge*

²² Evans (1982: 152) makes a similar point, although at the level of Here-thoughts.

²³ One further question about the *b*-replacement account is whether there is a difference between the way that *b* is individuated and the way that any other egocentric location is. One option would be to say that there is not any such difference: e.g. the reference rule for *ahead and to the left* will say that it refers to that location ahead and to the left of where the perceiver is spatially located, as that more or less extensive area which centres on the subject.

With the *b*-replacement account outlined, I now consider a challenge to it. A defender of the SLT might argue as follows:

Arg. from Action

- (1) A central constraint on how we specify the spatial contents of visual experience is that the perception-action link is preserved, such that given the experiences have the relevant content, they provide non-inferential reasons for action and so support intentional action explanations.
- (2) The *b*-replacement account of the spatial content of dynamic-reflexive visual experiences fails in this respect. It is not clear content of the form <x moving towards *b*> can provide a non-inferential reason for the appropriate behaviour.
- (3) A self-locating content fares better in this respect. It is because the subject represents their *location as their location* in such experiences that the relevant intentional action explanations make sense.
- (4) So, a *de se* reading of the relevant class of visual experiences is preferable to an *b*-replacement account which falls short of explaining a crucial explanandum.²⁴

The challenge for the *b*-replacement account is showing (2) to be false. The defender of the *b*-replacement account needs to defeat the presumption that contents of the form <x moving towards *b*> fail to provide the subject who takes those contents at face value with non-inferential reasons for action.

To begin, we can turn to what Christopher Peacocke says about ‘Degree 0’ subjects, that is subjects who lack the capacity for any form of self-representation and so whose visual experience *ex hypothesi* does not include a *de se* component:

²⁴ Bermudez 1999: 117-8 and Hurley 1998 gesture in this direction, claiming that only perceptions with first-person content can preserve the relevant perception-action links.

...this creature remains at Degree 0, however, because it never represents anything as standing in certain relations to itself. None of its perceptual states have *de se* contents of such forms as *that thing is that direction from me*. Instead, they have *here*-contents, such as *that thing is that direction from here*...there are no features of its actions that are explained by its content-involving states that can be accounted for only by attributing *de se* contents. The above indexical and demonstrative contents – here, now, that event, that thing – together with the predicative and relational notions with which they can be combined suffice for the explanations of its actions. (Peacocke 2014: 31)²⁵

The description we get here is opposed to premise (2) in the above argument. Visual-spatial contents which might be stated, by a theorist, with indexical and demonstrative notions (*that thing is moving towards here*), and so whose non-conceptual contents we could, in dynamic-reflexive cases, frame as $\langle x \text{ moving towards } b \rangle$ are supposed *sufficient* for explaining the relevant actions.

How might we defend this? Let's consider an example. As the ball flies through the air towards where the perceiver is located, and the subject enjoys a visual experience, it is plausible that they will engage in avoidance behaviour, such that taking the content of their visual experience at face value provides them with a non-inferential reason for *moving*. The content of the intentional action could approximate to $\langle \text{move left to avoid } x \rangle$. The SLT will claim that the content of the intentional action is more accurately given as $\langle \text{move to the left of } me \text{ to avoid } x \rangle$, on the basis that the content of the visual experience is $\langle x \text{ moving toward } me \rangle$. However, the *b*-replacement account would contest this, claiming that the content of the visual experience is $\langle x \text{ moving towards } b \rangle$, and that the content

²⁵ See also Perry 1986: 220 and Campbell 1994: 121 who make the same claim.

of the intentional action could be <move left of *b* to avoid *x*>. What is it about the *b*-replacement account that is supposed to fall short of providing the relevant action-explanation in this case, where we can understand the content of the visual experience as providing a non-inferential reason for action (such that I act in the way I do because of what I can see)?

The planned intentional action is agreed by both accounts to be that of moving to a different location, such that the ball is no longer on a trajectory to the perceiver's location. Why should this necessitate any *self*-locating content in the visual experience? Both 'me' and '*b*' seem equally well placed to support the relevant intentional action, such as to provide a non-inferential reason for moving. This is because in the spatial context both 'me' (as the *self*-locating component) and '*b*' (as the *here*-locating component) are spatial indexes that serve to index the subject relative to a location, specifically where the perceiver is. In terms of supporting intentional action-explanations of this basic kind (i.e. avoidance behaviour) it is unclear what explanatory work 'me' does that '*b*' does not. Note, '*b*' differs markedly, given its indexical status, from 'someone' or 'somewhere'. Insofar as 'me', in the context of dynamic-reflexive visual experience, serves in part to index the spatial location of the subject of the visual experience, and in virtue of this supports the relevant intentional action explanations (and so perception-action links), it is unclear why '*b*' (as individuated in terms of its reference rule) doesn't also do this. Simply put: to preserve the perception-action links in the cases we are considering the relevant indexes, figuring in the content of the experiences, need to be determinate spatial indexes – the impersonal 'h' index can play this role just as well as a *de se* index.

The difference, of course, is that on the *de se* reading which the SLT offers we implicate self-representational capacities, such that we have to add that the subject also represents their location *as their location* (as where 'I am' or 'me' is). Undoubtedly, in reflexive cases which go beyond the basic spatial cases, *de se* indexes serve more complex functions and

support more complex action explanations which the *b*-index may not. For example, perhaps reflexive emotional experiences (such as pride, shame, and embarrassment), and the kinds of actions they support, are only possible given that the subject who enjoys them can *represent their location as their location*. However, such cases arguably go beyond the minimal representational capacities involved in spatial contents and their perception-action links, at least for the case of dynamic-reflexive visual experience.

So, pace premise (2) in the argument from action, the *b*-replacement account preserves the relevant perception-action links, such that given the content of the visual experience is $\langle x \text{ moving towards } b \rangle$, subjects can, taking that content at face value, have non-inferential reasons for the relevant actions. Put otherwise, it is because the subject's visual experience has a representational content $\langle x \text{ moving towards } b \rangle$, and the subject takes that content at face value, that we can understand the avoidance behaviour as a rational response to *what they see*.

At this point, we are in a position to appeal to considerations of parsimony. The *b*-replacement account can explain the perception-action links just as well as the SLT, and so in this context, it is not clear that the *de se* reading explains something the *b*-replacement reading does not. Positing a self-locating representation of 'me' in the content of dynamic-reflexive visual experience is surplus to requirements. Moreover, given that the *de se* reading implicates complex self-representational capacities in what are basic visual-spatial scenarios, the *b*-replacement account should be preferred as an account of the contents of those experiences.²⁶

²⁶ Does the content of static visual experience, say of an object to the left also include the *b*-index? While more needs to be said on this, my sense is that it does not. In the case of *non-reflexive* static perception, the *b*-index is likely implicit (hence why the Campbell-Perry strategy of monadizing the contents gets a foothold in those cases), and so not part of the representational content of those visual experiences. The difference in dynamic-reflexive visual experience is that the *b*-index is explicit, and so is part of the content.

4. Visual Kinesthesia

This paper has focused on providing an alternative to a self-locating interpretation of the content of a particular class of visual experiences. In those cases, we were dealing with visual experiences of environmental movement, albeit a specific class therein, namely dynamic-reflexive visual experiences. This final section considers an extension of the *h*-replacement account by outlining how it deals with cases of putative *self-movement* exemplified by visual kinesthesia. Such cases are discussed by J. J. Gibson, and more recently by Schwenkler, as instances in which we need to posit self-locating content.

First, let me highlight the following passage from Gibson that brings the phenomena into focus:

Vision is kinaesthetic in that it registers movements of the body just as much as does the muscle-joint-skin system and inner ear system. Vision picks up both movements of the whole body relative to the ground and movement of a member of the body relative to the whole. Visual kinesthesia goes along with muscular kinesthesia. The doctrine that vision is exteroceptive, that it obtains “external” information only, is simply false. Vision obtains information about both the environment and the self. (Gibson, 1979: 183).

Framing this in terms of the SLT, Schwenkler glosses this as follows: ‘surely it is impossible visually to experience oneself as in motion without being able visually to experience one’s spatial location’.²⁷

²⁷ Schwenkler 2014: 142 (see also Bermudez 1998: 110-114).

To make the case for the SLT in this context, Schwenkler introduces a specific example. He appeals to the visual illusion of purported self-movement. The illusion goes as follows. A subject placed in a rotating drum (with the drum rotating clockwise) first has a veridical visual experience of the drum rotating (experience V). After a little while, it no longer seems to the subject that the drum is rotating clockwise, but that the subject themselves are rotating, and the drum is stationary, having stopped rotating (experience I). Yet this latter experience is an illusion; the subject is still stationary, and in fact, the drum has continued to rotate clockwise throughout. So, we have a visual illusion of movement.

I agree with Schwenkler that the Minimal View, encapsulated by the monadizing Campbell-Perry strategy considered at the start of section 2, will not easily be able to deal with this case on the model of replacing relational egocentric notions with monadic egocentric notions. Moreover, this is, at least in part, because we are dealing – akin in certain respects to dynamic-reflexive visual experiences – with experiences whose objects are *non-static*. As Schwenkler puts it, ‘they are ones in which the spatial content of visual experience is changing, and this aspect of change over time has to be accounted for in any adequate visual phenomenology’.²⁸ It is not clear how the Campbell-Perry strategy can do this.²⁹ In what follows, I consider how the *b*-replacement account can deal with such cases.

First, let’s note that we need an account of the content of V and I that respects their different phenomenologies, such that I includes an *apparent change* that is absent from V , which Schwenkler calls ‘a *new* aspect of apparent motion’, which was previously not present in V . What is required is capturing the distinctiveness of this new aspect of apparent motion, which (in part) distinguishes V from I . The SLT is arguably able to do so since it can claim that the content in V approximates to <drum is spinning, I am stationary>,

²⁸ Schwenkler 2014: 142.

²⁹ See Ibid: 142-6 for a detailed argument for this claim.

whereas in *I* the content approximates to <I am spinning, drum is stationary>. The *new aspect of apparent motion* in *I* is (purportedly) captured by the fact that the experience *represents* that I am spinning, and in so doing represents my spatial location *as my location*. So, consideration of this case seems to vindicate the SLT.

Before explaining how the *b*-replacement account can deal with these cases, I briefly reflect on features of the visual phenomenology of movement. Consider the following case. In the early 1990s, a video game console called the 3DO released a reasonably basic driving simulator called *Need for Speed*. In one of the in-game settings the player could choose to ‘change perspective’ to be inside the car while driving, so mimicking the first person visual experience of driving (as opposed to the typical third-person view of the car ‘from outside’ on the road). Although no body parts were visually represented a left input on the controller did alter the visual position of the steering wheel, which visually rotated toward the left. Occasionally, when driving on straight roads in this first-person mode, it appeared as if the car was stationary and the environment was moving. There was undoubtedly a phenomenological difference, in terms of what one saw on the screen, between the visual experience of the car moving through a stationary environment (the standard case) and the environment moving past what seemed like a stationary car (the odd case).

The import of this example (and similar ones) is to draw out a feature of illusions of vection which involve a switch between environmental movement and purported self-movement. First, let us note that in the case where the environment is visually experienced as moving this is not merely a case of dynamic visual experience as defined in section 2. In both the first veridical experience of the drum rotating clockwise, and in the ‘odd case’ from the gaming example, we do not just have a case of objects moving in what is experienced as a *static* environment (e.g. the aeroplane as moving from left to right across the horizon). Instead, what we have is *wholesale movement of one’s visual environment or visual*

world (i.e. the other objects in one's visual surroundings, which provides the frame of reference against which putative 'self-motion' is experienced). In the 'odd case' in the gaming example, it is everything external to the car, which seems to be moving past as the car is stationary. Likewise, in V in the barrel illusion, the entirety of one's 'visual world' is taken up by the rotating drum, and so it will (at first) seem as if the totality of one's 'visual world' is rotating.

This is important because it suggests that the phenomenology of *movement of one's visual world* (MVW hereafter), as the way we experience MVW, is as relative to the location of a (relatively) stationary object. This point is especially apparent in the case of MVW which takes the form of *spinning* (as in the barrel illusion). To enjoy a visual experience as of one's entire visual world as *spinning*, it is typically the case that one experiences that *spinning* relative to a stationary location, as *spinning* relative to something that is *not* experienced as in vection. This is brought out if one considers how different V is from a visual experience as of an external object, say a figure located in a stationary visual world, to appear as spinning (i.e. a visual experience with an approximate content <x is spinning over there>). Consider the difference between a visual experience of a ballerina spinning against a static background, and a visual experience as of the entire room spinning around one. So, we have good reason to think that MVW, when visually experienced, is necessarily relational, such that it requires *two terms* in the content of visual experiences, and that one of the terms will pick out a (relatively) stationary object.

In what follows, I explain how the analysis of the last couple of paragraphs allows us to formulate a *b*-replacement account for the cases under discussion. Let us first note that the general form of the content for the experiences we are considering should be <(A) spinning relative to (B)>. How, more precisely might we approximate the content of V ? In V , the (A) term needs to be the entire visual world, call this VW. Moreover, in V , the (B) term is the experienced *stationary position at which the perceiving subject is located* (the more

or less extensive area centred on the subject). So, it seems V is amenable to a b -replacement; we get the content of V approximating to $\langle VW$ spinning relative to $b \rangle$. When the illusion kicks in what we get is, as Schwenkler notes, a new aspect of apparent motion, which is present in I but not in V .³⁰ Crucially, the b -replacement account can accommodate this since it would approximate the content of I as follows: $\langle b$ spinning relative to $VW \rangle$. What occupies the (A) and (B) positions in V has flipped in I . The entire visual world (VW), in this case the barrel and its lines, is now stationary, and it is the location of the perceiving subject that seems to be spinning relative to the barrel. If this is plausible, then the b -replacement account has a ready-made explanation of the phenomenological contrast between V and I , since *what is experienced as spinning and what it as rest* has changed.³¹ So, if the challenge is to explain how the experience of visual kinesthesia, in case I , differs from the corresponding experience of a moving world, in case V , then the b -replacement account has the tools to do this.³²

³⁰ Note, this can't be presumed to be represented as *self*-movement since that would beg the question in favour of the SLT.

³¹ One objection to the account is as follows: locations or regions of space, which are the referents of 'here', do not move. As such, *a location itself* cannot spin. One response is to draw a *sense-reference* distinction in the following way. In V and I it is indeed the case that ' b ' refers to something, namely the spatial location at which the perceiver is located, which *does not in fact spin*, but in I there is undoubtedly an experience of b spinning (which is illusory). So, b can *seem like* it is spinning even if the referents of b cannot spin – put in technical language, a possible mode of presentation (a *sense*) for b can be one of spinning. Note also, a veridical experience or an illusion of *spinning on the same fixed spot* (as in the barrel illusion) should not be equated with a veridical experience or an illusion of *moving through* one's environment. While both are cases of visual kinesthesia, they are different in important respects, although more need to be said about this (see discussion in text).

³² Note it has the tools to this do in the way that, as Schwenkler (2014: 150) rightly highlights, the impersonal 'somewhere' or 'someone' replacements do not. I also agree with Schwenkler (fn.15 154), that a simple

Extending beyond this specific case, we can again turn to Peacocke who says the following concerning ‘Degree 0’ creatures:

If Degree 0 cases exist, then the sort of sensitivity to changing relations as one moves (or is moved) in the spatial world may be a sensitivity that does not require that one represent what is, in fact, one’s current location as one’s current location. The location may be represented simply as *here*. Not even a nonconceptual first person is required for representation of an objective spatial world (Peacocke 2014: 33).

The central move of the *b*-replacement account is to deal with cases of visual kinesthesia in precisely this way. The most typical content for an ordinary case of visual kinesthesia, say as one walks down the street or jogs around a park, will be $\langle b \text{ is moving through VW} \rangle$, where (remember) what occupies the (B) position will be experienced as stationary. Note, the move, as expressed by Peacocke and the *b*-placement account, is not to attribute movement in such cases to merely an arbitrary location or region of space *per se*, but rather to attribute motion to the perceiving subject as what in fact *occupies the relevant location* at any given moment or current time, and which is simply represented as *b*, without additionally representing where *b* is, one’s current location, *as one’s own current location*.

Undoubtedly more needs to be said about the way the *b*-replacement account has been extended to cover cases of visual kinesthesia. What has been offered here is not a conclusive argument in favour of its specifications of the relevant contents, but a plausible outline. And since, at least with respect to such cases, the SLT is presented as the only credible option (emerging from the elimination of implausible alternatives), then insofar

content of the form ‘here is moving’ in the case of an illusion of vection (or more broadly for cases of visual kinesthesia) is inadequate.

as the *b*-replacement is a live option then the SLT has more work to do in arguing for its interpretation.

Conclusion

This paper has argued that with respect to a specific class of visual experiences which seem amenable to the SLT, namely dynamic-reflexive visual experiences, an alternative can be offered by way of the *b*-replacement account. This account was also extended to deal with visual kinesthesia. While a full defence of the *b*-replacement account was not possible here, it is a promising alternative to the SLT in dealing with such cases of visual experience (further research needs to be conducted to see how it could be extended beyond the visual case, such as to auditory or olfactory experiences). It is certainly an alternative with which defenders of the SLT should engage.

References

- Bermúdez, J. (1998) *The Paradox of Self-Consciousness*. Cambridge, MA: MIT Press.
- (2002) ‘The Sources of Self-Consciousness’ in *Proceedings of the Aristotelian Society*, 102: 87–107.
- Campbell, J. (1994) *Past, Space, and Self*. Cambridge, MA: MIT Press.
- (2002) *Reference and Consciousness*. Oxford: Clarendon Press.
- Cassam, Q. (1994) *Self and World*. Oxford: Oxford University Press.
- Evans, G. (1982) *The Varieties of Reference*, ed. J. McDowell. Oxford: Clarendon Press.
- Gibson, J. J. (1979) *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin.
- Hurley, S. (1998) ‘Nonconceptual Self-Consciousness and Agency: Perspective and Access’ in *Communication and Cognition: An Interdisciplinary Quarterly Journal* 30 (3-4): 207-47.

- Husserl, E. (1989) *Ideas Pertaining to a Pure Phenomenology and to a Phenomenological Philosophy, Second Book: Studies in the Phenomenology of Constitution*, trans. R. Rojcewicz and A. Schuwer. Dordrecht: Kluwer Academic Publisher.
- Kelly, S. D. (2001a) 'Demonstrative concepts and experience' in *The Philosophical Review*, 110 (3): 397-420.
- McDowell, J. (1994) *Mind and World*. Cambridge: Harvard University Press.
- (2009) *Having the World in View: Essays on Kant, Hegel and Sellars*. Cambridge: Harvard University Press.
- Musholt, K. (2015) *Thinking about Oneself: From Nonconceptual Content to the Concept of a Self*, Cambridge, MA: MIT Press.
- Peacocke, C. (1992) *A Study of Concepts*. Cambridge MA: MIT Press.
- (2000) *Being Known*. Oxford: Clarendon Press.
- (2001) 'Does perception have a nonconceptual content?' in *Journal of Philosophy* 98: 239-64.
- (2014) *The Mirror of the World: Subjects, Consciousness, & Self-Consciousness*. Oxford: Oxford University Press.
- Perry, J. (1986) 'Thought Without Representation', in *Supplementary Proceedings of the Aristotelian Society*, 60: 137–152.
- Sartre, J. P. (1937) *The Transcendence of the Ego*, trans by Andrew Brown. London: Routledge.
- Schwenkler, J. (2014) 'Vision, Self-Location, and the Phenomenology of the Point of View' in *Nous* 48 (1):137-155.
- Shoemaker, S. (1968) 'Self-Reference and Self-Awareness' in *Journal of Philosophy* 65: 555-67.
- (1996) *The First-Person Perspective and Other Essays*. Cambridge: Cambridge University Press.
- Wittgenstein, L. (1921) *Tractatus Logico-Philosophicus*, translated by D.F. Pears and B.F. McGuinness, London: Routledge & Kegan Paul, 1961.
- (1958) *The Blue and Brown Books*. Oxford: Blackwell.