What do We Say When We Say How or What We Feel?

Berit Brogaard

February 7, 2009

"Like magic, she felt him getting nearer, felt it like a pull in the pit of her stomach. It felt like hunger but deeper, heavier. Like the best kind of expectation. Ice cream expectation. Chocolate expectation."
— Sarah Addison Allen (*The Sugar Queen*)

Contents

0. Introduction

Discourse containing the verb 'feel', almost without exception, purports to convey information about mental states. 'John feels that punishment is a good deterrent' is about the content of John's beliefs, 'I felt so certain' is about the epistemic properties of the speaker's beliefs, 'The entrance is so white that it feels as if you're walking into a huge iPod' is about the phenomenal properties of a typical visual experience of a certain entrance, 'Amy felt a sharp pain in her chest' is about Amy's pain sensations, and 'Kent felt the softness of the satin' is about Kent's tactile experiences.

Though it is evident that discourse containing the word 'feel' purports to convey information about mental states, the question remains as to what exactly is conveyed. Does discourse containing the word 'feel' actually succeed in describing the content and phenomenology of inner experience? If so, how does it reflect the phenomenology and content of the mental states it describes?

The aim of this paper is to answer both of these questions. I will begin by offering a linguistic analysis of discourse containing the word 'feel'. Following Roderick Chisholm's (1957) analogous distinction for 'look', I will argue that 'feel' can be used either as an epistemic verb or a perceptual verb. When it is used as an epistemic verb, it is (roughly) synonymous with 'believe'. When it is used as a perceptual verb, it functions either comparatively or non-comparatively. Comparative 'feels' reports, I argue, purport to describe aspects of perceptions, including bodily perceptions (proprioceptions). They have a distinctly comparative underlying syntax (e.g. 'Mary felt like she was on a roller roaster'). Non-comparative 'feels' reports, too, purport to describe aspects of perceptions. But they do not have a distinctly comparative syntax. There are three kinds of non-comparative 'feels' reports: those that result from combining 'feel' with an adjectival phrase (e.g. 'Mary feels sad'), those that result from combining 'feel' with an unsupported clause or a 'that'-clause (e.g. 'John felt Mary touch his leg', 'John felt that his heart was racing') and those that result from combining 'feel' with a noun phrase (e.g. 'Tom felt the softness of the satin').

Reflections on discourse about how and what we feel, I argue, yield some insight into the structure of emotions, tactile experiences and bodily sensations. Discourse about bodily

sensation purports to convey information about sensations *as of* the body behaving in certain ways. Discourse about emotions and tactile experiences purport to convey information about (i) perceptions of an external stimulus having certain emotive properties, (ii) perceptions of the body (or mind) behaving in certain ways, or (iii) perceptions of the body (or mind) responding to external stimuli. So, on the assumption that perceptual reports sometimes *correctly* describe what they purport to be about, bodily sensations are: (i) perceptions of the body behaving in certain ways, and emotions and tactile experiences are perceptions of an external stimulus having certain affective properties, (ii) perceptions of the body (or mind) behaving in certain ways, or (iii) perceptions of the body (or mind) responding to external stimuli.

It is sometimes said that an adequate account of perception cannot be given in terms of an analysis of ordinary language. I will address this concern at the end. I think that to the extent that ordinary language at least sometimes correctly describes bodily experience, ordinary language can shed light on its phenomenology and contents.

1. On the Semantics of 'Look'

Reports of feelings are linguistically similar to reports of visual appearances, though there are, as we will see, also important differences. Chisholm (1957) noticed that 'look' and other appear words can function either as an epistemic verb or a perceptual verb. When it functions as an epistemic verb, 'it looks as if/like' means, roughly, 'my total subjective evidence indicates that'. For example, if I look at a proof of Fermat's last theorem and say 'It looks like Fermat was right', 'look' is used as an epistemic verb. I am conveying that my total subjective evidence indicates that Fermat was right.

When 'look' functions as a perceptual verb, it has, as we will see, a distinct phenomenal meaning. There are two kinds of perceptual 'looks' reports: comparative and non-comparative (non-epistemic). I shall follow Chisholm in taking 'non-comparative' to be short-hand for 'non-comparative non-epistemic'. Comparative 'looks' reports tells us that two visual experiences share certain phenomenal properties but they need not tell us what the properties are. For

example, if I notice that Anna is visually similar to her sister, I may say 'Anna looks like her sister'. I have then said that a typical visual experience of Anna and a typical visual experience of her sister have certain phenomenal properties in common but I have not said what those properties are.

Comparative 'looks' reports differ from epistemic 'looks' reports. Comparative 'looks' reports always purport to describe aspects of visual experience, so they are true perceptual reports. Further, unlike epistemic 'looks' reports, comparative 'looks' reports can fail to be evidence-bearing, witness 'It looks like a cow but it's really a horse'.

Comparative 'looks' reports have a distinctly comparative structure. Hence, they are structurally related to more familiar comparative sentences. Here are some examples of comparative sentences:

(1)

(a) John is taller than every girl is.

(b) Peter writes like his brother

(c) Sonya eats like a bird

(d) Ingrid acts like a lady.

(e) Gertrud spams Facebook like a pro

A lot has been written on the semantics of comparative sentences. But on one influential account due to Irene Heim (2006), comparative sentences are to be analyzed as containing semantically vacuous 'wh'-items in the sentence structure. 1(a) can be read as 'John is taller than every girl is wh'. To a first approximation, 'every girl is wh' is to be read as: 'every girl x: x

is *this* tall'. This item scopes out of the comparative clause, and the 'wh'-item raises to a wide-scope position. Hence, 1(a) is of the following form:

$[wh_1[\text{every girl is } t_1]]_2$ [John is taller than t_2]

The truth-condition for 1(a) thus is: for every girl x there is a height y such that John's height is greater than y.

Likewise, 2(b) can be read as: 'Peter writes like his brother does wh', where 'Peter's brother does wh' is to be read, roughly, as: 'Peter's brother writes *this* way'. This item scopes out of the comparative clause. So, 2(b) has the following underlying structure: '$[wh_1[\text{Peter's brother writes } t_1]]_2$ [Peter writes t_2]'. We can assign the following truth-condition to 2(b): for some way w such that w is a way that Peter's brother writes, Peter writes that way too.

2(c) is to be read as: 'Sonya eats like a bird does wh', where 'a bird does wh' is to be read as: 'most birds x: x eats *this* much'. This item scopes out of the comparative clause. So, 2(c) has the following structure: '$[wh_1[\text{a bird eats } t_1]]_2$ [Mary eats t_2]'. We can assign the following truth-condition to 2(c): there is an amount x such that most birds eat x and Mary eats x too.

If we suppose that putatively comparative 'looks'-reports are truly comparative, which their grammar suggests, then it makes sense to apply Heim's analysis of comparative sentences to them too. The natural way to do so is as follows. Following Heim, ' X looks like Y' contains the implicit wh-clause 'wh_1[Y looks t_1]'. This item scopes out of the comparative clause. This then yields the underlying structure: '$[wh_1[\text{Y looks } t_1]]_2$ [X looks t_2]'. For example, 'Anna looks like her sister' is to be read as containing the implicit clause wh-clause: 'wh_1[Anna's sister looks t_1]'. This item scopes out of the comparative clause. So, the sentence has the following underlying structure: '$[wh_1[\text{Anna's sister looks } t_1]]_2$ [Anna looks t_2]'. We can assign the

following truth-condition to this structure: there is an x such that x is how Anna's sister looks, and Anna looks x.

Non-comparative 'looks' reports purport to describe the properties of visual experience directly (rather than comparatively). For example, if I look at my red chair and say 'the chair looks red', what I said is plausibly a non-comparative report. Some apparently non-comparative 'looks' reports are implicitly comparative. For example, 'John looks drunk' is plausibly a contraction of 'John looks like someone who is drunk'. Drunkenness, like hunger, is not a directly visually observable feature but can only be observed indirectly via, for example, visible, auditory, and olfactory evidence (e.g. slurred speech, loud voice, uncontrollable laughing, failure to walk straight, smelly breath, etc). The comparative report gets this right. If John looks drunk, then a visual experience of John shares certain phenomenal features in common with visual experiences of people who are drunk (e.g. a failure to walk straight, unfocused gaze, etc).

It may be argued that all apparent non-comparative reports are implicitly comparative. For example, 'The chair looks red' is plausibly a contraction of 'The chair looks like an object that is red'. The latter sentence could be true if the chair caused experiences as of greenness.

However, there are three reasons to resist this sort of pessimism about non-comparative 'looks' reports. First, as we have seen, comparative 'looks' reports plausibly just are existentially quantified non-comparative 'looks' reports.

Second, even if we were wrong about the linguistic analysis of comparative 'looks' reports, it is evident that we cannot successfully reduce all non-comparative 'looks' reports to comparative reports. 'That chair looks red' and 'that chair looks the way a red object looks' plausibly have the same truth-conditions. In this world, they are plausibly true just in case the demonstrated chair looks purely qualitatively red. Likewise, 'That chair looks purely qualitatively red' and 'That chair looks the way a purely qualitatively red object would look' plausibly have the same truth-conditions. However, the latter two reports are true just in case the demonstrated chair looks purely qualitatively red. Hence, the comparative report

presupposes a non-comparative use of 'looks'. So, not all non-comparative 'looks' reports reduce to comparative reports.

Third, as Chisholm (1957: 51) points out, if 'look red' is given a comparative reading, 'red things look red in daylight' is an analytic truth. It says 'things that are red look the way things that are red look', which is trivially true. If, on the other hand, 'look red' is given a non-comparative reading, then 'red things look red in daylight' is a synthetic truth. Even before she started studying neuroscience and physics Frank Jackson's Mary knew that in daylight red things look the way red things look. But she didn't know that red things looked non-comparatively red. So, 'red things look red' has two readings, one of which is a non-comparative (non-epistemic) reading.

It may be suggested that so-called non-comparative (non-epistemic) reports are really epistemic reports. For example, one might suggest that 'x looks red' can be analyzed as 'x is inducing in me the belief that x is red (Armstrong 1961). But it's implausible to think that all uses of 'x looks red' are epistemic. Epistemic reports are evidence-bearing. However, there is nothing incoherent in saying 'x looks red but it really is white'. But in this case 'x looks red' is not evidence-bearing and hence the sentence is not used epistemically (for discussion see also Maund 1986).[1]

Frank Jackson (1977) calls the non-comparative sense of 'look' 'the phenomenal sense'. However, as Alex Byrne (2009) points out, this terminology is somewhat misleading. If 'look' has a phenomenal sense just in case the sentence in which it occurs purports to describe the phenomenal properties of an experience, then both comparative and non-comparative 'looks' reports are phenomenal. The difference between comparative and non-comparative 'looks' reports is that comparative reports needn't make explicit mention of the phenomenological properties of the experience described.

[1] Byrne (1999) and Glüer (1999) take perceptual experiences to be beliefs. So, on their views, phenomenal 'looks' reports reflect the contents of beliefs. However, even on their views, it is clearly 'look' in the non-epistemic sense that figures in perceptual reports reporting the contents of perceptual experiences.

Elsewhere I have argued that Chisholm's distinction between an epistemic and a non-epistemic sense of 'look' is a genuine semantic distinction, and that perceptual reports reflect the content of perception. As for the semantic distinction, 'look' is systematically ambiguous (or polysemous) between an epistemic and a non-epistemic sense. One test of ambiguity is the coordination test, which makes use of the fact that most ambiguous and polysemous verb phrases impose incompatible requirements on the extension of the subject or predicate, depending on which sense they are taken to express. For instance, 'expired' combines with 'John' but not 'driver's license' when it means 'died' and combines with 'driver's license' but not 'John' when it means 'became invalid'. So, 'John and his driver's license expired' is infelicitous. There is no property of expiring that reasonably can be predicated of both John and his driver's license.

'Look' passes the coordination test. 'Premise 2 looks suspect and purple' is infelicitous, which suggests that 'look' is ambiguous between an epistemic and a non-epistemic sense. It may also be thought that 'look' is ambiguous between a comparative and a non-comparative sense. 'My chair looks purely qualitatively red and like a sofa' is indeed infelicitous. However, if comparative 'looks' reports are ultimately analyzable in terms of non-comparative existentially quantified 'looks' reports, as I have argued above, then the ambiguity probably is not an ambiguity in 'look' itself but rather an ambiguity in the phrases 'looks [adjectival phrase]' and 'looks like x'. So, while 'looks' *reports* are three-fold ambiguous (or polysemous), 'look' itself only has two senses.

2. The Comparative Sense of 'Feel'

Reports of feelings are linguistically similar to reports of appearances. 'Feel' can function as an epistemic verb or a perceptual verb. When 'feel' functions as an epistemic verb, it means, roughly, 'believe' or ' it is x's opinion' as in 'I feel that education at the college level should be focused on improving the critical thinking skills of students'. When 'feel' is given an epistemic reading, the sentence in which it occurs is evidence bearing.

When ‘feel’ functions as a perceptual verb it can occur either comparatively or non-comparatively. Comparative ‘feels’ reports have a distinctly comparative underlying syntax and purport to describe aspects of perceptual experience (broadly construed to include sensory perception from mixed modalities and proprioception). They tell us that a perceptual experience of an object or event is phenomenally similarly to another experience. The modality of the perception need not be made explicit, and in some cases what is described is the input from several sense-modalities. For example, I may say ‘It feels as if you are ignoring me, though I know perfectly well that you are not’. This is true only if perceptions of your behavior share certain phenomenal properties in common with typical perceptions of the behavior of someone who is ignoring me. Some further examples:

(2)

(a) This house feels like my childhood home.

(b) This country feels old-fashioned

(c) It feels as if my friends don’t respect my religion

(d) The entrance is so white that it feels as if you're walking into a huge iPod

2(a) is true if I have, say, a visual, auditory and olfactory experience as of this house instantiating certain properties, and a typical experience of my childhood home would be an experience as of the home instantiating some of these properties. 2(b) is true if I have an experience as of the country’s economy, traditions, music or architecture having certain properties, and experiences of old-fashioned things are experiences as of these things instantiating some of these properties. 2(c) is true if I have, say, a visual and auditory experience as of my friends behaving in a certain way, and that way is similar in its sensory properties (e.g. visual and auditory properties) to the ways of people who don’t respect my

religion. 2(d) is true just in case an experience of the entrance is phenomenally similar to an experience of a huge iPod.

Typical uses of the sentences in (2) have roughly the same meaning as typical uses of the sentences that result from substituting 'look' for 'feel'. But when 'look' is substituted for 'feel' and the sentences are read comparatively, the experience described is visual rather than auditory, olfactory or tactile. If this house looks like my childhood home, then it shares visible features in common with my childhood home. If this house feels like my childhood home, then it could share visible features in common with my childhood home but the features could also be tactile, olfactory or auditory or some of them could be visual and some of them olfactory.

Some comparative reports purport to describe aspects of bodily experiences (construed broadly to include emotions, tactile experience and proprioceptive experiences). They tell us that someone's bodily experience is phenomenally similar to an experience of something else but they need not indicate in which way they are similar. Here are some examples:

(3)

(a) I feel like my heart is going to burst

(b) It feels as if someone is jerking needles into the root of my tooth.

(c) This fabric feels like velvet

(d) It feels like I have to throw up

In 3(a) an experience is said to have a phenomenology similar to an experience as of a heart bursting. In 3(b) an experience is said to have a phenomenology similar to an experience of needles being jerked into the root of one's tooth. In 3(c) an experience of this fabric is said to have a phenomenology similar to a tactile experience of velvet. Finally, in 3(d) an experience is

said to have a phenomenology similar to the bodily sensation one has when one is about to throw up.

Both kinds of comparative reports have a distinctly comparative structure. Like the corresponding 'looks' reports, they are best analyzed as kinds of existentially quantified comparative sentences. If we apply Heim's analysis of comparative sentences to comparative 'feels' constructions, 3(c) is to be read as containing the implicit clause 'wh_1[velvet feels t_1]'. Or in English: 'this way x: velvet feels x'. This clause scopes out of the comparative claim. So, the sentence has the structure '[wh_1[velvet feels t_1]]$_2$ [this fabric feels t_2]'. 3(c) can thus be assigned the truth-condition: there is an x such that x is how velvet feels, and this fabric feels x. Or: there is an x such x is a phenomenal property of a bodily experience of velvet, and a bodily experience of this fabric has x. Similarly: 3(d) can be assigned the truth-condition: there is an x such that x is a phenomenal property of a bodily experience of having to throw up, and my bodily experience has x.

3. The Non-Comparative Sense of 'Feel'

'Feel' also has a non-comparative sense, as in 'John felt sick', 'Linda felt Michael pinch her arm', and 'Charlotte felt an itch'. Non-comparative reports can be used to describe (among other things) aspects of emotion, tactile experience or bodily sensation.[2] For example, 'John feels cold' can be used to describe John's bodily sensation of coldness or someone's tactile perception of John's body. 'John felt his heart beat very quickly' can be used to describe John's bodily sensation or tactile perception of the event: *John's heart beats very quickly* (in fact, it can also be used to describe a very quick bodily sensation or tactile experience of John's heart

[2] Non-comparative 'feels' reports do not just describe aspects of emotions, tactile experiences or bodily sensations but can also be used to describe other kinds of perception, intention, and feelings of doubt and certainty, as in 'He felt the anger of the crowd', 'I feel inclined to do it' and 'I feel doubt'.

beat), and 'John felt his leg' can be used to describe John's bodily sensation or tactile experience of his leg.

Non-comparative 'feels' reports are not reducible to comparative 'feels' reports. The reason for this is analogous to the reason non-comparative 'looks' reports are not reducible to comparative 'looks' reports. First, as we have seen, comparative 'feels' reports plausibly just are existentially quantified non-comparative reports.

Second, suppose we attempt to reduce 'x feels purely qualitatively sharp' to 'x feels a way in which an object would feel if it had instantiated purely qualitative sharpness'. One would then have to specify how an object that instantiates purely qualitative sharpness would feel. But such an object would feel purely qualitatively sharp. Hence, the reduction is unsuccessful.

Third, some sentences are either synthetic or analytic depending on whether they are given a comparative or non-comparative reading. 'Soft satin feels soft in normal conditions' is a synthetic report when read comparatively but is an analytic report when read non-comparatively.

Of course, it is plausible that many apparently non-comparative reports of feelings are implicitly comparative. 'I feel hungry', for example, plausibly is a contraction of 'I feel like someone who is hungry' or 'I feel the way that someone who is hungry normally feels'. There is good reason to think that many superficially non-comparative 'feels' reports are implicitly comparative. Because visual experience to a large extent is transparent we have a fairly rich vocabulary for describing the phenomenology of visual experience. To have a phenomenally red experience arguably just is to have an experience that represents something as (purely qualitatively) red. Because tactile experience to a large extent is transparent we also have a fairly rich vocabulary for describing tactile experience. To have a phenomenally sharp experience arguably just is to have an experience that represents something as (purely qualitatively) sharp.

When we describe emotions and bodily sensations we often borrow vocabulary from our descriptions of tactile and visual experience (including experience of space), as in 'I have a sharp pain in my chest' and 'I feel empty'. But these initially figurative ways of speaking often become integrated into the language and become conventionalized. 'Sharp' plausibly is an example of that. 'Sharp', it seems, can now be an equally good characterization of a pain and a tactually experienced object. But many phenomenal properties of emotions and bodily sensations still go conventionally unnamed, hence the need for comparative reports. If sufficiently rich, a comparative report can be as or more informative than a non-comparative report. Compare 'I feel lethargic, confused and choked up, I feel like I am in a thick hampering haze'.

As for the case of 'look' it may be suspected that non-comparative 'feel' reports are epistemic reports. For example, one may attempt to give truth-conditions for 'x is sharp' as follows:

(4)

'x feels sharp' is true iff x would normally induce in me the belief that x is sharp

The truth of (4), however, does not show that non-comparative 'feels' reports are epistemic reports. The normal conditions cited in the analysans of (4) are presumably meant to exclude conditions in which I feel that x is sharp yet fail to believe that it is sharp owing to irrationality or philosophical belief, for example concerning which properties objects instantiate. But this suggests that when I don't believe that x is sharp, there is a way that x feels if (4) is true. So, (4) reduces to:

(5)

‘x feels sharp’ is true iff X feels to me in a way that would normally induce me to believe that x is sharp

But (5) entails that there is a way in which x feels non-comparatively and non-epistemically.

3.1. ‘Feels’ Reports with Adjectival Phrases

We can divide non-comparative ‘feels’ reports into three kinds based on grammatical structure. One kind is of the form ‘x feels [adjectival phrase]’. Here are some examples.

(6)

(a) It feels sharp

(b) My face feels itchy

(c) Rosa feels tired

As already mentioned, many ‘feels’ constructions with adjectival phrases may be implicitly comparative claims. If Alice is asked after her marathon how she felt when she arrived sober but naturally high at the finish line, she might say ‘I felt drunk’. In this case her bodily experience may not be best described as representing Alice as drunk. It might be better described as representing Alice as instantiating certain internally felt properties that are characteristic of a state of drunkenness. This, of course, is not to rule out the possibility that if Alice’s bodily experience represents Alice as instantiating certain internally felt properties that are characteristic of a state of drunkenness, then it mistakenly represents Alice as drunk. This would be a case of a known illusion. I think a case could be made for either possibility. My argument below does not turn on how we settle this issue.

3.2. ‘Feels’ Reports with Unsupported Clauses

A second type of non-comparative report is one in which ‘feel’ combines with a so-called unsupported clause, as in ‘John felt his leg move’. As James Higginbotham points out, unsupported clauses are clauses that exhibit ‘none of the internal inflectional structure of a full sentence or a clausal complement: neither tense, nor infinitival *to*, nor progressive *–ing*.’ (1983: 102). ‘Feels’ constructions with unsupported clauses seem superficially related to ‘feels’ constructions with noun-phrase complements. For example, ‘John felt his leg move’ seems superficially related to ‘John felt his leg’. However, ‘feel’ constructions with noun-phrase complements must be kept apart from ‘feel’-construction with unsupported clauses. ‘Feel’ constructions with unsupported clauses are structurally similar to more familiar constructions with unsupported clauses. Consider:

(7)

(a) John saw Mary cry
(b) We like carrots raw
(c) I consider John smart

‘Mary cry’, ‘carrots raw’ and ‘John smart’ are unsupported clauses. Higginbotham (1983) argues that 7(a)-(c) cannot be paraphrased using ‘that’-clauses, as in:

(8)

(a) John saw that Mary was crying
(b) We like (it) that carrots are raw
(c) I consider that John is smart

8(a)-(c) appear to mean something quite different from 7(a)-(c). For 8(a) to be true, John needn't observe the crying event but merely needs to see signs that indicate that Mary had just been crying. So, unlike 7(a), 8(a) could be true if John saw Mary after she had been crying. Unlike 7(b), 8(b) entails that carrots (in general) are raw. Finally, unlike 7(c), 8(c) entails that John is smart.

Higginbotham suggests that 'seeing' constructions with unsupported clauses might be analyzed as follows (using Barwise's situation-semantics and 7(a) as an example):

There is an s, and John saw s, and $s \in [[\text{Mary cry}]]^{M}$

This is to be read as follows: There is an s which John saw and s is in the extension of event type: *Mary cries*. One reason given in favor of the event analysis is that (i) 'seeing' constructions with unsupported clauses are referentially transparent, that is, they do not admit of opaque readings, and (ii) the event analysis predicts that this is so.

However, I think it is questionable that 'seeing' constructions with unsupported clauses are referentially transparent. Suppose for the sake of argument that 'to cry' and 'to shiver and shed tears because of sadness, rage, or pain' are (necessarily) co-extensive. Now, consider:

(9)

(a) I saw Mary cry
(b) I saw Mary shiver and shed tears because of sadness, rage or pain
(c) I saw Mary shed tears

If 'to cry' and 'to shiver and shed tears because of sadness, rage or pain' are co-referential (as we have supposed for the sake of argument), and 'seeing' constructions with unsupported clauses are referentially transparent, then 9(a) and 9(b) have the same truth-value. As 'Mary shivered and shed tears because of sadness, rage or pain' entails 'Mary shed tears', 9(c) should be true if 9(a) is. But 9(a) could be true, while 9(c) is false. Suppose John sees Mary from behind, she is shivering, has a handkerchief in her hand and is making crying sounds. In this scenario John might correctly use 9(a) to report what he saw, but if he had used 9(b) or 9(c), he would at best have said something misleading. John did not see Mary shed tears. So, it seems that 'seeing' constructions with unsupported clauses are not referentially transparent. I think the lesson is that while 'seeing' constructions with unsupported clauses are referentially transparent in the noun-phrase position of the unsupported clause, it is not in general referentially transparent in verb-phrase position.

As it turns out, however, the fact that 'seeing' constructions with unsupported clauses are not referentially transparent does not undermine the event analysis. For, the event analysis does not predict that 'seeing' constructions with unsupported clauses are referentially transparent. If the event analysis predicted that 'seeing' constructions are referentially transparent, it would predict that 8(a) entails 8(c). But the traditional analysis does not predict this. When we see an event, we needn't see all of it. For example, 'John saw the car accident' does not entail 'John saw every part of the car accident'.[3] We see complex events, and other high-level properties,[4] in part by visually detecting what I will call 'manifestation properties'. I propose the following rough characterization of what it is to be a manifestation property:

Manifestation properties

The manifestation properties of an event are a range of detectable properties which reliably indicate and are causally impinged upon by the event.

[3] For discussion of what it is to see occluded parts of an object or event, see e.g. Nanay (forthcoming).

[4] On an argument for high-level properties entering into the content of perception see Siegel (2005).

The crucial concepts of reliable indication and causal impingement can be cashed out as follows:

Reliable Indication

E reliably indicate P iff in most of the closest situations in which E is present, P is present.

Causal Impingement

P impinges on E1, E2, … En iff in the closest situations in which P is present, a significant subset of {E1, E2, … En} are present.

The reliable indication requirement ensures that manifestation properties are distinct indicators.[5] Because the look of a drinking-water event is not distinct, it is not a manifestation property. It is not the case that in most (close) situations in which the look of a drinking-water event is present, a drinking-water event is present. A drinking-a-liquid event could easily look like a drinking-water event without being a drinking-water event.

The causal impingement requirement ensures that manifestation properties are relatively counterfactually invariant. Because the look of a mother sleeping is not invariant, it is not a manifestation property. The look of a mother sleeping is not counterfactually invariant. It is not the case that in all (close) situations in which a mother is sleeping, she looks like a mother that is sleeping. A sleeping-mother event could easily occur without looking anything like a sleeping-mother event.

The upshot is this: John can see a crying event in virtue of seeing various visual manifestation properties of crying, for instance, a shivering body, a handkerchief, runny mascara, and so on. Hence, even if all crying events essentially involve shedding tears, it can be true that John saw someone cry even if he didn't see them shed any tears.

[5] The terms 'distinctness' and 'invariance' are from Stalnaker (1997).

If, however, referential opacity is perfectly consistent with the event analysis, then the issue of whether 'seeing' reports with unsupported clauses are referentially transparent or opaque is not a factor in determining whether the event analysis is correct.

And, as it turns out, there is no particular problem in interpreting 'seeing' reports with unsupported clauses as seeing-*that* constructions. 9(a) can be interpreted as equivalent to (10):

(10) John saw that Mary was crying

(10) has two readings. On one reading, the past tense of the 'that'-clause makes a difference to the meaning of the whole. So, on this reading, (10) is best interpreted as meaning 'John saw that (it was the case that Mary is crying)'. On the second reading, the past tense is vacuous and makes no difference to the meaning of the whole. So, on this reading, (10) is best interpreted as meaning 'John saw that (Mary is crying)'. The same sort of ambiguity is apparent in belief reports. For example, 'In 1995 Mary believed that Nixon was president' can be read as saying that Mary believed that Nixon was president at some point prior to 1995, or as saying that Mary believed that Nixon was president at the time of her belief. Like the 'seeing' constructions with unsupported clauses 'seeing-that' constructions are referentially transparent in noun-phrase proposition but not in verb-phrase position.

The lack of substitutivity in verb-phrase position in 'seeing' constructions with unsupported clauses does not suggest that the narrow content of the visual experience described represents a property expressed by the verb. The narrow content of an experience is a content that supervenes on intrinsic properties of the perceiver. So a perceiver and an intrinsic duplicate of the perceiver cannot have experiences with different narrow contents. When narrow content is construed in this way, it could be true to say that John saw Mary cry even if John does not have a visual experience that narrowly represents Mary as crying. If John had a visual experience as of Mary shivering and shedding tears, and Mary in fact is crying, then at least in some contexts it is true to say that John saw Mary cry.

Seeing reports are similar to belief reports in this respect. There plausibly are contexts in which it is true to say that Lois Lane believes that the boring office worker named 'Clark Kent' can fly, even though she would not assent to 'the boring office worker named "Clark Kent" can fly', perhaps context in which it doesn't matter too much how the content of Lane's belief is described as long as it is clear that it concerns a certain man. Likewise, there plausibly are contexts in which it is true to say that Lois Lane believes that Superman can fly, even if she would not assent to 'Superman can fly' but only to 'Superman can travel through the air'. In general for a belief report to be acceptable it is required that the wide content of the 'that'-clause strictly matches the wide content of the belief (e.g. by being identical to it or by picking out the same state of affairs in the world). The narrow content of the belief report, on the other hand, just has to loosely match the narrow content of the belief, which it does if it is contextually appropriate (see e.g. Richard 1990, Chalmers 2009).

'Seeing' reports with unsupported clauses and 'that'-clauses appear to inherit this feature from the corresponding belief reports. 'John saw Mary cry' is true only if the wide content of 'Mary cries' strictly matches the wide content of the visual experience described. But its narrow content need only loosely match the narrow content of the visual experience described. How close the match has to be is determined by context.

What we have just said about 'seeing' reports with unsupported clauses carries over to 'feels' reports with unsupported clauses. Consider:

(11)

(a) Tom felt Mary pinch his arm

(b) John felt his leg move

(c) John felt Mary's remarks make a great impact on him.

Like ‘seeing’ reports with unsupported clauses, ‘feels’ reports with unsupported clauses needn’t be referentially transparent. Suppose again that ‘Mary cried’ entails ‘Mary shed tears’. The traditional analysis, then, predicts that 12(a) entails 12(b):

(12)

(a) John felt Mary cry on his shoulder

(b) John felt Mary shed tears on his shoulder

To feel an event is to feel some of its bodily manifestation properties. The event of crying and the event of shedding tears have different manifestation properties. So, even if crying entails shedding tears, 12(a) can be true without 12(b) being true. For example, John may feel someone cry in virtue of feeling someone’s body shiver, but feeling someone’s body shiver does not suffice for feeling someone shed tears.

These observations are consistent with both the event analysis and the propositional analysis. On the event analysis, 12(a) is to be read as saying that there is an s which John felt and s is in the extension of event type: *Mary cries on John’s shoulder*. On the propositional analysis, 12(a) is to be read as saying: ‘John felt that (Mary is crying on his shoulder)’. John can feel an event by feeling some but not all of its manifestation properties. So, even if crying essentially involves shedding tears, John can feel a crying event without feeling the shedding of tears. Similarly, a ‘that’-clause can express a semantic content c and c can be an accurate description of a mental content m, even if c is not identical to m. So, it can be true that John felt that (x is crying on his shoulder), even if the mental content of John’s bodily sensation is not the content of ‘x is crying on my shoulder’.

3.3. ‘Feels’ Reports with Noun-Phrase Complements

A third kind of non-comparative 'feels' report is the 'feels' construction with a noun-phrase complement. Here are some examples:

(13)

(a) Tom felt a quickening of his heartbeat

(b) John felt the softness of the satin

(c) Mary felt a pain in her leg

'Feels' constructions with a noun-phrase complement are syntactically different from 'feels' reports with unsupported clauses. They result from combining 'feel' with a noun phrase rather than an unsupported clause. But the two kinds of construction are semantically similar. 's feels NP' is true in virtue of s undergoing a bodily perception of the bodily manifestation properties of an object or event. So, they do not require that the literal semantic content of the description of the perceived object be a mental content of the bodily perception. For example, one can feel a quickening of one's heartbeat by feeling the bodily manifestation properties of a quickening of one's heartbeat.

For this reason 'feels' constructions with unsupported clauses are not referentially transparent. Even if 'x cries' analytically entails 'x sheds tears', 'Tom felt a crying woman who was leaning against his shoulder' does not imply 'Tom felt the tears of a woman who was leaning against his shoulder'. Tom can feel a woman who is crying without feeling her tears.

3.4. Tests for Tactile Experience vs. Bodily Experience

Non-comparative 'feels' reports purport to describe (among other things) emotional experience, tactile experience and bodily sensation. One might wonder whether there is a way

to determine if a report describes bodily sensation (including emotional experience) or tactile experience. Some have argued that there is. Anthony Kenny (2003), for example, argues that one can identify what a 'feels' report of the form 's feels NP' describes by looking at its semantic implications. He divides 'feels' reports into three grammatical types:

> Type 1: 'Feel' may be followed by noun phrase: I feel you, the smoothness of the satin, sadness, an itch on my back.
>
> Type 2: 'Feel' may be followed by an adjective: I feel sad, itchy, feverish, hot.
>
> Type 3: 'Feel' may be followed by a 'that' clause or an 'accusative or infinite clause': I feel that the satin is smooth, that I am sad. I feel the satin to be smooth.

According to Kenny, when a 'feels' report of type 1 describes a (tactile) perception, it implies a proposition of type 3 but not type 2. When a 'feels' report of type 1 describes an emotion or a sensation it implies a proposition of type 2 but not type 3. For example, 'I feel the coldness' implies that I feel that the something is cold but it does not imply that I feel cold. 'I feel anger' implies that I feel angry but it does not imply that I feel that I am angry.

Though initially plausible, Kenny's test is not surefire. It is certainly often true that when 's feels NP' is used to describe an aspect of an emotion or bodily sensation, we can infer a type 2 proposition but not a type 3 proposition. If I feel anger, I feel angry but I need not feel that I am angry, in the epistemic sense of 'feel'. 'I feel angry' is perhaps best understood as a contraction of 'there is a way an angry person feels, and I feel that way', which does not entail that I judge that I am angry. I could feel angry but not know or believe that I am angry.

However, it is not quite right that, when a description of the form 's feels NP' is used to describe a tactile (or other) perception, we can infer a type 3 proposition. As noted above, 'feel' followed by a 'that'-clause can, on one reading, express an opinion or judgment (and hence not an aspect of a tactile experience). On another reading, it is equivalent to a 'feels'

construction with an unsupported clause. Though we can sometimes infer a 'feels' construction with an unsupported clause from 's feels NP', we cannot infer an epistemic 'feels' report. For example, from the fact that I feel a lump in the mattress it doesn't follow that I judge that there is a lump in the mattress. Joyce Trebilcot (1970) offers the following counterexample. I may feel the lump in the mattress but mistake the lump for a sock and judge that there is a sock under the sheet. Of course, if I feel an explosion, then I feel something explode and hence that something is exploding. But here 'feel' must be given a perceptual reading rather than an epistemic reading.

A further problem is this. While it is often true that when a type 1 report purports to describe a tactile experience, then it does not entail a type 2 proposition, this is not true in general. Suppose I run a fever and feel cold, I am shivering, yet when I touch my body it feels hot. In that case 'I feel the hotness of my skin' entails 'I feel hot' but the latter does not entail that I have a bodily sensation as of being hot.

Kenny's test can, however, be modified to serve as an approximate test of whether 'feels' constructions with noun-phrase complements purport to describe emotions/bodily sensations or tactile sensations as follows. When type 1 reports describe emotions or bodily sensations (as in 'John felt the slimness of his body'), they tend to entail a type 2 proposition (e.g. 'John felt slim'). When they describe tactile perceptions ('John felt the explosion'), they tend not to entail a type 2 proposition ('John felt the explosion' does not entail 'John felt explosive').

Kenny's test can also be employed as an approximate test for 'feels' constructions with adjectival phrases. If 'John feels cold' is entailed by a 'feels' construction with a noun-phrase complement, e.g. 'John feels the coldness of his body', then 'John feels cold' purports to describe a sensory experience. If it doesn't, then it purports to describe tactile experience.

There is no good test of whether a 'feels' construction with an unsupported clause purports to describe a tactile experience or a bodily/emotional experience. 'John felt his quick heart-beat' entails 'John felt his heart beat quickly', regardless of whether the latter report purports to describe a tactile or bodily experience. The same goes for 'John felt something

explode', which entails 'John felt the explosion', regardless of whether the latter is used to describe a tactile or bodily experience.

4. Is 'Feel' Ambiguous?

If there are three uses of 'feel', then one might suspect that 'feel' is ambiguous. Here we need to distinguish between unsystematic and systematic ambiguity, also known as lexical ambiguity and polysemy, respectively. If a string of letters is lexically ambiguous, then the fact that the same string of letters spells two different words is a linguistic coincidence. 'Bank' is a good example of this. If, on the other hand, a string of letters is polysemous, then it spells a single word with different but related meanings. A good example of this is 'healthy'. 'Healthy' as it occurs in 'a healthy person' and 'healthy' as it occurs in 'a healthy appetizer' have different but related meanings.

Polysemy can be explained semantically or pragmatically (see Bach 1999). A polysemous word is a semantically underspecified lexical entry. Because the lexical entry is underspecified, linguistic or extra-linguistic context is required to determine which proposition is conventionally conveyed by a sentence containing it. Whether the full proposition conventionally conveyed is best treated as semantically expressed or pragmatically conveyed by the sentence will depend on further theoretical assumptions. When a lexical entry is semantically underspecified but its completion is not conventionalized, we do not have a case of polysemy but a case of conversational implicature, conversational impliciture or contextual completion.

As an example of the latter phenomenon, consider the two comparative uses of 'feel': one is used in descriptions of over-all experience and the other is used in descriptions of bodily experience. But this phenomenon is not likely to be a case of polysemy. Whether a 'feels' report describes a bodily experience or an over-all experience is contextually determined. In this respect the comparative use is like other cases of unconventionalized completion. For

example, ‘I saw a man’ does not specify the time or location of the seeing or the features perceived. But the sentence is not systematically or unsystematically ambiguous.

The questions that remain are whether ‘feel’ is ambiguous between an epistemic and a perceptual sense, and whether ‘feel is ambiguous between a comparative and a non-comparative sense. There are various linguistic tests of ambiguity. One is the translation test. If a word is ambiguous, it will tend to translate into different words in other languages. The translation test doesn’t strongly differentiate between unstructured (or lexical) ambiguity and structured ambiguity (or polysemy). If it is a linguistic coincidence that a single string of letters spells two different words, then we should expect the term to translate into different words in most other languages, and if a word has two closely related and contextually determined meanings, we should expect the different meanings to be lexically manifest in at least some other languages.

We could distinguish between the unsystematic and the systematic translation tests. If a word is lexically ambiguous, we should expect the word to translate into different words in most other languages. If a word has different closely related contextual meanings, we should expect it to translate into different words in some other languages in which different lexical items are used for different meanings despite their closeness.

‘Feel’ passes the systematic translation test but not the unsystematic translation test. In many languages the translation of ‘feel’ can be used as widely as it can in English. The German word ‘fühlen’, for instance, has, roughly, the same range of uses of the English word ‘feel’. The fact that variation across languages is not wide-spread suggests that the ambiguity of ‘feel’ is not genuinely lexical but is a case of polysemy.

A second test of ambiguity is the coordination test discussed above. Most ambiguous and polysemous verb phrases impose incompatible requirements on the extension of the subject or object. We can use this observation to test for polysemy. If the different senses of a polysemous verb put different requirements on the subject or object, then it should be possible to construct an infelicitous sentence with a conjoined subject or object. To take our earlier example, ‘John and his driver’s license expired on Tuesday’ is infelicitous, because living things

die whereas legal documents cease to be valid. So, regardless of how 'expire' is interpreted, the sentence is infelicitous.

'Feel' passes the coordination test. 'I feel that the education in this country is first-rate and that my leg is itchy' is infelicitous, regardless of which sense is assigned to 'feel', which suggests that there is a genuine (systematic) ambiguity in 'feel' between its epistemic and perceptual meanings.

It may seem that 'feel' is also ambiguous between a comparative and a non-comparative sense. The occurrence of 'feel' in 'my leg feels soft' is most naturally read non-comparatively. The occurrence of 'feel' in 'my leg feels like somebody else's leg' must be read comparatively. So, if 'feels' reports are ambiguous between a comparative and a non-comparative sense, we should expect 'feels' reports containing the conjunctive phrase 'like somebody else's leg and soft' to be infelicitous.

This is indeed the result we get. 'My leg feels like somebody else's leg and soft' is infelicitous. However, if comparative 'feels' reports are existentially quantified non-comparative reports, then the ambiguity is not likely to be an ambiguity in the verb itself but is rather an ambiguity in the phrases 'feels like x' and 'feels [adjectival phrase]'. These phrases, then, are the proper bearers of ambiguity not 'feel'.

There is further evidence for this hypothesis. 'John felt something explode and an itch' is no less fortuitous than 'my leg feels like somebody else's leg and soft', which suggests that there is an ambiguity in 'feels [unsupported clause]' and 'feels [noun phrase]' but not that 'feels', when used comparatively, has several different senses.

5. Do Our Descriptions of How and What We Feel Reflect How and What We Feel?

Let us now turn to the question of whether 'feels' reports reflect the representational content of bodily experiences (including emotions). This question must be kept apart from the question of what 'feels' reports purport to describe or even succeed in describing. Someone who has

had an injury in his right foot and who now has an injury in his left foot might say 'My toe feels like it's about to explode'. Though this is a comparative report and hence needn't name a property that the speaker's toe actually has, it can be a successful description nonetheless. But intuitively, a sensation of pain in one's foot does not have as its representational content *the sensation of my toe is phenomenally similar to the sensation of a body part that is about to explode*. This latter describes the phenomenal properties of the pain sensation but not its representational content.

To answer the question of whether 'feels' reports reflect the representational content of bodily experiences, we must say a bit more about what we mean by 'reflect' and 'representational content'. I propose the following definitions of 'S reflects property P' and 'S reflects content p':

Phenomenal Property Reflection

A report that describes experience e reflects a phenomenal property P iff [necessarily, the report is true iff P is a phenomenal property of e]

Content Reflection

A report that describes experience e reflects a content p iff [necessarily, the report is true iff p is a content of e].

Phenomenal Property Reflection and Content Reflection are meant to be restricted to reports that are tokens of sentences that can have true tokens when uttered by us. Without this restriction, my report 'John, a human, looks exactly like an egg inside and out to me in good viewing conditions' reflects phenomenal redness. For, my report is necessarily false and it is necessarily false that my current visual experience instantiates phenomenal redness (I am

looking at a black and white screen). So, the right-hand side is true. So, it follows that my report 'John, a human looks exactly like an egg inside and out to me in good viewing conditions reflects phenomenal redness. We can avoid this counterexample by restricting the definition to reports that are tokens of sentences that can have true tokens when uttered by us. Sentences that do not have true tokens do not reflect any phenomenal properties.

How then should we understand 'the representational content of bodily experiences'? The content of a bodily experience is conveyable information, something which, at least in principle, could become the content of a belief or judgment grounded in the experience. If the content is representational, it must furthermore be what is represented by the experience.

'My toe feels as if it's about to explode' reflects a certain quality of the bodily sensation described, and that quality can become part of the content of a phenomenal belief directly about the sensation's phenomenology. But this content is not a *representational* content of the pain sensation. Though a phenomenal belief represents phenomenal properties of experience, an experience does not represent its own phenomenal properties.

However, 'my toe feels as if it's about to explode' doesn't just reflect the phenomenology of a pain sensation, it also reflects of a content represented by the pain sensation described. There are two sorts of phenomenal properties an experience might reflect: representational phenomenal properties and non-representational phenomenal properties. Representational properties are properties which experiences have just in case they represent a particular content, for instance, the property of representing an object as round. Non-representational phenomenal properties are properties of experiences that contribute to their phenomenology but which do not correspond to a representational content or a property in the representational content.

If we focus on the case of vision, there is good reason to think that imprecise experiences do not represent objects as imprecise. There is a correlation between representational phenomenal properties and phenomenal looks. Phenomenally red experiences represent objects as red and objects sometimes look red. Imprecise experiences do not represent objects as imprecise, and no object (phenomenally) looks imprecise. If we

take this correlation at face-value, we have evidence that 'looks' reports do not reflect non-representational phenomenal properties but reflect only representational phenomenal properties.

Like visual experience, bodily experiences have non-representational phenomenal properties. Unclear or imprecise bodily experiences do not represent objects as unclear or imprecise. My pain sensation might represent my chest as sore, and my chest might feel sore, but my pain sensation does not represent my chest as unclear, and my chest does not feel unclear. If we take this correlation between how things feel and what bodily experiences represent at face-value, then we have reason to think that 'feels' reports do not reflect non-representational phenomenal properties but reflect only representational phenomenal properties.

There is good reason to think, then, that if 'my toe feels like it's about to explode' reflects a phenomenal quality of my pain experience, then it reflects a representational phenomenal property of my pain experience. So, by reflecting the quality of the experience the report also reflects a representational content of the experience. It reflects the existentially quantified content *my toe has a salient property in common with body parts that are about to explode*.

It may seem that pain sensations do not have existentially quantified contents of the form 'my toe has a salient property in common with body parts that are about to explode'. However, if we construe representational contents of bodily experiences as conveyable information about what the experience represents, as suggested above, and conveyable information is information which, at least in principle, could come to constitute the content of a belief, then the content *my toe has a salient property in common with body parts that are about to explode* can indeed be *a* content of a pain sensation.

Unlike linguistic entities, non-linguistic bearers of content do not have content that can be compositionally determined from grammatical structure and the semantic content of lexical items. Some non-linguistic bearers of content are plausibly individuated by their content. Belief may be a case of that. It is plausible that *there is a dog in front of me* and *there is an*

animal in front of me cannot both be contents of a single belief. But perceptual experience is different from belief in this respect. If perceptual experience is individuated by its phenomenology rather than by its content, which is highly plausible, then there cannot be a single content that is *the* content of a perceptual experience. The phenomenology of perceptual experience is not fine-grained enough to narrow down the possible contents to a single one. For example, it is plausible that if one's phenomenology represents a dog, it also represents an animal. If this is so, then one cannot have a perceptual experience as of there being a dog in front of one without having a perceptual experience as of there being an animal in front of one.

But even if the phenomenology of perceptual experience could represent a dog without representing an animal, surely it cannot structurally represent in the way that language structurally represents. So, while the sentences 'Lisa gave the package to John' and 'John received the package from Lisa' can be construed as expressing different propositions, a perceptual experience as of Lisa giving a package to John just is a perceptual experience as of John receiving a package from Lisa. These sorts of considerations give us some reason for talking about the contents of perception [in the plural] rather than the content of perception [in the singular].

Of course, the report 'my toe feels like it's about to explode' need not reflect any representational content other than *my toe has a salient property in common with body parts that are about to explode*, but it may. I take *reflects* to be a transitive relation. So, if x reflects y, and y reflects z, then x reflects z.

Whether x reflects y will depend in part on the speaker's internally justified beliefs about the common causes of, constitutive elements of, or normal salient properties of x. So, what else the report 'my toe feels like it's about to explode' reflects will depend in part on the speaker internally justified beliefs about the common causes of, constitutive elements of, or normal salient properties of body parts that are about to explode. If it is known that a typical salient property of body parts that are about to explode is that of providing discomfort, then the report reflects that the experience represents a discomfort in the toe.

What I just said about pain sensations carries over to other descriptions of bodily experiences. For example, 'I feel sad', if accurate, reflects that I have the property of being sad. I may be completely ignorant of what sorts of sensations sad people tend to have. For me, then, 'I feel sad' reflects that the speaker being sad. If I am justified in believing that sad people often have a feeling of pressure in their chest, of being choked up, of tears coming on, of life being meaningless, and so on, then 'I feel sad' also reflects that the emotive state described represents the body as undergoing some of these changes.

6. Do Descriptions of Bodily Experiences Tell Us Anything about their Structure?

We can draw some interesting conclusions from the foregoing considerations about the structure of emotions, tactile experiences and bodily sensations. Suppose 'I feel sad' is an accurate report of an aspect of an emotional state, and that it reflects a representational content of that state. The report reflects the content that the speaker is sad. But being sad is likely to reflect that certain bodily and mental conditions are present, for instance, diminished interest in one's surroundings, decreased appetite, irritability, depressed mood, feeling choked up, feeling confused, feeling lethargic, and so on. 'I am sad' then reflects not only the representational content that the speaker is sad but also the content that speaker is in some of the mental and bodily states characteristic of sadness.

But emotional reports need not just describe bodily sensations. They often describe external objects and state of affairs as well (see e.g. Solomon 2004). For example, we often say things like:

(13)

(a) John feared the angry dog

(b) Mary was angry at John

(c) Tom was happy that Amy was back

How are reports like those in (13) related to 'feels' reports? It might be thought that the reports in (13) reflect a different kind of representational content of emotional experiences: what is sometimes called a 'formal content'. For example, one could argue that my fear has the formal content that the dog is a threat to my well-being. 13(a) then might be thought to reflect this content. 'John feels his heart race' and 'John feels afraid', then, would reflect one kind of content of John's state of fear, and 'John fears the angry dog' would reflect a different kind of content of his state of fear. This sort of view is consistent with a conjunctive theory of emotions, according to which emotions consists in bodily sensations and a perception or judgment that attributes an emotive property to an external object.

However, the conjunctive theory cannot be quite right. First, 13(a) could be false even if John has a sensation of a bodily reaction typically associated with fear, and he perceives the dog as a threat to his well-being. For example, John could perceive Jill's anger (while they are being attacked by a mad dog) as causing his bodily reaction as of fear but judge only the dog to be a threat to his well-being. In that case John judges that the dog is a threat to his well-being, and he has a sensation of a bodily reaction typically associated with fear. But he does not fear the dog. He fears Jill's anger. So the conjunctive theory makes the wrong predictions here.

Second, 13(a) can be true even if John's fear does not attribute the property of being a threat to his well-being to anything. John might just feel that *that* thing over there causes his body to shiver, his heart to race, and the hairs on his arms to rise.

What is required for the truth of 13(a), it seems, is that John *feels* the angry dog cause or give rise to his fear response. So, 13(a-c) could be true partially in virtue of the following sentences being true:

(14)

(a) John felt the dog cause his body to shiver, his heart to race, and the hairs on his arms to rise.

(b) Mary felt John cause her face to flush, her fists to tighten and her jaw to clench.

(c) Tom felt Amy's being back cause a tickle in his tummy, a rush of blood through his body and a smile on his lips.

If 13(a-c) are true partially in virtue of the truth of 14(a-c), then at least some of the content of emotions can be accurately reported using a 'feel' construction with an unsupported clause (or a 'that'-clause). Emotional experiences thus have causal contents of the form 'x causes r', where x is an object, event or state of affairs and r is a set of bodily reactions. As emotions involve perceptions of external objects, perceptions of bodily reactions and perceptions of causal processes, they are partially sensory (e.g. visual) and partially proprioceptive.

The sentences in (14) are perhaps a bit stilted. However, there are other more natural 'feels' reports with unsupported clauses which purport to describe aspects of emotional experiences. Here are a few examples:

(15)

(a) She felt him getting nearer, felt it like a pull in the pit of her stomach

(b) Mary felt his grip make her breathing ragged and shallow.

(c) She felt the dark shadow make her heart beat faster.

The account of the structure of emotions just presented suggests a unified theory of objectual and propositional emotions. It is common to draw a distinction between objectual and non-

objectual emotions. If John fears the dog, the dog is an object of his fear. If John fears that the dog will kill him, the proposition that the dog will kill John is an object of his fear.

On the present account, the difference between objectual and propositional emotions just is a difference in whether John perceives the dog or the dog's being able to kill him as a cause of his bodily reactions. Of course, one does not rule out the other. John could perceive the dog as causing a fear response in virtue of perceiving the dog's being able to kill him as causing a fear response.

The unified account of objectual and propositional emotions avoids a well-known problem for objectual attitudes. It is possible that one can fear Satan, admire Sherlock Holmes, and love Odin. Theories that take these attitudes to be attitudinal relations to objects are committed to objects that do not exist. On the present account, to fear Satan is to experience Satan give rise to feelings of fear, to admire Sherlock Holmes is to have experience Sherlock Holmes give rise to feelings of admiration, and to love Odin is to experience Buddha give rise to feelings of worship. These experiences are non-veridical, as the relevant objects do not exist, but they are genuine emotional experiences nonetheless.

In short: There is good reason to think that at least some emotional reports reflect a causal content of the emotion described. On the assumption that some of these reports reflect representational contents of emotional states, emotions have causal contents.[6]

The derived causal account of emotional experience may seem problematic. It may be argued that the sentences in (14) and (15) describe the causing of an emotion rather than the emotion itself. However, it is widely held that *the cause* of an emotion is different from *the object* of the emotion. G. E. M. Anscombe offers the following example:

> A child saw a bit of red stuff on a turn in a stairway and asked what it was. He thought his nurse told him it was a bit of Satan and felt dreadful fear of it. (No

[6] This view, that emotions are perceptions as of the world acting upon us, may be seen as a development of Solomon's (2004) view of emotions as engagements with the world and Goldie's (2004) view of emotions as attributions of affective properties to external objects.

> doubt she said it was a bit of satin.) What he was frightened of was the bit of stuff; the cause of his fright was his nurse's remark. The object of fear may be the cause of fear, but, as Wittgenstein remarks, is not *as such* the cause of fear. … Or again, you may be angry *at* someone's action, when what *makes* you angry is some reminder of it, or someone's telling you of it. …'This sort of cause of a feeling or reaction may be reported by the person himself, as well as recognized by someone else, even when it is not the same as the object. (1957: 16)

In other words, one can be afraid of an object that didn't cause the fear, one can be angry at someone who didn't cause the anger, and one can be happy that x is the case, even though x's being the case didn't cause one's happiness.

However, upon further reflection, the view Anscombe expresses is perfectly consistent with the view that accurate descriptions of emotions can have the form 'I feel x cause bodily reaction r'. The child experiences some red stuff cause her bodily reactions, even though the actual cause of her bodily reactions is the nurse's remark that the red stuff is a bit of satin. The nurse might later convince the child that it was her remark that caused the child's fear, not the red satin. This may put an end to the child's fear, or the child may continue to be afraid of the red satin. In the latter scenario the child correctly judges that the nurse's remark was a cause of her fear but she continues to experience the red satin as fearsome and hence as a cause of her fear. Her emotional experience then is a case of a known illusion.

Likewise, if you are angry at someone's action, you may experience an action cause your bodily reactions, even though the actual cause of the bodily reaction is the reminder. If this discrepancy is brought to your attention but you continue to feel angry at the person's action, the person's action is perceived as that which makes you angry. This too is a case of a known illusion.

If at least some descriptions of emotions have the form 's feels x cause bodily reaction r', then we needn't conceive of emotions as conjunctive states consisting of a bodily sensation and a judgment. Rather, emotions are perceptions of external objects giving rise to certain bodily reactions.

Emotions are often thought to be appraisals. For example, Goldie (2004) argues that emotions attribute emotion-proper properties to objects. My fear of the dog represents the dog as *frightening*, and my love for John represents John as adorable. However, this thesis is consistent with the view that emotions are perceptions of causal processes. To be frightening to x plausibly just is to cause certain familiar bodily reactions in x. So, if my fear represents the dog as frightening, then it represents the dog as causing certain familiar bodily reactions. So, acknowledging that emotions attribute emotion-proper properties to objects does not require accepting a conjunctive theory of emotions.

I take at least some descriptions of emotions to reflect both narrow and wide contents. A description of my fear may reflect the narrow contents *my mouth is dry*, *my voice is faltering*, *my body is shivering*, and *there is something with dog-like (or perhaps even Fido-like) properties in front of me, and it is causing the bodily changes*. But by reflecting these narrow contents the description may also reflect corresponding wide contents that consist of Fido himself, certain physiological changes in my body and a causal reaction between them.

Some of the above considerations carry over to tactile experiences. Suppose 'The rock feels hard' is an accurate report of an aspect of a tactile experience. Which contents does it reflect? It evidently reflects that the rock is hard. But tactile experiences can reasonably be thought to involve not just representations of properties of objects but also properties of the body (see e.g. de Vignemont et al 2005). Plausibly I cannot have a tactile experience as of an object being hard without experiencing pressure to the part of my body that does the touching. This is reflected in the way we describe tactile experiences. We will say things like:

(16)

(a) The bottle feels cold to the touch

(b) John felt the rock press against his palm

(c) Mary felt the fire make her body warm

How are the sentences in (16) related to the tactual reports cited earlier? It may be thought that they reflect a different kind of mental state, a perception of the body responding causally to an object, as opposed to a tactile experience. On this view, then, tactile experiences, bodily sensations and perceptions of the body responding causally to an object are different kinds of mental states.

However, there is good reason to think that tactile experiences cannot be states entirely distinct from bodily sensations and perceptions of the body responding causally to an external object. If a rock feels hard when I put my hand around it, it feels the way it does partially in virtue of its giving rise to the sensation of pressure to my palm. To feel hard to someone plausibly just is to add pressure to the touching part of the body.

'The rock feels hard to me' then may be true in virtue of 'I feel the rock press in a certain way against my palm' and hence in virtue of me perceiving the rock press in a certain way against my palm.

If this is right, then the content of tactile experiences can be accurately reported using a 'feels' report with an unsupported clause. So, at least some accurate descriptions of tactile experiences are structurally similar to accurate causal descriptions of emotions. Like emotions, tactile experiences are not conjunctions of bodily sensations and a perception of an external object having a certain property. They are perceptions of an external object causing certain bodily reactions.

Of course, just as most of our everyday descriptions of emotions do not describe every salient aspect of emotions, so descriptions of tactile experiences do not describe every salient aspect of tactile experiences. Most descriptions are partial. If 'I feel the rock cause pressure to my palm' is an accurate description of the speaker's tactile experience, then 'I feel the rock', 'I feel the pressure of the rock', 'I feel the rock in my palm' and 'the rock feels hard', too, are accurate descriptions of aspects of the speaker's tactile experience.

The language we use to describe pain sensations and other bodily sensations suggests that bodily sensations are akin structurally to emotional and tactile experiences. Suppose 'I feel pain in my arm' is an accurate report of an aspect of a pain state. Then the report reflects the content that there is a pain in the speaker's arm, where the concept of pain can be understood as a mentalistic concept which, in normal cases, picks out a physiological bodily disturbance. If this connection between the narrow and the wide content of the concept of pain is understood by the speaker, then 'I feel pain in my arm' may reflect not only the content that there is a pain in the speaker's arm but also the content that there is a physiological disturbance to the speaker's body.

Pain sensations do not ordinarily represent external objects but they do represent parts of the body. 'I feel a pain spiraling down my arm' describes an experience that represents not just a pain (and hence, in good cases, a physiological disturbance) but also my arm and a relation between my arm and the pain. Descriptions of pain sensations can thus be descriptions of perceptions of events, as in 'my arm feels like needles are piercing their way through it', 'I feel an intense throbbing pain that shoots through the back of my left, upper arm', and 'I feel pain that throbs for a second and then goes away'. So pain sensations sometimes have a structure that is superficially similar to that of tactile experiences and emotions. Though pain experiences are not typically perceptions of external objects causing bodily reactions, they are sensations of pains moving through or being located in body parts.

There are thus good reasons why the word 'feel' occurs in our descriptions of emotions, tactile experiences and bodily sensations. First, there is good reason to think that emotions and tactile experiences have contents that are structurally similar. Both represent an external object as causing a bodily reaction. Second, descriptions of how and what we feel are interwoven. Descriptions of pain sensations, for example, can themselves be partial descriptions of emotions, as in 'I feel pain in my chest'. Descriptions of pain experiences, bodily sensations and perceptions of external objects, then, can function as partial descriptions of emotions, and descriptions of bodily sensations and descriptions of perceptions of external objects can function as partial descriptions of tactile experiences.

7. Concluding Remarks: Skepticism about Introspection

I have offered a linguistic analysis of 'feels' reports and argued that 'feels' reports, when accurate, reflect the content of mental states. 'Feels' reports most commonly purport to describe emotions, tactile experiences and bodily sensations. Because 'feels' reports reflect the representational content of mental states, our descriptions of bodily experience can give us some insight into its structure. I have argued that our descriptions of bodily experience, on the assumption that they are sometimes accurate, indicate that emotions and tactile experiences are experiences of bodily reactions to objects, whereas bodily sensations are partial descriptions of emotions and tactile experiences or other events of the body.

Before concluding I would like to briefly address a possible concern about the main project of this paper. It is widely held that attempts to read off metaphysical structure from ordinary language are doomed to fail. For example, from the fact that the English language is a tensed language, we cannot conclude that there is an ontologically fundamental present, from the fact that English makes reference to both mental and physical entities we cannot conclude that mind-body dualism is true, and from the fact that English makes reference to moral facts, we cannot infer that moral realism is true. So, some may be doubtful about a project that attempts to draw conclusions about the structure of bodily experiences based on how we describe them in ordinary language. However, while I am generally sympathetic to the limits of language as a method for gaining insight into extra-linguistic phenomena, I think there is less reason for pessimism in this particular case.

Even if we are not infallible reporters of the contents of our minds, it does seem that we have a sort of privileged access to these contents. So, at least first-person reports ought to have some degree of reliability. To completely reject the project of gaining insight into the structure of emotions based on their descriptions, one would probably need to reject the possibility of accurately describing bodily experience.

It is not uncommon to encounter this sort of skepticism with respect to descriptions of bodily experience (see e.g. Schwitzgebel 2007, 2008). As reasons for skepticism it is often pointed out that mistakes about our own emotions are uncovered repeatedly, and that we often are unable to pinpoint which emotion or bodily experience we are undergoing. People frequently retract their claims about their emotional states. Retractions such as 'I thought I was angry, but I now realize that I was really deeply saddened' in no way strike us as pathetic.

Identifying bodily perceptions is complicated further by the fact that they needn't be attended to or even conscious. For these reasons descriptions of bodily perceptions are thought to be highly unreliable or perhaps even largely inaccurate.

Though I am sympathetic to these sorts of concerns, I do find myself firmly located in the optimistic camp. This is not the place to offer full justification for an optimistic stance, but I would like to add a word or two to motivate it.

Some visual processing is unconscious. Some visual processing is conscious but not attended to. Yet we do not doubt the accuracy of descriptions of visual experience on these grounds. Similarly, we should not doubt the accuracy of descriptions of bodily experiences on the grounds that information processing about a state of the body needn't be conscious and can fail to be attended to.

A more reasonable ground of doubt is the fact that we frequently retract descriptions of emotions. We may say that we are angry at a person but later we realize that we were experiencing jealousy. However, in these kinds of cases I am tempted to doubt either the retraction (as opposed to the initial judgment) or our ability to name emotions accurately.

It is well-known that memories of perceptions are less reliable than the perceptions themselves. And some retractions are no doubt grounded in difficulties recognizing and naming particular bodily experiences. The perception of the body responding in certain ways to the thought of an upcoming job interview may be interpreted as excitement when it really ought to be classified as fear.

But there is an analogy to vision here as well. We do not linguistically classify visual experiences depending on their phenomenal properties. But suppose we did. We might then employ one name for visual experiences in which objects in the center of the visual field are attended to, a different name for visual experiences in which objects in the periphery are attended to, a third name for visual experiences in which the objects attended to are familiar, and so on. If indeed we had linguistically classified visual experiences in this way, I suspect recognizing and naming visual experiences would be as difficult as naming and recognizing emotional experiences. However, these sorts of difficulties should not lead us to doubt the accuracy of descriptions which do not name the experience described. So, I think there are plenty of grounds for optimism.

There is one further reason for optimism which I would like to mention here. It is hard to see how bodily experiences could have contents if our descriptions of bodily experiences were systematically inaccurate. If our descriptions of bodily experiences were systematically inaccurate, then presumably our judgments about our bodily experiences would be systematically false. But then information carried by our bodily experiences would hardly ever come to constitute the content of a judgment about our bodily experiences. But information that is not passed on to a conscious level of processing does not constitute experiential content, at best it would be purely computational content. Hence, to the extent that our bodily experiences have experiential contents, these contents must at least in some cases constitute the contents of true conscious beliefs (and hence also descriptions) about our bodily experiences, which gives us good grounds for optimism.

References

Anscombe, G. E. M. (1957/2000). *Intention*. Harvard University Press.

Armstrong, D. M. (1961). *Perception and the Physical World,* London: Routledge.

Bach, K. (1999). The semantics–pragmatics distinction: What it is and why it matters. In Ken Turner, ed., *The Semantics–Pragmatics Interface from Different Points of View*, 65–84. Oxford: Elsevier.

Byrne, A. (2009). "Experience and Content". *Philosophical Quarterly* 59: 429-451.

Chalmers, D. (2004). "The Representational Character of Experience", In B. Leiter (ed.), *The Future for Philosophy* (153-81). Oxford: Oxford University Press.

Chalmers, D. (2009). "Propositions and attitude ascriptions: A Fregean account", *Nous*, forthcoming.

Chisholm, R. M. (1957). *Perceiving: A Philosophical Study*, Ithaca: Cornell University Press.

de Vignemont, F., Erhsson, H., Haggard, P. (2005), "Bodily Illusions Modulate Tactile Perception", *Current Biology* 15:1286-90.

Goldie, P. (2004). "Emotion, Feeling, and Knowledge of the World", in *Thinking about Feeling: Contemporary Philosophers on Emotions*, Solomon, R. C. ed., Oxford: Oxford University Press, 91-106.

Glüer, K. (2009). "In Defence of a Doxastic Account of Experience", *Mind and Language* 24: 297-327.

Heim, I. (2006). "Remarks on comparative clauses as generalized quantifiers", Ms, MIT.

Higginbotham, J. T. (1983). "The Logic of Perceptual Reports: An Extensional Alternative to Situation Semantics", *Journal of Philosophy* 80: 100-127.

Jackson, F. (1977). *Perception: A Representative Theory*, Cambridge: Cambridge University Press.

Kenny, A. (2003/1963), *Action, Emotion and Will*, Routledge.

Maund, B. J. (1986). "The Phenomenal and Other Uses of 'Looks' ", *Australasian Journal of Philosophy* 64: 170-180.

Nanay, B. Forthcoming. “Perception and Imagination: Amodal Perception as Mental Imagery”. *Philosophical Studies* (forthcoming)

Richard, M. E. (1990). *Propositional Attitudes: An Essay on Thoughts and How We Ascribe Them*, New York: Cambridge University Press.

Siegel, S. (2005). “Which Properties Are Represented in Perception?”. In *Perceptual Experience*, eds. T. Szabo Gendler and J. Hawthorne, Oxford: Oxford University Press.

Solomon, R. C. (2004). “Emotions, Thoughts, and Feelings: Emotions as Engagements with the World”, in *Thinking about Feeling: Contemporary Philosophers on Emotions*, Solomon, R. C. ed., Oxford: Oxford University Press, 76-90.

Stalnaker, R. (1997). “What Might Non-conceptual Content Be?”, presented by R. Stalnaker at the SOFIA conference, Barcelona, June 1997.

Schwitzgebel, E. (2007). “No Unchallengeable Epistemic Authority, of Any Sort, Regarding Our Own Conscious Experience”, Phenomenology and the Cognitive Sciences 6.

Schwitzgebel, E. (2008). “The Unreliability of Naïve Reflection”. *The Philosophical Review*, *117*, 245–273.

Trebilcot, J. (1970). “Dr. Kenny's Perceptions”, *Mind* 79: 142-143.