

Visual Attention Fixes Demonstrative Reference by Eliminating Referential Luck^{*}

Suppose you are looking at an ordinary material thing in front of you. You think thoughts you would express by saying ‘That’s orange. It’s spherical. It’s moving to the left.’ In this case you are thinking what I shall call ‘perceptual demonstrative thoughts’: thoughts of the kind ordinarily made available by current perceptual contact with the objects they are about. This paper is about how perceptual contact with objects enables us to have perceptual demonstrative thoughts about them.

The paper has four parts. §1 motivates two constraints on right accounts of perceptual demonstrative thought. §2 summarises some empirical results about perceptual attention. §3 uses these results to develop an account of perceptual demonstrative thought that meets the constraints. §4 gives two reasons to prefer the resulting account to its nearest rivals.

§1 Two constraints on accounts of perceptual demonstrative thought

This section proposes two constraints on accounts of perceptual demonstrative thought. The first part of the section states the constraints. The second and third parts motivate them.

§1.1 Stating the constraints

I begin with some necessary terminology. I shall use ‘concept’ to abbreviate ‘representation deployable in thought’. I shall say that concept v ‘refers to’ object o iff, for all Φ (where Φ is a concept of property Φ), $\langle v \text{ is } \Phi \rangle$ is true iff o has Φ ¹. So I am treating reference as a relation between thought components and things thoughts are about. I shall use ‘perceptual demonstrative’ to abbreviate ‘concept referring to a particular, and of the kind deployed in thinking perceptual demonstrative thoughts’.

In addition, I shall say that a concept’s ‘basic’ deployments are the deployments you must be disposed to make to count as grasping it. For example, classical inferential moves are often held to be basic for concepts of classical logical constants². And I shall say that a basic deployment is ‘canonically justifiable’ iff to count as grasping the concept you must be disposed not just to make the deployment, but to make it on the grounds of a specific (‘canonical’) justification. (I am going to argue that perceptual demonstratives have canonically justifiable basic deployments. I leave it open which other concepts do too.)

Finally, I want to introduce the notion of justification that is ‘luck eliminating with respect to an object’ as follows.

Start with this account of what it is for a justified true belief to be ‘merely luckily’ true:

S’s justified true belief that p is ‘merely luckily’ true iff S’s justification fails to exclude not- p situations that it must exclude if S is to be rationally entitled to the belief (‘relevant’ not- p situations).³

For example, suppose I believe that X lives in a particular house on the basis of twice having seen X leaving it. Suppose there is no special background in place to make these observations good evidence for the belief: I have simply leapt to a conclusion. Then if my belief is true it is merely luckily true. For, though there is a hard question about which situations are ‘relevant’ to rational entitlement, any plausible view will recognize the relevance of alternatives the subject knows to be commonplace. And I know that there are commonplace situations where people emerge from houses other than their own. So my justification fails to exclude relevant contrary alternatives. Therefore, I am merely lucky if my belief turns out true.

In contrast, many philosophers regard global sceptical situations as irrelevant to our rational entitlement to everyday beliefs in ordinary contexts. Our ordinary justification for everyday beliefs does not exclude the possibility that we are brains in vats. But this possibility is not usually relevant to our entitlement to these beliefs. So, as long as your justification for an everyday belief excludes all relevant alternatives, if your belief is true it is not merely luckily true. And if your belief is false because you actually are a brain in a vat, it is unluckily false: you are ‘unlucky’ in that a situation you are not rationally required to exclude happens to obtain.⁴

It is a short step from this definition of ‘merely lucky’ true belief to a definition of ‘luck eliminating’ justification:

S’s justification for the belief that p is ‘luck eliminating’ iff it excludes all relevant not- p situations.

Justification for a belief is ‘luck-eliminating’ iff it tends towards truth, so that the subject is unlucky if the belief is false, and not merely lucky if it is true. Similarly, justification for a belief is ‘luck-eliminating with respect to an object’ iff it tends towards getting the object’s properties right:

S’s justification for the belief that $\langle v \text{ is } \Phi \rangle$ is ‘luck-eliminating with respect to o ’ iff it excludes all relevant situations in which o is not Φ (so that S is unlucky if o is not Φ , and not merely lucky if o is Φ).⁵

Finally, we can define ‘luck-elimination with respect to o ’ for the justification underwriting a specific move towards a belief in terms of how this justification contributes to the belief’s total justificatory package:

S’s justification for a specific move towards the belief that $\langle v \text{ is } \Phi \rangle$ is ‘luck-eliminating with respect to o ’ iff, if S makes the move on the basis of the justification, but ends up with a belief whose total justification is not luck-eliminating with respect to o , some luck-introducing fault elsewhere in S’s total justification for the belief is to blame.

For example, consider an astronomer (hereafter ‘A’) taking telescope readings. Suppose A is justified in believing that the readings derive from a single object. A introduces a name, ‘Xena’, and forms beliefs she would express using ‘Xena’ by inferences like

The telescope is detecting iron-richness
Xena is iron-rich

The telescope is detecting γ -radiation
Xena emits γ -radiation

A's justification for these moves derives from her justified belief that the telescope is delivering information about a single object. This justification is luck-eliminating with respect to o iff the telescope is focused on o . For suppose the telescope is focussed on o . Then if A is merely lucky to end up with Xena-beliefs that match o 's properties, or is not unlucky to end up with Xena beliefs that do not match o 's properties, there must be a luck-introducing fault elsewhere in the justificatory package for these beliefs (perhaps A should have, but has not, checked the telescope's reliability). Conversely, suppose the telescope is not focussed on o (it is focussed on o^* or on nothing). Then if A ends up with Xena-beliefs that match what o is like this will be a mere matter of luck.

In these terms, the constraints I want to introduce can be put like this:

NON-LUCKINESS – if v belongs to a class of concepts that have canonically justifiable basic deployments, v refers to o iff canonical justification for basic deployments of v is luck-eliminating with respect to o .

GROUNDING – Canonical justification for basic deployments of a perceptual demonstrative referring to o requires perceptual attention to o .

I shall suppose that perceptual attention to an object involves both selection and awareness. It involves selection in that attention to o both recruits information processing power to the extraction of information from o , and imposes a structure on the extracted information, for example, binding together features as features of a single thing. It involves awareness in that attending to an object is a way of being aware of it: there is something that a state of attending to an object is like. (This allows for perceptual contact involving awareness without selection: this is the kind of perceptual contact we have with objects and properties in unattended peripheral vision. It also allows for perceptual contact that involves selection without awareness. I shall say more about both kinds of case in motivating GROUNDING below.)⁶

NON-LUCKINESS and GROUNDING entail that perceptual attention enables perceptual demonstrative reference because it secures canonical justification for basic uses of a perceptual demonstrative that is luck-eliminating with respect to the attended object.

The rest of this section motivates the two constraints.

§1.2 Motivating NON-LUCKINESS

NON-LUCKINESS states both a necessity claim:

Where v has a range of canonically justifiable basic deployments, v refers to o only if canonical justification for basic deployments of v is luck-eliminating with respect to o .

and a sufficiency claim

If canonical justification for basic deployments of v is luck-eliminating with respect to o , v refers to o .

I shall motivate each in turn.

First consider the necessity claim. Suppose this claim is false. Then v might refer to o even though canonical justification for basic deployments of v is not luck-eliminating with respect to o . But to allow this is to allow the following combination:

- i) S arrives at the belief that $\langle v \text{ is } \Phi \rangle$ by a canonically justified basic move.
- ii) There are no luck-introducing faults elsewhere in the belief-forming process.
- iii) S 's justification fails to exclude all relevant situations where $\langle v \text{ is } \Phi \rangle$ is false.

And this combination is incoherent. For suppose (iii). Then there are situations where $\langle v \text{ is } \Phi \rangle$ is false that must be excluded for rational entitlement to the belief, but which S 's justification fails to exclude. So S would not be irrational to doubt that $\langle v \text{ is } \Phi \rangle$. Now suppose (ii). Then the possibility of rational doubt is not generated by elements in the belief-forming process prior to the move at (i): (ii) says that these other elements are in good order. Rather, the possible rational doubt concerns the belief-forming move itself. But (i) says that this is a canonically justified basic move. And canonically justified basic moves cannot rationally be doubted. Contradiction.

So we have the 'necessity' part of NON-LUCKINESS: given that v has canonically justifiable basic deployments, if v refers to o then S 's canonical justification for basic deployments of v is luck-eliminating with respect to o .

Now consider the 'sufficiency' part: if canonical justification for basic deployments of v is luck-eliminating with respect to o , v refers to o . Suppose not. Then it is possible that

- (i) S has rational entitlement generating justification for the belief that $\langle v \text{ is } \Phi \rangle$; this justification is luck-eliminating with respect to o .

but

- (ii) Something other than o 's being Φ is required for the truth of $\langle v \text{ is } \Phi \rangle$.

And this combination is at best extremely implausible. For consider

- (iii) S ought to believe that p only if p is true.

and

- (iv) S ought to believe that p only if S has justification that secures rational entitlement to the belief.

(iii) states the so-called 'objective norm' for belief. To deny (iii) is to deny that a belief that turns out to be false turns out to be (objectively) wrong. (iv) states the 'subjective norm'. We ought (subjectively) to form only beliefs to which we are rationally entitled because only these beliefs are in good order with respect to our mental economies.⁷

Now consider the following distinction between ‘direct’ and ‘indirect’ regulation of mechanisms by norms⁸:

Norm N regulates a mechanism ‘directly’ iff right functioning of the mechanism consists in responding to detection of the condition for N’s application by doing what N calls for or allows. (So if N regulates a mechanism directly, a violation of N is a malfunction of the mechanism.)

Norm N regulates a mechanism ‘indirectly’ iff the mechanism is directly regulated by some other norm that is followed partly as a means of conforming to N.

It is plausible that (iv) regulates our belief-forming mechanisms directly. If you end up with beliefs to which you are not rationally entitled, your belief-forming mechanisms have misfired.

It is also plausible that (iii) does not regulate our belief-forming mechanisms directly. For it is plausible that there are cases of unlucky false belief: cases where a subject forms a belief on the basis of justification that secures rational entitlement, but the situation is unlucky and the belief false. And if there are such cases, they do not involve malfunction of our belief-forming mechanisms. Rather, they are cases where our belief-forming mechanisms function as they should, but the world does not oblige.

If (iii) does not regulate our belief forming mechanisms directly, it must do so indirectly. So there must be a norm that regulates our belief-forming mechanisms directly and that we follow as a means of conforming to (iii). But we conform to (iv) directly. And there are no other candidates to be direct norms for belief. It follows that we obey the subjective norm for belief (iv) partly as a means of conforming to the objective norm (iii).⁹

But this relation between (iii) and (iv) requires

(v) Justification that secures rational entitlement to a belief rules out situations relevant to whether you ought (subjectively) to form the belief in which the belief is not true.

Without (v), direct conformity to (iv) is not a means of conformity to (iii). And (i) and (ii) entail that (v) is false. Given (i), S has rational entitlement generating justification that rules out all relevant situations in which *o* is not Φ . But given (ii), this is not to rule out the relevant situations in which S’s belief is not true.

So the story about how belief is regulated – directly by (iv) and indirectly by (iii) - generates a case for the sufficiency part of NON-LUCKINESS: if canonical justification for basic deployments of *v* is luck-eliminating with respect to *o*, *v* refers to *o*.

§1.3 Motivating *GROUNDING*

GROUNDING says that canonical justification for basic deployments of a perceptual demonstrative referring to *o* requires perceptual attention to *o*.

As a first step towards motivating *GROUNDING*, I need to say more about which concepts count as perceptual demonstratives. So consider again the case where you are looking at ordinary thing *o* in front of you and thinking <that>-thoughts about it: thoughts

like <That is orange> or <That is spherical>. Here are some observations about this simple case. (I intend these observations as data that will help precisify what counts as a perceptual demonstrative, and that a right account of perceptual demonstratives should be able to explain.)

Firstly, your <that> thoughts are made available solely by your current perceptual link with *o*. If you were not currently perceiving *o*, some extra relation would be required for you to think about it. For example, you might remember it, know a proper name for it, or grasp a description that it satisfies. No such extra relation is required for aboutness in the present case.

Secondly, your <that>-thoughts are existentially committing: you are disposed to move from <That is orange> to <Something is orange>. They are also ‘unified’ in the following sense: you are prepared to move from <That is orange> and <That is spherical> to <That is both orange and spherical> without needing extra reasoning to show that your tokens of <that> co-refer.¹⁰

Thirdly, there is a close relation between your taking it that you are rationally entitled to make these moves and your preparedness to essay your <that>-thoughts. If you think you are not entitled to move from <That is orange> to <Something is orange> or from <That is orange> and <That is spherical> to <That is both orange and spherical>, you will not be prepared to essay the initial atomic <that> thoughts in the first place. It seems that you are rationally entitled to these moves just in virtue of the availability of your <that>-thoughts; the rational entitlement is secured by whatever makes the <that>-thoughts available.

So reflection on this simple case generates a candidate delineation of a kind of conceptual representation. The kind’s members are conceptual representations that are made available by perceptual contact with objects, where the perceptual contact with the object also justifies moves from <*v* is Φ > to <There is something Φ there>; from <*v* is Φ > and <*v* is Ψ > to <*v* is both Φ and Ψ >; and perhaps other moves that would be revealed by further reflection. I suggest that perceptual demonstratives are the members of this kind. (This precisifies my earlier claim that perceptual demonstratives are thoughts of the kind ‘ordinarily made available’ by current perceptual contact with the objects they are about. Note that I am allowing that there is perception-based reference that is not perceptual demonstrative reference. For example, suppose I think <He must be really tall>, looking at the enormous shoe of somebody I have never met. Then perceptual contact with one thing is enabling me to think about another. But this kind of reference – reference secured by ‘deferred ostension’¹¹ – is not perceptual demonstrative reference.)

This account of what perceptual demonstratives are entails that they have canonically justifiable basic deployments. The basic deployments are the deployments that characterize the kind. Canonical justification for these basic deployments is justification in virtue of perceptual contact.

But why think that the perceptual contact that secures canonical justification for basic deployments of perceptual demonstratives must be attentive perceptual contact – perceptual contact involving both selection and awareness? Why will awareness without selection or selection without awareness not suffice?

Consider awareness without selection first. The central case of awareness without selection is unattended peripheral vision. There is an ordinary use of ‘that’ associated with this case: ‘I’m going to ignore that’ you say, resisting the temptation to let your

attention be drawn by some change in your unattended visual field. But the moves that are automatically justified for ordinary perceptual demonstrative <that> are not justified for unattended peripheral <that> (hereafter ‘<pthat>’). Unattended peripheral vision does not bind features as features of single objects¹². So there is no automatic justification for the move from <Pthat is orange> and <Pthat is spherical> to <Pthat is orange and spherical>. Similarly, <pthat> is not existentially committing. We do not move from <Pthat is orange> to <There is an orange object there>. Your expectation when you think <Pthat is orange> is that it may turn out that, though there is orangeness in a region of your visual field, there is no orange object there at all: you anticipate finding that your experience is caused by a trick of the light, or a streak of orange in the evening sky.

Now consider selection without awareness. Here are two kinds of case where selection without awareness occurs.

Firstly, there are cases where perceptual contact is too brief for awareness, but the mind nevertheless registers information derived from the thing that was (briefly) perceived, and registers this information with the structure characteristic of selective processing (for example, binding features as features of a single thing). The fact that information has been registered and has selective structure is shown by the ‘priming’ effects of the brief perceptual contact. For example, perceptual contact with a famous face can decrease the time it takes to recognize an associated famous face even when contact with the first face was too brief for awareness.¹³

Secondly, there are cases of ‘blindsight’. Blindsight subjects have a type of brain damage that leaves them without conscious awareness of all or part of their visual fields (the ‘blank’ part is the ‘blind field’). These people are, nevertheless, able to make reliable forced-choice guesses about the orientation, direction, and kind of an object in the blind field, and successfully reach out and grab these objects when prompted to do so. So in blindsight there seems to be selective processing of information from objects of which the subject is not aware.¹⁴

Could selection without awareness generate the justification characteristic of basic deployments of perceptual demonstrative <that>? To see what is involved in a ‘Yes’ answer to this question, let us focus on the case of blindsight. Suppose that a blindsight subject (hereafter ‘B’) introduces a term, ‘bthat’, to use in responding to forced-choice questions about the object in the blindfield. So in response to ‘Is it square or round?’ B says ‘Bthat is round’, and so on. Then to allow that the perceptual link underlying B’s ‘Bthat’-reports secures the justification characteristic of basic deployments of perceptual demonstrative <that> is to endorse the following view (the ‘<bthat> hypothesis’). B’s perceptual link with the object in the blindfield makes available a demonstrative (<bthat>) whose uses it justifies. B’s responses to forced-choice questions are expressions of thoughts of form <Bthat is Φ >. B is aware of his reports of <bthat> thoughts, but is not aware of the thoughts themselves. (Blindsight subjects claim that they are ‘just guessing’, or ‘just find themselves’ making the forced-choice responses that they do.) Nevertheless, B’s perceptual contact with the object in the blindfield justifies both his <bthat>-thoughts and his ‘bthat’ utterances. So B is not really guessing. He just thinks he is because his condition leaves him unaware of his grasp of <bthat> and his justification for deploying it the way he does.

Here is what I take to be the main argument against the <bthat> hypothesis (this is a special case of an old line of thought¹⁵):

(i) The intentionality (or ‘aboutness’) characteristic of thought is essentially a phenomenon of subjective consciousness: no mental state has conceptual content (the content characteristic of thought) unless the subject has or can have direct subjective awareness (awareness not based on inference) of being in the state.

(ii) According to the <bthat> hypothesis, B’s forced choice reports express justified deployments of conceptual representations of which B has and can have no direct subjective awareness. So the <bthat> hypothesis violates (i). (B might gain access to his <bthat> thoughts by reasoning as follows: ‘I am disposed to say “Bthat is spherical”’; when I am disposed to say “Bthat is spherical” there is a spherical thing in my blindfield; when there is a spherical thing in my blindfield I am thinking <Bthat is spherical>; so I am thinking <Bthat is spherical>.’ But this is not direct subjective awareness.)

So

(iii) A right account of perceptual demonstrative thought must rule out the <bthat> hypothesis.

I hasten to acknowledge that (i) is controversial. Rejection of anything like (i) is central to ‘naturalist’ views of intentional content. On a ‘naturalist’ view, all relations of ‘aboutness’ are reducible to relations that can be stated in purely non-intentional, non-mentalistic terms.¹⁶ For example, the claim might be that ‘aboutness’ is purely a matter of causal relations between physical and non-intentional states of a person’s brain and things and properties in the world.¹⁷ Or it might be that ‘aboutness’ is to be explained in evolutionary terms, where intentional and mentalistic notions (like belief, justification, and awareness) play no irreducible role in the explanation.¹⁸ The notion of ‘subjective awareness of being in a mental state’ is both mentalistic and intentional. So on a naturalist view, the suggestion that there is a kind of intentionality (the intentionality characteristic of thought) that essentially involves subjective awareness of being in a mental state is just wrong.

It is obviously not possible to address the naturalism/ non-naturalism debate in a paper whose central focus lies elsewhere. So in what follows I shall just suppose (i). But to meet naturalists half-way, I point out that a central part of the proposal I am going to make can be excised from my own preferred non-naturalist setting and translated into a naturalist account of how perception secures perceptual demonstrative reference. I shall indicate the opportunity for the excision when it arises.

§1.4 Conclusion to §1

I conclude that there are at least good initial reasons to accept both NON-LUCKINESS and (for non-naturalists) GROUNDING. Given these constraints, the problem of explaining how perception enables perceptual demonstrative reference reduces to the following question: How does perceptual attention secure justification for basic deployments of perceptual demonstratives that is luck-eliminating with respect to the attended object? The rest of the paper answers this question.


2 Perceptual attention and visual objects

This section summarises some empirical results about perceptual attention. I should stress that I am not aiming for empirical exhaustiveness. I shall describe only (some of) the empirical data most relevant to the account of perceptual demonstratives I want to propose.

Recall that as I am using the term here ‘perceptual attention’ involves both awareness and selection. The notion of ‘selection’ itself has two components. Selection recruits extra information-processing to attended parts of the visual field (so that we typically have more information from, and are quicker to detect changes in attended regions). And selection structures the input from an attended region, for example, binding features together as features of a single object. Some instances of perceptual attention involve only the first type of selection: when you attend to a region of a blank wall your perceptual processing is devoting extra processing power to the region, but not binding features. Since mere enhanced processing does not bind features, it does not support the inferences characteristic of perceptual demonstrative <that>. So it is the richer, structure-involving kind of selection that must underlie perceptual demonstrative reference. The empirical results I am going to summarise concern how the relevant structuring works.

Let us start with a widespread assumption: the assumption that pre-conceptual processing occurs, as psychologists say ‘in parallel’, while conceptual processing occurs ‘in series’. A parallel processing system can perform multiple tasks at once without significant drop-off in performance, up to a point where its capacity is exceeded and performance collapses. In contrast, a serial processing system can perform only one task at once. If assigned more than one, it must cycle between tasks, with a consequent decline in performance. The assumption is plausible because pre-conceptual processing seems to have the first performance profile – level performance across increasing demands up to capacity – while conceptual processing seems to have the second: the more conceptual-level tasks you attempt simultaneously, the worse your performance gets¹⁹.

Now consider the following experiment (a ‘multiple object tracking’ experiment). Subjects are shown a display containing a number of qualitatively identical dots. Some dots are ‘flashed’ to distinguish them from others ((a) in the diagram). Then they all move randomly around the display (b). When motion stops the subject is asked whether a given dot flashed at the outset (c).



(from Scholl ‘Objects and Attention’)

To succeed at this task, you must pick out the initially cued dots and keep track of them over time. If either picking out or keeping track involved conceptual processing, our performance should be worse at tracking two dots (so when two dots flash at (a)) than one, worse at three than two, and worse again at four. But in fact our performance does not follow this pattern. We are as good at tracking two, three, or four things as we are at tracking one. At five the capacity seems to hit overload and fall off suddenly. This pattern

– uniform success up to a point where the system’s capacity is reached – is characteristic of pre-conceptual processing. So we can conclude that the processing involved in both stages of the task is pre-conceptual.²⁰

Other experiments²¹ reveal that whether performance follows the efficiency-and-falling-off pattern characteristic of preconceptual processing depends on how the dots behave. For example, the pattern does not arise for dots whose motion is spatio-temporally discontinuous (so that dots seem to leap from one point to another without traversing the ground between). We can track one dot or ‘thing’ behaving in this way. (Imagine watching an aeroplane’s shadow flicker across an uneven landscape.) But the capacity to track a ‘thing’ moving like this lacks the performance profile characteristic of non-conceptual capacities. We are much worse at tracking two jumping-around dots than one, and much worse again at tracking three. So tracking of dots whose spatio-temporal paths appear discontinuous is conceptually aided. Similarly, the non-conceptual efficiency pattern does not arise for dots whose motion appears to violate the boundaries of initially cued objects and the relations between their parts. (In the relevant experiments, dots seem to dissolve and flow to new locations where they reassemble.) We can easily track one ‘thing’ behaving like this. But we are much worse at tracking more than one. So when you track a ‘thing’ whose motion does not preserve boundaries and relations between parts, you are engaged in conceptually aided tracking. Parallel claims hold for ‘things’ that engage in apparent fission or fusion or apparent non-independent motion (motion as a part of a larger whole), and ‘things’ whose sizes appear to fluctuate wildly over short periods.

The visual cuing at (a) works by grabbing attention: detection of the cue recruits extra information-processing resources to the cued element of the visual field. And additional experiments show that the tracking in (b) involves maintaining attention to the initially cued dots. (We are much quicker at detecting gradual changes in tracked than in untracked dots; quicker detection indicates attention because speed of detection is proportional to dedication of information-processing resources.²²)

So the conclusion that locking on to cued dots and tracking them over time are pre-conceptual can be restated as a conclusion about perceptual attention. Pre-conceptual processing parses the visual field into units of coherence (‘visual objects’) that can draw and hold attention without conceptual assistance.

Similarly, the results about how dots must behave for conceptually unaided tracking translate into a conclusion about the structuring of perceptual input involved in perceptual attention to visual objects.

To state this conclusion, let us first introduce a distinction between what I shall call the ‘features detected’ and the ‘property information delivered’ through an attentional channel. Imagine that you are watching a coin spin slowly in front of you. There is an intuitive sense in which the coin’s shape ‘looks different’ as it turns: seen heads-on, the coin looks round; as it turns, it looks like a flatter and flatter ellipse, then a skinny rectangle, then an ellipse again. But there is also an intuitive sense in which the coin’s shape appears unchanging: your experience seems like an experience of a thing with a stable shape that is presenting different angles to you, not like an experience of a thing whose shape is unstable. I shall use ‘features detected’ to capture what appears to change in this kind of case, and ‘property information delivered’ to capture what appears to remain the same.²³

Given this distinction, the results about the conditions for conceptually unaided attentional tracking emerge as results about combinations of features that can be detected through a conceptually unaided attentional channel: the features detected must be completable by the visual system into a characteristic kind of property information. The location features detected ($\langle \text{at } p \text{ at } t \rangle$; $\langle \text{at } p^* \text{ at } t^* \rangle$) must be completable by the visual system into the appearance of a continuous spatio-temporal path. The size-features must be completable into the appearance of relatively stable size given motion at detectable speeds in three dimensions. The features detected when the visual object overlaps in the visual field with other visual objects must demand resolution by the visual system into the appearance of interactions between self-standing objects: appearances of collision, rebounding, or one thing passing behind another, but not fission or fusion. And the features detected must not demand completion that results in the appearance of motion that violates object boundaries, or the appearance that the attended thing is part of something larger. (This is probably an incomplete list. Further experimentation may reveal extra conditions.)

The next section uses these empirical conclusions to provide an account of how perceptual attention fixes perceptual demonstrative reference that meets the constraints from §1.

§3 Perceptual attention, luck-eliminating justification, and perceptual demonstrative reference

The first step towards the account I want to propose is to note something that is surely not coincidental: there is a match between the conditions for sustaining conceptually unaided attentional tracking and the traditional ontological category of ordinary objects. Here is a working definition of the traditional category ‘ordinary object’:

An ‘ordinary object’ traces out a continuous spatio-temporal path. It moves and acts in ways that respect its boundaries and relations between its parts (so that usually parts that are adjoined at t are still adjoined at $t + 1$). It does not survive fission or fusion. It is not ‘part’ of any other object in the same way its parts are parts of it. Its macroscopic properties are relatively stable.²⁴

And I repeat the conditions for sustaining conceptually unaided attentional tracking from the end of §3:

We can sustain conceptually unaided attentional tracking only as long as the stream of features detected demands completion into the appearance of an object that traces out a continuous spatio-temporal path; has boundaries that remain intact during motion; moves in ways that respect relations between its parts; moves independently (rather than as a part of something larger); does not appear to engage in fission or fusion; and has a relatively stable size.

The second step towards the account I want to propose is to import a notion from the wider literature on thought and speech about particulars: the notion of a ‘file’ of

beliefs organized in such a way that it is a candidate to be a file about a particular (I shall abbreviate this to ‘object file’). An object file’s developmental path is determined by the incoming information directed to it, and how this incoming information is marshalled (the rejection of combinations taken to be incompatible; the drawing of conclusions from information the file already contains). In a ‘files’ framework, conceptual representations potentially ‘of’ objects just are files of information with appropriate organisational structure. And reference (a relation between conceptual representations and objects) is a relation between a file and the object it is about.²⁵

Suppose that some of our object files (‘attention-based’ object files) inherit their contents directly from conceptually unaided attentional channels. So an attention-based object file develops like this (imagine that you are watching an orange ball roll away from you down a corridor at constant speed; at $t + 1$ it passes through a patch of bright sunlight; at $t + 2$ it comes to rest partially obscured)²⁶:

Time	t	t + 1	t + 2
<i>Features detected through attentional channel</i>	orange (given lighting conditions x), round, occupying a z-sized patch of the visual field, located at p, moving with velocity v	brighter orange (given lighting conditions x*), round, occupying a z* sized patch of the visual field (for z* smaller than z), at p*, moving with velocity v	orange (given lighting conditions x), half-round (viewed from <i>here</i>), occupying a z** sized patch of the visual field (z** smaller than z*), at p**, stationary
<i>Property information delivered</i>	is orange, is spherical, is size s, is at p, is moving with velocity v	is the same shade of orange, is spherical, is size s, is at p*, is moving with velocity v	is the same shade of orange, is spherical, is size s, is at p**, is stationary
<i>Contents of attention-based object file</i>	<is orange>, <is spherical>, <is size s>, <is at p*>, <is moving with velocity v>	<is orange>, <is spherical>, <is size s>, <has maintained stable colour, shape, velocity and size since t>, <has moved from p to p*>	<is orange>, <is spherical>, <is size s>, <has maintained stable colour, shape and size since t>, <is at p**>, <moved to p** from p by way of p*>, <is stationary>

Then I propose the following conjecture:

Conjecture – Perceptual demonstratives are attention-based object-files.

(This conjecture will be turned into something more respectable by the account of perceptual demonstratives it enables us to provide.)

And now consider the following toy scenario. You are trying to construct a full account of someone’s journey from snippets of information (‘X was in Chicago on the 17th, and Butte on the 24th’; ‘X spent the night of the 20th in Laramie’). You are using a railway schedule that you know to be accurate as a guide to which routes and journey times are possible. As you construct your account, you use the railway schedule to justify your attempts to fill in missing details. (‘The only train out of Laramie on the morning of the 21st was going to Denver, so X went from Laramie to Denver’; and so on.) Your justification for these moves (justification by appeal to the railway schedule) is luck-eliminating with respect to X’s journey iff X in fact travelled only by train. For suppose X did travel only by train. And suppose you are careful to draw only conclusions that really are justified by the railway schedule. Then if you end up with a wrong account of X’s journey, or are merely lucky to end up with a right account, the snippets of information you started with will be to blame, not the moves you made by marshalling these snippets according to the railway schedule. Conversely, suppose X actually drove or rode a mule. Then even if the snippets you start with are accurate, you would have to be extremely lucky to end up right about the path X followed.

Given the story about attention-based object-files so far, we almost have a parallel account of how perceptual attention secures luck-eliminating justification for deployments of perceptual demonstratives. For attentional information processing completes incoming feature-snippets into property information as if the attended object is an ordinary object. And an attention-based object file inherits its contents from the property information delivered through the attentional channel that feeds it. So if you are attending to an ordinary object, and end up with an attention-based object-file that gets the object's properties wrong, or are merely lucky to end up with a file that gets its properties right, it will be your feature-detection that is to blame, not the marshalling of features to give an attention-based object-file. Conversely, if the attended object is not an ordinary object, it will be as if you were using a railway schedule to construct an account of someone's journey from snippets of information when in fact the person drove. If the attended object is not an ordinary object, even if the feature information you are detecting is accurate, you are just lucky if you end up with an attention-based object file that gets the object's properties right. (I shall say more about attention to visual objects that are not ordinary objects in §4.2 below.)

This takes us as far as

Non-lucky correctness – Suppose S has an attention-based object-file fed by an attentional link with ordinary object *o*. Then, provided that S's perceptual information processing is accurate with respect to detection of features, S will be unlucky to end up with an attention-based object-file that does not match *o*'s properties, and not merely lucky to end up with an attention-based object-file that does match *o*'s properties.

Now recall

NON-LUCKINESS – if *v* belongs to a class of concepts with canonically justifiable basic deployments, *v* refers to *o* iff canonical justification for basic deployments of *v* is luck-eliminating with respect to *o*.

NON-LUCKINESS entails that two further steps will deliver an account of how perceptual attention enables perceptual demonstrative reference. The first is to show that justification for file-maintaining moves is canonical justification for basic deployments of perceptual demonstratives. The second is the step from the purely causal *non-lucky correctness* to a normative parallel: the claim that justification for the moves involved in maintaining an attention-based object-file is luck-eliminating with respect to the attended object iff it is an ordinary object.

In fact, there is little room for argument about the first of these steps. For NON-LUCKINESS entails that if canonical justification is undermined, reference is undermined. But perceptual demonstrative reference is robust in the face of factors that undermine rational entitlement to treat perceptual input as accurate. You can refer to a thing using a perceptual demonstrative even though you know that your visual system is befuddled by a distorting mirror and bad light, so you know that your detection of colour, shape, size, and location features is unreliable.²⁷ (Note that what is undermined in this case is your entitlement to believe <That is orange> or <That is spherical> on the basis of apparent orangeness or sphericity. Your entitlement to move from <That is orange> to

<Something is orange> or from <That is orange> and <That is spherical> to <Something is orange and spherical> is not undermined: entitlement to these moves is constitutive of grasp of perceptual demonstrative <that>.) Since reference for perceptual demonstratives is not undermined by factors that undermine the normative status of incoming feature information, canonical justification for basic uses of attention-based object-files is not the justification that this incoming information carries. It is justification for the marshalling of features that is involved in maintenance of an attention-based object file.

So to deliver an account of how perceptual attention enables perceptual demonstrative reference it remains to show how *non-lucky correctness* can be turned into a normative parallel.

It is at this point that the opportunity arises for a naturalist variant on the proposal of this paper. From a naturalist point of view, justification just is a matter of non-lucky correctness. So a naturalist who is convinced by the argument for NON-LUCKINESS will be envisaging an account of justification on which there is no distinction between justification that is luck-eliminating with respect to an object, and causal factors that secure non-lucky correctness with respect to the object. From this point of view, if NON-LUCKINESS is true, *non-lucky correctness* is all that an account of perceptual demonstrative reference requires.

But suppose we want a non-naturalist account of how perception enables perceptual demonstrative reference. How might the step from the merely causal (so, from a non-naturalist point of view, non-normative) *non-lucky correctness* to a normative parallel go?

There is a tactic for answering this question that is thrust to the fore by traditional discussions of justification. These discussions have focussed on justification for belief. And they have shared the following starting assumption:

Justification for belief – A belief is justified iff there is an appropriate dependence of the way the belief is formed on whether the state of affairs required for its truth obtains (or will obtain).

Given this assumption, the debate about how beliefs are justified has been about what constitutes an ‘appropriate dependence’. For example, philosophers who think that justification involves reliability say that ‘appropriate dependence’ is at least partly a matter of causal dependence: justification requires that a belief be caused by a mechanism that will generally operate only if its truthmaker does or will obtain. Those who think justification involves counterfactual dependence say that a belief is justified only if the subject’s holding the belief depends counterfactually on whether or not it is (or will turn out to be) true²⁸. Given *justification for belief* as a model, the obvious move in attempting to get from a claim about causal dependence to a claim about justification is to argue that the causal dependence, if properly described, will reveal itself as a justificatorily ‘appropriate’ dependence. I shall discuss a view of perceptual demonstratives that makes this kind of move in §4.

But I want to suggest an alternative. I suggest that the relation between causation and justification that an account of perceptual demonstratives requires should not be modelled on *justification for belief*. It should be modelled on justification in virtue of intention.

Here is an example to bring out how justification in virtue of intention works.

Example: A skilled archer shoots at a target

Robin Hood intends to shoot a willow wand stuck into the ground 200 yards away. Robin's skills are such that he is easily able to hit this kind of target at this kind of distance unless conditions are abnormally windy. He takes careful aim and fires.

Consider what justifies Robin's aiming and firing the way he does. A right account must connect whatever justification Robin has with the facts that he intends to hit the target, and that, given his skill, he would be unlucky to fail in this intention, and not merely lucky to succeed. Here is one plausible proposal²⁹:

- i) Robin intends to hit the willow wand.
- ii) Because Robin is a skilled archer, this intention selects (causes in an appropriate way) activation of a means of implementation that is a reliable generator of its fulfilment.
- iii) Robin's aiming and firing the way he does are justified because they are non-luckily selected non-lucky generators of fulfilment of his intention (so he will be unlucky to fail and not merely lucky to succeed).

According to i-iii, Robin's movements as he aims and fires are justified because they are selected by his intention as a non-lucky route to its fulfilment. But this normative story incorporates a purely causal story about visuo-motor information-processing. Robin's visuo-motor system is enacting a routine of input and output (perceptual inputs x determine parameters for action y). The normative story depends on this purely causal story (it is the parameter-setting mechanism that gives Robin a reliable method of hitting the target). But the information processing story is not itself normative. It is just a story about causal input and caused output. The whole package becomes normative only when we factor in Robin's intention to hit the target. Robin intends to hit the target; this intention selects a perceptual-input-action-output routine; given the intention, the movements produced by this routine are justified because they are non-luckily selected non-lucky generators of fulfilment of the intention.

I suggest that the relation between the causal story about formation of attention-based object files (the story that gets us as far as *non-lucky correctness*) and justification for uses of perceptual demonstratives should be understood along these same lines. On its own, the story about generation of attention-based object files is merely causal. The package becomes normative when we factor in what I shall call our 'representational intentions'. The mind intends to reach out to the world in thought. The moves involved in maintenance of attention-based object-files are 'justified' in that they are non-luckily selected non-lucky generators of fulfilment of representational intention.

The rest of the section provides the main details of this view. The last section shows why I think it should be preferred to its nearest rivals.

First let me clarify the notion of 'representational intention'. There is an obvious initial objection to accounts of perceptual demonstratives that appeal to an 'intention to represent'. For it is plausible that perceptual demonstratives are 'basic' in the following sense: grasp of a perceptual demonstrative referring to o does not depend on grasp of any other conceptual representation of o , or on any conceptual specification of a domain of which o is a member³⁰. But intentions are normally taken to be propositional attitudes.

And if a representational intention is a propositional attitude, its content must be specified either by <I intend to represent...> (where the gap is filled by a conceptual representation of a thing) or <I intend to represent something (some thing outside the mind)>.

Assuming the first option, grasp of a perceptual demonstrative rests on grasp of some other conceptual representation of its referent. Assuming the second, grasp of a perceptual demonstrative rests on the capacity for conceptual specification of the domain of mind-independent things. So neither option respects the basic status of perceptual demonstratives.

However, this objection arises only on the assumption that representational intentions have propositional content. And I want to suggest that they do not. I suggest that representational intentions belong to a class of non-propositional motivational states. This is the class of basic psychological needs or compulsions. Some other members of this class (not all shared, or shared to the same extent, by all agents) are the need to avoid loneliness; the compulsion to act aggressively; the need to conform; the need not to conform; the need to feel secure.

Here are three reasons to think that these needs or compulsions do not have propositional content.

Firstly, having a need does not require grasp of the concepts that figure in a canonical account of what it is to have it. For example, you can need to feel secure without grasping the concept of security. This is a defining characteristic of non-conceptual or non-propositional content.³¹

Secondly, there is no pressure towards saying that needs have propositional content to explain how we come by them. Beliefs are often arrived at by rational deliberation. And it is hard to explain this fact unless we suppose they have propositional content. But we do not arrive at basic psychological needs as a result of rational deliberation. So the parallel pressure does not arise.

Thirdly, there is no pressure towards saying that a need has propositional content to explain what it would take to fulfil it. A desire has determinate propositional content insofar as it is a desire that some specific state of affairs (or perhaps some simple disjunction of states of affairs) come about. But a basic psychological need does not determine a specific action, event or state of affairs that would fulfil or relieve it (there is no telling in advance what will relieve X's loneliness, or what outlet Y's aggression will find).³²

The basic psychological needs I have listed so far are emotional needs: needs we have in virtue of the fact that we are emotional beings. But having recognized emotional needs, it is hard to see grounds for resisting the claim that there are also rational needs: needs we have in virtue of the fact that we are rational beings (and the having of some of which is constitutive of rationality). Some philosophers' accounts of our rational lives already appeal, explicitly or implicitly, to rational needs or compulsions: the need to understand oneself³³; the need to avoid epistemic blame³⁴; the compulsion to make basic inferential moves when presented with simple premisses³⁵. I suggest that representational intention is a basic rational need: the mind has a basic, non-propositional, need to represent things outside itself. (To avoid quibbling about whether a mental state of this kind is an intention, from now on I shall say 'need to represent' rather than 'representational intention'.)

With the notion of the basic need to represent in place, the account I want to propose of how perceptual attention enables perceptual demonstrative reference can be set out as follows.

- (i) The mind has a basic need to represent things outside itself.
- (ii) This need selects formation of attention-based object-files as a means to its fulfilment.
- (iii) The moves involved in maintaining an attention-based object-file are (weakly) justified in that they are selected as a means to fulfilment of a basic cognitive need.
- (iv) Given *non-lucky correctness*, the justification at (iii) is luck-eliminating with respect to the attended object iff it is an ordinary object.

So

- (v) Given NON-LUCKINESS, an attention-based object-file refers to the attended object iff it is an ordinary object.

But

- (vi) Most of the time if you form an attention-based object-file the attended object is an ordinary object. (I defend this claim immediately below.)

v and vi entail

- (vii) Formation of an attention-based object-file is a reliable fulfiller of representational intention.
- (viii) Given vii, we can recognise a richer sense in which the moves involved in maintaining an attention-based object-file are justified: these moves are (richly) justified because they are non-luckily selected non-lucky generators of fulfilment of the basic need to represent.

To see why (vi) should be accepted, think of what it is like keeping visual track of a ripple across the surface of a lake or a raindrop running down a window. Attention can be drawn to such things without conceptual guidance (detection of motion is a very common attentional ‘cue’). But most ripples and raindrops do not move in ways that respect stable boundaries. And ripples and raindrops routinely split apart and fuse with other ripples and raindrops. Tracking a ripple or a raindrop as it behaves like this involves a host of (little, unarticulated) decisions about where the boundaries of the thing are, or which fission product to keep attending to. These decisions involve conceptual processing. So, in most cases, attentional tracking of a ripple or raindrop (or other visual object that is not an ordinary object) is conceptually aided attentional tracking. Now suppose that your attention is caught by a visual object that it is not an ordinary object. Most of the time, the visual object will display its non-ordinary-objecthood almost

immediately. And when non-ordinary-objecthood is displayed, you will either lose track of the thing (your attention will wander elsewhere) or you will begin to devote conceptual resources to tracking it. In the first kind of case, the attentional link is too short-lived for an attention-based object-file to be formed. In the second, the resulting object file is not an attention-based object file. For, by definition, an attention-based object file inherits its contents from the information delivered by a conceptually unaided attentional link.

i-viii give us what the GROUNDING + NON-LUCKINESS framework from §1 requires: an account of how perceptual attention secures justification for basic uses of a perceptual demonstrative that is luck-eliminating with respect to the attended object.

§4 Good consequences

I shall close with two points of comparison between the account I have proposed and nearby alternatives. The comparisons concern the elements of the account that I think others are most likely to find problematic: the appeal to representational needs, and the treatment of attention to visual objects other than ordinary objects.

§4.1 Classical normativity, practical normativity, and selection without awareness

Consider the contrast between justification for belief and justification in virtue of intention:

Justification for belief – A belief is justified iff there is an appropriate dependence of formation of the belief on whether the state of affairs required for its truth obtains.

Justification in virtue of intention – A behaviour is justified iff it depends in an appropriate way on an intention (it is a non-luckily selected non-lucky generator of the intention's fulfilment).

Both kinds of justification involve a non-lucky relation between a mental state and a way the world is. But these relations have opposite directions of orientation. In the belief case, the orientation is, as I shall say, 'classical': it runs from a way the world might be (namely, such that the belief's truthmaker obtains) to the formation of the mental state. In the intention case the relation is 'practical': it runs from the mental state (the intention) to the state of affairs that fulfils it.³⁶

I have suggested that the justification for basic uses of perceptual demonstrative <that> secured by perceptual attention has practical orientation: basic deployments of a perceptual demonstrative are canonically justified in virtue of the fact that they are selected by the basic need to represent as a means to its fulfilment.

Though it is not possible to attempt a survey here, I think it is fair to say that every other extant view of perceptual demonstratives that respects the kind of concerns about normativity made explicit by NON-LUCKINESS favours a classical orientation. For example, consider the following recent proposal.³⁷

Suppose you are attending to an ordinary object in front of you. Your attentional link with the thing provides a stream of information derived from it. Because you are receiving this information through an attentional channel, you respond by forming

perceptual demonstrative beliefs. Your attentional channel provides input information about location, shape, and so on. You respond by forming beliefs like <That is about a foot in front of me> and <That is spherical>. You might also respond by forming a perceptual demonstrative intention (<I am going to pick that up>). If you do form such an intention, your attentional link determines how you will reach out in the attempt to execute it. So your attentional link underpins a basic pattern of informational input to cognition, and output from cognition to action. This ‘underpinning’ is both causal and normative. It is causal because your attentional link causes you to form the beliefs you do, and combines with your intention to cause you to reach out distance x with your grip adjusted in manner y . It is normative because your attentional link determines standards of right functioning for your input-to-cognition and output-to-action information processing. If you are attending to o , right input-to-cognition information processing results in accurate beliefs about what o is like, and right output-to-action information processing generates actions on o . (To attend to o just is to stand in a relation to o that generates these standards of right functioning.) So perceptual attention provides normative underpinning for the input-to-cognition/output-to-action pattern that it causes: it provides standards for assessing this information processing pattern as going right or wrong.³⁸ (Recall the telescope example from §1. The fact that a telescope is focused on o provides a standard for assessing whether it is delivering accurate information: it is delivering accurate information iff it is getting o ’s properties right. If the telescope is not focused, there is no standard for assessing the accuracy of the information it delivers. This view treats the role of perceptual attention in securing perceptual demonstrative reference as a parallel to the role of ‘focus’ in securing the ‘aboutness’ of beliefs formed on the basis of information delivered through a telescope.)

Though these are not the terms in which it was initially stated, this proposal fits the GROUNDING + NON-LUCKINESS framework. For the proposed account of why attention suffices to fix reference – attention plays a role parallel to the role played by focus in the telescope case – also entails that attention secures justification that is luck-eliminating with respect to the attended object. This justification has a classical direction of orientation because justification in virtue of focus is justification in virtue of tending towards getting an object’s properties right.

Why prefer the (practical) view I have proposed to the ‘focus’ view and other classical accounts? A full discussion of this question is not possible here. But I want to suggest one argument in favour of the practical view: it offers a better account of the status of merely selective perceptual processing. (Recall that attentive processing is both selective and awareness-generating. Merely selective processing shares the structure of attentive processing but does not generate awareness.)

To see how the practical view treats merely selective processing, it is helpful to add another layer to the parallel between essaying perceptual demonstrative reference and shooting at a target. Consider the following example

Example: Robin Hood shoots under hypnosis

Robin has been hypnotized to shoot at something red every time he hears a whistle. Robin’s skills as an archer remain intact. So each time he hears a whistle he shoots an arrow, and this arrow usually hits a red thing.

On the face of things, Robin's aiming-and-firing motions here are not justified. The story about why he moves as he does is a mere causal story: the hypnosis and the whistle cause activation of a perceptual-input/action-output routine. This is the same information processing routine as is selected when Robin shoots intentionally. But in the hypnotized case the routine is not embedded in a wider normativity-conferring setting: Robin's movements are not justified because they are not part of a causal routine that has been selected as a non-lucky generator of fulfilment of his intention.³⁹

On the practical view, the reason perceptual attention suffices for perceptual demonstrative reference while mere selective processing does not should be understood along the same lines. The causal routine for formation of attention-based object files may be intact in a mere selective case. But this causal routine has no normative status on its own. It gains normative status when it is harnessed by representational intention as a means to its fulfilment. Attention is essential to canonical justification for basic deployments of perceptual demonstratives because mere selective information processing (as opposed to attentive information processing) does not deliver awareness, so is not harnessable by representational intention.⁴⁰

Now compare this with the account of the status of merely selective processing generated by the classical 'focus' view. On this view, perceptual attention suffices to fix perceptual demonstrative reference because an attentional link is a focused link. So someone taking this view who wants to uphold GROUNDING must maintain that perceptual information processing that does not generate awareness is not genuinely focused. For example, you might argue as follows. Perceptual attention to *o* generates a rich range of capacities with respect to *o*: capacities to report *o*'s (apparent) properties; to try to find out what *o* is like; and to act on *o*. In contrast, a merely selective perceptual link, like the link a blindsight subject has with an object in the blindfield, generates a very impoverished range of capacities. But it is the richness found in the case of perceptual attention that warrants the claim that perceptually attending to *o* is 'focusing' on *o*, so that perceptual attention to *o* can be said to set standards of correctness for the information processing associated with it. Given the comparatively tiny range of capacities generated by the blindsighter's perceptual contact with an object in the blindfield, there are no grounds for saying that the blindsighter has a perceptual link that is 'focused on' the object: this would be like saying that a photo-electric cell is 'focused on' a light-source that activates it.⁴¹

One problem with this move is that it is vulnerable to an as yet unexcluded empirical possibility: there might turn out to be a form of merely selective processing that mirrors the richness in capacities generated by perceptual attention.

A second problem is that many cases of attentive processing do not generate richer capacities with respect to an object than mere selective processing might. Suppose you are attending to a thing visible merely as a speck in the sky. In this case, your capacities to make justified reports on the basis of your attentional link are very limited. You might be able to report only the thing's apparent location – something a blindsight subject can also do. Similarly, your attentional link does not put you in a position to try to find out what the thing is like. And it does not enable you to act on the thing in any sense in which a blindsight-link cannot: you can point at the object, or walk towards it, but so can a blindsight subject. So the suggestion would have to be that perceptual attention to *o* counts as focus on *o* in virtue of the fact that attention to a thing 'potentially' generates a

rich range of capacities. But now there is a hard problem about why a mere selective link does not ‘potentially’ generate a rich range of capacities too.

I do not want to deny that an advocate of the ‘focus’ model might find a more plausible account of why perceptual processing that does not generate awareness is not really focussed.⁴² Rather, my point is that on the practical view the need for such an account does not arise.

To bring out the importance of this point, let us go back to the claim that motivated GROUNDING in §1.3:

(i) The intentionality characteristic of thought (conceptual intentionality) is essentially a phenomenon of subjective consciousness.

It is an old observation⁴³ that this claim is consistent with

(ii) Intentionality has a physical basis (a mental state with intentional content just is a physical state with intentional content).

All that is required to maintain both (i) and (ii) is

(iii) Any physical state that has conceptual content is a subjectively conscious state.

Now suppose

(iv) Perceptual contact with objects and properties in the world secures conceptual content because it secures justification for deployment of perception-based concepts. (NON-LUCKINESS is a special case of this claim.)

Finally, suppose a classical direction of orientation for this justification:

(v) Perception secures justification for deployments of perception-based concepts because it generates appropriate dependence of the way these concepts are deployed on what the world is like.

(v) and (iii) entail

(vi) A perceptual link that secures ‘appropriate dependence’ of deployments of perception-based concepts on what the world is like also generates awareness.

So to keep (i), (ii), (iv) and the classical orientation (v) you must find a difference between the ‘appropriate dependence’ at (vi), and whatever dependence merely selective processing might secure for tokenings of sub-personal states. And it is going to be hard to find this difference. For it is hard to see what might count as a ‘way of securing dependence’ that might not be present at the sub-personal level.

The practical view I have proposed offers a different and, I think, better account of why merely selective processing cannot secure conceptual intentionality. On this view, there is no need to look to structural differences between perceptual processing with

awareness and perceptual processing without for an explanation of the normative difference between perceptual attention and merely selective perceptual contact. Rather, the difference lies with whether the information processing involved in the perceptual link can be harnessed to the mind's basic need to represent.

§4.2 Perceptual demonstrative reference failure and perception as the source of the basic subject matter of thought

The second point of comparison with nearby proposals that I shall consider concerns essayed perceptual demonstrative reference where the attended visual object is not an ordinary object but a ripple, a shadow, a mereological griffin consisting in two ordinary objects that the visual system is treating as one, or some other spurious 'thing'.

The practical view entails that in such cases reference fails. Essaying perceptual demonstrative reference in these cases is like using a railway schedule to construct an account of a journey given snippets of information, when in fact the person drove. In this case, your justification for your operations on the incoming snippets is not luck-eliminating with respect to the journey. Similarly, on the practical view, when you essay perceptual demonstrative reference to a visual object that is not an ordinary object, your file-maintaining moves are justified (because they are selected by representational need as a generator of its fulfilment). But this justification is not luck-eliminating with respect to the attended object. So (given NON-LUCKINESS) your attempt to refer does not succeed.

Of course, we do refer to ripples, shadows, and so on on the basis of perceptual links. But on the practical view this is not (pure) perceptual demonstrative reference. It is conceptually aided or 'complex' demonstrative reference based on perception. I have argued that almost all attentional tracking of non-ordinary-objects is conceptually aided. And an object-file formed on the basis of a conceptually aided attentional link is not an attention-based object file. It is a conceptual representation whose basic deployments are justified partly by the relevant conceptual processing, and which has its own reference-fixing story (which I cannot consider here). It is only where your visual system is taken in by a non-ordinary object behaving with enough ordinary-object-like coherence for conceptually unaided tracking that you will essay a pure perceptual demonstrative that fails to refer.

There is an old⁴⁴ and intuitive model of the importance of perceptual demonstrative reference against the background of which this consequence about reference failure emerges as a virtue. To motivate the model, consider the following argument.

(i) Our grasp of what it takes for a general proposition to be true is derived from our grasp of what is required for the truth of its potential particular instances: knowledge what is required for $\langle \text{Something is } \Phi \rangle$ to be true rests on knowledge what is required for the truth of $\langle v \text{ is } \Phi \rangle$.

(ii) Chains of 'derivation' of the kind mentioned in (i) must come to an end in a range of 'fundamental' propositions, where if p is a proposition in this range, knowledge what it takes for p to be true is not derived from knowledge what it takes for any other proposition to be true. (Otherwise our knowledge what it takes for propositions to be true

will never bottom out in knowing that p is true iff the world is such-and-such way. And in that case our thought will have no subject matter.)

But (i) entails that these fundamental propositions cannot be general propositions. So

(iii) There must be a fundamental range of particular propositions: a range of propositions of form $\langle v \text{ is } \Phi \rangle$ such that grasp of v does not require grasp of any other conceptual representation of v 's referent, and does not require grasp of any general proposition.

This argument leaves us with three options: reject (i); reject (ii); or identify the fundamental range of propositions and explain how we are able to grasp them. The only serious contender with respect to the third option is the claim that the fundamental range of particular propositions is the range of propositions grasp of which would involve thinking a perceptual demonstrative thought. Pursuing this option, we get

(iv.a) For perceptual demonstrative $\langle \text{that} \rangle$, and observational concept $\langle \Phi \rangle$, grasp of $\langle \text{That is } \Phi \rangle$ is basic: it does not rest on grasp of any more fundamental way of identifying the referent of $\langle \text{that} \rangle$, or on grasp of any general proposition.

(iv.b) Grasp of the simplest propositions of form $\langle \text{Something is } \Phi \rangle$ consists in knowing that $\langle \text{Something is } \Phi \rangle$ is true iff $\langle \text{That is } \Phi \rangle$ is true for some potential instance of perceptual demonstrative $\langle \text{that} \rangle$.

(iv.a) and (iv.b) state the old model of the importance of perceptual demonstratives: perceptual demonstratives provide the basic subject matter of thought.

Now add a consequence of the classical 'focus' view:

(v) The reach of potential demonstrative reference is the reach of potential attentional tracking (so that if o can draw and hold attention, o can be the object of perceptual demonstrative reference).

Finally, add an apparent consequence of the empirical data from §2:

(vi) We can keep attentional track of any 'thing' that appears to behave with ordinary-object-like coherence.

Given (v) and (vi), the class of potential objects of demonstrative reference includes any 'thing' that can appear to behave with ordinary-object-like coherence for long enough to engage attentional tracking. This class includes at least the following: ordinary objects, shadows; ripples; reflections; dots or marks on surfaces; parts of ordinary objects; parts of ripples and reflections; mereological griffins. But according to (iv.b) our grasp of a general proposition like $\langle \text{There is something square in this room} \rangle$ is to be explained in terms of our grasp of its potential perceptual demonstrative instances. So (v) and (vi) entail that it is far too easy for $\langle \text{There is something square in the room} \rangle$ to be true. All it takes is for there to be some perspective from which some square visual object can grab attention. For example, $\langle \text{There is something square in the room} \rangle$ will come out true if no

ordinary object is square, but there is some perspective from which a subject would be able to attend to a square mereological griffin.

To retain (i) and (ii) while avoiding this absurd result we must reject either (v) or (vi).

The classical ‘focus’ view entails (v). So a proponent of this view who wants to keep (i) and (ii) must reject (vi). But it is hard to see how this move could be sustained. To reject (vi) you must deny that we really do keep attentional track of non-ordinary objects (so that when your visual system is locked to a ‘thing’ that is in fact a shadow or a ripple, the information processing story is different from the story for a visual lock on an ordinary object). But our perceptual systems are sensitive to how objects behave in the actual world, not how they behave in nearby possible worlds. So as long as a ‘thing’ is in fact behaving with ordinary-object-like coherence, it will be treated by the perceptual system as an ordinary object. There is no plausibility to the claim that two processes (attentional tracking of ordinary objects on the one hand, whatever you want to call the process on the other) lock onto the same kind of coherence in the same ways.⁴⁵

In contrast, the view I have proposed entails that if (vi) is true (v) is false. So it allows us to reclaim the old model of perceptual demonstrative reference as providing the basic subject matter of thought: grasp of perceptual demonstrative <that> does not rest on grasp of any other conceptual representation of the thing you are thinking about; the domain of our most basic quantifiers is defined by the reach of potential perceptual demonstrative reference.

Anscombe, G. E. M. (2000) *Intention*. Cambridge, MA: Harvard University Press 2000.

Ayers, Michael. (1991) ‘Substance: Prolegomena to a Realist Theory of Identity’ *Journal of Philosophy*.

Boghossian, Paul. (2008) *Content and Justification*. Oxford: Oxford University Press.

Boghossian, Paul. (2008a) ‘Is Meaning Normative?’ in Boghossian 2008.

Boghossian, Paul. (2008b) ‘Blind Reasoning’ in Boghossian 2008.

Brewer, Bill. (1999) *Perception and Reason*. Oxford: Oxford University Press.

Campbell, John. (2002) *Reference and Consciousness*. Oxford: Oxford University Press.

Dehaene, Stanislas, Jean-Pierre Changeux, Lionel Naccache, Jerome Sackur, and Claire Sergent. (2006) ‘Conscious, preconscious, and subliminal processing: a testable taxonomy’. *Trends in Cognitive Science*.

Dickie, Imogen. (2010) ‘We Are Acquainted With Ordinary Things’ in Robin Jeshion (ed.) *New Essays on Singular Thought*. Oxford: Oxford University Press.

Dienes, Zoltan and Josef Perner. (2007) ‘Executive Control without Conscious Awareness: the Cold Control Theory of Hypnosis’. In Graham Jamieson (ed.) *Hypnosis and Conscious States: the Cognitive Neuroscience Perspective*. Oxford: Oxford University Press.

Evans, Gareth. (1982) *The Varieties of Reference* Oxford: Oxford University Press.

Fine, Kit. (2007) *Semantic Relationalism*. Oxford: Blackwell.

- Fodor, Jerry. (1990) 'A Theory of Content' in Fodor *A Theory of Content and Other Essays*. Cambridge MA: MIT Press.
- Hiller, Avram and Ram Neta. (2007) 'Safety and Epistemic Luck' *Synthese*.
- Jeshion, Robin. (2001) 'Acquaintanceless *De Re* Belief' in J. Campbell, M. O'Rourke and D. Shier (eds.) *Topics in Contemporary Philosophy: Truth and Meaning*. New York: Seven Bridges Press.
- Lawlor, Krista. (2001) *New Thoughts about Old Things: Cognitive Policies as the Ground of Singular Concepts*. New York: Taylor and Francis.
- Lewis, David. (1996) 'Elusive Knowledge'. *Australasian Journal of Philosophy*.
- Nozick, Robert. (1981) *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
- Palmer, Stephen E. (1999) *Vision Science: Photons to Phenomenology*. Cambridge, MA: MIT Press.
- Peacocke, Christopher. (1992) *A Study of Concepts*. Cambridge, MA: MIT Press.
- Pritchard, Duncan. (2007) 'Anti-Luck Epistemology' *Synthese*.
- Pylyshyn, Zenon. (2003) *Seeing and Visualizing*. Cambridge MA: MIT Press.
- Pylyshyn, Zenon. (2007) *Things and Places: How the Mind Connects with the World*. Cambridge MA: MIT Press.
- Recanati, Francois. (1993) *Direct Reference*. Oxford: Blackwell.
- Russell, Bertrand. (1956) 'The Philosophy of Logical Atomism' in R. C. Marsh (ed.) *Logic and Knowledge*. London: Allen and Unwin.
- Scholl, Brian.(ed.) (2001) *Objects and Attention*. Cambridge MA: MIT Press.
- Scholl, Brian. (2001a) 'Objects and Attention: the State of the Art' in Scholl 2001.
- Scholl, Brian, Zenon Pylyshyn and Jacob Feldman (2001) 'What is a Visual Object? Evidence from Target Merging in Multiple Object Tracking' in Scholl 2001.
- Scholl, Brian. (2007) 'What have we learned about attention from multiple object tracking (and vice versa)?' in D. Dedrick and L. Trick (eds.) *Computation, Cognition, and Pylyshyn*. Cambridge MA: MIT Press.
- Searle, John. (1980) 'Minds, Brains, and Programs' *Behavioral and Brain Sciences*.
- Searle, John. (1994) *The Rediscovery of the Mind*. Cambridge MA: MIT Press.
- Siegel, Susanna. (2004) Review of John Campbell *Reference and Consciousness* in *Philosophical Review* vol. 113 no. 3.
- Smithies, Declan. 'What is the Role of Consciousness in Demonstrative Thought?' forthcoming, *Journal of Philosophy*.
- Velleman, J. David. (2000) *The Possibility of Practical Reason*. Oxford: Oxford University Press.
- Velleman, J. David. (2000a) 'The Aim of Belief' in Velleman 2000.

Velleman, J. David. (2006) 'Motivation by Ideal' in Velleman *Self to Self*. New York: Cambridge University Press.

Velleman, J. David. (2007a) *Practical Reflection* David Hume Series Edition. Stanford: CSLI.

Velleman, J. David. (2007b) 'What Good is a Will?' in Anton Leist (ed.) *Action in Context*. Berlin: De Gruyter.

Weiskrantz, Lawrence. (2009) *Blindsight – A Case Study Spanning 35 Years and New Developments*. Oxford: Oxford University Press.

Wittgenstein, Ludwig. (1961) *Tractatus Logico-Philosophicus* trans. D.F. Pears and B. F. McGuinness. London: Routledge and Kegan Paul. 1961.

* This paper is descended distantly from talks at Harvard, Barcelona, the Ohio State Dubrovnik Conference, and Warwick, and more closely from talks at Columbia (with a half-CUNY audience) and UMass Amherst. Thanks to all members of the audiences on those occasions. Particular thanks to Chris Hill and Barry Stroud, who commented at Harvard and Warwick respectively. For conversations that improved the paper thanks to Louise Anthony; Ned Block; Bill Brewer; John Campbell; Gary Ebbs; Benj Hellie; Robin Jeshion; Hilary Kornblith; Hakwan Lau; Joe Levine; Chris Meacham; David Rosenthal; James Shaw; Susanna Siegel; Matt Soteriou; David Velleman; Jay Wallace; and Sebastian Watzl. Finally, many thanks to the editors of this volume for two rounds of extremely helpful comments.

¹ '< v is Φ >' abbreviates 'The proposition constructed from v and Φ '.

² Peacocke 1992 ch. 1; Campbell 2002 ch. 5 §5.

³ Compare Pritchard 2007 §2; Hiller and Neta 2007 §1.

⁴ This parallels Lewis's response to scepticism in his 1996 (see esp. 559). But Lewis's concern is knowledge, not rational entitlement. So he is dealing with a different notion of 'relevance'. For example the actual situation is always relevant to knowledge, but is sometimes irrelevant to rational entitlement (otherwise there could be no unluckily false beliefs).

⁵ I am supposing that justification can 'exclude' only situations that are potentially relevant to rational entitlement. So my rational entitlement generating justification for believing <Jack is not tall> is luck-eliminating with respect to Jack, but not with respect to the number 7: given my justification, there are no unexcluded relevant situations in which 7 is tall, but my justification does not exclude any such situations because there were none to begin with.

⁶ This classification of perceptual processes into attentive, merely selective, and merely awareness-generating owes its structure to Dehaene et. al. 2006. Some philosophers and psychologists use 'perceptual attention' more broadly, counting all selective perceptual processing as attentive processing. Nothing in this paper depends on how this classificatory matter should be resolved.

⁷ Compare Boghossian 2008a 99-101.

⁸ This distinction derives from Boghossian 2008a 101.

⁹ For the same conclusion established by different arguments, see Boghossian 2008a 101; Velleman 2000a 245-246, 252-255.

¹⁰ This is not to deny that there can be informative identities involving perceptual demonstrative <that>. The point is just that in an ordinary case where you use two tokens of <that> to refer to a thing you are keeping track of in perception the identity is not informative. I discuss informative identities involving perceptual demonstratives in my 2010 §3.2.

¹¹ So-called by Evans 1982 150-151.

¹² See Palmer 1999 §11.2.6 and compare Campbell 2002 31-33.

¹³ See Palmer, 1999 §13.3.2 for a textbook account and references.

¹⁴ Weiskrantz 2009.

¹⁵ Searle 1994 contains an extended defence of a stronger version of (i): the claim that all intentionality is a phenomenon of subjective consciousness. See especially 156-159. See also Searle 1980. Unlike Searle, I allow that sub-personal states might have a kind of intentionality distinct from the intentionality characteristic of thought. The (i)-(iii) argument is indebted to Campbell 2002 7-10, 142-145.

¹⁶ For this definition of 'naturalism' see Searle 1994 49-52; Fodor 1990 51-52 and note 3.

¹⁷ Fodor 1990.

¹⁸ Kornblith 2002.

¹⁹ See for example Palmer 1999 554.

²⁰ Pylyshyn 2003 §5.3.1, 2007 34-58; Scholl 2001a §2.5, 2007.

²¹ Scholl 2001a §6, 2007 §4; Scholl et. al. 2001.

²² Scholl 2001a 10 gives a summary and references.

²³ For a basic account of this topic see Palmer 1999 ch. 7. In the terms Palmer uses at 313, I am using 'features detected' to describe 'proximal' aspects of perceptual experience and 'property information delivered' to describe 'distal' aspects.

²⁴ Compare Ayers 1991. The proposal that follows in fact provides a way to refine the account of the category of ordinary objects. I discuss this further in my 2010 §2.3.

²⁵ For recent uses of the 'files' framework see Evans 1982; Recanati 1993; Lawlor 2001; Jeshion 2001. Nothing I say here supposes that the notion of a file is prior in order of explanation to the notion of grasp of coreference (see Fine 2007 67-68 for an objection on this score; Lawlor 2001 63, 79-80 for what I take to be the right reply).

²⁶ The proposal I want to make can accommodate some adjustments with respect to detail. For example, nothing I say here depends on whether shape features detected are three-dimensional or (as I suggest in row 1 of the table) two-dimensional. And nothing depends on exactly how much of the contents of an attention-based object file at t is incorporated into the file at $t + 1$.

²⁷ For similar claims see Campbell 2002 §5.7; Evans 1982 179.

²⁸ A counterfactual model of justification combines Nozick's account of knowledge as true belief that tracks the truth (1981 172-178) with the claim that the justification is whatever a true belief must have to count as knowledge.

²⁹ This is modelled on Anscombe's 2000 account of the relation between an intention and its fulfilment. Anscombe says that if I intend to Ψ I believe that I will Ψ , and if I actually do Ψ my intention counts as knowledge iff it is appropriately connected to the action that fulfils it. But she distinguishes 'practical' from 'speculative' knowledge. A true belief

counts as ‘speculative’ knowledge iff it depends appropriately on its truthmaker; a fulfilled intention counts as ‘practical’ knowledge iff it plays an appropriate role in bringing its truthmaker about: ‘Practical knowledge is “the cause of what it understands”, unlike “speculative” knowledge, which is “derived from the objects known”’(87). See also 51-53; 57. This reading of Anscombe is indebted to Velleman 2007a xxi-xxv and 2007b .

³⁰ I argue for this ‘basicness’ claim in §4.2. The claim also follows from the ‘no extra relation’ datum from §1.3.

³¹ See, for example, Peacocke 1992 ch. 3.

³² Compare Velleman 2006 315-316.

³³ Velleman 2000 20-24; 2007a xx.

³⁴ Boghossian says that the most basic inferences transmit warrant in virtue of instantiating epistemically ‘blameless’ patterns (2008b esp. 278-279). This account requires that basic inference is not itself a matter of thinking ‘Pattern X is blameless; this step is in accord with pattern X; so I may take this step’.

³⁵ Peacocke 1992 6-8.

³⁶ Compare Velleman 2007 xii-xv.

³⁷ This is Campbell’s view in his 2002. For other classical proposals see Brewer 1999; Smithies *forthcoming*.

³⁸ See Campbell 2002 13-19 for the initial claim about causal underpinning; 3, 16, 26, 34-38, 41-43, 55-56, 87-88 for the claim that conscious attention defines the target of the perceptual information-processing it causes; 22-26, 84-90 for the relation between causal and normative underpinning and reference.

³⁹ This is in step with ‘cold control’ accounts of hypnosis, according to which a hypnotic ‘trigger’ activates sub-personal routines without mediation by personal level mechanisms for intentional action. See Dienes and Perner 2007.

⁴⁰ The question of the exact relation between awareness and harnessability by representational intention is a topic for future work.

⁴¹ Compare Campbell 2002 10: ‘...experiential highlighting...affects the functional role of ... experience of [an object]...It means that you are in a position to keep track of [it] deliberately over time, you are in a position to answer questions about [it] on the basis of vision, and you are now able to act with respect to [it]’; similar at 3, 19, 28, 42-43. Note that this is a point about difference in richness of capacities, not difference in richness of information delivered: 9-10. For the claim that, there is selection without attention, but it cannot underpin demonstrative reference see 4, 31-34, 51-57, 142-145.

⁴² For example, Campbell says that attention but not mere selection enables ‘intentional interrogation’ of the environment (3, 27, 33-34, 89); that attention but not mere selection enables intentional action on an object (27, 48-49, 89); and that attention but not mere selection makes available the ‘categorical thing itself’ (10, ch. 7 esp. 137-145, ch. 8 esp. 250-254). Though it is not possible to argue this point here, I think these differences are part of what an account of why mere selection cannot secure perceptual demonstrative reference must explain, rather than resources for the explanation. Objections to Campbell on the ground that merely selective processing can be focussed are also raised by Siegel 2004 429; Smithies *forthcoming* §2.

⁴³ Searle 1994 especially 14-15, 28-29, 54-55, 89-93.

⁴⁴ I suggest that the best interpretation of *Tractatus* 3.2-3.261 is as providing something like the argument for iii below. Russell states a similar argument at xiii of his introduction to the *Tractatus*, and the argument's ingredients occur throughout his logical atomist writings (see Russell 1956 230 for a version of (i); 197 for a version of (ii)).

⁴⁵ Campbell proposes that different 'styles' of attention (involving different ways of binding features) enable reference to different categories of object. For example, he says that you are attending in different styles when you attend to a person, a cloud or a valley. (2002 ch. 4 especially 61-63; 70; 74-75; 82). In these terms, the current problem concerns cases where a thing's behaviour enables you to attend to it in a style inappropriate to its category.