JOURNAL PRE-PROOF

From immediate perception to perceptual belief

Derek H. Brown

DOI: 10.1387/theoria.26036

Received: 14/02/2024

Final version: 17/06/2025

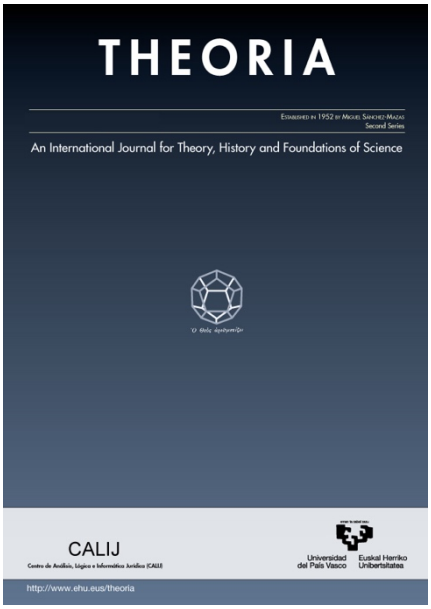

FROM IMMEDIATE PERCEPTION TO PERCEPTUAL BELIEF

(De la percepción inmediata a la creencia perceptiva)

Derek H. Brown
University of Glasgow

ABSTRACT: Pretheoretically, our senses give us immediate access to the world around us. Perceptual demonstratives (e.g. uttering "that is a bottle" while pointing at a bottle) seem to directly refer to their intended referents. Should we explain such referential directness by perceptual immediacy? Bermúdez (2000) argues that we should, and offers a powerful theory of immediate and mediate perception (Naturalized sense-datum theory) to explain how. He also argues that a different but influential theory of immediate and mediate perception (that found in traditional sense-datum theory) cannot explain referential directness by perceptual immediacy. I argue that his criticisms fail in instructive ways. Following this I assess the extent to which Bermúdez's theory can explain a critical perceptual challenge, that of explaining the experiential similarities between perceptual experiences with varying degrees of veridicality. I argue that Bermúdez's theory can only meet this challenge by embracing a veil of perception akin to that of his opponent – something he sought to avoid. Before concluding, I present the underlying framework of this debate devoid from any appeal to sense-datum theories. The reason is because that framework has considerable power and relevance to a number of contemporary issues in philosophy of perception.

Keywords: direct perception; indirect perception; immediate perception; mediate perception; demonstratives; perceptual demonstratives; deferred demonstration; ostension; perceptual belief; sense-datum theory; naturalized sense-datum theory.

RESUMEN: Preteóricamente, nuestros sentidos nos proporcionan acceso inmediato al mundo que nos rodea. Los demostrativos perceptuales (e.g. proferir "Eso es una botella" mientras se señala una botella) parecen referir directamente a sus referentes. ¿Deberíamos explicar este carácter directo de su referencia por su inmediatez perceptiva? Bermúdez (2011) defiende que sí, y ofrece una poderosa teoría de la percepción inmediata y mediata para explicar cómo (la teoría de los datos sensoriales naturalizados). También sostiene que otra teoría distinta, pero influyente, de la percepción inmediata y mediata (la que se encuentra en la teoría tradicional de datos sensoriales) no puede explicar el carácter directo de la referencia por su inmediatez perceptiva. En este trabajo, argumento que sus críticas fallan de manera instructiva. A continuación, evalúo en qué medida la teoría de Bermúdez puede explicar un desafío perceptual crítico: explicar las semejanzas experienciales entre experiencias perceptivas con diferentes grados de veracidad. Argumento que la teoría de Bermúdez solo puede abordar este desafía adoptando un velo de la percepción similar al de su oponente – algo que

intentaba evitar. Antes de concluir, presento el marco subyacente a este debate, desprovisto de menciones a las teorías de datos sensoriales. La razón para esto es porque este marco tiene un poder considerable y resulta relevante para una serie de cuestiones contemporáneas en filosofía de la percepción.

Palabras clave: percepción directa; percepción indirecta; percepción inmediata; percepción mediata; demostrativos; demostrativos perceptuales; demostración diferida; ostensión; creencia perceptiva; teoría de datos sensoriales; teoría naturalizada de datos sensoriales.

SHORT SUMMARY: Bermúdez (2000) offers a theory that explains referential directness via perceptual immediacy, and argues that an opposing theory cannot do the same and thus should be rejected. I reply to his criticism, and in turn argue that his theory can only explain The Problem of Perception via postulation a veil of perception, something he sought to avoid.

How direct or immediate is our access to the physical world? It seems undeniable that perception is our most basic epistemic and informational route to the physical world. Thoughts and speech can be about the physical world, but when they are, their accuracy relies on what we see, hear, and touch in the world. Perception is our most basic route. How direct or immediate is our perceptual access to the physical world? Twentieth century was largely dominated by a view, sense-datum theory, according to which our perceptual access to the world is mediated by the perception of non-physical mind-dependent objects. Many today find this unpalatable. This also seems to conflict with a kind of directness in thought and speech that we arguably can have about physical objects. When I look at a bottle and think or say 'That is a bottle', 'that' seems to directly refer to that object. How can such referential directness be explained, if perception of the physical world is mediated? We can turn this question on its head: ideally, shouldn't we explain such referential directness by our ability to immediately perceive the physical world, particularly given that perception affords our most basic access to that world? If so, what might such an account look like?

That is the topic of Bermúdez (2000).[1] His proposed solution posits a tight explanatory connection between perceptual immediacy and direct demonstrative thought or belief. However, he argues, while we cannot immediately perceive physical objects, we can immediately perceive their front-facing surfaces. He calls the view *Naturalized sense datum* theory (NSD). My interest in this topic is twofold. First, I think the issues just outline are foundational, and have not received enough attention in recent years. I thus wish to demonstrate their contemporary interest. Second, while there are many merits to Bermúdez's account, I think his argument fails, and that NSD is ultimately forced to accept a kind of veil of perception that he sought to avoid. Thus, overall, I defend a veil of perception, but in a way that explains how we can directly refer to physical objects via perceptual demonstratives.

[1] All references to Bermúdez are to his (2000) unless otherwise specified.

I proceed as follows. I first describe and defend the proposed connection between immediate perception and direct demonstrative refence (section 1) and then outline Bermúdez's overall argument for NSD (section 2). Following this, I explain and motivate Bermúdez's NSD (section 3), and why he thinks traditional sense datum theory (SDT) cannot explain direct demonstrative reference (section 4). I then respond to his critique of SDT (section 5). That completes my critique of Bermúdez's overall argument, the upshot of which is that there are different, viable conceptions of immediate perception that can be used to explain how we directly demonstrate physical objects. I take this to be of value on its own. In the rest of this paper, I examine the extent to which NSD can address The Problem of Perception. The reason is because a balanced debate between NSD and SDT theorists requires consideration, not only of a NSD-friendly phenomenon like perceptual demonstratives, but also of a SDT-friendly one like The Problem of Perception.

The Problem of Perception is the challenge of explaining the differences between veridical, illusory, and hallucinatory perceptions while respecting the fact that these disparate types of perceptions can be subjectively indistinguishable. This is a foundational perceptual problem in its own right. It is particularly important in a discussion about perceptual immediacy and different versions of sense datum theory, because sense datum theory is (rightly) regarded as offering a powerful, though controversial, solution to this problem. I argue, however, that NSD cannot solve The Problem of Perception, and that The Problem creates an internal tension for NSD that is difficult to resolve (section 6). I then consider a resolution by adding nonconceptual content to NSD, nonconceptual content being something that Bermúdez (alongside others) has done much to articulate and defend (section 7). Bringing nonconceptual content to bear on these issues is also of interest because nonconceptual content has many current advocates. I argue that this addition can indeed solve The Problem of Perception, but at the cost of reverting to a kind of perceptual veil that NSD sought to avoid. The upshot of this part of the paper is therefore that NSD cannot, in the end, avoid a veil of perception. I briefly connect this discussion to the recent interest in naïve realism, which offers yet another conception of immediate perception (section 8). In the penultimate section (section 9), I lay bare the underlying framework of this discussion, divorced from the constraints of sense-datum theories (SDT and NSD). I indicate how this framework has numerous unexplored applications, and then offer a brief conclusion.

What I hope to convince readers of is the importance of examining perceptual immediacy and perceptual mediation, both on their own and in relation to directness in thought and speech. Bermúdez's Naturalized sense datum theory provides a fascinating means of facilitating this discussion. Following Bermúdez, I focus on visual perception throughout.

1. The Reference Constraint

The notion of "indirect" or "mediated" perception has a long, distinguished history, with numerous philosophers arguing that we "directly" or "immediately" perceive ideas or sensations and "indirectly" or "mediately" perceive the surrounding physical world. Following Bermúdez (see also

e.g. Jackson, 1977; Snowdon, 1992), I distinguish between an epistemic and a perceptual version of this distinction. To start:

> **Perceptual distinction**: When one perceives something O by virtue of perceiving something else, one mediately perceives O. When a perception of a thing isn't mediated, it is immediate.
>
> **Epistemological distinction**: Beliefs about perceptions that are fully justified by those perceptions are direct perceptual beliefs, and those that are not indirect perceptual beliefs.[2]

The perceptual distinction applies to perceptual states on their own, and the epistemological distinction is centred on the kinds of justification that perceptual states can provide for perceptual beliefs. The former is about the structure or metaphysics of perceptual states themselves, and the latter is about how perceptual states feed into associated epistemic states.[3] The kinds of perceptual states I focus on are perceptual experiences (or instances of perceptual consciousness), as I take these to be what philosophers of perception are most often concerned with. Let me first elaborate on the epistemological distinction.

One might argue that if I have a perceptual experience as of a black cup of coffee before me, that experience is not sufficient to justify the belief that there is a black cup of coffee before me. This is because the experience could be veridical but could also be illusory or hallucinatory. If one accepts this line of thought, then it seems that perceptual experience alone is not sufficient to justify perceptual beliefs about ordinary physical objects. From here one might, for example, argue that the experience does justify the belief that there *seems* or *appears* to be a black of coffee before me. This latter belief would then be a direct (or basic) perceptual belief, and the former belief (that there is a black of coffee before me) would be an indirect (or non-basic) one. This general debate in perceptual epistemology is not my concern.

There is, however, a specific type of perceptual belief that is important to this discussion, namely *perceptual demonstrative beliefs*. These are beliefs involving a demonstrative that purport to refer to something that is (by hypothesis) currently perceived. To simplify, they are beliefs of the form 'That is P', where 'that' purportedly denotes some perceived object (or property or event or whatever) and 'P' is a property term of some sort. For example, the belief "That is blue" or "That is a laptop" which is held in response to a perception of something (and purports to be about that thing). When demonstratives are employed in this kind of way I will call them *perceptual demonstratives*. Perceptual demonstrative beliefs of this sort are particularly simple perceptual beliefs that rely quite heavily on what is perceived. They rely heavily on what is perceived because demonstratives are perhaps the most compelling example of a directly referential term, where the meaning or content of an employed demonstrative is directly and arguably wholly given by the referent.[4] In this case the demonstrative is a perceptual demonstrative, and thus the intended referent is something

---

[2] One could formulate this in terms of knowledge instead of beliefs if one preferred.

[3] Lyons (2023) identifies some other relevant notions of 'directness'. For reference my 'immediate perception' is akin of his 'perceptual directness' and my 'direct perception' to his 'epistemological directness'.

[4] The character of the demonstrative, in Kaplan's sense, is something like "the thing being demonstrated" or "the thing one is intending to demonstrate."

purportedly perceived. The subject of the belief is in this way "directly given" by the perception. This makes simple perceptual demonstrative beliefs strong candidates for beliefs that can be directly and fully justified by perceptions – that is, they are candidates for one type of basic belief. While I am tempted by this view, for present purposes I will set aside debates about basic beliefs and focus on perceptual demonstratives and their associated beliefs as phenomena of interest in themselves. At minimum, interest in perceptual demonstratives stems from the fact that their successes seem to hinge precisely on the perceptions that they are parasitic on.

Most believe that we can demonstratively refer to ordinary physical objects like computers, cars, and trees. If a successfully employed perceptual demonstrative owes its success to the perception that it is parasitic on, then one might suppose that our account of perception, and in particular of immediate perception, should explain how we can perceptually demonstrate ordinary physical objects. Put another way, successful perceptual demonstratives of ordinary physical objects involve direct perception (in the epistemic sense), and this directness derives from the associated current perception. Perhaps we can and should explain this epistemic directness in terms of what one is immediately perceiving (in the perceptual sense). This is Bermúdez's Reference Constraint:

> **Reference Constraint**: If it is indeed the case that we make demonstrative reference to ordinary physical objects, then our account of the immediate object of perception must explain how this is possible. (Bermúdez, p. 365)

The Reference Constraint asserts that we should explain successful perceptual demonstratives in terms of immediate perception, setting the ground for explaining an important type of perceptual belief in terms of immediate perception. I will grant the Reference Constraint as a working assumption.[5] The Reference Constraint prompts us to analyse immediate and mediate perception (the perceptual distinction). It also sets the stage for Bermúdez's main argument. Let me turn to those now.

2. Bermúdez's main argument

Bermúdez considers two views of immediate perception: what we usually mean when we speak of sense-datum theory (SDT) and his proposed alternative, naturalized sense-datum theory (NSD). Stated most generally, the term 'sense-datum' means what is given to the senses. Studies of sense-data aim to give a theory of what it is that is given to the senses, of how what is given is related to the ordinary physical world and to ourselves, of how we should conceive of perception in this context, and of how all of this feeds into perceptual beliefs, knowledge and other aspects of epistemology. In this regard SDT and NSD are two theories about the nature of sense-data.[6] Here are the core commitments of each.

---

[5] Lyons (2023) provides an overview of perceptual epistemology that contains some options that may not require the Reference Constraint. I cannot delve into this terrain in the present work, though I will conclude that the Reference Constraint is important and worth preserving.

[6] See Hatfield (2021) for an overview of sense-data. The history of the topic is complex, as is the relation between that history and the issues discussed in this work. For example, NSD bears an important relation to Moore's (1918-19) original version of sense-datum theory. What I will refer to as 'SDT' is the version that became more popular in

According to NSD, "The immediate objects of visual perception, or what are often called sense-data, are parts of the facing surfaces of physical objects" (353). We nonetheless can and do perceive ordinary physical objects. However, according to NSD, we perceive ordinary physical objects by perceiving their front-facing surfaces. In this sense we mediately perceive those objects. The idea that we immediately perceive only the front-facing surfaces of objects is familiar from discussions of sense-data. I will unpack it below.

According to SDT, the immediate objects of perception are mind-dependent objects. Within SDT these mind-dependent objects are themselves typically called 'sense-data'. We thus have two uses of 'sense-data'. To help us keep track of this, I will use the term 'SD' for the mind-dependent objects that are central to SDT and continue to use 'sense-data' for the immediate objects of perception (whatever they may be). According to SDT, although the immediate objects of perception are these peculiar mind-dependent objects, namely SD, we can nonetheless perceive ordinary physical objects by immediately perceiving the mind-dependent ones. This is because SD *represent* ordinary physical objects, either intrinsically or in virtue of the way perceivers interpret them. In this way we mediately perceive ordinary physical objects. There are numerous arguments in favour of SDT, and at least as many criticisms. In my view the most powerful argument in favour of SDT stems from perceptual error, namely illusion and hallucination (see Section 6), and other important ones stem from causation and perception (see esp. Robinson, 1994), the scientific conception of the world in contrast to the phenomenal one (see esp. Jackson, 1977), and so on.

We thus see that, according to both NSD and SDT, we mediately perceive ordinary physical objects by immediately perceiving sense-data. However, what counts as sense-data (the immediate objects of perception) is strikingly different on the two accounts. According to NSD, sense-data are the front-facing surfaces of physical objects, and according to SDT sense-data are SD. We therefore mediately perceive physical objects for very different reasons on these views. This groundwork is sufficient to explain Bermúdez's overall argument, which I have distilled from Section III of his 2000 (pp. 365-372):

(1) We do make demonstrative reference to ordinary physical objects.

(2) The Reference Constraint is true.

(3) There are only two viable theories of immediate perception: NSD and SDT.

(4) SDT fails to adequately explain perceptual demonstratives about ordinary physical objects.

(5) NSD successfully explains perceptual demonstratives about ordinary physical objects.

(6) Therefore, we should endorse NSD (and reject SDT).

Premise (1) is put forward as a reasonable, intuitively justified hypothesis. One could object to it, including a defender of SDT. In this case the argument is stalled. However, I regard (1) as plausible and find the debate that ensues from assuming its truth fascinating. Premise (2) is also a hypothesis, but it is justified by the kind of reasoning given in Section 1. Premise (3) is what Bermúdez takes to

subsequent decades, and is associated with Russell (1912), Jackson (1977), and Robinson (1994) among others. In the interests of space, I suppress these many important historical details.

be the relevant views for this debate. I will remark on this in Section 8 but for now wish to note that even if one rejects the 'only' in premise (3), the argument could be modified so that (3) says that two of the important theories relevant to this debate are NSD and SDT, and the conclusion (6) asserts that NSD is preferable to SDT'. This would retain the validity of the overall argument, and (as I hope to show) retain the interest of the substance of Bermúdez's insights. Premises (4) and (5) are clearly substantive claims in need of justification. Assuming they are justified, the argument is on solid footing.

At this point one might ask whether we need another argument against SDT, given how out of favour the view is in our current climate? I think this worry misses Bermúdez's point and the interest of his paper. First, there is enduring interest in direct perception (epistemological notion) and perceptual demonstratives. In addition, the Reference Constraint is *prima facie* plausible and interesting. This prompts us to examine viable notions of immediate perception (and mediate perception) so that we can adequately understand how perceptual demonstratives can and by hypothesis do apply to ordinary physical objects. Second, the idea of trying to capture what discussions about sense-data were after (the immediate objects of perception, whatever they may be), put it to good use (via the Reference Constraint), and in the process show that SDT isn't a good a way of doing this, is fascinating on its own. If Bermúdez's argument works, then he effectively shows that SDT failed at its own game (giving a viable account of sense-data), but that we can nonetheless recover what the debate was after by embracing NSD. Third, the conceptions of immediate and mediate perception embodied in NSD and SDT are distinct, and stand on their own, independently of these theories. I will briefly remark on this in the conclusion, as I believe these disparate conceptions of immediate and mediate perception have fruitful, underexplored applications. Let me move on to justifying (5), and then critiquing (4).

3. Why NSD?

Bermúdez claims that (5) NSD successfully explains perceptual demonstratives about ordinary physical objects. I agree. Let me explain the idea. In NSD, we immediately perceive unoccluded front-facing surfaces of ordinary physical objects and mediately perceive the occluded parts of these objects and the objects as wholes. This is definitive of the view. To get a sense of the significance of the NSD, consider an everyday example and three ways of interpreting it.

> **Person Perception:** When you look at a person who is facing you, there are parts of the person that are "presented" to you in a way that other parts are not. For example, their face and the front of their torso are oriented toward you and "in view" in way that their backsides, insides and body as a whole are not.

Consider three interpretations of this case.

One might propose an austere "parts only" theory of perception according to which all we really see are the unoccluded front-facing parts of physical objects. Judgements about whole objects and about occluded object parts are constructed in post-perceptual cognition, based on inferences from perceived front-facing parts of those objects. Thus, in Person Perception, you perceive only the

front of the person's face and torso. Insofar as you are "committed" to them having backsides and insides, and to them being a discrete whole physical object, such commitments arise only in post-perceptual cognitive judgement. This might be how to conceive of perceptual experience if something like Marr's (1982) 2.5D sketch or Peacocke's (1992) base level scenario content constituted perceptual experience.[7]

Second, one might propose that we do experience whole objects, but that there is nothing to the idea that we "immediately perceive" the unoccluded front-facing parts or that there is any interesting difference between perceiving whole objects and perceiving their front-facing parts. This is an "objects only" view. Thus, in Person Perception, you perceive the person as a whole object and there is nothing distinctive about the difference between perceiving them as a whole object and perceiving their front-facing parts. On this view, we can safely set aside this whole discussion and talk simply about the objects that we perceive. This interpretation might be how to conceive of perceptual experience in contemporary naïve realism (e.g. Martin, 2017) or in some 4E theories of perception (e.g. Shapiro and Spaulding, 2024).

NSD proposes a different interpretation. According to NSD, and against the "objects only" view, there is an important sense in which we more immediately perceive the unoccluded front-facing parts of physical objects than their occluded parts, and than the objects as wholes. But NSD also asserts that we nonetheless generally perceive whole objects, in contrast to the "parts only" view. To hold these claims simultaneously, NSD proposes that we perceive all of these things but in different senses of 'perceive'. Since our perceptions of the unoccluded front-facing parts of objects seem more immediate (or direct or present or epistemically available in perception), and our perceptions of whole objects and of their occluded parts occurs *by virtue of* our perceptions of the unoccluded front-facing parts of objects, the NSD theorist proposes that we mediately perceive whole physical objects and their occluded parts by immediately perceiving their unoccluded front-facing parts. Thus, in Person Perception, you immediately perceive the front of the person's face and torso and mediately perceive them as whole objects and mediately perceive their occluded parts.

I take the NSD interpretation to be quite intuitive in part because it charts a well-motivated middle course between the more extreme "parts only" and "object only" views. Here are two further arguments for the approach. As Bermúdez argues (pp. 363-4), to perceive something we must discriminate it from its surrounds (see also e.g. Dretske, 1979). We generally visually discriminate objects by means of the front-facing parts that are open to view. Front-facing parts are thus serving a privileged perceptual role that is captured in the NSD conception of immediate perception and absent from the "objects only" view. Second, we can draw support for NSD from the phenomenon of amodal completion (see e.g. Briscoe, 2011). In psychology it is typical to distinguish between the front-facing parts of perceived objects and the occluded parts, and posit a special mechanism – amodal completion – to explain how we manage to perceive whole, complete objects despite in some sense only "immediately" perceiving their unoccluded front-facing parts. One important part

[7] Neither Marr nor Peacocke restrict perceptual experience to the front facing parts of things. Instead, Marr posits the 2.5D sketch as present in an early stage of visual processing, and Peacocke posits scenario content as a key component of perceptual experience. Perceptual experience itself is for both theorists much richer than this.

of research in this area concerns illusions that involve amodal completion. Such illusions provide strong evidence for this kind of "filling in" being distinctive (against the "objects only" view), and being perceptual and not merely something arising in post-perceptual cognition (against the "parts only" view). Therefore, the NSD idea that we immediately perceive front-facing surfaces and mediately perceive ordinary physical objects is credible and a worthy foundation for a perceptual theory.[8] In a broader discussion, I would delve deeper into this debate. However, for present purposes it suffices to situate NSD against some of its rivals, and explain why NSD advocates take the view to be compelling. Let me move to the substance of (5).

Premise (5) asserts that NSD can explain how we demonstrate ordinary physical objects (on the assumption that we do). The argument for (5) is as follows. According to NSD, we immediately perceive unoccluded front-facing parts of ordinary physical objects. It seems reasonable to assert that we can demonstrate what we immediately perceive: if there is anything that we can demonstrate, we can demonstrate what we can immediately perceive (in whatever sense of 'immediately perceive' one finds acceptable). One may deny this claim, for example by rejecting the very idea of immediate perception. However, I take it this wouldn't be charitable to Bermúdez, and I've given adequate motivation for the existence of immediate perception in any case. One may also deny this by offering some theory of demonstratives that doesn't rely heavily on what is perceived, but I will assume that perceptual demonstratives not only exist but are a foundational kind of demonstrative.[9] Given the important link between demonstratives and perception, the claim that we can demonstrate what we immediately perceive is an attempt to carve out a basic and safe subset of perceptions and say that surely demonstratives, if they are applicable at all, must be applicable to this subset. I agree. It follows that we can demonstrate unoccluded front-facing surfaces of ordinary physical objects.

Bermúdez's key move at this point is to appeal to *deferred demonstration*. This is where a demonstration carries beyond what is before one in the following sense. For example, if one says "That man is the greatest!" while pointing to a picture of Messi, one is referring to Messi the person with 'that man' and not the picture. 'That' can thus go beyond a picture being perceived and refer to what the picture represents.[10] To be sure, this doesn't prevent us from demonstrating the picture, for example by saying "That is a picture of the greatest footballer." Such a demonstration is *non-deferred* relative to the former demonstration.[11] Similarly, you might point to an oversized jacket on a chair and say "That's a large person" (deferred demonstration) or "That's a large jacket" (non-deferred demonstration). You may also hear a soundbite from an advert and exclaim "This is a catchy tune" (non-deferred demonstration) or "This is my favourite restaurant!" (deferred demonstration). I think that a full, adequate theory of deferred and non-deferred demonstration is difficult to develop, but

[8] Note that I haven't gone through Bermúdez's potent analysis of some critiques of the NSD idea (see esp. pp. 359-364).

[9] For comparison, Brovold and Grush argue that "the fashionable philosophical notion of the 'perceptual demonstrative' is an aberration" (2012, 39). There is a fascinating debate to be had here, but it falls outside my current purview.

[10] Bermúdez (370) helpfully reminds us that this idea is found in Quine (1968, pp. 40-44).

[11] I will stick with the terminology 'deferred' and 'non-deferred' when discussing demonstratives so as not to get confused with the other terminology already in place.

the existence of the phenomenon is straightforward, and is appreciable independently of the issues we have discussed to this point.

Deferred demonstration is relevant in the following way. I have argued that we can demonstrate unoccluded front-facing surfaces of ordinary physical objects. These surfaces are parts of those objects. It seems plausible that deferred demonstration can "carry over the part-whole relation" (Bermúdez, p. 370). This is arguably true in general and in any case is very compelling when the part at issue is being immediately perceived and the whole at issue is being mediately perceived in the NSD sense of these terms. NSD therefore provides a compelling explanation of how demonstrative reference to ordinary physical objects can occur. I think we should accept this argument and hence (5).

4. Why reject SDT?

Premise (4) asserts that SDT cannot do what I, following Bermúdez, just argued NSD can, namely explain demonstrative reference to ordinary physical objects. Recall that according to SDT, we immediately perceive the mind-dependent objects SD, and by virtue of these representing ordinary physical objects (either intrinsically or by virtue of how we interpret them[12]) we mediately perceive the latter. Bermúdez's argument for (4) is to pose a dilemma for SDT. There are two possible ways to demonstratively refer to ordinary physical objects: via deferred and non-deferred demonstration. Neither model, he argues, can be applied within the constraints of SDT. Thus, (4) is true – SDT cannot explain demonstrative reference to ordinary physical objects. Let me unpack this.

Let's consider deferred demonstration first. Within SDT, we immediately perceive SD. Bermúdez argues that deferred demonstration to ordinary physical objects cannot be appealed to by SDT advocates. To see why, suppose that a perceiver immediately perceives a given SD_m[13], and mentally demonstrates that SD_m while saying or thinking "That is a laptop" with the intention of demonstrating the laptop that by hypothesis is represented by that SD_m. If this demonstrative succeeds in denoting the laptop, this would involve deferred demonstration. The reason Bermúdez believes this isn't credible is because:

> linguists and philosophers of language are agreed that deferred ostension [demonstration] can only work in a communicative context when speaker and hearer share knowledge of an individuating principle linking the physically present ostended [demonstrated] object with the intended object of reference. (p. 372)

However, Bermúdez argues, within SDT this requires shared knowledge of the speaker's private SD and the individuating principle that links those SD and ordinary physical objects. "It seems obvious, however, that no such principle is implicated in everyday demonstrative reference to material objects" (p. 372). Thus, Bermúdez concludes, SDT cannot appeal to deferred demonstration to explain demonstrative reference to material objects.

[12] I suppress this qualification below.

[13] I will use 'SD_m' to refer to a singular sense-datum, and continue to use 'SD' to refer to the plural sense-data.

In other words, deferred demonstration works with a picture of Messi because I and my interlocutors understand the link (or individuating principle) between the picture and Messi and can implicitly appeal to this to correctly interpret what I'm trying – and by hypothesis succeeding – in saying when I utter "That is a great football player." By contrast, when I mentally point to my SD_m and say "That is a laptop", my interlocutors cannot have analogous knowledge of the requisite individuating principle because my SD_m and my demonstration of it are private. Thus, deferred demonstration is unavailable within SDT to explain demonstrative reference to ordinary physical objects.

This leaves SDT to explain demonstrative reference to ordinary physical objects via non-deferred demonstration. For Bermúdez this yields the other horn of the dilemma, for here demonstrative reference must somehow bypass SD and latch onto the ordinary physical object. The problem, however, is that this violates "the Reference Constraint because their account of the immediate object of perception has no part to play in their explanation of how demonstrative reference to three-dimensional material objects is achieved. The sense datum is otiose in the explanation of direct perception" (p. 372).

Bermúdez concludes that, whether SDT utilizes deferred or non-deferred demonstration, it cannot adequately explain how we demonstratively refer to ordinary physical objects. (4) is thus true.

5. In defense of SDT

I reject both horns of the dilemma Bermúdez tries to impose on SDT. Before addressing these claims directly, let me remark on two relevant issues about SDT. First, there is a general worry about SDT that stems from its appeal to a private relationship between a subject and her SD. One might hold that these kinds of private relations are problematic, perhaps by appeal to Wittgenstein's private language argument. I cannot engage in a detailed discussion of this topic. However, I can point to an important work like Fodor (1975) which provides a detailed account of how a private language of thought can coherently exist and arguably does exist in humans. This is sufficient to counter any blanket rejections of my proposal by appeal to worries about private languages. I therefore assume that private languages are possible. Among other things, this means that the Reference Constraint can in principle be interpreted to apply to private languages.

Second, one might worry that SDT leads to skepticism about the external world. How could we come to know anything about a mind-independent, physical world if all we immediately perceive are mind-dependent SD? Russell (1912) addressed this worry by appeal to the various patterns and stabilities that one can observe between SD. For example, every evening the SD I associate with my bed and its many features are like the ones I experienced when I left for work in the morning. Russell argued that the best explanation of the observed patterns and stabilities of an individual's SD is that there is a distal physical world that has similar patterns and stabilities and is the cause of one's SD. In principle, this reasoning is analogous to the reasoning that justifies our beliefs about things that we cannot "immediately observe" in nature, including subatomic particles and black holes. This includes justified beliefs about causal relations between those unobservables, and causal relations

between those unobservables and our experiences of data observed on (e.g.) computer monitors that has been gathered by measurement devices like particle accelerators and the James Webb telescope. The full rational story of how we acquire and justify beliefs in unobservables in nature is complex, as is the full rational story of how, given SDT, one could come to believe in a distal physical world as the cause of one's SD. Regardless, I submit that SDT can avoid collapsing into a kind of idealism by an inference to the best explanation of what is the likely cause of various patterns of one's SD. I thus assume that SDT is consistent with knowledge of a mediately perceived physical world.[14] With these two assumptions in hand, let me return to Bermúdez's claims.

5.1 The First horn

The first horn of Bermúdez's dilemma concerns using deferred demonstration, where the SD are demonstrated and ordinary physical objects are referred to via deferred demonstration. Bermúdez's objection rests on the hypothesis that deferred demonstration requires shared knowledge by speaker and hearer of an individuating principle linking the demonstrated object (e.g. the picture or the SD_m) to the intended referent (e.g. Messi or the laptop). He then claims that such shared knowledge is "obviously" not present in everyday demonstrative reference to physical objects. The first claim should be rejected and the second claim, while true, doesn't prevent the advocate of SDT from giving an account of how we might – in principle – use deferred demonstration to demonstrate ordinary physical objects.

Bermúdez claims that shared knowledge of a principle that links the demonstrated object to the intended object is required for deferred demonstration. It is true that deferred demonstrations expressed via public languages rest on shared knowledge of this sort. It is also true that our actual practice of deferred demonstrations developed via shared knowledge of this sort. The relevant question, however, is a principled one: if SDT were true, could someone demonstrate one of their own SD in an effort to generate a deferred demonstration of an ordinary physical object, and succeed in referring to that object? Let me explain why we should answer positively.

Suppose that I am alone and possess adequate cognitive and perceptual capacities (setting aside the practical questions of how these things came to be). Suppose also that I correctly believe that SDT is true. Currently, I immediately perceive a blue, oval SD_m and believe that it is caused by something in the physical world. I suggest that I could internally demonstrate the SD_m with the intention to refer to what I believe is the distal physical cause of the SD_m. Such an internal demonstration might be accomplished via a focused attention on the SD_m, and this could be combined with the cognitive intention to refer to what I believe is the distal physical cause of the SD_m. In fact, the SD_m is caused by a robin egg. Thus, the egg is the relevant cause of the SD_m, the SD_m resembles the egg in relevant ways (i.e. both are blue and oval), and my intentions are straightforward. I see no reason why my efforts to generate a deferred reference to the egg would fail, anymore than my efforts to refer to Messi via a picture of him would fail. It does not matter whether there is a shared or public principle that links the demonstrated object to the intended object.

[14] Russell is well known for having (1927) argued that knowledge of the physical world is limited to knowledge of the world's "structural" features. This limitation of worldly knowledge is sufficient for my purposes.

The other claim in the first horn of Bermúdez's dilemma is that everyday public demonstrative reference to physical objects does not seem to involve shared knowledge of deferred reference through SD. This is doubtless true. Regardless of the truth of SDT, common sense does not operate as though anything like SDT true, and thus doesn't embody shared knowledge of the kind of deferred reference to the robin egg described above. The same is true of linguistics, and thus we wouldn't expect linguistic analyses of demonstrative reference to generally reflect anything like the kind of private demonstration I am appealing to. The relevant question, however, is not how best to interpret everyday practice. The relevant question is whether deferred reference to the robin egg can, in principle, be shared among a community of speakers. I see no reason why it can't. Here is a rough sketch of how this could come about.

Suppose a group of people come to believe that SDT is true and they are in fact correct in this belief. Given my initial assumptions, this means that they also believe that there is a surrounding physical world, and that these physical things are the typical causes of each person's private SD. At the beginning, they may speak in a cumbersome technical way about SD and the physical world. For example, suppose A says:

"The physical thing that is causing me to have this blue, oval SD_m..."

Here the demonstrative 'this' purports to refer to A's SD_m. A community member B can understand this, even if, for example, B doesn't rule out the skeptical possibility that A's SD are (e.g.) inverted with respect to B's. Furthermore, A can say:

"That robin egg is causing me to experience this blue, oval SD_m."

Here, as before, 'this' demonstrates A's SD_m and 'that' demonstrates the hypothesized egg. B has no difficulty understanding A's claim. This would, for example, involve deferred demonstration of the egg if A's employment of 'that' involves A mentally pointing to her SD_m with the intention of referring to the egg that A's blue oval SD_m represents. If needed, A could explain to B that this is what she is trying to do, and then do it. B would thus understand that deferred demonstration was used to refer to the hypothesized egg with 'that' when A utters "That egg is causing me to experience this blue, oval SD_m."

One may doubt whether B could come to know which physical object A is referring to – that is, can B come to know that A is talking about the specific robin egg in question? Yes. To see why, assume B is standing next to A and the egg is in front of A. This means that B is experiencing private SD that represent A and A's location relative to B, and vice versa. When B's gaze is directed in front of where A is represented to be, B also experiences a blue, oval SD_m (or least what B would describe as a blue, oval SD_m). In this way, B and A could come to single out the same object – the egg – as the referent of A's deferred demonstrative. While this scenario is underspecified in many ways, it strikes me as entirely coherent. Thus, while a shared principle of individuation of the right sort is required for shared communication of demonstrated objects, it is not required for deferred demonstration itself.

Consider, finally, a more evolved stage of communication where, as a community, deferred demonstration of this sort is the presumed interpretation of demonstrative expressions of this sort.

Thus, A simply says "That is my favourite egg" and B correctly understands A to be demonstrating her blue, oval SD_m and intending to, with 'that', refer to the hypothesized egg. This would also involve deferred demonstration of the sort that Bermúdez rejects, but the shared knowledge of the individuating principle would be implicit in their communicative practice, much as it is implicit in our practice when I point to the picture of Messi and say "That is a great football player."

To be sure, Bermúdez's objection is that it "seems obvious, however, that no such principle is implicated in everyday demonstrative reference to material objects" (p. 372). As a point about how to understand our everyday linguistic practice, I agree. The way we tend to understand our everyday actions is not in accordance with the SDT community just described. But this isn't an objection to the ability of SDT to explain demonstrative reference to ordinary physical objects. Further, the SDT advocate can argue that, in the end, their description is *in fact* an accurate description of our everyday practice, whether we realize it or not. Regardless, I conclude that SDT can explain demonstrative reference to ordinary physical objects via deferred demonstration.

5.2 The Second Horn

The second horn of the dilemma asserts that SDT cannot use non-deferred demonstration to explain demonstrative reference to ordinary physical objects. This is because to do this, the demonstration would have to bypass the SD, in which case the demonstration doesn't appeal to the immediate objects of perception (i.e. the SD), violating the Reference Constraint. I disagree. Return to speaker A in a community that believes SDT and in which SDT is true. Suppose A says "That egg is my favourite" and with 'that' succeeds in bypassing her blue, oval SD_m and directly referring to the purported egg. It doesn't follow that the explanation of how this works can avoid appeal to A's SD. Assume A's demonstration consists entirely of her publicly pointing with her finger to the cup, and thus avoids the kind of inner demonstration by A to her own SD discussed in 5.1. Now suppose that A's act of demonstrating occurs in significant part because of the immediate objects of her perception: she is able to perceptually connect to the egg, form her thought about the egg, and move her arm so as to point at the egg *because* of the SD_m that she is experiencing and the fact that the SD_m represents the egg. On this account, A's demonstration of the egg is non-deferred in the relevant sense, but its success depends critically on the immediate objects of her perception, namely her SD.[15] The Reference Constraint is therefore preserved. This strikes me as viable within SDT.[16] This practice is also much closer to our everyday practice.

In conclusion, premise (4) of Bermúdez's overall argument, that SDT fails to adequately explain perceptual demonstratives about ordinary physical objects, should be rejected. Since this is Bermúdez's primary motive for rejecting SDT in favour of NSD, SDT can withstand his assault. However, there is much more of interest to be said about these issues. At minimum, we should

[15] Compare with the "epistemic" role played by what is immediately perceived in Brown (2009, pp. 381-386).

[16] One alternative possibility is to suppose that A demonstrates the cup solely by appeal to unconscious perceptual states. The perceptual experience would thus overdetermine what is needed for reference and by hypothesis for perceptual demonstrative belief. While this is possible, I do not believe this is an accurate model for understanding our perceptual lives (though I recognize that some, like Jack Lyons (personal communication), believe this model is largely correct). I also suspect that, to the extent that unconscious perception undermines SDT it also undermines NSD. These details, however, would take us too far afield.

consider whether NSD can explain the phenomena that originally motivated SDT (section 6), and, if not, whether it might be supplemented to do so (section 7).

6. The problem of perception & NSD

Suppose I am correct that Bermúdez has not given us a good reason to reject SDT. As mentioned earlier, the opponent of SDT may not be stressed, believing that there are several other reasons to reject SDT. Fair enough, and this is not the place to itemize and attempt to reply to each of them. However, I do wish to highlight another part of this dialectic. The SDT advocate needn't strictly play defense. Instead, they can go on offense and argue that NSD cannot adequately explain what is often called The Problem of Perception. This is the problem of (a) explaining the differences between veridical, illusory and hallucinatory perceptual experiences, while (b) explaining the phenomenal similarities that can occur between them.[17] Let me explain how The Problem of Perception is particularly challenging for NSD.

Here is a simple characterization of the core notions that is sufficient for our purposes.[18] A veridical perceptual experience is fully accurate of the physical world. For example, you have an experience of a blue cylindrical thing, and you are in fact looking at a blue cylindrical thing. An illusion occurs when you experience a physical object, but experience at least one of its features incorrectly. For example, you have an experience of a blue cylindrical thing, and you are in fact looking at a purple cylindrical thing. An hallucination occurs when your experience doesn't have a corresponding physical object at all. For example, you have an experience of a blue cylindrical thing and there is nothing corresponding to this in the physical environment.

Veridical, illusory and hallucinatory experiences are "latching" onto the physical world to different degrees, and in this regard these experiences are different types of perceptual states. Nonetheless, while we can at times subjectively distinguish between veridical, illusory and hallucinatory experiences, at times we cannot. In principle, these different types of experiences can be subjectively indistinguishable. For many, the most powerful justification for SDT is its ability to explain this conundrum. SDT postulates a type of mind-dependent object – SD – as a perceptual intermediary between the perceiver and the physical world. This intermediary is present regardless of whether one's experience is veridical, illusory or hallucinatory – SDT is a "common factor" theory of perception. This explains why experiences of these types can be subjectively similar and even indistinguishable. However, the difference between veridical, illusory and hallucinatory experiences does not collapse, since which type of experience one is having depends on how well one's current SD latches onto (i.e. represents and is caused by) the physical world.

It is straightforward that, on its own, NSD doesn't solve The Problem of Perception. NSD is a theory according to which the immediate objects of perception are the front-facing surfaces of ordinary physical objects, and ordinary objects themselves are mediate objects of perceptions. As it

[17] The reasoning that leads from here to SDT was first clearly articulated by Robinson (1994) and has since been codified in Crane and French (2021). See again Hatfield (2021) for an overview of various additional motives for SDT.

[18] For a powerful new theory of illusion and hallucination see Macpherson and Batty (2016).

stands, this framework provides no resources to explain the difference between veridical perception, illusion and hallucination, nor the respect in which experiences of each sort can be subjectively indistinguishable. Bermúdez concedes one part of this challenge when he notes that NSD is "is only an account of non-hallucinatory perception" (p. 354). He thus recognizes that NSD needs to be bolstered to explain hallucination. However, he fails to make a similar concession about illusion, and about the subjective similarities that occur across experiences of all three sorts. He also fails to explicitly appreciate how acute the problem is.

Consider, briefly, illusion. Suppose you look at a purple cylindrical bottle and have an experience of that bottle as a blue cylindrical bottle. That is, the blue that you experience isn't experienced as a feature of just any cylindrical bottle, it is experienced as a feature of the very physical cylindrical bottle that sits before you (a bottle which is in fact purple). This is what makes this an illusion: you are actually experiencing the object in question (i.e. there is successful object perception), but you are experiencing it to have a feature that it doesn't have (i.e. there is an error in property perception). How should NSD be applied here? In line with NSD, the front-facing surface of the bottle is present before you in a way in which the backsides and insides are not and in which the bottle as a whole is not. However, the front-facing surface is purple, and you aren't experiencing that surface as purple, you are experiencing it as blue. In what sense, then, are you "immediately perceiving" the front-facing purple surface? If anything, it seems like the front-facing surface, although physically available, is perceptually occluded by the experienced blue. So, what is immediately perceived? Perceptually, it seems like a blue cylindrical bottle is immediately perceived, but there is nothing in the NSD view that captures this. If the NSD theorist adds something new to explain the illusion, whatever is "seemingly blue" arguably has to (a) capture the sense in which one is experiencing blue, (b) do so in a way such that blue is occluding the purple of the object, while (c) preserving the core NSD thesis that the front-facing surface of the perceived purple object is immediately perceived.

This is a problem for NSD because NSD is built around the idea of immediate perception and in particular the idea that front-facing surfaces are what is immediately perceived. There is no easy way for the NSD theorist to address these worries.[19] To see why, consider a possible reply.

7. Nonconceptual content & NSD

Bermúdez is a well-known defender of nonconceptual content (e.g. Bermúdez, 1995, 2007; for an overview see Bermúdez & Cahan, 2024). It would be natural for him to appeal to nonconceptual content to address The Problem of Perception.[20] To what extent can nonconceptual content solve The Problem of Perception? How can this purported solution be integrated with NSD? These are the questions that I address in this section. My conclusions are as follows. Nonconceptual content can adequately address the semantic part of The Problem of Perception. However, it can only positively address the phenomenological part by appeal to postulates that are no less extreme than

[19] Compare with Hatfield's (2021) worry that NSD may have to postulate "appearances" to accommodate perceptual constancy.

[20] Thank you to an anonymous referee for suggesting that I address this possibility.

those of SDT. Further, the resulting framework is either incoherent, or asserts that in veridical perceptual we are mediately aware of the physical world. The latter commits NSD to a perceptual veil that is similar to the one embraced by SDT, and which NSD is designed to avoid. In effect, via nonconceptual content NSD can solve the Problem of Perception, but only by embracing a veil of perception. Let me unpack this more slowly.

A full discussion of nonconceptual content falls outside the scope of this work. For our purposes, it is instructive to recall what nonconceptual content is and what kinds of reasons are offered for its existence. Contents are expressed by representations.[21] Nonconceptual content is a form of content that is in principle distinct from conceptual content, that is, distinct from contents found in thought. Bermúdez's "Master Argument" for nonconceptual content in perception rests on two principles:

1. "[T]he content of perception is linked to capacities for perceptual discrimination";
2. "[C]oncept possession can never wholly be explained in terms of capacities for perceptual discrimination." (Bermúdez, 2007, 59)

From these principles it follows that there is perceptual content that is nonconceptual.

There are important case studies that help justify this conclusion. Many animals can perceptually discriminate things that their cognitive systems are not sufficiently complex to generate concepts for. This point carries over to humans as well. We can perceptually discriminate detailed, nuanced features like highly specific colours and shapes despite not having concepts for such specific properties. These cases suggest that perceptual discrimination is operating, in a fundamental way, independently of cognition, and thus that perceptual contents are nonconceptual.

Contents, by their nature, can have varying degrees of accuracy to their subject. Thus, Bermúdez identifies the possibility of misrepresentation as one of the core criteria of representations (1995, p. 351). Attributing nonconceptual content to perceptual experience thus permits, indeed guarantees, the possibility of perceptual error. This makes it poised to help address The Problem of Perception.[22]

Recall that, as standardly conceived, in illusion there is successful object perception but erroneous property perception. For example, one experiences a purple bottle as blue. By contrast, successful object perception does not occur in hallucinations. For example, one has an experience as of purple bottle and there is nothing in the physical world that corresponds to this. Let me unpack what a theory of perceptual error must explain, focusing again on illusion. The key distinction for our discussion is between the *semantic* and *phenomenological* challenges posed by perceptual error.

One fundamental puzzle raised by illusion is a *semantic* one of how to model perceptual experience such that it can be said to involve the perception of an object but the misperception of one or more of its features. It is straightforward that adding nonconceptual content to NSD affords an account

[21] Some of the terminology that follows I will employ without argument (e.g. the term 'expressed' in above sentence). Hopefully, context is sufficient to capture what is intended. In all cases I intend the terminology to be as neutral as possible.

[22] What follows is a discussion of a contemporary form of representationalism about perceptual experience. As is well known, there are now numerous approaches to representationalism. I cannot survey them all, but what follows is representative of a very influential approach.

of the semantic challenge of illusion that is unavailable to NSD proper. In the bottle case, I am perceiving a bottle that is in fact purple but erroneously experiencing it as blue. By appeal to nonconceptual content, we can assert, for example, that my perceptual system is expressing something like the content <there is a blue bottle>[23]. Thus, my perceptual system is correctly representing that there is a bottle but incorrectly representing it as blue. For our purposes, this is an adequate explanation of the semantic challenge of illusion. However, it is important to see that this account of illusion is incomplete, and incomplete in a manner that is at the heart of our discussion.

It is one thing to give a model of the semantics of illusion. It is quite another to give a model of the *phenomenology* of illusion, a model of what it is like for the perceiver to undergo an illusion. On its own, nonconceptual content yields no positive account of the phenomenology of illusion. One way to see this is to recognize that nonconceptual content is consistent with SDT. As I've emphasized throughout, according to SDT, SD represent the physical world, either intrinsically or via an interpretation of SD. Either way, this is consistent with SD expressing nonconceptual content. For example, a charitable interpretation of SDT holds that humans can perceive many specific colours and shapes in the physical world that they don't have concepts for *because* SD can represent specific colours and shapes that humans don't have concepts for. This permits humans to mediately perceive those specific physical features. The same holds for animals who can discriminate things they don't possess concepts for.

Regarding illusion, SDT already embodies the semantic solution offered by nonconceptual content: when one misperceives the purple bottle as blue, one's blue, bottle-shaped SD correctly represents that there is a bottle in the physical world but incorrectly represents it as blue. What SDT adds to the semantic account of illusion afforded by nonconceptual content is an account of the perceptual phenomenology: the bottle perceptually seems blue (in the phenomenal sense of 'seems') because one is *immediately perceiving* a blue bottle-shaped SD. This is a contentious commitment of SDT. It is also a positive account of the phenomenal challenge posed by illusion and, by extension, by hallucination. What can the NSD advocate say?

On its own, appeal to nonconceptual content does not provide a positive account of the phenomenology of perceptual error. I, for one, cannot find such a positive account in Bermúdez's work. Thus, to provide a concrete analysis I will make two simplifying, charitable assumptions. First, other well-known defenders of nonconceptual content have offered substantive and similar proposals. Here are two examples:

> Along with (most) other representationalists, I am happy to say that, in the hallucinatory case, the perceiver is conscious of an un-instantiated property. (Tye, 2014, p. 304, fn. 20)
>
> In hallucinating…We are aware of pure universals, uninstantiated properties (Dretske, 2003, p. 73)

[23] There are of course various ways to unpack the semantics of perceptual contents. The existential format in the text is reminiscent of Tye (1995, 2002). These details don't substantively impact the argument that follows.

Dretske and Tye postulate uninstantiated properties (e.g. Platonic universals[24]) and a capacity to be perceptually conscious or aware of them. These tools are then used to explain the phenomenology of hallucination. For ease of reference call this first assumption *the Dretske-Tye account of hallucination*. This proposal has the advantage of providing a positive account of the phenomenological puzzle. However, it raises two challenges. It relies on metaphysical entities that are not straightforwardly friendly to naturalism (e.g. universals are generally regarded as being outside space and time) and it explains perceptual phenomenology via a novel perceptual capacity to become perceptually aware of the proposed metaphysical entities. More concretely, it explains perceptual phenomenology via awareness of entities outside space and time.[25] One reason this explanation is questionable is because there is a decided "here and now" aspect to perceptual experience that seems at odds with universals. One would not blame the SDT advocate for arguing that explaining the phenomenology of hallucinations via a capacity to immediately perceive SD is no more troubling than explaining it via awareness of universals outside space and time.[26]

Let me first conclude that (a) there are influential defenders of nonconceptual content who recognize a need to provide not only a semantic account of perceptual error but also a phenomenological one, and (b) their positive phenomenological account appeals to postulates that are, on their surface, no less worrying that the postulates of SDT.

To extend this account to veridical perception and illusion we need to make another assumption. For ease of discussion, I will confine myself to *common-factor views*, which are arguably endorsed by Bermúdez, Dretske, and Tye. According to common-factor views, the nature of perceptual experience is fundamentally the same across hallucination, illusion and veridical perception – what varies is how those experiences connect to (i.e. causally and semantically) the physical world.[27] As mentioned, SDT is a common-factor view. Applying the common-factor view to the Dretske-Tye account of hallucination arguably yields the conclusion that one is aware of universals during hallucinations, illusions and veridical perceptions. What varies is how those experiences causally and semantically connect to the physical world.

One odd consequence of this view is that it isn't explicit if or when we are aware of the physical world. Instead, to this point, awareness is a relation we bear to universals, and the physical world is only entering the theory at the semantic and causal levels (i.e. where one explains how well the physical world satisfies a given experience/ content). Yet, surely, we are aware of the physical world in some sense. One natural suggestion is that our awareness of the physical world, in some sense, is mediated

[24] While Dretske appeals to universals to understand uninstantiated properties, Tye (2014) appeals to sets of possible worlds. For simplicity, I will stick to universals in the text. I don't believe this affects the substance of my argument.

[25] In the case of Tye's preferred metaphysical postulated, sets of possible worlds, the result is that perceptual phenomenology is explained via awareness of sets possible worlds, i.e. awareness of a collection of entities that constitute different spacetimes.

[26] Compare with chapter 5 of Schellenberg (2018).

[27] By contrast, disjunctivist views hold that the nature of perceptual experience is fundamentally different, depending on whether the experience is veridical or erroneous. Disjunctivism is standardly aligned with nonrepresentational views of perception, though there are exceptions (e.g. Johnston, 2004; Schellenberg, 2018). Crane and French (2021) provide an overview. For the interested reader, Tye (2014) critiques so-called "gappy content" views (which e.g. Schellenberg defends) and defends two niche common-factor content views, surveying numerous alternatives in his discussion.

by our awareness of universals. That is, if experience has the same nature across hallucinations, illusion and veridical perceptions, and involves awareness of universals in all cases, then, when one is aware of the physical world in perceptual experience, such awareness seems to be mediated by one's awareness of universals. One's awareness of the physical world arises through that world satisfying the structure of universals that explains the phenomenology of one's experience in a given case. For example, if one has an experience as of a blue bottle, then one is aware of the structured universal <there is a blue, bottle-shaped object>. In the veridical case, to the extent that one is aware of the relevant blue bottle, one is aware of it by virtue of that bottle satisfying this structure of universals.

The problem is that this is starting to sound suspiciously like the framework espoused by SDT. While this doesn't commit to the *immediate perception* of universals and the *mediate perception* of the physical world, it is committed to the *immediate awareness* of universals and the *mediate awareness* of the physical world, in a sense of these terms that looks very friendly to the SDT framework. Suppose, for example, that SDT were altered so that, instead of perceivers immediately perceiving SD, perceivers are immediately *aware* of SD, and instead of perceivers mediately perceiving the physical world, perceivers are mediately *aware* of the physical world. I am not sure what of substance this would change in SDT. In this case, both the revised SDT and the generalized Dretske-Tye views would then be committed to the thesis that perceivers are immediately *aware* of some entity type X and mediately *aware* of the physical world (where, for all x, x ≠ the physical world). Put another way, to explain The Problem of Perception the union of NSD and nonconceptual content is committed to a distinction between immediate and mediate perceptual consciousness of a sort that NSD was designed to avoid, i.e. one that is akin to the sort utilized in SDT. That is my second claim.[28]

My third claim, which is a corollary of the second, is that the case of illusion brings out this problem in a particularly forceful way. I sketched this idea in the last section, but can state the problem with more precision given the resources that are now available. Illusions are challenging because they are semantically-hybrid states involving successful object perception and at least one unsuccessful property perception. When I experience the purple bottle as a blue bottle, I am successfully perceiving the bottle and the bottle shape but not successfully perceiving the bottle colour. Pretheoretically, the purple of the bottle is occluded by the experience of blue (however one wants to unpack the latter). According to NSD, during perceptual success I immediately perceive the front-facing surfaces of objects, and mediately perceive the objects themselves. Given that I am successfully perceiving the bottle, its shape and its location, it seems to follow (within NSD) that I am immediately perceiving the front-facing surface of the bottle. However, given that the purple of the bottle is perceptually occluded, and that the bottle's front facing surface is imbued with that purple, it seems hard to assert that I am immediately perceiving the front-facing surface of the bottle. This is a problem for NSD, even when supplemented with nonconceptual content.

[28] It is instructive to compare the above argument with that found in Kriegel (2011). Like me, Kriegel argues that views of the sort just discussed are committed to a veil of abstracta. However, he then argues that this must be avoided, and can be avoided by adverbialism, concluding that we should adopt adverbialism. I, unsurprisingly, reject the second stage of his reasoning. In the interest of space and focus, I cannot work through the details here.

I see no compelling way to resolve this tension without collapsing NSD into an SDT-friendly view. According to the solution outlined above, although I am immediately perceiving the front-facing surface of the bottle in the NSD sense of the term, I am nonetheless mediately aware of that surface in the sense that my awareness of the surface is mediated by my awareness of universals. That is, NSD seems to collapse into a view akin to SDT in the sense that NSD is committed to the claim that we are not immediately aware of the physical world, even during veridical perception.

In conclusion, this argument rests on two assumptions, the Dretske-Tye account of the phenomenology of hallucination, and a common-factor account of nonconceptual content. These assumptions are justified because they are defended by leading advocates of nonconceptual content. Nonetheless, the NSD advocate is welcome to suggest alternative assumptions.[29] For our purposes, I submit that I've adequately justified my conclusion: when extended to the Problem of Perception via an appeal to nonconceptual content, NSD is either rendered incoherent or collapses into a view that is akin to SDT in that both views deny that we are immediately aware of the physical world in perceptual experience.

8. Is it SDT, NSD or bust?

NSD and SDT are viable views, despite facing some challenges. As I hope to have showed, examining how they conceive of immediate and mediate perception, and how they explain demonstrative reference to ordinary physical objects, is incredibly instructive. Nonetheless, one might wonder whether we should embrace premise (3), that NSD and SDT are the only two viable theories of immediate perception. It is thus worth briefly considering an alternative that has received recent attention: naïve realism.

According to this view, we immediately perceive not the front-facing surfaces of physical objects but physical objects themselves. That is, naïve realism is built around the idea that we immediately perceive ordinary objects like trees and cars. In this regard the view differs from NSD. However, Bermúdez dismisses naïve realism because it is "no longer in play" (p. 368). While that was true in 2000, when Bermúdez's article was published, in our current climate naïve realism is very much in play.[30] Naïve realism satisfies the Reference Constraint, as the view can explain perceptual demonstratives by the immediate perception of ordinary objects. In this regard, were one to defend NSD in today's climate, one would arguably be obliged to provide reason to prefer NSD to naïve realism.

A defense of NSD over naïve realism can begin with the arguments NSD advocates offer in favour of treating the front-facing surfaces of physical objects, and not the whole objects themselves, as the immediate objects of perception (Section 3). Recall, one argument for NSD asserts that we visually

[29] For example, Pautz (2021) defends a quietist account of perceptual phenomenology. According to his view, in erroneous perceptual experience it indeed seems like one is aware of uninstantiated properties, but no positive account of this is needed. Why? Roughly, because a quietist view of this sort is superior to all non-quietist alternatives. To me, as a solution this has all the advantages of theft over honest toil.

[30] In fact, the recent interest in naïve realism arguably began very soon after Bermúdez's (2000). For example, two influential works came out in 2002: Campbell (2002) and Martin (2002). See Crane and French (2021) for an overview.

discriminate objects by their front-facing surfaces and thus perceive front-facing surface more immediately than we perceive objects themselves (Bermúdez, 2000, pp. 363-4). The other argument appeals to the psychological capacity for amodal completion (e.g. Briscoe 2011), which arguably entails that there is a primacy to the perception of the front-facing surfaces of objects over the perception of objects themselves.

There is much interesting work to do here, but this work sadly falls outside my present purview. I will emphasize that, however the debate between NSD and naïve realism is resolved, the challenge from SDT remains. SDT theorists demand an explanation of The Problem of Perception. Providing such an explanation is not only a difficulty for NSD, but also a well-known difficulty for naïve realism.[31] In this regard SDT again sits in a comfortable position.

There are, of course, other options that should be considered in a wider discussion. For example, a distinctive, recent account is Hill's (2022). On Hill's view, we do not directly experience objective properties of physical objects (like their colour and size). Instead, a special class of abstract properties – Thouless properties – are the "most immediate objects" of perceptual experience (p. 50).[32] Thus, Hill argues that "perceptual experience presents us with Thouless properties rather than their objective counterparts. This undercuts naïve or direct realism" (p. 57). This arguably also undercuts NSD and SDT.

There is thus much more to be said about premise (3). Nonetheless, even restricting our discussion to NSD and SDT (as premise (3) does) yields enormous fruit that can be applied to a broader suite of options in future work.

9. The underlying framework, without SDT and NSD

SDT and NSD appeal to different conceptions of the distinction between immediate and mediate perception. What has not been made explicit is that these conceptions do not depend on either theory. Here is each conception on its own:

- **Front-Facing Surface Conception**: the immediate$_1$ object of perception is a (front-facing) part of the mediate$_1$ object of perception. For example, you mediately$_1$ perceive the car by immediately$_1$ perceiving the front side of it.
- **Representation Conception**: the immediate$_2$ object of perception represents (and is distinct from) the mediate$_2$ object of perception. For example, you mediately$_2$ perceive a football match by immediately$_2$ perceiving a television.

[31] Naïve realists have offered numerous proposals for how they might account for illusions. See e.g. Brewer (2008), Genone (2014) and French & Phillips (2020). See Brown & Macpherson (2025) for a critical discussion.

[32] Thouless properties are a kind of appearance property that is postulated to explain deviations from perfect perceptual constancy found in vision. More precisely, regarding size, "The Thouless size of an object x with respect to viewpoint y is F(v, d), where v is the visual angle x subtends with respect to y and d is information pertinent to the distance from x to y. The function F is a "partial- constancy" function, in the sense that its values are more stable than ever-fluctuating retinal images but are nonetheless always characterized by under- constancy" (2022, p. 45).

Although NSD is built around the first conception and SDT around the second, each conception can stand on its own. Further, each can and should play an important explanatory role in perceptual theory. Let me briefly explain why.

Regarding the Front-Facing Surface Conception, it is hard to deny that we have more immediate perceptual access to the front-facing parts of things than we do to their backsides, insides or to the things as wholes. The difference can, for example, be sidelined in psychophysical studies of two-dimensional figures on computer monitors. But this restriction is obviously an artefact of experimental design, and ultimately three-dimensional physical objects are a, if not the, fundamental topic of perceptual psychology and of everyday perception. With three-dimensional object perception comes the Front-Facing Surface Conception of immediate and mediate perception. We thus can and should debate about how to understand the Front-Facing Surface Conception and its significance for issues like object perception and perceptual demonstratives. However, at the end of this debate some form of this conception of immediate and mediate perception will remain intact.

Assuming SDT is false, the Representation Conception arguably played a small role in our evolutionary history. Independently of this, it is playing a decidedly outsized role in our current lives, with the proliferation of televisions, smartphones, computers, and the like. It is straightforward to me that my son is not merely growing up with a blurred understanding of the reality of his immediate physical environment versus the reality represented by various screens to which he has access. He is arguably growing up without a sharp division between these realities. I won't pause to try to assess the value of this. My point is that what is mediately$_2$ perceived is arguably as real to him as what is immediately$_2$ perceived. Our theories of perception should pay much closer attention to the Representation Conception than they have. This is additionally true given the rapid development of virtual and augmented reality technology. The kinds of experiences found in these settings test the limits of our understanding of what is "immediately" versus "mediately" perceived, and of what we are referring to when we think and talk about what is experienced. The solution isn't to abandon the topic of immediate and mediate perception, it is to develop it in a way that permits it to illuminate our burgeoning perceptual reality.

It is also worth making explicit that both the Representation Conception and the Front-Facing Surface Conception can be applied to a single perceptual experience. For example, one can immediately$_1$ perceive a front-facing surface of an object, and mediately$_1$ perceive the object itself, and, simultaneously, if that object represents some further object, then one can mediately$_2$ perceive the further object by immediately$_2$ perceiving the former. In my view this is precisely what happens when one is watching a live football match on a smartphone.

Above I argued that both deferred and non-deferred demonstration can be applied, in various ways, to cases to which the Front-Facing Surface Conception can be applied, and to cases to which the Representation Conception can be applied. A full account of these matters should appeal to all of these resources, regardless of whether one adheres to SDT, NSD or some other view. We need both conceptions of immediate and mediate perception, we need both deferred and non-deferred perceptual demonstratives, and, although the details are complex, many instances of the latter should be explained via combinations of the former. In this regard the Reference Constraint, or something

in the vicinity, remains in place. I therefore see the underlying conceptual framework found in Bermúdez's 'Naturalized Sense Data', when properly understood, to hold enduring value.

10. Conclusion

Throughout the twentieth century, philosophers of perception often discussed how immediately we perceive ordinary physical objects, and, if not immediately, then how what is immediately perceived feeds into and justifies basic perceptual beliefs about ordinary physical objects. In recent years, philosophy of perception has moved away from this suite of issues. In this regard, Bermúdez's 'Naturalized Sense Data' may seem old fashioned. To my mind, viewing it as such is a mistake. There are different and legitimate conceptions of the distinction between mediate and immediate perception, and it is instructive to examine how each feeds into perceptual beliefs, such as those involving perceptual demonstratives about ordinary physical objects, regardless of whether one views these as basic beliefs. I argued that Bermúdez was wrong to contend that SDT could not adequately explain such demonstratives. Instead, I argued, there are multiple theories that can explain perceptual demonstratives. His preferred view, NSD, would also benefit from an account of The Problem of Perception: perceptual errors and their relation to veridical perceptions have always been important to the distinction between mediate and immediate perception. NSD is right to apply the distinction to veridical perceptions, but wrong to not also address perceptual error, and in particular illusion, given the forceful challenge illusion poses to the view. We can supplement NSD with nonconceptual content. However, in so doing we must be careful to not merely address the semantic component of The Problem of Perception. Addressing the phenomenological component requires an explicit account of perceptual phenomenology. I considered a proposal of Dretske and Tye, influential defenders of nonconceptual content. I argued that, in the end, the union of NSD with nonconceptual content yields either an incoherent view, or a view according to which we are not immediately aware of the physical world during veridical perceptions, but instead mediately aware of it. This introduces a kind of perceptual veil that is akin to that embraced by SDT – an outcome Bermúdez sought to avoid.

In conclusion, I suggest that the underlying framework Bermúdez articulates to guide his discussion is essential to perceptual theory, regardless of which specific theory you endorse. His distinction between immediate/mediate perception and direct/indirect perceptual beliefs is critical (see also Jackson, 1977; Snowdon, 1992). Further, his distinction between two types of immediate-mediate perception – one embodied in what I've called SDT (reminiscent of Russell, 1912) and one in his NSD (reminiscent of Moore, 1918-1919) – is under-appreciated and, to my mind, foundational to perceptual theory. These distinctions can and should be abstracted from SDT and NSD and applied to numerous case studies and other perceptual theories. That landscape is at present largely unexplored. I hope it won't be for long.[33]

[33] Project SENSOR (2024-2027), an AHRC-DFG funded project on "sensory engineering" has this as one of its main topics. Sensory engineering refers to ways perceptual experiences can be intentionally manipulated – engineered – via technologies of different sorts. See https://www.gla.ac.uk/research/az/cspe/projects/sensorsensoryengineering/.

Acknowledgements

This paper was written for the 30th SIUCC conference in honour of the work of José Luis Bermúdez, held in Valencia from 26-28 April 2023. I wish to thank the audience for helpful comments, and the conference organizers Victor Verdejo and Matheus Valente. Thanks are also owed to my colleagues David Bain, Robert Cowan, Jack Lyons and Fiona Macpherson, all of whom facilitated a discussion on this material that resulted in a dramatic improvement in my thinking. Thanks are also due to Robert Briscoe. Finally, thanks to the audience at the University of St. Andrews Philosophical Society. This work was supported by the Arts and Humanities Research Council AH/Y007638/1.

References

Bermúdez, J.L. (1995). Nonconceptual content: from perceptual experience to subpersonal computational states. *Mind and Language 10* (4), 333-369.

Bermúdez, J.L. (2000). Naturalized sense data. *Philosophy and Phenomenological Research 61*(2), 353-374.

Bermúdez, J.L. (2007). What is at stake in the debate on nonconceptual content? *Philosophical Perspectives 21* (1), 55–72.

Bermúdez, J. & Cahen, A. (2024). Nonconceptual mental content. In Edward N. Zalta & Uri Nodelman (eds.), *The Stanford Encyclopedia of Philosophy* (Fall 2024 Edition). URL = <https://plato.stanford.edu/archives/fall2024/entries/content-nonconceptual/>.

Brewer, B. (2008). How to account for illusion. In A. Haddock and F. Macpherson (eds.) *Disjunctivism: Perception, action, knowledge* (pp. 168–80). Oxford University Press.

Briscoe, R. E. (2011). Mental imagery and the varieties of amodal perception. *Pacific Philosophical Quarterly 92* (2), 153-173.

Brovold A. & R. Grush. (2012). Towards an (improved) interdisciplinary investigation of demonstrative reference. In A. Raftopoulos & P. Machamer (eds.), *Perception, realism, and the problem of reference* (pp. 11-42). Cambridge University Press.

Brown, D. H. (2009). Indirect perceptual realism and demonstratives. *Philosophical Studies, 145,* 377-94.

Brown, D. H. & Macpherson, F. (2025). *Colour illusion*. Routledge.

Campbell, J. (2002). *Reference and consciousness*. Oxford University Press.

Crane, T. & French, C. (2021). The problem of perception. *The Stanford Encyclopedia of Philosophy* (Fall 2021 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/fall2021/entries/perception-problem/>.

Dretske, F. (1979). Simple seeing. Gustafson, D.F., Tapscott, B.L. (eds) *Body, mind, and method. Synthese Library, 138* (pp. 1-15). Springer, Dordrecht. https://doi.org/10.1007/978-94-009-9479-9_1

Dretske, F. (1995). *Naturalizing the mind*. Cambridge, MA: MIT Press.

French, C. & Phillips, I. (2020). Austerity and illusion. *Philosophers' Imprint, 20*(15), 1-19.

Fodor, J. (1975). *The language of thought.* Harvard University Press.

Genone, J. (2014). Appearance and illusion. *Mind 123*(490), 339–376.

Hatfield, G. (2021). Sense data. In Edward N. Zalta & Uri Nodelman (eds.), *The Stanford Encyclopedia of Philosophy* (Fall 2021 Edition). URL = <https://plato.stanford.edu/archives/fall2021/entries/sense-data/>.

Hill, C. (2022). *Perceptual experience.* Oxford University Press.

Jackson, F. (1977). *Perception.* Cambridge University Press.

Johnston, J. (2004). The obscure object of hallucination. *Philosophical Studies, 120*(1), 113-183.

Kriegel, U. (2011). The veil of abstracta. *Philosophical Issues 21*(1), 245-267.

Lyons, J. (2023). Epistemological Problems of Perception. In Edward N. Zalta & Uri Nodelman (eds.), *The Stanford Encyclopedia of Philosophy* (Fall 2023 Edition). URL = <https://plato.stanford.edu/archives/sum2023/entries/perception-episprob/>.

Macpherson, F & Batty, C. (2016). Redefining illusion and hallucination in light of new cases. *Philosophical Issues 26*(1), 263-296.

Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information.* W. H. Freeman.

Martin, M.G.F. (2002). The transparency of experience. *Mind and Language, 17,* 376–425.

Martin, M.G.F. (2017). Elusive objects. *Topoi 36*(2), 247-271.

Moore, G. E. (1918-1919). Some judgements of perception. *Proceedings of the Aristotelian Society, 19,* 1-29.

Pautz, A. (2021). *Perception.* Routledge.

Peacocke, C. (1992). *A study of concepts.* MIT Press.

Quine, W.V.O. (1968). Ontological relativity. In W.V.O. Quine (ed.) *Ontological relativity and other essays* (pp. 26-68). Columbia University Press, 1969.

Robinson, H. (1994). *Perception.* Routledge.

Russell, B. (1912). *The problems of philosophy.* Williams and Norgate.

Russell, B. (1927). *Analysis of matter.* Kegan Paul.

Schellenberg, S. (2018). *The unity of perception: content, consciousness, evidence.* Oxford University Press.

Shapiro, L. & Spaulding, S. (2024). Embodied cognition. In Edward N. Zalta & Uri Nodelman (eds.), *The Stanford Encyclopedia of Philosophy* (Fall 2024 Edition). URL = <https://plato.stanford.edu/archives/fall2024/entries/embodied-cognition/>.

Snowdon, P. (1992). How to interpret 'direct perception'. In T. Crane (ed.), *The contents of experience* (pp. 48-78). Oxford University Press.

Tye, M. (1995). *Ten problems of consciousness.* MIT Press.

Tye, M. (2002). *Consciousness, color, and content.* MIT Press.

Tye, M. (2006). Nonconceptual content, richness, and fineness of grain. In Tamar Szabo Gendler & John Hawthorne (eds.), *Perceptual experience* (pp. 504-530). Oxford University Press.

Tye, M. (2014). What is the content of a hallucinatory experience? In Berit Brogaard (ed.), *Does perception have content?* (pp. 291-330). Oxford University Press.

DEREK H. BROWN worked at Brandon University (Canada) prior to joining the Centre for the Study of Perceptual Experience and Philosophy at University of Glasgow in 2017. He works on philosophy of perception, with particular interest in philosophy of colour and mediated perception. He is currently the UK-PI of the AHRC-DFG funded project *Sensory Engineering* (2024-2027), which is examining ways that sensory experiences can be intentionally manipulated via technologies.

ADDRESS: School of Humanities, University of Glasgow, G12 8QQ, UK. E-mail: derek.brown.2@glasgow.ac.uk – ORCID: https://orcid.org/0000-0002-5765-2443