

Naïve realism about unconscious perception

Paweł Jakub Zięba¹ 

Received: 20 March 2017 / Accepted: 13 September 2017 / Published online: 6 October 2017
© The Author(s) 2017. This article is an open access publication

Abstract Recently, it has been objected that naïve realism is inconsistent with an empirically well-supported claim that mental states of the same fundamental kind as ordinary conscious seeing can occur unconsciously (SFK). The main aim of this paper is to establish the following conditional claim: if SFK turns out to be true, the naïve realist can and should accommodate it into her theory. Regarding the antecedent of this conditional, I suggest that empirical evidence renders SFK plausible but not obvious. For it is possible that what is currently advocated as unconscious perception of the stimulus is in fact momentaneous perceptual awareness (or residual perceptual awareness) of the stimulus making the subject prone to judge in some way rather than another, or to act in some way rather than another. As to the apodosis, I show that neither the core of naïve realism nor any of its main motivations is undermined if SFK is assumed. On the contrary, certain incentives for endorsing naïve realism become more tempting on this assumption. Since the main motivations for naïve realism retain force under SFK, intentionalism is neither compulsory nor the best available explanation of unconscious perception.

Keywords Unconscious perception · Naïve realism · Intentionalism · Perceptual relation · Consciousness

Recently, a number of philosophers have pressed that naïve realism (a.k.a. relationalism) cannot deliver a plausible account of unconscious perception (see e.g. [Berger and Nanay 2016](#); [Block 2017](#); [Nanay 2014](#)). The objection assumes that naïve realism is inconsistent with the following claim:

✉ Paweł Jakub Zięba
pawelek.zieba@uj.edu.pl

¹ Institute of Philosophy, Jagiellonian University, Cracow, Poland

(SFK) Mental states of the *same fundamental kind* as ordinary conscious seeing can occur unconsciously.

Considering the large amount of empirical evidence in favour of SFK, the prospects of rejecting this claim are not promising. Although Phillips (2016) has raised some substantial doubts whether empirical evidence for SFK renders it unquestionable, he admits that this is far from enough to reject SFK on empirical grounds. And a rejection of SFK on purely theoretical grounds would be taken as ad hoc by the critics. Does it all mean that naïve realism hinges on a controversial interpretation of empirical data? Not necessarily. The naïve realist is not compelled to make the risky bet on the falsity of SFK insofar as it can be shown that naïve realism is compatible with SFK. The main goal of this paper is to establish precisely this compatibility. In particular, I demonstrate that the core of naïve realism, as well as its main motivations, retain force even if it is assumed that SFK is true. On top of that, I show that some incentives for embracing naïve realism become even more tempting on this assumption. And if motivations for naïve realism work under SFK, intentionalism is neither compulsory nor the best available account of unconscious perception.

In Sect. 1, I briefly outline the basics of the naïve realism versus intentionalism debate in philosophy of perception (Sect. 1.1), and spell out the intentionalist objections to naïve realism which rest on the assumption that this view is incompatible with SFK (Sect. 1.2). In Sect. 2, I assess the plausibility of SFK against the background of the ongoing debate between Phillips and Ned Block. Even though I am nowhere near as convinced of SFK as Block is, I think that Phillips' objections to SFK are assailable inasmuch as they rely on the distinction between personal and subpersonal levels of explanation. In Sect. 3, I show that the core of naïve realism (Sect. 3.1), as well as each of its main motivations (Sect. 3.2), maintain their argumentative force even if SFK is true. Section 4 gives a response to objections mentioned in Sect. 1.2 by sketching a naïve realist account of unconscious perception.

1 Naïve realism and the challenge of unconscious perception

1.1 Naïve realism versus intentionalism

Because I want to stay neutral with respect to internal debates within the naïve realist camp, it will be useful to extract the core of naïve realism, which is arguably common to all of particular formulations of this claim. I take it that the core of naïve realism comprises the following claims: (i) perception, understood here as a successful sensory encounter with the mind-independent world, is a direct relation between the subject and the mind-independent physical object; (ii) the relata of this relation, the perceiver and the perceived object, stand to each other as components of a unified state called 'perception' (this point distinguishes naïve realism from traditional direct realism); (iii) perceptual relation provides some cognitive import for the subject, which normally consists in creating an opportunity to gain knowledge about mind-independent reality.

Typically, motivating naïve realism takes form of a transcendental argument, i.e. it is argued that naïve realism is either the only or the best available way of accounting

for X.¹ This variable is filled with one or more of the following: (A) the possibility of demonstrative thought (Campbell 2002); (B) the transparency of sense experience (Martin 2002b); (C) the fine-grainedness of experience (Brewer 2011); (D) the particularity of experience (Martin 2004, 2006); (E) the primacy of veridical perception in the theory of perception (Martin 2006); (F) the possibility of thought about the external world (Travis 2013b); (G) epistemic humility regarding the judgments about the kind of experience one is having (Martin 2004); (H) relationality of experience without the commitment to sense data (Brewer 2011).²

Although this list is not exhaustive, it tenably encompasses the main motivations of naïve realism. By the same token, it can be seen as indicating the main points of disagreement between naïve realists and intentionalists. Intentionalism³ is a group of theories sharing a common commitment that all perceptual experiences are representational, i.e. that perceptual experiences have contents which represent the world as being a certain way. Naïve realists usually discard intentionalism on the grounds that positing perceptual representations is tantamount to introducing problematic intermediaries between the subject and the object of perception (see e.g. Travis 2013a), even though intentionalists usually do not consider themselves as doing so (see e.g. Burge 2005, pp. 30, 49).⁴

Intentionalists remain unconvinced by the (A–H) arguments. In some cases, they simply disagree with a claim of which naïve realism is advocated to be the best account. A clear example is rejection of (E) by Burge (2011). In other cases, however, the apple of discord is not whether a particular desideratum of naïve realism is in itself legitimate but whether accounting for it necessarily requires embracing naïve realism, and whether it allows or forbids positing perceptual representations. For instance, the naïve realist charge that only naïve realism can explain the particularity of experience faces the intentionalist response that this idea can be incorporated into the framework of perceptual representations (see e.g. Nanay 2012; Schellenberg 2014; Tye 2007). I think that the unconscious perception controversy is similar: while intentionalists argue that positing perceptual representations is indispensable to account for unconscious perception, it is in fact open for naïve realists to respond that their view has the resources necessary to explain this phenomenon. In the remainder of this section, I outline the intentionalist perspective. The naïve realist rejoinder is offered in Sects. 3 and 4.

¹ Stern (2007) has pointed out that transcendental arguments can be thought of as either ‘ambitious’ or ‘modest’. The former conclude how things are, the latter—how one should think about things in order to secure X. Whether the authors I am referring to take their arguments for naïve realism to be ‘ambitious’ or ‘modest’ would have to be determined on a case-by-case basis. To avoid unnecessary complications, I leave this issue aside.

² This list draws on (but is not identical to) a similar list in Locatelli and Wilson (2017, p. 201).

³ The idea that all perceptual experiences are representational is sometimes called ‘representationalism’ (see e.g. Nanay 2015). Following (Fish 2010, pp. 66–69), I call it ‘intentionalism’ instead, and reserve ‘representationalism’ for a specific version of this idea, according to which the phenomenal character of experience depends on the intentional content of experience.

⁴ Intentionalists often think of themselves as direct realists. They consider perceptual representations as means of directly perceiving the environment, not as intermediaries between the subject and the object of perception. One exception is the so-called phenomenology-first intentionalism (a.k.a. qualia theory) (see e.g. Block 2010).

1.2 The challenge of unconscious perception

According to the naïve realist orthodoxy, the notion of perceptual relation explains the phenomenology of perception. For the subject, being in perceptual relation is just having things in view. Thinking about perception in this way provides a simple yet comprehensive account of both qualitative character and cognitive import of the phenomenology of perception. Consequently, naïve realists often define perception as a specific kind of awareness, which, in turn, makes it seem obvious that their view precludes unconscious perception. And if naïve realism is understood as entailing that perception is conscious by definition, each bit of empirical evidence for unconscious perception puts this view into question. Thus, the problem for naïve realism arises when SFK is juxtaposed with the following putative commitment of naïve realism:

(PC) Perception is conscious by definition.

If PC is true, SFK is false. For PC entails the following claim:

(anti-SFK) Mental states of the *same fundamental kind* as ordinary conscious seeing *cannot* occur unconsciously.

In Sect. 3.1, I argue that naïve realism is not committed to PC, and thus consistent with SFK. To say that perceptual relation explains the phenomenal character of perception is not equivalent to saying that there could be no perceptual relation without phenomenal character. Nor does the possibility of perceptual relation without phenomenal character entail that perceptual relation does not explain the phenomenal character of conscious perception. While it is expectable that not every naïve realist is going to agree with me on this,⁵ I am not alone in the naïve realist camp (broadly construed) in thinking that perception can be unconscious. The position I advance here has been closely anticipated in the work of Hellie (2014a, pp. 248–250).⁶

Still, even if my proposal is accepted, unconscious perception does not cease to be problematic for naïve realism. That is to say, even if the core of naïve realism can be rescued, unconscious perception can still be used to undermine the reasons for endorsing it. In fact, the unconscious perception-based objections to naïve realism usually concern the motivations of the view, not the view itself. They are problematic for naïve realism whether it is committed to PC or not. For they purport to show that naïve realism fares worse than intentionalism at explaining certain perceptual phenomena.

For example, Block argues that accounting for unconscious perception must consist in introducing unconscious perceptual representations, and ‘if unconscious seeing is unconsciously representing, there is a strong case that conscious seeing is consciously

⁵ Note that my proposal rules out another stance the naïve realist could take towards SFK, namely that the naïve realist analysis does not apply to unconscious perception. In this respect, no less is my proposal a rejoinder to naïve realism’s critics than it is an offer to the proponents of naïve realism which they may or may not accept.

⁶ Hellie (2011, 2013a, b, 2014a, 2018), who draws on works of Rosenthal (1986) and Ryle (1949), uses the possibility of unconscious perception to argue that questions about perception should be sharply distinguished from questions about consciousness, and that a ‘non-objective’ conception of consciousness should be embraced. While I acknowledge that a lot of what I say below can be used to reinforce these desiderata, I want to stress that my argumentation in this paper does not presuppose them.

representing' Block (2017, p. 169). Simply put, the possibility of unconscious perception supports intentionalism.

Berger and Nanay (2016, p. 431) and Nanay (2014, pp. 42–45) raise a bit more intricate objection to naïve realism. It starts with the observation that visual system in humans and other mammals has two subsystems (dorsal and ventral). The former is responsible for visual control of motor actions, whereas the latter carries out identification and recognition. The problem for naïve realism arises in cases where the cooperation between the subsystems malfunctions. For example, in Ebbinghaus and Müller-Lyer illusions, the dorsal subsystem is much less prone to be fooled by the misleading nature of the stimulus than the ventral system. In effect, the subject can visually misrecognise the size of the stimulus, and yet retain the ability to manipulate it efficiently, as if the recognition was correct. The intentionalist can explain this by ascribing two different perceptual representations to the two visual subsystems. But how is the naïve realist supposed to account for this duality with a single perceptual relation? Moreover, an analogous problem arises in regard to multimodality of perception. How to explain the integration of information delivered by numerous sense modalities with a single perceptual relation?

Here (Nanay 2014, pp. 44–45) presents the naïve realist with a dilemma. On the first horn, the naïve realist might respond that both property ascription by the two visual subsystems and multimodal integration are unconscious subpersonal processes, while the term 'perception' should be reserved for conscious effect of those. However, the reply entails that there is no such thing as unconscious perception, which amounts to contradicting a well-established empirical claim. On the second horn, the naïve realist might allow for unconscious perceptual representations in order to accommodate the phenomena in question. This response is also unsatisfactory, since it renders naïve realism as a claim about conscious perceptual experience only, not about perception in general.

Berger and Nanay (2016) indicate three possible lines of defence the naïve realist can choose from. Needless to say, Berger and Nanay find them all unpromising.⁷ The first option is to persist with PC and follow Phillips in questioning the empirical evidence for unconscious perception. I evaluate this strategy in Sect. 2. Although I am not entirely convinced that unconscious perception is an unquestionable empirical fact, I think that Berger and Nanay are right that 'such an approach would hold relationalism hostage to forthcoming experimental results' (Berger and Nanay 2016, p. 2), which speaks against embracing this strategy. Naïve realism is a metaphysical claim, not an empirical one. As such, it ought to be evaluated on both empirical and theoretical (as opposed to solely empirical) grounds (cf. Paul 2012, p. 18).

The second option consists in rejecting PC while simultaneously retaining the commitment to anti-SFK (PC entails anti-SFK, but not *vice versa*). According to this view, conscious perception is relational, whereas unconscious perception is not. PC is false because perception can be unconscious, but anti-SFK is true because conscious and unconscious perception differ with respect to relationality. Berger and Nanay predict that this general idea might be developed in two directions. On the first version,

⁷ For a defence of the view that all three routes are promising ways of defending naïve realism, see Anaya and Clarke (2017).

consciousness transforms unconscious perceptual representations into conscious perceptual relations. But what would such transformation consist in? And what theory of consciousness should one endorse to make such view palatable? One might expect that an answer to these questions would be a philosophical equivalent of a Heath Robinson contraption. Thus, I concede to Berger and Nanay that such view would not be as convincing as its simple and intuitive intentionalist counterparts: representationalism and second-order intentionalism. On the second version, conscious and unconscious perception are radically different phenomena. This view is also unsatisfactory, as it could be undermined by empirical evidence that the difference between conscious and unconscious perception is a difference of degree, not of kind (see Sect. 2.1). Moreover, as I understand naïve realism, it entails that two experiences are of different fundamental kinds if they have different objects. On the assumption that physical objects can be perceived both consciously and unconsciously, the idea that two experiences of physical objects are radically different seems to be in tension with naïve realism.

The last alternative is to jettison both PC and anti-SFK by embracing the claim that both conscious and unconscious perception are relational phenomena. I think this is precisely the way the naïve realist should go. Berger and Nanay rightly notice that opting for this view pledges one to explain the difference between perceptual relations with phenomenal character and perceptual relations without phenomenal character in terms of the difference between the relata of these relations. However, they discard this strategy by suggesting that naïve realism remains unable to explain the aforementioned possibility of discord between the two visual subsystems even after being disburdened of PC and anti-SFK.

I respond to all these objections in Sects. 3 and 4. Before that, let us see what makes certain intentionalists so confident about the possibility of unconscious perception, and what reasons might be given to resist it.

2 Empirical evidence for unconscious perception

In this section, I assess SFK against the background of the main themes in the ongoing debate between Block and Phillips. Block (2017) contends that unconscious perception is an unquestionable empirical fact. As it turns out, this is a precipitant diagnosis (Sect. 2.1). However, Phillips' criticism of empirical evidence for SFK is also flawed (Sect. 2.2), and does not constitute a promising strategy for defending naïve realism (Sect. 2.3).

2.1 Jumping to conclusions

On the one hand, there is a very large body of empirical evidence in favour of the claim that mental states of the *same fundamental kind* as ordinary sense perception can occur unconsciously (SFK). Unconscious perception is claimed to occur in certain clinical cases (e.g. type-1 blindsight,⁸ unilateral neglect, visual agnosia), as well as in many experimentally induced cases (most notably subliminal priming and continuous flash

⁸ Psychological literature distinguishes two types of blindsight. Type-1 blindsight involves no awareness of the stimulus. Type-2 blindsight involves residual awareness of the stimulus.

suppression) (see e.g. [Block 2012](#); [Breitmeyer 2015](#)). In all these cases, subjects report no consciousness of the stimulus despite performing with above-chance accuracy in experimental tasks which require perception of the stimulus. As a result, SFK is widely considered as unquestionable empirical datum.

On the other hand, there are some good reasons to suspect that the inference from all this data to the conclusion that perception can be unconscious is too hasty. According to [Phillips \(2016, 2017, manuscript\)](#) for each putative case of unconscious perception, it is hard to rule out that the inference from empirical data to SFK rests on some methodological artefact. There is always a worry whether (i) the subject is really unconscious of the stimulus, or whether (ii) the mental state in question is genuinely perceptual.

In experimental practice, two thresholds are used to determine whether the subject is conscious: subjective and objective. The former is met if the subject reports that she is conscious of the stimulus. The latter is met if the subject is able to discriminate the stimulus. [Phillips \(2016\)](#) points out that there can be cases where the objective threshold is met, while the subjective is not. Namely, there can be cases where the stimulus is discriminable, albeit it barely stands out from the noise. In such cases, the subjects are prone to conservative response bias, i.e. to report that they did not perceive anything (when asked a ‘yes/no’ question), or that the stimuli are the same (when asked a ‘the same/different’ question). Phillips concludes that such bias is very likely to occur in experiments on patients who suffer from blindsight and neglect. The fact that they report no consciousness of the stimulus can be explained as due to residual conscious perception of the stimulus in conjunction with conservative response bias. This explanation becomes particularly convincing when two further facts are taken into account: (i) blindsight patients’ reports regarding the phenomenal character of their experiences are sometimes contradictory; (ii) both type-1 blindsight, which involves no awareness of the stimulus, and type-2 blindsight, which involves residual awareness of the stimulus, occur in the same patients, and thus stem from the same brain damage ([Overgaard and Mogensen 2015](#), pp. 34–35). Therefore, positing unconscious perception is not the only plausible explanation of blindsight and neglect, which in turn casts doubt on the inference from these cases to SFK.

Block does not persist on blindsight and neglect, however. He grants that reports of people with brain damage should be taken with a grain of salt ([Block 2016](#), p. 457). According to Block’s Anna Karenina theory of unconscious perception, ‘all conscious perceptions are alike but each unconscious perception is unconscious in its own way’ ([Block 2016](#), p. 452). If each kind of unconscious perception is different, and, apart from the clinical cases, there are 24 other kinds concerning people with healthy brains, it is improbable that they all can be undermined with a single objection. Here Block refers to a paper by Bruno Breitmeyer, who enumerates 24 methods of non-invasive suppression of phenomenal awareness of visual stimuli ([Breitmeyer 2015](#)).

The number is impressive indeed, which makes it quite disappointing to discover that the differences between these methods are so nuanced that the list can be easily boiled down to 10 categories ([Breitmeyer 2015](#), p. 239). These are further systematized into a functional hierarchy, which indexes each blinding technique to one of 8 various levels of unconscious processing. Moreover, Breitmeyer makes it clear that visual

crowding, Gestalt-switching effects, and attentional effects suppress access consciousness of the stimulus only, leaving phenomenal consciousness intact (Breitmeyer 2015, pp. 236, 239). Thus, they do not demonstrate unconscious perception because sustaining phenomenal consciousness entails that the objective threshold is met. Simply put, not every method of suppressing consciousness yields unconscious perception. As a result, the list of 24 kinds of unconscious perception is reduced to 6 various kinds of priming and flash suppression (which have a lot in common) plus transcranial magnetic stimulation (which arguably yields hallucination rather than perception). Out of 24, we came down to 2.

Of course, it is not my intention to suggest that distinguishing 24 methods is just splitting hairs. Taking the differences between the methods into consideration is essential for determining which one is the most appropriate for a given explanatory task. From this perspective, it certainly does matter what masks should be used, whether they should flicker and/or be swapped between the eyes, and how long the resulting suppression can last. What I find dubious, however, is whether these nuance differences make SFK any more plausible.

The experimental blinding techniques provide solid evidence for SFK only if they can suppress consciousness of the stimulus *completely*. As the recent study by Gelbard-Sagiv et al. (2016) shows, this is not necessarily the case. The results suggest that unconscious processing of high-level features of the stimulus (e.g. face identity, fame) requires conscious perception of low-level features (e.g. colour, location). Furthermore, even if participants were asked to perform tasks on low-level features only, it would still be possible that consciousness of certain features (e.g. shape, colour, size) is suppressed, while consciousness of other features (e.g. motion, gray level, spatial location) is not. For this reason, Breitmeyer finds ‘calls for *complete absence* of phenomenological content to be *the* criterion for invisibility as misdirected’ (Breitmeyer 2015, p. 240, *emphases in original*). Breitmeyer’s point is that specific methods render specific features unconscious, leaving consciousness of other features sustained. If this is correct, then it is possible that conscious perception of non-suppressed features of the stimulus enables or at least facilitates perception of the suppressed features.

There are couple reasons for taking this possibility very seriously. First, there are good reasons to think that consciousness of low-level features is graded (see Windey and Cleeremans 2015). If this is correct, the distinction between the conscious and the unconscious is so blurry that it becomes unfeasible to sort out the exact threshold of consciousness. There will always be a possibility of there being a degree of minimal consciousness below the threshold. Second, unconscious sensory representations might be insufficient for creating stable associations (see Raio et al. 2012), which severely limits the presumptive cognitive import of unconscious perception. Third, perception is not an isolated state. Usually, it is subject to top-down influences, and it is multimodal. Although the subject might not be conscious of the target stimulus visually, she can still become aware of it through other channels, e.g. other sense modalities, or inferences from perceptual judgments based on conscious perception of other stimuli and beliefs stored in her memory. It seems unfeasible to isolate a per-

ceptual episode from all such influences. And even if it were possible, such situation would be extremely artificial.⁹

In the light of all that, the question whether perception can be unconscious becomes much more difficult to answer. Gelbard-Sagiv et al. (2016) say that it is too early to generalise their results. But suppose it is confirmed by future studies that conscious perception of low-level properties of the stimulus is necessary for, or correlates with, unconscious perception of its high-level properties. Now it might be asked: is the stimulus consciously or unconsciously perceived? It seems that giving a clear-cut response to this question is impossible without adopting one of the competing philosophical theories of perception. For naïve realism emphasises the primacy of objects (as opposed to properties) in perception, whereas intentionalism takes exactly the opposite stance. Under naïve realism, the so-called ‘perception of high-level properties’ is basically bringing what is perceived under a category, i.e. perceptual judgment (cf. Travis 2013d, p. 31). The naïve realist interpretation of such findings would therefore have it that momentaneous perceptual awareness (or residual perceptual awareness) of the stimulus makes one prone to judge in some way rather than another, or to act in some way rather than another.¹⁰ The intentionalist, on the other hand, would probably say that perception of high-level properties is unconscious. Under intentionalist lights, then, the situation in question could be described as an integration of conscious and unconscious perceptual contents (cf. Block 2012, pp. 181–182). At any rate, the mere possibility of the naïve realist interpretation being true suffices to regard Block’s claim that 24 blinding techniques somehow cumulatively render SFK obvious as inconclusive.

2.2 A dead end

Although currently available empirical evidence for unconscious perception is inconclusive, some method of suppressing consciousness could eventually turn out sufficient to establish SFK. In comparison to the clinical cases, subliminal priming and continuous flash suppression constitute evidence for SFK that is much harder to resist. On the standard interpretation of these methods, the stimulus influences the subject’s behaviour despite not being discriminable by the subject (i.e. the objective threshold is not met). The hypothesis that the subject is not completely unconscious of the stimulus is much less credible here than it is in the case of blindsight and neglect. Presumably for this reason, Phillips puts more pressure on questioning whether the effects of subliminal priming and continuous flash suppression qualify as genuine cases of perception. He contends that, even under the assumption of Burge’s intentionalist theory of perception (which is endorsed by Block), they do not. According to Burge, perception is a sensory objective representation by the individual (Burge 2010, p. 368). The main point in Phillips’ criticism is this: because the representation

⁹ I thank Sebastian T. Kołodziejczyk for drawing my attention to this.

¹⁰ One might think that the naïve realist would have to posit unconscious perceptual judgment, i.e. unconscious recognition, to explain this situation. This, in turn, could be used as an objection to John Campbell’s claim that conscious experience of X is necessary to form a demonstrative judgment about X, which would amount to undercutting one of the main motivations of naïve realism. However, this is not the only route the naïve realist could go. A plausible alternative is offered in Sect. 3.2.1.

in question is not available to central agency, there is no reason to ascribe it to the individual. Consequently, such representation should be considered subpersonal, as opposed to personal. It is not perception because it is not a state of the subject. It is merely a state of the perceptual system (Phillips 2017, p. 181).

I find this reference to Dennett (1969) famous distinction between personal and subpersonal level of explanation deeply problematic. To the extent to which Phillips' criticism of SFK depends on this distinction, it is doomed to fail. I base this diagnosis on three observations.

The first observation is more of a clarification than actual objection to Phillips. It concerns what the personal/subpersonal distinction is a distinction of, namely, two types of explanatory account. Drayson (2012) has rightly noticed that the difference between such semantically evaluable items does not necessarily map onto corresponding difference between their truthmakers. Put differently, distinguishing between personal and subpersonal explanations does not have to correspond to any particular distinction between personal and subpersonal states. Introducing the latter distinction would require making some additional metaphysical claims, albeit at the cost of blurring the distinction. For, at least under the naturalist framework, 'personal' and 'subpersonal' would often pick out the same psychological state. This is because personal and subpersonal explanations often pertain to the same *explanandum*; they just explain it in different ways. There are, of course, subpersonal explanations of phenomena which do not correspond to any personal explanation, but this is not true of all subpersonal explanations. In this regard, it does not make sense to say that a representation is not perceptual because it is subpersonal. Following (Drayson 2012, pp. 12–14), I suggest that a different distinction is appropriate here: the doxastic/subdoxastic distinction coined by Stich (1978). Employing it allows to differentiate those subpersonal explanations that share their *explananda* with personal explanations from those that do not. The latter kind of subpersonal explanations is concerned with subdoxastic states, i.e. states carrying information that is not consciously accessible to the subject and 'largely inferentially isolated from the large body of inferentially integrated beliefs to which a subject has access' (Stich 1978, p. 507).

Therefore, when Phillips claims that representation of the stimulus in subliminal priming and continuous flash suppression is not ascribable to the individual, he can be charitably interpreted as claiming that the representation in question is a subdoxastic state, not a subpersonal state. Still, is Phillips right that subdoxastic states are not ascribable to individuals? This brings up the second observation I have foreshown above. Let me illustrate it with an example.

One of the main themes in the Block versus Phillips debate on unconscious perception concerns a study by Jiang et al. (2006). The study uses the paradigm of continuous flash suppression to demonstrate that unconsciously perceived erotic stimuli can guide spatial attention. Such stimuli are shown to either attract or repel the subjects' attention in accordance with their gender and sexual orientation. Block claims that the subjects in this study unconsciously perceive the stimuli because the content of the relevant representations reflects 'personal-level values and understanding' Block (2017, p. 182). He also raises the significance of the data in question by cautioning that this is the only study he knows of 'in which personal level preferences are so strongly revealed in unconscious perception' Block (2017, p. 187).

Phillips responds that this phenomenon can be explained as a reflexive, automatic response, unavailable to central agency. As such, it is not ascribable to the individual (Phillips 2017, p. 181). Even if the availability to central agency is not a necessary condition for attribution to the individual, its presence would be sufficient for such attribution. Conversely, ‘when a representation is unavailable to central agency, we lack a positive ground for attribution’ (Phillips 2017, p. 181). If an *explanandum* is not ascribable to the individual, its explanation is not a personal level explanation. It only explains components of persons, i.e. states of perceptual systems of persons. Since perception is supposed to be a personal state (i.e. a state of the individual), an explanation of perception has to be personal level explanation. From this perspective, the bone of contention between Block and Phillips is whether unconscious representations of the stimuli in Jiang’s study belong to the personal, or the subpersonal level of explanation, i.e. whether they explain states of persons, or merely states of their cognitive systems.

But as we have already seen, the real issue is whether the representations in question are doxastic, or subdoxastic states. Block opts for the former, Phillips—for the latter. However, this terminological refinement does not make the personal level explanations autonomous from the subpersonal level explanations. Nor does it make the doxastic states autonomous from the subdoxastic states. Firstly, Stich allows for unconscious inferential connections between doxastic and subdoxastic states (Stich 1978, pp. 507–511). Secondly, Phillips’ argument from unavailability to central agency is ill-founded. As the classic (and, to my knowledge, still not outdated) paper by Nisbett and Wilson (1977) informs us, people have no reliable introspective access to their higher-order psychological processes. When asked about reasons for their choices and actions, they either apply implicit, a priori causal theories, or simply confabulate. In this regard, it does not make sense to say that the stimulus does not guide action because it is unconscious. For it is a common phenomenon, not an exception, that what guides people’s behaviour is not consciously accessible to them.¹¹ Thirdly, as (Bermúdez 2000) and (Wong 2014) have more recently pointed out, there is no good candidate for a determinant of the autonomy of personal level explanation.¹² Although both personal/subpersonal and doxastic/subdoxastic are genuine distinctions, there is no convincing theoretical obstacle to cross-level explanations.

2.3 How not to defend naïve realism

Of course, it is still open for enthusiasts of the personal level’s autonomy to defend the way Phillips uses the personal/subpersonal distinction. Nevertheless, my third issue with this strategy is that it marks out the bounds of perception in a way that is inconsistent with naïve realism, which means that it does not lend itself to defend this view.

¹¹ I am grateful to Marcin Miłkowski for pointing this out to me. For more recent advocacy of the view that introspection is highly untrustworthy, see (Schwitzgebel 2008).

¹² Note that the original exposition of the personal/subpersonal distinction by Dennett did not involve the claim that the personal level is autonomous from the subpersonal level.

Those of naïve realists who want to stick by PC could try to advance further arguments for the distinction between personal and subpersonal states, and in this way unlock the dead end I have indicated in the previous subsection. But such attempts would be futile because naïve realism construes explanation of perception as neither personal nor subpersonal. On this view, the perceived object is literally a proper part of perception. This entails that any comprehensive explanation of perception has to be superpersonal, in that its *explanandum* reaches beyond person. Under naïve realism, even a personal explanation of what happens to the perceiver is not an exhaustive explanation of perception: it only explains the subperceptual state of the perceiver, i.e. only one relatum of perceptual relation.¹³

Let me illustrate this point against the background of Tyler Burge's rejection of one of the implications of naïve realism known as 'disjunctivism'. According to Burge (2005, 2011), since vision science groups perception and hallucination under the same fundamental category, denying that perception and hallucination are of the same fundamental kind (i.e. accepting disjunctivism) amounts to contradicting well-established science. There is no contradiction here, however, because naïve realism and disjunctivism pertain to perception as a realised state, whereas vision science is primarily concerned with the capacity of the subject to get into that state. Because the capacity to perceive is fallible, some exercises of this capacity fall short of perception. Burge's point is that all exercises of the capacity fall under the same fundamental category because they are exercises of the same capacity. If perception is understood as the successful exercise of the capacity to perceive, it is of the same fundamental kind as unsuccessful exercises of this capacity. But the naïve realist does not conceive of perception as the successful exercise of the capacity to perceive. She thinks that perception is *the result* of the successful exercise of the capacity to perceive. That the capacity is fallible does not entail that the result of its successful exercise is of the same fundamental kind as the result of its unsuccessful exercise. Just as a chick is not of fundamentally the same kind as a bad egg just because it could have failed to hatch out. Various exercises of the same capacity can result in fundamentally different mental states.

This is precisely why naïve realism renders the exhaustive explanation of perception superpersonal. To explain the successful exercise of the capacity to perceive, whether on personal or subpersonal level, or in a way that links both levels together, is to explain only one relatum of perceptual relation. If naïve realism is correct, this is not yet a complete account of perception.

For the naïve realist, SFK means that the result of the successful exercise of the capacity to perceive can occur unconsciously. Although it is not obvious that SFK is true, the prospects of rejecting this claim are uncertain. Even though Phillips raises some further worries concerning both Jiang's experiment and continuous flash suppression in general (see e.g. Phillips 2017, pp. 174–175, 180–181), invoking the personal/subpersonal distinction on subliminal priming and continuous flash suppres-

¹³ Failing to recognise that perception is a superpersonal phenomenon results in claiming that 'our perceptions are fantasies that coincide with reality' (Frith 2007, p. 135) and 'perception is controlled hallucination' (Clark 2016, p. 14). To the naïve realist ear, such slogans sound utterly preposterous. For the naïve realist, the only way to make them palatable is to replace the term 'perception' with 'perceptual judgment'.

sion is one of the main strategies he uses against SFK,¹⁴ and to this extent his criticism of SFK remains unsatisfactory. Thus, questioning SFK is not a promising way of defending naïve realism.

3 Naïve realism about unconscious perception

Let us put all the doubts aside and assume, for the sake of argument, that SFK is true. Why think that naïve realism is still a viable option? Because under naïve realism, perception is first and foremost a direct relation between the subject and the mind-independent object. That is to say, perception is relational no matter whether it is conscious or unconscious. In what follows, I spell this out in detail (Sect. 3.1), and consider the ways in which SFK affects the main motivations of naïve realism (Sect. 3.2).

3.1 Is naïve realism inconsistent with SFK?

I suggest that the conclusion should be drawn from the conjunction of naïve realism and SFK, namely that there can be perceptual relations such that the subject is not conscious of the object. For short, I will call such perceptual relations *unconscious perceptual relations*, and distinguish them from *conscious perceptual relations*, i.e. perceptual relations such that the subject is conscious of the object. These two are, of course, terms of art; I do not mean it literally that relations are conscious or unconscious. Instead, I claim that the occurrence of a perceptual relation may or may not result in the subject becoming conscious of the object. If consciousness is graded, there is a whole spectrum of partially conscious perceptual relations between these two extremes. On this exposition of naïve realism, whether a given perceptual relation is conscious (or to what extent is it conscious) turns on the state of its relata, i.e. the subject and the object. By ‘the subject’ I mean ‘a living organism capable of sense perception’. Paradigm exemplars of the subject thus understood are participants of empirical experiments conducted by vision science. By ‘the object’ I mean whatever the subject is perceptually related to. In particular, ‘the object’ is constituted by the totality of the perceived scene.

While this formulation of naïve realism is consistent with SFK, it is also consistent with the view that conscious perception is a paradigmatic case of perception. That view regards cases of unconscious perception as non-paradigmatic. They can be classified as either defective (as in clinical cases) or immature/not fully developed (as in subliminal priming, continuous flash suppression) cases of perception. Therefore, allowing for the possibility of unconscious perceptual relations *does not* entail that conscious and unconscious perception, though different regarding consciousness, are on a par in every other respect.

But even with this caveat in place, the idea of unconscious perceptual relation will be rejected by those who insist that perceptual relation is necessarily a direct

¹⁴ For instance, Phillips (2017, p. 174) adopts this strategy to deny that the Raio et al. (2012) study constitutes evidence for unconscious perception of threatening stimuli.

conscious acquaintance with the mind-independent physical object (i.e. the proponents of PC). For example, they might object that the only ‘subject’ that can be a relatum of perceptual relation is ‘the subject of conscious experience’. To purport that the subject may feature as a relatum in both conscious and unconscious perceptual relations is to equivocate on ‘the subject’.

To this, one might respond that viewing the perceptual relation as necessarily conscious is not essential to naïve realism. Rather, it is a consequence of arguing for naïve realism in a specific way. As Nanay has rightly suggested, ‘the difference between the relational and the representational view of perception could be framed in such a way that it does not concern the nature of perception, but rather the individuation of perceptual states’ (Nanay 2015, p. 325). Nanay’s insight is in line with Tim Crane’s remark that the main bone of contention between intentionalism and naïve realism concerns whether perception is essentially a relation (Crane 2006, p. 133). Intentionalists say “no” because the core of their view is that perception is essentially a representation which may or may not relate the subject to what it represents. Naïve realists say “yes” because the core of their view is that perception is a relation between the subject and the mind-independent physical object. Whether this relation is necessarily conscious is a further question.

This response is correct, yet unsatisfactory. For what is really at stake is whether perceptual relation can be both unconscious and direct. The naïve realist will be reluctant to the idea of unconscious perceptual relation insofar as accepting it might be seen as tantamount to revival of the so-called screening-off problem.

The ‘screening-off’ is a specific relationship between three variables: X, Y, and Z. If X is present in a situation S, then whether Y is present in S is probabilistically independent of the presence of Z in S (Hellie 2014b, p. 153). To see how the screening-off problem for naïve realism arises, consider the following pair of experiences: a perception and an hallucination that is phenomenally indistinguishable from that perception. It seems reasonable that both experiences have qualitatively identical phenomenal characters and neural bases. These two features constitute a common factor between the two experiences. Whether an experience is a perception, or an indistinguishable hallucination, the common factor is present. The occurrence of the common factor suffices for the subject to formulate a perceptual thought of the sort “That is O.” It is sufficient in the sense that whether the subject is perceiving or hallucinating, as far as the subject is concerned, the subject has an experience of O. This is all that is needed to explain how the thought “That is O.” came into existence. Consequently, it is irrelevant to this explanation whether O is really present and experienced by the subject. For the occurrence of the thought “That is O.” is probabilistically independent of the occurrence of the perceptual relation with the actual O. Perceptual relation is therefore screened-off by the common factor, where ‘screened-off’ amounts to ‘rendered explanatorily irrelevant’ (cf. Martin 2004, pp. 61–62).

Naïve realists typically respond to the screening-off problem by adopting a kind of disjunctivism about phenomenology, i.e. the view that the phenomenal character of a perception cannot be qualitatively identical to the phenomenal character of an hallucination even if those experiences are indistinguishable from the first-person perspective. Various ways of specifying this general idea include: (i) negative disjunctivism, i.e. the view that the only positive characteristics that can be given to an hallucination is

that it is indiscriminable from the first-person perspective from a perception (Martin 2004); (ii) eliminativism about the phenomenal character of hallucination (Fish 2009); (iii) multidisjunctivism, i.e. the view that no perception shares its fundamental kind with any hallucination, and some hallucinations fail to share fundamental kinds with other hallucinations (Hellie 2014b).

Whether successful or not, all these strategies presuppose that the phenomenal character of perception is at least partially constituted by the mind-independent physical object. This assumption is necessary to remove the common factor, i.e. to show that the phenomenal character of perception and the phenomenal character of subjectively indistinguishable hallucination are not qualitatively identical, even though they might appear identical from the first person perspective. But does this rejoinder still work if perception can be unconscious?

As long as it is presupposed that PC is true, the naïve realist can maintain that perception consists in fragments of mind-independent world ‘shaping the contours of perceptual experience’ (Martin 2004, p. 64). Less ornately, the phenomenal character of perception can be explained in terms of direct relation between the subject and the objective properties of the mind independent object. But if it is assumed that perception can be unconscious, it seems to follow that perception cannot be direct in this way. The assumption seems to favour the view that the mind-independent object can affect the perceptual system whether the subject is directly aware of it or not. Unconscious perception can modify consciousness only indirectly, by affecting thoughts and actions of the subject. Accounting for unconscious perception requires positing perceptual representations, which of course brings the common factor back again.

However, there is a good reason to think that granting SFK does not resurrect the screening-off problem. The objection in the previous paragraph presupposes that:

(T) A mental state X is ascribable to the subject only if X is conscious.

Nevertheless, T has been out of the picture at least since SFK has been assumed. For SFK is antithetical to T. Simply assuming T would thus amount to begging the question against SFK.¹⁵ Although PC is consistent with denial of T, I surmise the latter might constitute an important incentive to endorse the former for many PC-naïve realists. The way Phillips employs the personal/subpersonal distinction provides an illustration of how this might work.¹⁶ Provided that T is ruled out, there can be unconscious mental states of the subject. Crucially, if some of these mental states are partially constituted

¹⁵ If T were proven indispensable, Phillips’ personal/subpersonal objection to unconscious perception would be validated, and the naïve realist would not have to worry about SFK at all.

¹⁶ This is not to say that Phillips himself is a PC-naïve realist. The following passage suggests that he regards the kind of naïve realism about unconscious perception that I am advancing here as at least a viable option: ‘Could a hypothetical naïve realist think of unconscious perception as involving just such relations? It might be objected that such relations are intended to explain phenomenal character. Yet consider two imperfect analogies. Suppose a good life constitutively involves personal relationships, relationships which partly explain why that life is good. It does not follow that personal relationships of the same kind cannot occur within a bad life. Suppose (with Russellians) that true propositions have mind-independent objects amongst their constituents, objects whose identities partly explain such propositions’ truth values. It does not follow that false propositions lack mind-independent objects as constituents’ (Phillips 2017, p. 176).

by the mind-independent physical object, it follows that the subject is capable of becoming sensorily acquainted with such objects both unconsciously and directly.

Block (2017) motivates SFK with the following alternative to T:

(T*) A mental state X is ascribable to the subject if X has a proper content and plays a proper functional role for the subject.

He uses empirical evidence for unconscious perception (e.g. Jiang et al. 2006) to illustrate how unconscious perception can have the relevant content and play the relevant role. Contrary to his own declarations, however, the stronger the evidence he provides, the more viable it is for the naïve realist to argue that unconscious perception fulfils the T* condition precisely because unconscious perception is relational.

Unconscious perception is perception in virtue of having cognitive import for the subject.¹⁷ On the one hand, this import is severely limited, and can be assessed only *post hoc*, by showing that the subject's behaviour was different than it would have been had no perception happened. On the other hand, unconscious perception can be regarded as having such import for the subject only if it puts the subject in direct relation with the mind-independent reality. And as long as it is assumed that unconscious perception is relational and by the subject, it can also be regarded as direct in a way that averts the screening-off problem.

3.2 Are the motivations for naïve realism undermined by SFK?

As indicated in Sect. 1.1, naïve realism is substantiated by a set of arguments (A–H). In what follows, I argue that they all retain force under the assumption of SFK (which of course entails that none of them hinges on PC). They are sorted into three groups: those which concern conscious perception only (Sect. 3.2.1); those which apply equally to conscious and unconscious perception (Sect. 3.2.2); those which would gain in plausibility if SFK turned out to be true (Sect. 3.2.3).

Because my aim is confined to showing that these arguments work under SFK, I try to say only so much that is necessary to achieve this particular goal. For this reason, the arguments are introduced very crudely. Although each of the arguments is subject to criticism unrelated to SFK, responding to that criticism is beyond the scope of this paper.

3.2.1 Arguments which apply to conscious perception only (A–C)

(A) SFK does not entail that there are no differences between conscious and unconscious perception. All it says is that both conscious and unconscious perception are mental states of the same fundamental kind. And all that this entails under naïve realism (at least under my exposition of this view), in turn, is that both conscious

¹⁷ The cognitive import of unconscious perception is what enables the subject to perform an experimental task with above-chance accuracy despite not being conscious of the stimulus. In another paper (Zięba 2017), I argue that contrary to cognitive import of conscious perception, cognitive import of unconscious perception does not provide the subject with an opportunity to acquire perceptual knowledge. It only makes the occurrence of such opportunity more probable.

and unconscious perception are direct relations between the subject and the mind-independent physical object. Since SFK does not entail that perceptual consciousness is explanatorily idle, the fact that unconscious perception is not as cognitively efficient as conscious perception does not undermine the grounds of naïve realism. Even though some of those grounds stipulate that perception has to be conscious in order to play certain explanatory roles, none of it is undermined by simply allowing that perception can sometimes occur unconsciously.

This is why SFK does not undermine John Campbell's claim that the possibility of demonstrative judgment is conditional on direct conscious acquaintance with the mind-independent object (Campbell 2002, pp. 114–115). Some intentionalists have tried to account for the role of perception in grounding demonstrative thought and reference by binding perceptual content down to the mind-independent object. For example, it has been argued that the propositional content of perception contains an indexical term which picks some particular object (Tye 2014), or that the properties perceptual content represents the perceived object as having are tropes, i.e. abstract particulars (Nanay 2012). For Campbell, however, this is not sufficient to accredit demonstrative judgment to perceptual consciousness because under intentionalism experience with the same phenomenal character as perception can occur when the represented object is not present (Campbell 2002, p. 117).

The idea of unconscious perceptual relation is consistent with Campbell's view. Empirical evidence suggests that unconscious perception can direct one's attention, facilitate behaviour, and create short-term, unstable associations. Perhaps it makes the subject prone to think about things in some ways rather than another. This, however, does not suffice for grounding demonstrative thought and reference.¹⁸ Subjects do not form demonstrative judgments about what they were primed with [unless they actually become conscious of the primes, which often ends up in excluding them from the experiment by the researchers (see e.g. Block 2017, pp. 178, 183; Phillips 2017, p. 180)]. Nor do the blindsight and neglect patients judge demonstratively about what they have unconsciously seen (again, unless their condition involves some residual awareness of the environment). This is not to say that unconscious perception (or residual, or very short-lasting perception) cannot influence demonstrative judgments. Just as there can be unconscious top-down influences on one's conscious perceptual judgments (e.g. one's emotions, desires, or beliefs can bias one's perceptual judgment without one being aware of it), by the same token, there can also be bottom-up influences on one's conscious perceptual judgment (i.e. unconscious, or residual, or very short-lasting perception can alter one's perceptual judgment without one being aware of it). If this is right, then SFK poses no threat to naïve realism as a theory of conscious perception.

¹⁸ (Cf. Campbell 2002, p. 223: 'Suppose that, without anyone having any glimmering as to how this came about, our visual systems, at some level remote from consciousness, became attuned to the presence of certain items in our surroundings—suppose we call them "spooks". Then we might find that we could achieve intersubjective agreement on when to assent to "Spooks!" on the basis of visual information, without any of us ever having experienced these things. But this would evidently be a case in which none of us had the slightest idea what we were talking about. The hard facts about meaning are hard facts about our knowledge of meaning, and our knowledge of meaning depends on our capacities specifically for conscious attention to objects.').

(B–C) The same point applies to two other features of conscious perception which naïve realists use to argue for their view, namely transparency and fine-grainedness. Transparency has it that when one introspectively focuses on one's experience, one becomes aware of nothing but the features of the experienced object. Some intentionalists also endorse transparency, albeit they think that the features in question are represented to the subject, whereas naïve realists think of them as presented to the subject. According to [Martin \(2002b\)](#), the intentionalist account of transparency in visual imagination (a.k.a. visualising) has perplexing consequences, which constitutes a reason to prefer the naïve realist explanation.

In a nutshell, Martin's argument goes as follows: (i) visualising an object consists in imagining a visual experience of that object; (ii) the phenomenology of visualising is committal on the question whether the object is present in the imagined scene; (iii) on naïve realism, the phenomenal character of experience is committal with respect to the presence of the experienced object, which is not the case under intentionalism; (iv) under naïve realism, to say that visualising X amounts to imagining an experience of X is sufficient to account for (ii), whereas the intentionalist has to supplement this with a non-imagistic supposition that specifies whether the content of the imagined experience is veridical in order to accommodate (ii); (v) the need for non-imagistic suppositions to explain visual imagination makes intentionalism fail to account for transparency of visual imagination; (vi) the naïve realist account is superior because it does not generate any such requirement.

The argument from fine-grainedness, in turn, trades on the difference between experience and thought. It is based on the observation that the phenomenal character of sense experience is extremely rich. No conception of perceptual content can convey this richness, since there is no way to construe such content to be as determinate as the phenomenal character of experience. Accordingly, Bill Brewer argues that no currently available conception of perceptual content is able to account for the fine-grainedness of experience ([Brewer 2011](#)). The problem does not arise for naïve realism, since this view is not committed to the idea that experience has content.

One important difference between Campbell's argument on the one hand, and Martin's and Brewer's arguments on the other, is that the former takes consciousness to be the *explanans* of certain phenomena, whereas in the latter two consciousness figures as an *explanandum*. Consequently, Campbell's argument could be undermined by showing that unconscious perception can enable demonstrative thought, which would mean that consciousness either does not explain this phenomenon at all, or provides a superfluous explanation. This strategy could not be employed against Martin's and Brewer's arguments, however, since these two concern what phenomenology of experience is like. For example, even if it was shown that unconscious visualising is possible, this datum could not be used to undermine Martin's argument because the latter concerns phenomenology of such experiences, i.e. something that unconscious visualising would lack by definition.

Provided that unconscious perception does not enable demonstrative thought, I conclude that neither of the three arguments I have considered in this subsection is impaired by SFK, since they are all concerned with conscious perception only. For the same reason, however, they are insufficient to establish naïve realism as a general theory of perception if SFK is granted.

3.2.2 Arguments which apply equally to both conscious and unconscious perception (D–F)

(D) Perhaps the most important motivation for naïve realism consists in acknowledging that the object of perception is particular. For instance, when I see a prawn, what I am aware of is not the fact that there is at least one prawn around, or that there is some prawn in the vicinity. Rather, I am aware of the presence of *this* particular prawn. Had it been a different prawn, the object of my perceptual awareness would have been different as well. Now, is seeing a prawn and, say, being primed with a prawn (most likely a picture of one) different in terms of particularity? It is not, since both experiences put the subject in perceptual relation with a particular object.

Although naïve realism is arguably the most straightforward way of explaining the particularity of perception, this is not enough to regard it as the best explanation (Martin 2002a, p. 200). The conclusion that naïve realism is indeed the best explanation of particularity might be drawn indirectly, from the aforementioned arguments about transparency and demonstrative thought, as well as Charles Travis' criticism of the idea that perception has content (see point F below).

(E) One consideration against defining perception as direct relation with something particular concerns the possibility of illusion and hallucination phenomenally indiscernible from perception. For example, there can be cases in which it seems to one from the first person perspective that one is presented with a prawn, where in reality no prawn is around. This is often used to support the following picture. Perception, Illusion, and Hallucination are species of a genus—Sense Experience (a.k.a. Perceptual Experience).¹⁹ Whether X is a member of the genus might turn on X's having the phenomenal character specific to sense modalities (e.g. X conveys a certain look, or smell, etc.). Alternatively, it might turn on X's being related in some way to a specific kind of brain activity (e.g. X's neural basis is in visual cortex, or rhinencephalon, etc.). Yet another possible criterion has it that X belongs to the genus if it is an exercise of the subject's capacity to perceive certain kinds of stimuli, e.g. visual, auditory, and so on. On all three alternatives, Perception is, on a par with Illusion and Hallucination, an instance of a more fundamental kind of mental state called Sense Experience.

Naïve realism, at least as I understand it (but see Martin 2006, p. 363), recommends a different picture. First, Perception is more fundamental than Sense Experience even though the latter is more general. This is because Perception is the prototype of Sense Experience. Kinds of Sense Experience other than Perception (e.g. Illusion,²⁰ Hallucination) are considered as kinds of Sense Experience *because* they sufficiently resemble Perception. Second, the fundamental criterion for membership in a particular kind of Sense Experience concerns what a given experience relates the subject to (i.e. what it is an experience of). On this picture, a given experience counts as an instance of Perception only if it relates the subject to the mind-independent macroscopic physical object (cf. Brewer 2011, pp. 92, 94). Thus understood, naïve realism entails metaphysical disjunctivism, i.e. the idea that Perception and Hallucination are

¹⁹ Here I use capital letters to indicate where I speak about kinds, and where not.

²⁰ Naïve realists typically regard Illusion as either a sub-kind of Perception, or a sub-kind of Hallucination.

fundamentally different kinds of phenomena even if there can be instances of Hallucination subjectively indistinguishable from instances of Perception. Furthermore, should it be granted that Sense Experience of mind-independent non-physical objects (e.g. ghosts) is possible, as incredible as this may seem, that would call for introducing another kind of Sense Experience. Experiences of this kind would not belong to Perception because their objects would be different from the objects of Perception.

The crucial question for present purposes is whether this view holds water under the assumption of SFK. I think that the positive answer follows straightforwardly from the conclusion of Sect. 3.1. Namely, as long as it is granted that perceptual relation can be unconscious, everything I have said in the previous paragraph remains true. An experience belongs to Perception if it relates the subject to the mind-independent physical object, irrespective of whether the subject is conscious of the object or not. Accordingly, it is open for the naïve realist to classify Unconscious Perception as a sub-kind of Perception. Still, the intentionalist might reply that SFK always comes together with perceptual representations, and thus there is no way to explain unconscious perception without positing unconscious representations (cf. Block 2017, p. 169; Nanay 2014). In the next subsection, I explain why this objection fails.

(F) Another incentive for joining the naïve realist camp comes from Travis' insight that the idea of perceptual content renders thought about the world unintelligible (Travis 2013d). According to Travis, senses are *silent*, i.e. they merely bring things in view for us, without specifying how things are. It is up to us to decide how things are based on what we perceive, which we do by way of forming perceptual judgments. Perception cannot do it for us, since there is no conception of perceptual content which would explain how perception could do that.

Under intentionalism, perceptual representation has a face value, i.e. it univocally represents things as being a certain way. This is consistent with things not being the way experience represents them to be. Accordingly, the subject can either take her experience at its face value (i.e. consider it veridical), or reject it as misleading. So far, so good. But how does the subject recognise the face value of her experience? The face value cannot be fixed by how things visually look because the way things visually look is equivocal between contents. For any given look, there is more than one way things can actually be. For instance, if something visually looks like a filleted prawn, it can either be a filleted prawn, or a prawn-shaped candy, or countless other things. Nor can the face value be fixed by what the subject takes the experience to represent, since in that case the face value would not be *for the subject* to take a stance on. Rather, it would be ascribed to the experience *by the subject*. That I take my perceptual experience to be of a filleted prawn is hardly the face value at which I could later accept it or reject it. Yet another option would consist in something most intentionalists reject, namely regarding perceptual content as representing how things subjectively seem to the subject, as opposed to how things are. Even if this manoeuvre allows the intentionalist to keep the perceptual content univocal, it does so at a high cost of hiding the world beyond the so-called veil of perception. For this assumption

renders perceptual judgment as concerned with how things subjectively appear to the subject, not how they really are.²¹

Travis deems it impossible to delineate perceptual content in a way that would avoid these problems. He concludes that the notion is incoherent, and intentionalism is false. However, the argument was designed to cover conscious perception. Does it work for unconscious perception as well? Suppose that the subjects in [Jiang et al. \(2006\)](#) study unconsciously perceived naked bodies, and this is why their attention moved in accordance with their sexual orientation. For ([Block 2017](#); [Nanay 2014](#)), positing perceptual representations is indispensable to explain this phenomenon. But if Travis is right about the ambiguity of visual looks, it seems perfectly reasonable to ask what disambiguates the content of unconscious perceptual representations the subjects had during the Jiang's experiment. Why did the relevant perceptual contents specify naked bodies rather than any of all the things one could mistake them for? Considering that, as Travis himself notes, such disambiguation would not require that the subject 'can characterize such representational content accurately, or formulate it explicitly' ([Travis 2013d](#), p. 28), the question does not seem even a bit less puzzling than when it is asked in relation to conscious perception. This indicates that the argument applies to both conscious and unconscious perception. Therefore, contrary to what Block and Nanay have suggested, accepting SFK does not make it imperative that perceptual representations be invoked.

3.2.3 Arguments reinforced by SFK (G–H)

Naïve realism is not optimistic about the cognitive import of consciousness all along the line. In particular, naïve realism is sometimes motivated with a certain sort of scepticism about introspection. To this extent, it actually gains support from SFK. There are two incentives for adopting naïve realism that can be corroborated in this way.

(G) The first one is the idea that the subject is not infallible about the phenomenal character of her own experiences (see e.g. [Hawthorne 2007](#); [Hellie 2011](#); [Martin 2004](#); [Williamson 2000](#)). As I have mentioned above, naïve realism can be defined as the claim that the phenomenal character of perception is literally co-constituted by the mind-independent physical object. An obvious objection to this view arises from the possibility of hallucinations which are phenomenally indistinguishable from perceptions. Fallibilism about phenomenal character arms the naïve realist with a neat rejoinder to this objection: when the subject undergoes a hallucination phenomenally indistinguishable from perception, the phenomenal character of hallucination merely seems to be the same as the phenomenal character of the corresponding perception, whilst in fact it is not the same (some go as far as to say that hallucinations have no phenomenal character at all). What fallibilism about phenomenal character buys the

²¹ One problem with this consequence is that it is unclear what are the truth conditions for such judgments. On the conception of judgment ([Travis 2013a, e](#)) draws from Frege, something can be an object of judgment only if it exists independently of being judged, and can be judged by more than one subject. Otherwise it would be unclear what makes the judgment true or false. Because claims about inner appearances fail to fulfil this condition, Travis concludes that they are not judgments at all.

naïve realist is the reason to reject the inference from the possibility of hallucinations phenomenally indistinguishable from perceptions to the claim that such hallucinations have the same phenomenal characters as corresponding perceptions.

Now, I claim that fallibilism about phenomenal character acquires additional support from SFK. Empirical findings in favour of SFK consist of situations in which subjects make false introspection-based judgments about their own mental states. Subjects claim that they have not perceived anything, whereas in fact they did. This constitutes empirical support for the claim that introspection does not supply the subject with reliable information as to what kind of mental state she is in. This, in turn, lends colour to fallibilism about phenomenal character, and so to naïve realism as well.

(H) The second motivation might seem to contradict the first. It concerns a certain similarity between naïve realism and sense data theory (a.k.a. indirect realism). Namely, both views share an assumption that perception is fundamentally a relational phenomenon (contrary to intentionalism, which denies it). The idea is that in perception we are consciously related to various entities. Such entities have to exist in order for us to be consciously aware of them. This is supported by phenomenology of experience: the mere introspection of one's sense experience reveals its relational nature (cf. [Soteriou 2016](#), pp. 6–8, 56–57).

Wait a moment. Is it not a clear example of double standard to rely on introspection when it supports the naïve realist view and discard it as unreliable when it bolsters the opponent's position? Not necessarily. Firstly, fallibilism about phenomenal character and the appreciation of relationality of experience are perfectly consistent with each other. Secondly, and more importantly, the objection works both ways, since the intentionalist relies on introspection in relation to phenomenal character, and simultaneously discards it with respect to the relationality of experience. As ([Martin 2002b](#), pp. 420–421, [2004](#), pp. 84–85) rightly points out, each theory of perception is in some sense an error theory, and determining which theory is correct ultimately boils down to figuring out what the error is.

Anyway, the difference between naïve realist relationalism and sense data relationalism concerns the nature of direct objects of perception. The former view renders these objects as mind-independent (i.e. their natures do not depend in any way on them being perceived), whereas the latter construes them as mind-dependent²² (i.e. their natures in some way depend on them being perceived). A lot has been written about why positing mind-dependent objects of perception (a.k.a. sense data) is problematic. From the naïve realist perspective, one problem with sense data theory is that it is 'inconsistent with the claim that perceptual presentation provides us with a conception of what mind-independent physical objects are' ([Brewer 2011](#), p. 52). Suppose that the sense data theorist replies that the correct theory of perception does not have to provide any such conception. Suppose further that the sense data theorist rejects any other objection the naïve realist has in her arsenal. In that case, allowing for the possibility of unconscious perceptual relation arms the naïve realist with yet another, and most likely deciding, way to show that her view is superior. For, contrary to sense data theory, naïve realism can secure the intuitive idea that perception

²² I have in mind contemporary versions of sense data theory (see e.g. [Jackson 1977](#)). The original proponents of this view (e.g. C.D. Broad, G.E. Moore, Bertrand Russell) held that sense data are mind-independent.

is fundamentally relational even if mental states of fundamentally the same kind as perception can occur unconsciously.²³

To sum up, both issues covered in this subsection add up to a substantial and tempting motivation for naïve realists to keep their fingers crossed that SFK gets established beyond doubt. At this point, it can be seen that neither the core nor any of the main motivations of naïve realism turns out to be false if it is assumed that SFK is true.²⁴

4 How to account for unconscious perception?

Even if it is granted that the SFK-naïve realism advanced in the previous section is consistent with the core of naïve realism and all of its main motivations, one might still doubt whether it can provide satisfactory responses to all of the objections mentioned in Sect. 1.2. The goal of this section is to show that it can.

If considerations in Sect. 3.2.2(F) are correct, positing perceptual representations is not necessary to account for unconscious perception. More specifically, SFK does not require positing the kind of representation that intentionalism propounds. Under intentionalism, perception *is* a representational state that can represent the world accurately or inaccurately, which means that it is a kind of representation that can mislead. Importantly, rejecting this kind of representation is consistent with allowing for what Travis calls *effect-representation* (Travis 2013c, p. 315). The latter is a relation between an effect and any of its causes. For instance, if the subject perceives naked bodies, the activation in her visual cortex is an *effect-representation* of her perceiving the naked bodies, whether she perceives the bodies consciously, or not. Importantly, this sort of representation cannot mislead. Had no naked body been perceived, there would have

²³ Perhaps the sense data theorist could respond with introducing unconscious sense data. However, adopting this strategy does not seem promising, since the arguments for sense data rely very strongly on introspective access, which is obviously unavailable in unconscious perception.

²⁴ There is an objection to naïve realism which does not trade on unconscious perception but could be used to argue that the above considerations do not secure naïve realism after all. According to Brogaard (2015), type-2 blindsight constitutes a counterexample to naïve realism because it is a case of veridical visual perception which nonetheless does not exhibit particularity, transparency, nor fine-grainedness. Type-2 blindsight differs from type-1 blindsight in that it involves residual awareness of the stimulus (hence it is not unconscious perception). Because the subjects with this syndrome report experiencing ‘something’ or ‘that something is happening’ in a way that it seems to them ‘foggy’ or ‘behind a screen’, whilst denying that they see any particular object, Brogaard claims that (i) none of the three characteristics applies to such experiences. And yet there are good reasons to think that (ii) such experiences are veridical perceptions (apparently due to above-chance performance by type-2 blindsight patients in relevant experimental tasks). From the conjunction of (i) and (ii), Brogaard concludes that particularity, transparency, and fine-grainedness are not essential features of perception. This, in turn, ‘give[s] us reason to reject direct realism as a general theory of visual experience’ (Brogaard 2015, p. 101).

Had naïve realism been committed to the claim that sameness and difference in phenomenal character solely turns on sameness and difference of the perceived object, the objection would have reached the target. But naïve realism is not so committed. The idea that the phenomenal character of perception has the mind-independent object among its constituents does not entail that the object is the only determinant of the phenomenal character. Thus, the commitment does not stick even if one takes naïve realism as entailing that perception is conscious by definition (cf. Phillips 2017, p. 176). Besides, it is disputable whether experiences in type-2 blindsight are visual (see e.g. Overgaard and Mogensen 2015).

been no corresponding *effect-representation* in the brain. If the scientist mistakenly takes a brain activity to be an *effect-representation* of a perception where in fact no perception happened, there is no misleading *effect-representation* involved. The brain activity in question is just *effect-representing* something else.

The distinction between perceptual representation and *effect-representation* arms the naïve realist with rejoinders to some of the objections I have wheeled out in Sect. 1.2. As we have already seen, Block's argument that SFK supports intentionalism is a *non sequitur*. Because Travis' argument applies to both conscious and unconscious perception, no need for perceptual representations follows straightforwardly from SFK.

According to Berger and Nanay (2016, p. 431) and Nanay (2014, pp. 42–45), where the naïve realist speaks of a single perceptual relation, there are in fact two distinct perceptual episodes, i.e. vision for recognition, which is executed by the ventral system, and vision for action, which is carried out by the dorsal system. Similarly, what the naïve realist describes as a single perceptual relation is in many cases a complex integration of sensory information from various sense modalities. Consequently, the critics argue, it is unclear how could this multitude of perceptual episodes be explained with a single relation. To account for these phenomena, the naïve realist would have to either (i) confine perception to conscious experience and regard all the rest as subpersonal unconscious processes enabling conscious experience (i.e. deny SFK), or (ii) allow for unconscious perception and take naïve realism to be the theory of conscious perception only.

This is a false dilemma, however. The naïve realist can explain the phenomena in question by stipulating that perception results in many distinct *effect-representations* in the brain. Consider the following example. When I look at a prawn on my plate and simultaneously reach for it with a fork, this single event (I have looked only once) is *effect-represented* by many various happenings in my brain, including the activation of my dorsal and ventral systems, as well as multimodal integration (the prawns have a certain look, they smell like sea, and make noises when I prick them with my fork). What exactly is wrong with this picture?

It seems that Nanay evaluates naïve realism through the lens of intentionalism. Under intentionalism, there has to be a match between the means by which the subject perceives (i.e. perceptual content) and its *effect-representation* in the brain. For example, if a perceptual representation activates two different systems in the brain (e.g. dorsal and ventral), this has to be reflected in perceptual content this representation carries. If one thinks of perceptual relation in the same way (i.e. as a means by which perception happens), one is likely to expect a similar correspondence in the naïve realist explanation. But perceptual relation is not a means by which perception happens. It just *is* perception, i.e. having things in view. On this picture, the only correspondence one might expect is the one between the objective properties of the stimulus and the *effect-representations* in the brain. For this reason, the cases in which dorsal and ventral visual subsystems fail to cooperate (e.g. Ebbinghaus and Müller-Lyer illusions, see Sect. 1.2) do not constitute a problem for naïve realism. If an experience results in two conflicting *effect-representations* in the brain, all that this reflects is the nature of the stimulus. It certainly does not follow that there had to be two perceptual rela-

tions instead of one. The naïve realist does not even have to accept SFK to rebut this objection.

This brings us to Berger and Nanay's rejection of the idea of unconscious perceptual relation. Berger and Nanay (2016) argue that the distinction between conscious and unconscious perceptual relation would have to be explained in terms of difference between the relata of these relations. This, they claim, is unattainable, as none of the relata seems appropriate for playing that role. The object of perception does not make the difference because it could be perceived both consciously and unconsciously. The perceiver does not make the difference because it is unclear which proper part of the perceiver is supposed to make it. Neither does the sometimes added third relatum (e.g. perceiver's point of view, sense modality, lighting conditions, background knowledge, etc.) because it would have to be simultaneously consciousness-apt and consciousness-inapt in order to explain the alleged multitude of perceptual episodes I have rejected in the two preceding paragraphs.²⁵

Let us consider this objection against the background of particular examples of unconscious perception. Suppose that type-1 blindsight is unconscious perception. Suppose further that a blindsight patient who has earlier declared having taste in seafood is now shown a prawn in her blind field. As the increased saliva secretion indicates, the prawn is perceived, although further tests show that neither subjective, nor objective thresholds for consciousness are met. What would have to change in order for the subject to consciously perceive the prawn? Obviously, the subject: the damage in her visual cortex would have to cease to exist.

Just to make it clear, I do not claim that in order to change a case of unconscious perception into a case of conscious perception we would have to replace the blindsighter with a numerically distinct subject who happens to have a healthy brain. Instead, I claim that to make that particular blindsighter consciously see the stimulus in her blind field, we would have to 'repair' her brain. The case of conscious perception we would thereby enable would have a *numerically the same yet qualitatively different* subject.

Is it always the case that if the subject in two perceptual relations is numerically the same yet qualitatively different then those perceptual relations are different? This seems too strong, as it entails that the tiniest change in the state of the subject renders a change in the experience the subject undergoes. If, for example, the subject's legs are cut off, that does amount to a qualitative change in the subject, and yet does not prevent the subject from becoming perceptually related to a prawn in exactly the same way as before. However, the naïve realist can posit some restrictions on what kind of qualitative difference in the subject would be necessary to yield a difference in perceptual relation. The difference concerning presence/absence of lesions in visual cortex definitely does constitute the relevant contrast. A detailed story about this difference should be left for empirical vision science to tell.

To take another example, suppose that continuous flash suppression induces unconscious perception in subjects with healthy brains. This method consists in presenting the target stimulus to one eye and a "Mondrian" to another. The "Mondrian" is a high contrast and rapidly changing stimulus, which somewhat resembles paintings of Piet

²⁵ Note that denying that vision for recognition and vision for action should be treated as two distinct perceptual episodes *does not* amount to rejecting the empirical hypothesis about two visual subsystems.

Mondrian. “Mondrians” are known to suppress the awareness of the competing stimuli for ‘dozens of seconds’ (Yang et al. 2014, p. 1). Suppose a subject is presented with a prawn to one eye and the “Mondrian” to another. The tests show that, before the suppression abates, the prawn is perceived unconsciously. What would have to change in order for the subject to perceive it consciously? Obviously, the stimulus. Had something else instead of “Mondrian” been presented to the other eye, the prawn would have been consciously perceived much quicker. Had the prawn been presented to both eyes, it would have been perceived consciously.

One might object that the stimulus is just the prawn, which remains the same in both cases, while the Mondrian stands for something like circumstances of perception. If this was the case, there would be no difference in the stimulus after all. But why should the prawn be considered the proper object of perception while the Mondrian merely count as the circumstances of perception? Even if we had to choose which of those two is the proper object, why not opt for the opposite? That is, why not consider the Mondrian the proper object and the prawn the circumstances? In experimental practice, it is the prawn that is more likely to play the role of ‘the target stimulus’, but the identification of the target stimulus is purely functional, namely relative to the goal of a given experiment. By contrast, the object of perception is whatever the subject is perceptually related to. In this connection, I see no reason to deny that both the prawn and the Mondrian equally qualify as elements of the perceived scene. They both co-constitute what the naïve realist calls the object of perceptual relation.

Such examples could be multiplied, but my point should be clear by now. Whether the stimulus is consciously or unconsciously perceived turns on the state of the subject and/or the nature of the stimulus. As to the third relatum, I think it reduces to either the state of the subject or the nature of the stimulus, but I will not argue for it here.

Describing unconscious perception in terms of the state of the subject and the nature of the stimulus amounts to the most basic explanation of unconscious perception. Had these factors been irrelevant for unconscious perception, Breitmeyer’s considerations described in Sect. 2.1 would have been useless for explaining this phenomenon. This basic explanation is of course far from satisfactory, hence both relata of unconscious perception have to be put under further empirical scrutiny. That said, I fail to see how could the results of that further examination cast any doubt on the following naïve realist definition of unconscious perception:

(UP) Unconscious perception is (i) the relation between the subject and the mind-independent physical object such that (ii) it has measurable cognitive import for the subject despite the fact that (iii) the state of the subject and/or the nature of the object precludes the subject from becoming conscious of the object.

Simple as it is, it is hard to deny. By comparison, explaining unconscious perception by way of stipulating some modification in a contentious theoretical object called ‘perceptual representation’ (not to be mistaken with *effect-representation*) does not seem to help us understand the phenomenon any better. (Nanay 2015, p. 322) advocates intentionalism as a compelling explanation of how (i) experience can be misleading, and (ii) perception plays the role in justifying beliefs. However, both claims fall prey to Travis’ objections, and neither gains support from SFK.

Perhaps it could still be shown, on some other grounds, that intentionalism is the preferable theory of perception (e.g. by offering a definition of perceptual content that would surmount Travis' criticism). But if we focus on unconscious perception exclusively, intentionalism does not seem to have any advantage over naïve realism. Considering everything I have said above, nothing seems to prevent the naïve realist from adopting the definition I have offered.

5 Conclusion

I have argued for the following conditional claim: if SFK turns out to be true, the naïve realist can and should accommodate it into her theory. Regarding the antecedent of this conditional, I have suggested that empirical evidence renders SFK plausible but not obvious. For it is possible that what is currently advocated as unconscious perception of the stimulus is in fact momentaneous perceptual awareness (or residual perceptual awareness) of the stimulus making the subject prone to judge in some way rather than another, or to act in some way rather than another. As to the apodosis, I have shown that neither the core of naïve realism nor any of the main motivations for endorsing it is undermined if SFK is assumed. On the contrary, certain incentives for endorsing naïve realism become more tempting on this assumption. This is because the stronger evidence for unconscious perception is provided, the more viable it is for the naïve realist to argue that unconscious perception is perception in virtue of being a direct relation between the subject and the mind-independent physical object. Since the main motivations for naïve realism work under SFK, intentionalism is neither compulsory nor the best available explanation of unconscious perception. Although the naïve realist alternative to intentionalism might still be rejected on some other grounds, simply establishing SFK does not suffice to do so.

Acknowledgements I am deeply grateful to Benj Hellie, Michał Klincewicz, Sebastian T. Kołodziejczyk, Marcin Miłkowski, and an anonymous referee for this journal for helpful comments and suggestions.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Anaya, A., & Clarke, S. (2017). Naïve realism and unconscious perception: A reply to Berger and Nanay. *Analysis*, 77(2), 267–273.
- Berger, J., & Nanay, B. (2016). Relationalism and unconscious perception. *Analysis*, 76(4), 426–433.
- Bermúdez, J. L. (2000). Personal and sub-personal: A difference without a distinction. *Philosophical Explorations*, 3(1), 63–82.
- Block, N. (2010). Attention and mental paint. *Philosophical Issues. A Supplement to Noûs*, 20, 23–63.
- Block, N. (2012). The grain of vision and the grain of attention. *Thought: A Journal of Philosophy*, 1(3), 170–184.
- Block, N. (2016). The Anna Karenina principle and skepticism about unconscious perception. *Philosophy and Phenomenological Research*, 93(2), 452–459.

- Block, N., & Phillips, I. (2017). Debate on unconscious perception. In B. Nanay (Ed.), *Current controversies in philosophy of perception* (pp. 165–192). New York, London: Routledge.
- Breitmeyer, B. G. (2015). Psychophysical “blinding” methods reveal a functional hierarchy of unconscious visual processing. *Consciousness and Cognition*, 35, 234–250.
- Brewer, B. (2011). *Perception and its objects*. Oxford, New York: Oxford University Press.
- Brogaard, B. (2015). Type 2 blindsight and the nature of visual experience. *Consciousness and Cognition*, 32, 92–103.
- Burge, T. (2005). Disjunctivism and perceptual psychology. *Philosophical Topics*, 33(1), 1–78.
- Burge, T. (2010). *The origins of objectivity*. Oxford: Oxford University Press.
- Burge, T. (2011). Disjunctivism again. *Philosophical Explorations*, 14(1), 43–80.
- Campbell, J. (2002). *Reference and consciousness*. Oxford: Oxford University Press.
- Clark, A. (2016). *Surfing uncertainty: prediction, action, and the embodied mind*. Oxford: Oxford University Press.
- Crane, T. (2006). Is there a perceptual relation? In T. Gendler & J. Hawthorne (Eds.), *Perceptual experience* (pp. 126–146). Oxford: Oxford University Press.
- Dennett, D. (1969). *Content and consciousness*. London: Routledge & Kegan Paul.
- Drayson, Z. (2012). The uses and abuses of the personal/subpersonal distinction. *Philosophical Perspectives*, 26(1), 1–18.
- Fish, W. (2009). *Perception, hallucination, and illusion*. Oxford: Oxford University Press.
- Fish, W. (2010). *Philosophy of perception. A contemporary introduction*. London: Routledge.
- Frith, C. (2007). *Making up the mind. How the brain creates our mental world*. London: Blackwell Publishing.
- Gelbard-Sagiv, H., Faivre, N., Mudrik, L., & Koch, C. (2016). Low-level awareness accompanies “unconscious” high-level processing during continuous flash suppression. *Journal of Vision*, 16(1), 1–16.
- Hawthorne, J. (2007). Direct reference and dancing qualia. In T. Alter & S. Walter (Eds.), *Phenomenal concepts and phenomenal knowledge: New essays on consciousness and physicalism* (pp. 195–209). Oxford: Oxford University Press.
- Hellie, B. (2011). There it is. *Philosophical issues*, 21(1), 110–164.
- Hellie, B. (2013a). Against egalitarianism. *Analysis*, 73(2), 304–320.
- Hellie, B. (2013b). It’s still there!. In R. Brown (Ed.), *Consciousness inside and out: Phenomenology, neuroscience, and the nature of experience* (pp. 127–136). Berlin: Springer.
- Hellie, B. (2014a). Love in the time of cholera. In B. Brogaard (Ed.), *Does perception have content?* (pp. 241–261). New York: Oxford University Press.
- Hellie, B. (2014b). The multidisjunctive conception of hallucination. In F. MacPherson & D. Platchias (Eds.), *Hallucination, philosophy and psychology* (pp. 149–173). Cambridge, London: MIT Press.
- Hellie, B. (2018). An analytic-hermeneutic history of consciousness. In K. M. Becker & I. Thompson (Eds.), *Cambridge companion to history of philosophy 1945–2015*. Cambridge: Cambridge University Press.
- Jackson, F. C. (1977). *Perception: A representative theory*. Cambridge: Cambridge University Press.
- Jiang, Y. V., Costello, P., Fang, F., Huang, M., & He, S. (2006). A gender- and sexual orientation-dependent spatial attentional effect of invisible images. *Proceedings of the National Academy of Sciences*, 103(45), 17048–17052.
- Locatelli, R., & Wilson, K. A. (2017). Introduction: Perception without representation. *Topoi*, 36, 197–212.
- Martin, M. G. F. (2002a). Particular thoughts and singular thought. In A. O’Hear (Ed.), *Logic, thought and language* (pp. 173–214). Cambridge: Cambridge University Press.
- Martin, M. G. F. (2002b). The transparency of experience. *Mind and Language*, 4(4), 376–425.
- Martin, M. G. F. (2004). The limits of self-awareness. *Philosophical Studies*, 120(1–3), 37–89.
- Martin, M. G. F. (2006). On being alienated. In T. Gendler & J. Hawthorne (Eds.), *Perceptual experience* (pp. 354–410). Oxford: Oxford University Press.
- Nanay, B. (2012). Perceiving tropes. *Erkenntnis*, 77(1), 1–14.
- Nanay, B. (2014). Empirical problems with anti-representationalism. In B. Brogaard (Ed.), *Does perception have content?* (pp. 39–50). New York: Oxford University Press.
- Nanay, B. (2015). The representationalism versus relationalism debate: Explanatory contextualism about perception. *European Journal of Philosophy*, 23(2), 321–336.
- Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, 84(3), 231–259.
- Overgaard, M., & Mogensen, J. (2015). Reconciling current approaches to blindsight. *Consciousness and Cognition*, 32, 33–40.

- Paul, L. A. (2012). Metaphysics as modeling: The handmaiden's tale. *Philosophical Studies*, 160(1), 1–29.
- Phillips, I. (2016). Consciousness and criterion: On block's case for unconscious seeing. *Philosophy and Phenomenological Research*, 93(2), 419–451.
- Phillips, I. (2017). *Unconscious perception reconsidered*. Retrieved from <http://www.ianbphillips.com/uploads/2/2/9/4/22946642/antwerp.pdf>.
- Raio, C., Carmel, D., Carrasco, M., & Phelps, E. A. (2012). Nonconscious fear is quickly acquired but swiftly forgotten. *Current Biology*, 22(12), 477–479.
- Rosenthal, D. M. (1986). Two concepts of consciousness. *Philosophical Studies*, 49(3), 329–359.
- Ryle, G. (1949). *The concept of mind*. London: Hutchinson and Co.
- Schellenberg, S. (2014). The relational and representational character of perceptual experience. In B. Brogaard (Ed.), *Does perception have content* (pp. 199–220). Oxford: Oxford University Press.
- Schwitzgebel, E. (2008). The unreliability of naive introspection. *The Philosophical Review*, 117(2), 245–273.
- Soteriou, M. (2016). *Disjunctivism*. London: Routledge.
- Stern, R. (2007). Transcendental arguments: A plea for modesty. *Grazer Philosophische Studien*, 74, 143–161.
- Stich, S. P. (1978). Beliefs and subdoxastic states. *Philosophy of Science*, 45(4), 499–518.
- Travis, C. (2013a). Frege, father of disjunctivism. In: *Perception. Essays after Frege* (pp. 59–89). Oxford: Oxford University Press.
- Travis, C. (2013b). Reason's reach. In *Perception. Essays after Frege* (pp. 118–143). Oxford: Oxford University Press.
- Travis, C. (2013c). The preserve of thinkers. In *Perception. Essays after Frege* (pp. 313–363). Oxford: Oxford University Press.
- Travis, C. (2013d). The silences of the senses. In *Perception. Essays after Frege* (pp. 23–58). Oxford: Oxford University Press.
- Travis, C. (2013e). Viewing the inner. In *Perception. Essays after Frege* (pp. 90–117). Oxford: Oxford University Press.
- Tye, M. (2007). Intentionalism and the argument from no common content. *Noûs*, 41(1), 589–613.
- Tye, M. (2014). What is the content of a hallucinatory experience? In B. Brogaard (Ed.), *Does perception have content?* (pp. 291–310). New York: Oxford University Press.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.
- Windey, B., & Cleeremans, A. (2015). Consciousness as a graded and an all-or-none phenomenon: A conceptual analysis. *Consciousness and Cognition*, 35, 185–191.
- Wong, H. Y. (2014). Personal and sub-personal: Overcoming explanatory apartheid. In T.-W. Hung (Ed.), *Communicative action. Selected papers of the 2013 IEAS conference on language and action*. Berlin: Springer.
- Yang, E., Brascamp, J., Kang, M.-S., & Blake, R. (2014). On the use of continuous flash suppression for the study of visual processing outside of awareness. *Frontiers in Psychology*, 5(724), 1–17.
- Zięba, P. J. (2017). Unconscious perception and perceptual knowledge. In C. Limbeck-Lilienau & F. Stadler (Eds.), *The philosophy of perception and observation. Contributions of the 40th international Wittgenstein symposium August 6–12, 2017 Kirchberg am Wechsel* (pp. 301–303). Kirchberg am Wechsel: ALWS.