

CONSCIOUS AND UNCONSCIOUS MENTALITY

Examining their Nature, Similarities
and Differences

*Edited by Juraj Hvorecký, Tomáš Marvan and
Michal Polák*

First published 2024

ISBN: 978-1-032-52979-0 (hbk)

ISBN: 978-1-032-52974-5 (pbk)

ISBN: 978-1-003-40952-6 (ebk)

10

TROUBLES WITH THE ORTHOGONALITY THESIS

Juraj Hvorecký

(CC-BY-NC-ND) 4.0

DOI: 10.4324/9781003409526-13

This Chapter was funded by EU grant No. CZ.02.2.69/0.0/0.0/18_054/0014626



Routledge
Taylor & Francis Group
LONDON AND NEW YORK

10

TROUBLES WITH THE ORTHOGONALITY THESIS

Juraj Hvorecký

10.1 Introduction

This chapter evaluates some basic assumptions behind the ‘dual’ theories of phenomenal consciousness and finds them problematic for at least some examples of conscious episodes. Because the dual theories claim to be universalist in nature, even a small number of counterexamples damages their status. I will present evidence, both empirical and conceptual, that demonstrates the untenability of the dualist assumptions that conscious phenomenality and its unconscious counterpart possess essentially the same qualities. In doing so, I will primarily attack the claim that is shared by many dual theorists that the procedure responsible for bringing phenomenal content into consciousness serves this sole function and does not influence phenomenal qualities of content. I use the term orthogonality, first introduced in Vosgerau et al. (2008) as a convenient shortcut for the conception in which consciousness and content come unproblematically apart. I will conclude by hinting at an alternative proposal that explains the emergence of conscious phenomenality as a single step operation.

Let me start with a general observation. Theories of consciousness come in immense variety and have flooded the literature both in philosophy and cognitive sciences. There are so many options to choose, from the empirically based to the purely speculative, from metaphysically conservative to ultra-liberal. Whatever one’s preference, there is likely a theory that matches it. Compare this abundance with an interest into unconsciousness and its theoretical reflection. There is no real competition between various theories of unconsciousness and no strong stances on various central issues. In fact, there is hardly any theory that would be worthy of its name. It is also interesting to

observe that hardly any books are devoted to unconsciousness (and the titles that are about the topic, at least superficially, do not cover it from an analytic perspective and use no empirical data). Tellingly, the *Stanford Encyclopaedia* has no entry on the subject.

Regardless of unconsciousness being such a neglected topic, many theorists have advocated a view of the unconscious that constitutes a proto-theory. This proto-theory has been often evoked by various defenders of the dual theory of phenomenal consciousness, but despite its frequent usage, it does not attain a status of a full theory. Even those who make use of it do not have a full story on what is the scope of the unconscious remain unanswered. What rules govern unconscious content? How is the unconscious connected to the neurobiological substrate? And many other central questions. In the field of philosophy of mind, the proto-theory was popularized by John Searle, who attributes it to Sigmund Freud.¹ Searle (1992) portrays Freudian unconscious mental states as being akin to conscious ones minus the consciousness (p. 152). He employs two metaphors to illustrate his exposition. In the first one, unconscious states are like fish in a pond. When swimming close to the surface, fish are visible, and all their qualities are grasped easily. When they dive deeper into the pond, we lose sight of them and of their properties, yet nothing substantial changes with the fish themselves. In the second metaphor, unconscious items are imagined to be stored in a dark attic, and “these objects have their shapes all along, even when you can’t see them” (ibid.) Both metaphors are aimed to support the thesis that unconscious mental states possess the very same properties that are familiar to us from conscious experiences. Importantly (using vocabulary not employed by Searle), conscious and unconscious states share the same phenomenal properties.

Recently, the presupposition of the sameness of both conscious and unconscious content has been utilized by so-called two-factor theories (I will sometimes use the term ‘dual theory’ to designate the same set of claims). These theories (examples include Rosenthal 2010; Marvan and Polák 2017, Fink et al. 2021) argue that it would be beneficial to theories of consciousness to split their task into at least two subtasks to achieve the final goal of explaining its target phenomenon. Specifically, the task of explaining contents of consciousness should be separated from tackling issues of the emergence of consciousness itself. According to the two-factor theorists, the first factor processes ensure that contents are prepared for the uptake by the eventual second factor. The first factor’s processes operate solely at the unconscious level, endowing the content with all the qualities that will eventually be experienced consciously. There are many properties that content possesses, including intentional, temporal, and phenomenal. By the time the first factor processes complete their respective jobs, the content is fully prepared—*the cake is fully baked*, as I like to put it.

The second factor is identical to the consciousness procedure. The procedure serves one and only purpose: it makes conscious the content, readied by the first factor procedures. Two-factor theories should be applauded for making the problem of consciousness more tractable. In fact, the dual strategies help to focus the attention on essence of the problem—the consciousness itself. In fact, one might say that the dual theories show that the problem of consciousness is (a relatively) easy problem, contrary to a long tradition that labels it as hard.²

A great illustration of some of the essential claims behind the two-factor theories is found in the paper by Vosgerau and colleagues (Vosgerau et al. 2008). While the terminology of the authors is a bit out of line with some other defenders of the dual strategy, their argumentative strategy is illustrative of some of the essential points that I want to focus on. The authors argue that “representations with one and the same content can be conscious or unconscious” (314). This unproblematic thesis is, however, made less defensible when a stronger reading is offered: “Various types of content can be phenomenally conscious or unconscious, but there is no type of content which has the exclusive right to become phenomenally conscious”³ (7). What follows is the orthogonality thesis in its essence: Whatever is conscious can be unconscious. Consciousness and content are orthogonal.

There are two main lines of arguments used in the paper of Vosgerau et al. First, authors refute several competing theories of consciousness that stress the role of particular content in consciousness production. Specifically, ventral-stream hypothesis, enactivist approaches, and theories based on non-conceptual content are addressed and rejected. The orthogonality thesis is then presented as a viable alternative, not vulnerable to problems the alternatives are facing. Second, the orthogonality thesis is supported by additional empirical arguments that come from cases of impoverished conscious experiences or unconscious perception. Blindsight, agnosia, hemineglect, or failures of attention are discussed, and the orthogonality thesis is supported by repeated claims that content in clear-cut conscious cases and content in the impoverished scenarios are identical. When it comes to a measure of the equality of content in both scenarios, behavioural measures are offered. It is claimed that subjects in the unconscious or impoverished scenarios are prone to the very same behaviour that we witness in analogous conscious situations. This strategy of equating behaviour under both conscious and unconscious scenarios is often invoked by the defenders of the two-factor theories, yet in recent years an additional support has been harnessed by neuroscientific evidence that compares neural processes in both situations (see Marvan, this volume).

Sadly, the paper by Vosgerau and colleagues bears some clear marks of simplification that are to be found in other papers on this topic. First, it only discusses perceptual examples, as if consciousness did not have other contents. Yet on any common-sense account, consciousness is full of non-perceptual

items: thoughts, emotions, desires, feelings, and so forth. This simplification leads authors to state that “the phenomenality is the same for all conscious experiences” (19). One can (roughly) see what that amounts to for conscious experiences within a single-modality, say vision. It is sufficiently difficult to assess what the statement amounts to in cases across modalities. In what sense is phenomenality of touch and olfaction ‘the same’? Adding non-perceptual cases makes the claim virtually impossible to comprehend. The idea of sameness of phenomenality of an occurrent thought, an experienced high pitched sound and a toothache is, at least to me, impossible to evaluate. The only property these various episodes share is their existence in consciousness. Phenomenally they could not be more different (and there is a legitimate worry that occurrent cognitive states might possess no phenomenality at all, which makes the comparison utterly impossible).

Second, equating behavioural responses in standard conscious and lesioned or otherwise non-standard unconscious conditions is problematic as well. While analogous statements that put broadly conceived behavioural responses in two scenarios on the same level can be found elsewhere in the literature, on critical reading there is no identity of responses worthy of discussion.⁴ In cases of neurological conditions, brought about by injuries, lesions or concussions, patients have very weak responses that often have to be elicited by external prompts. In virtually all tested subjects, we find a high percentage of cases when the exposition to stimuli leads to no response whatsoever. These cases are glossed over while we concentrate on cases with detectable responses, yet such cases are crucial for rejecting the claims of identity of responses to conscious and unconscious scenarios.

Notice that even if both of my concerns could be dismissed, the orthogonality thesis is not conclusively defended by the authors. Suppose for a moment that ignoring non-perceptual cases does not influence the thesis and sameness of behavior in conscious and unconscious scenarios can be established. The claim of orthogonality is a general statement that covers relations between content and phenomenality. As such, it cannot be defended by citing individual cases. One has to have some principled way to argue for the conclusion. No such move is found in the article. It therefore opens the door to a substantial criticism of the thesis. Any general statement is rejected by just a single counterexample. In what follows we will search for it.

10.2 Counterexamples

In a recent article, Skrzypulec (2022) discusses differences between conscious and unconscious perception of colours. He concentrates on the way conscious vision represents surface colours and, referring to literature on the topic, distinguishes between two ways in which they can be represented. On one hand, conscious vision represents surface colours relationally. In relational

representations, consciously perceived colours are ordered and judged on the similarity scale. Orange is closer to red than it is to green. On the other hand, conscious colour percepts can be represented categorically: Burgundy, ruby and crimson are all types of red. Upon reviewing large amounts of available evidence, Skrzypulec comes to the following conclusion: “There are certain sensory qualities, in particular representing division of colors into categories or representing sameness of colors despite illumination changes, which are possessed only by some of the conscious states” (670). The categorical colour representations only exist at the conscious level and that should be disturbing news to any defender of orthogonality.

While my real target is the orthogonality thesis and I should be sympathetic to any argument that attempts to undermine it, I am not strongly convinced that Skrzypulec’s observations pose a serious threat to the thesis. The main reason for my scepticism lies with a specific nature of the domain from which the supposed counterexample originates. As we have seen, Skrzypulec is pointing out the absence of categorical perception in unconsciousness. However, there are some reasons to think cases of categorical perception are artefacts of cognitive penetration. Linguistic categories influence judgments that are based on perceptual inputs (Dubova and Goldstone 2021). As such, categorical perception is not necessarily a perceptual phenomenon, but likely a result of an interaction between perception and higher cognitive processes of categorization. It is therefore hard to see what corresponding *perceptual unconscious* phenomenon should serve as its counterpart to judge validity of the orthogonality thesis. Note that this remark does not weaken the overall argument of Skrzypulec about mental phenomena that exist on the conscious level only. It just indicates that the target mental state might not be perceptual and therefore we should not be looking for its unconscious counterpart in the perceptual domain. I will have more to say about similar kinds of examples shortly.

My own examples that undermine the orthogonality thesis are coming from varied sources. Some of them are supported by empirical evidence, others are conceptual in nature. Empirical ones are obviously the strongest, yet even conceptual ones should make us think harder about the kinds of examples that are usually utilized in debates over the status of the orthogonality thesis.

Before I introduce and discuss the relevant examples, let me comment on a repeated property of many debates on consciousness in general and perceptual consciousness in particular. It has been noted many times before that when the topic of consciousness is discussed, plentiful examples from the domain of vision are used, and other modalities and mental phenomena are largely neglected. This overemphasis on vision gives us a very simplified view of what consciousness is like and what its properties might be.

Additionally, when a debate concentrates solely on perceptual consciousness, cited examples present mostly simple single-modality cases. Yet there is so much

more in perceptual consciousness and infinitely more in all of consciousness! It is my shift to the more complicated cases of multisensory (conscious) perception that will unveil some basic difficulties for the orthogonality thesis.

There are many multisensory mental phenomena and especially intriguing are those in which the resulting percepts are not simple sums of single-modality inputs. These multisensory illusions are often little known and somewhat hard to imagine. Take parchment skin illusion: Subjects are asked to describe the tactile sensation on their hand that is accompanied by sound experience. Even though identical pressure is applied to the same place on their skin, subjects report a distinct tactile feeling in accordance with the frequency of co-occurring sound. When a high-frequency sound is played out, the tactile sensation is rough. With the low frequency it changes to a smooth one. This result is surprising and does not seem to answer to any prior intuitions that we have about tactile or auditory sensations. Even more counterintuitive is an observation that deficits in a third, seemingly unrelated modality, have strong influence on susceptibility to this illusion. Champoux and colleagues have found out that blind people, with both early and late onset visual impairment, do not experience the parchment skin illusion (Champoux et al. 2010). The result confirms a crucial contribution of a modality that is not even directly involved in this multisensory illusion.

The resulting multisensory percept is certainly not a sum of two co-occurring experiences (with the implicit role of the third modality) that we are familiar with in single modality scenarios. In order for the orthogonality thesis to hold up, its defenders have to explain this (and other similar) multisensory illusion in the following way: The resulting multisensory percept must have been completed ('fully baked') already at the unconscious level and then this complex percept was elevated to consciousness via the dedicated consciousness-conferring procedure. This is a daring statement for which we have no empirical evidence. Given how understudied similar phenomena are, there are, to my knowledge, no studies that look at the underpinnings of this phenomenon. Fortunately, we are luckier with other multisensory illusions.

McGurk Effect (McGurk and MacDonald 1976, Tiippana 2014) is better known than parchment skin illusion, partly because it was discovered much earlier. In this illusion, subjects hear sounds and visually perceive a face articulating incongruent sound and the resulting percept corresponds to neither the heard nor the seen sounds. In its most famous version, a dubbing voice utters /b/ over the face articulating /g/ and the subject reports hearing /d/. This is once again a very surprising result, and while we have some models that explain why the effect arises for some combinations of auditory and visual percepts and not for others (Magnoti and Beauchamp 2017), additional evidence is needed to account for the processes that elevate the resulting McGurk compound into consciousness. Fortunately, this evidence is available, and it is not very welcoming news to the orthogonality

thesis defenders. Ramsey and Palmer (2012) tested various conditions under which the multisensory integration takes place and discovered that, “For the McGurk effect to occur, ... the lip-streams must be perceived consciously, even though participants might very well be unaware of integrating them with the speech-streams” (ibid., 362). While the integration might be taking place unconsciously (it is hard to imagine that the integration itself would be conscious!), at least one of the compounding experiences must be conscious. This point is generalized by the same authors in the conclusion of their paper: “cross-modal effects can occur in the absence of consciousness, but the influencing modality must be consciously perceived for its information to cross modalities” (ibid., 363).

This finding threatens the status of the orthogonality thesis. We have found at least one case where the simple transformation of content from the unconscious level to the conscious one or, more specifically, the transformation takes place because some of its compounding states are conscious and would not exist otherwise. This is a direct refutation of an approach that simply assumes that states are first formed at the unconscious level and then elevated to consciousness by a dedicated process. I can envision at least two ways in which the thesis of orthogonality could be saved. First, one can allow for iterations of the consciousness-conferring process that can take in a variety of contents, regardless of whether they were conscious beforehand. Second, the defender of orthogonality might claim that she is not at all interested in the precise etiology of the content, and the only fact that matters is that just before the final elevation to consciousness, the designated content was briefly unconscious. I will leave it to the defenders of the view to improve on these suggestions or devise even better responses to the perceived threat. My aim is to undermine their thesis even further, with the use of several conceptual examples.

Take the case of an extremely emotionally charged state, say a state of terror. Experiencing terror consciously has many manifestations. While very undesirable, we have all witnessed these situations first hand: Your mind is focused on a single terrifying event, unable to grasp anything else. Rational judgment is ruled out, your body trembles and fears the consequences of what you are witnessing. You scream, swear and concentrate all powers within to get out of the situation. Now, detach yourself from this dramatic scenario and ask yourself a simple question. Could you undergo the state of terror unconsciously? If so, would it have the same behavioural, psychological, motivational effects on you? If you tend to answer in the positive, think again. Would your mind be focused on a single stimulus, unable to reason, while your body shakes and your muscles are stretched to the utmost, without you knowing what is going on, because the core mental state that is causing all this commotion remains unconscious? I very much doubt that. And I remain sceptical when you claim that your responses might not be so manifest, yet

the unconscious mental state is identical to the conscious state of terror. Undoubtedly, there might be unconscious states that are analogous to the experience of terror, but with less similarities between their manifestations, and those of the genuine conscious terror, we should feel less justified to make the analogy. Similar examples can be found across the emotional domain, especially with intense emotional episodes and their manifestations.

Another conceptual example from a different mental domain is also available, though I am ready to admit that this one is slightly less persuasive. Take (very frequent) cases of cognitive states with very detailed content: I hear steps in the corridor as that of my boss, I smell the rose as Mount Shasta variety, I hope that shares of my favourite company will not lose more than 15 percent of their value today, and so forth. These states are heavily impacted by conceptual content that makes them particular and context sensitive. On my conception, it is hard to see what their unconscious counterpart would be like. While I am ready to admit that some conceptual penetration is detectable at the unconscious level, to my knowledge there is no indication in the literature that it reaches such depth and influence. We have yet to see clear cases of such states at the unconscious level. This line of argumentation leads us back to my discussion of Skrzypulec's colour categorization cases above. To me it looks like his example falls into a wider category of a cognitive penetration with a detailed content. In his case the content roughly bears the meaning, *this shade called X falls under a more general category Y*, and once again might represent the type of state that is just not available unconsciously.

Finally, there are persuasive arguments out there that, despite an obsession of philosophers and psychologists concentrating on single modality cases, multimodal integration and conceptual penetration are much more widespread than we are ready to admit. Casey O'Callaghan (2017; 2019) has been arguing for years that even seemingly simple cases of perception are fused with intermodal and conceptual content, making single modal cases a rarity, not a norm. If that is the case (and I believe it is), there might be even further space for scepticism about a simple orthogonal view in which conscious and unconscious content are on a par, and a simple procedure suffices to elevate a fully composed unconscious state into consciousness. My worry about the simple consciousness-conferring mechanism is that with more integration of various building blocks from perceptual and non-perceptual sources, the process of content preparation takes places later in the cognitive system. And some (many?) of these detailed content preparatory processes are simply appearing in consciousness only, in what is their final stage of integration. There is no need for a heavily redacted content, enriched by multisensory input and concepts to appear twice—once unconsciously and then again at the conscious level. It would be perfectly sufficient for some content to just appear consciously, and the cognitive system would not be burdened with an additional conscious making mechanism.

I have presented some empirical and conceptual arguments that undermine the status of the orthogonality thesis. If one thinks outside of the box of simple single-modality cases, the picture of unconscious mental states that are ‘fully baked’ and ready to be delivered to consciousness loses its traction. There are good reasons to think that two-factor approaches are suitable for simple mental states but cannot be applied to more complicated ones that rely on several sources and their integration. I could stop here and simply claim that my tasks of indicating troubles with the orthogonality is over. Yet I want to go a bit further. Instead of dwelling on the negative result, I want to briefly introduce an alternative scenario in which two-factor is replaced with a single factor, significantly simplifying an overall cognitive architecture of conscious systems.

10.3 A speculative outline of the one-factor view

An attentive reader might have noticed that there is a single guiding idea behind my various criticisms of the orthogonality thesis. As content gets its final form that is to be presented consciously, many steps have been taken to integrate its various aspects into a single whole. This might be a long and complicated process, but one that has a finishing line—the content being conscious. On my reading of the dichotomy of conscious and unconscious states, the conscious ones are always, in some respects, distinct from the unconscious ones. Their resulting integration of various properties that characterize them *are necessarily* different from integration that takes place unconsciously. If you are sceptical and believe that many states might be identical on both sides of the divide, carefully consider *all* properties that conscious and unconscious states possess: their intentionality, temporality, mineness, subjectivity, valence, perspectivalness, vividness, motivational force and many others. I have yet to see examples of unconscious states that match their conscious counterparts along all these dimensions.

My thought behind the one-factor theory is as follows. A cognitive system processes some content and enriches it with further and further properties. Once these processes of enrichment reach a certain threshold, the content just becomes conscious. There is no need for a dedicated mechanism that elevates content into consciousness. Instead, consciousness simply consists of contents that were sufficiently enriched. Importantly, the threshold is not uniform for all types of mental states. Rather, one can envision the threshold as a bundle of properties that are set with different thresholds. For conscious emotions, it is their intentionality, motivational force and vividness that might be crucial. Perceptual states need to be endowed with perspectivalness, temporality and subjectivity to count as conscious. Conscious cognitive states like thoughts and hopes depend on intentionality and subjectivity. It is a varied pack; one whose outline needs careful considerations for each and every type

and maybe even for each and every individual token. Overall, the one-factor theory looks at what characteristics individual conscious states possess, and all thresholds particular states need to exceed in order to become conscious. No single conscious-conferring mechanism is needed, just a careful processing assessment of individual states or state types. On this story, consciousness is a result of specific content-filling processes that nature did not orchestrate as a simple on/off feature, dependent on a single mechanism, but as a complex and mutually interconnected procedure that takes lots of variables for every one of its contents.

Acknowledgements

Work on this chapter was supported by The Czech Science Foundation (GAČR), project no. 20-14445S ('Dual models of Phenomenal Consciousness'), realized at the Institute of Philosophy, Czech Academy of Sciences.

Notes

- 1 It is beyond the scope of this chapter to evaluate whether Searle got Freud right. What matters to us is that Searle made that view prominent in his well-known 1992 book (though chapter 7, *The Unconscious and Its Relation to Consciousness* that introduces the topic is not very widely cited). It also needs to be pointed out that Searle strongly rejects the Freudian view. He believes there is no such thing as unconsciousness, only neural states with the potential to become conscious.
- 2 On my reading, what makes the problem of consciousness the hard problem is an insistence that consciousness makes the content phenomenal. 2-factor theories reject that notion and argue that phenomenality precedes consciousness. This is, obviously, a buck-passing claim, but an important one as it frees theories of consciousness of some of the heaviest baggage that might prevent us from moving forward.
- 3 Although not explicitly mentioned by the authors, I assume the modal verb 'can' is used in empirical or nomological sense. Logical or metaphysical possibilities have little bearing on how to account for consciousness in this world.
- 4 To be more precise, in induced impoverished conditions (masking, short duration of exposition, etc.), responses might be more like those in standard conditions. However, the full identity is rarely, if ever the case. Vosgerau et al. concentrate on neuropsychological cases such as blindsight or agnosia and I thereby concentrate on these as well.

References

- Champoux, François, Olivier Collignon, Benoit A. Bacon, Franco Lepore, Robert J. Zatorre, and Hugo Théoret. 2010. "Early- and Late-Onset Blindness Both Curb Audiotactile Integration on the Parchment-Skin Illusion." *Psychological Science* 22 (1): 19–25. <https://doi.org/10.1177/0956797610391099>

- Dubova, Marina, and Robert Goldstone. 2021. "The Influences of Category Learning on Perceptual Reconstructions." *Cognitive Science* 45 (5): 1–24. <https://doi.org/10.1111/cogs.12981>
- Fink, Sascha Benjamin, Lukas Kob, and Holger Lyre. 2021. "A Structural Constraint on Neural Correlates of Consciousness." *Philosophy and the Mind Sciences* 2: 1–23. <https://doi.org/10.33735/phimisci.2021.79>
- Magnotti, John, F., and Michael S. Beauchamp. 2017. "A Causal Inference Model Explains Perception of the McGurk Effect and Other Incongruent Audiovisual Speech." *PLoS Computational Biology* 13 (2): e1005229. <https://doi.org/10.1371/journal.pcbi.1005229>
- Marvan, Tomáš, and Michal Polák. 2017. "Unitary and Dual Models of Phenomenal Consciousness." *Consciousness and Cognition* 56: 1–12. <https://doi.org/10.1016/j.concog.2017.09.006>
- McGurk, Harry, and John MacDonald. 1976. "Hearing Lips and Seeing Voices." *Nature* 264 (5588): 746–748. <https://doi.org/10.1038/264746a0>
- O’Callaghan, Casey. 2017. *Beyond Vision: Philosophical Essays*. Oxford: Oxford University Press.
- O’Callaghan, Casey. 2019. *A Multisensory Philosophy of Perception*. Oxford: Oxford University Press.
- Palmer, Terry D., and Ashley K. Ramsey. 2012. "The Function of Consciousness in Multisensory Integration." *Cognition* 125 (3): 353–364. <https://doi.org/10.1016/j.cognition.2012.08.003>
- Rosenthal, David M. 2010. "How to Think about Mental Qualities." *Philosophical Issues* 20: 368–393. <https://doi.org/10.2307/1522921>
- Searle, John. 1992. *The Rediscovery of the Mind*. Cambridge, Mass.: The MIT Press.
- Skrzypulec, Błażej. 2022. "Contents of Unconscious Color Perception." *Review of Philosophy and Psychology* 13: 665–681. <https://doi.org/10.1007/s13164-021-00552-7>
- Tiippana, Kaisa. 2014. "What Is the McGurk Effect?" *Frontiers in Psychology* 5: 75. <https://doi.org/10.3389/fpsyg.2014.00725>
- Vosgerau, Gottfried, Tobias Schlicht, and Albert Newen. 2008. "Orthogonality of Phenomenality and Content." *American Philosophical Quarterly* 45 (4): 309–328.