Rational relations between perception and belief: the case of color

invited contribution for *Review of Philosophy and Psychology* (resubmitted)

Dr. Peter Brössel
Director of the Emmy Noether Research Group —
From Perception to Belief and Back Again
Department of Philosophy II
Ruhr-University of Bochum, Germany

Contents

1 Introduction

It is a longstanding philosophical idea that we form our perceptual beliefs (PBs) on the basis of perceptual experiences (PEs), and that a PB can be justified by the corresponding PE. This procedure corresponds to the first step of rational, or justified, belief acquisition. In a second step, we use these PBs to justify all our other a posteriori beliefs. Concerning the second step, traditional and formal epistemology and philosophy of science have progressed considerably in explaining how we proceed rationally from observational data to theories about the world (Carnap 1950, 1971a,b; Hartman & Sprenger 2010; Leitgeb 2015; Popper 1935). Concerning the first step, some progress has been achieved in traditional epistemology and in the philosophy of perception (Evans 1982; Gaucker 2011; Heck 2000; Macpherson 2012, 2015; Peacocke 1992, 1998, 2001; Pryor 2000; Siegel 2010; Stalnaker 1998). However, in formal epistemology virtually no progress has been made concerning the first step, and indeed some of the published works present negative results (Bradley 2015; White 2006). Only very recently have formal philosophers started to make progress in this area, most of whom are working at the intersection of formal epistemology and the philosophy of cognition (e.g., Douven & Decock 2014; Churchland 1992, 1998; Gärdenfors 2000, 2014). The present paper investigates the first step of rational belief acquisition. Thus, it focuses on justificatory relations between PEs and PBs, and between their contents, respectively. In particular, the paper aims at outlining how it is possible to reason from the content of PEs to the content of PBs. The paper approaches this aim by combining a formal epistemology perspective with an eye towards recent advances in the philosophy of cognition. Furthermore, the paper restricts its focus, concentrating on the case of color perception and perceptual beliefs about color.

The paper is structured as follows: Section 2 outlines the state of the art concerning how to reason from PEs to PBs and it draws a rather bleak picture. Starting from Section 3 we outline a new approach concerning how to recast the traditional empiricist picture of the relation of PEs and PBs in a new formal framework. The paper concentrates on the example of color perception and beliefs about the color of objects. Thus, Section 3 discusses how to represent the content of color perception and the content of corresponding beliefs—more specifically, the content of PEs of color and the content of the main building block of our PBs about the color of objects, i.e. color concepts. In particular, in subsection 3.1 it is argued that PEs have non-conceptual content and that that non-conceptual content of color experience is best represented as points in a phenomenal similarity space. Section 3.2 introduces Gärdenfors's idea to use geometrical spaces to model the content of concepts as regions in that space. (Since we focus on the case of color, we restrict ourselves to color space instead of using the full-blown apparatus of Gärdenfors's Conceptual Spaces framework and we argue that an essential semantic aspect of perceptual color concepts can be represented as regions in a phenomenal similarity space.) Section 4 outlines what a rational relation between the non-conceptual content of PEs of color and the conceptual content of PBs about color might look like. In particular, we define a rational degree of belief function that shows how one could reason from such PEs to PBs about color. In Section 5 we conclude the paper and outline the directions of future research.

2 The State of the Art

According to the traditional picture advocated by British and Logical Empiricists the following three claims are correct: (i) The content of PBs is conceptual. (ii) The content of PEs is independent of our beliefs, desires, and of our concepts or conceptual schemas (and independent of all of cognition in general). Paraphrasing (ii) in modern terminology: the content of PE is non-conceptual (Evans 1982; Heck 2000; Peacocke 2001) and cognitively impenetrable (Pylyshyn 1999). Finally, (iii) we form and we can justify our PBs (i.e., our empirically most basic beliefs) on the basis of our PEs.[1] This traditional idea is also very entrenched in the cognitive sciences.

---

[1]Indeed, some empiricists have even argued that in virtue of the independence of the non-conceptual content of PEs and the conceptual content of beliefs in general and PBs in particular, PEs can serve as a secure fundament for justifying

Focussing on visual perception, adopting terminology from Marr (1982) and following Pylyshyn (1999), defenders of the traditional picture in cognitive science distinguish between, on the one hand, early and cognition-independent visual processes and, on the other, late recognitional processes essential only for categorization tasks that work on the basis of the output of early vision. The output of early visual processes is thereby considered not to involve any conceptual representation and to be independent of cognition. Thus, according to this classical picture, the output of early visual processes is non-conceptual and cognitively impenetrable. The output of late vision on the other hand is considered to involve conceptual representation and therefore to be cognitively penetrable to at least some extent. Among other things, the conceptual content of late perceptual processes certainly depends on having the relevant concept in the first place.

This traditional picture possesses many features that epistemologists as well as—and especially— cognitive scientists consider to be advantageous. From an epistemic point of view, this is an attractive picture; one cognition-independent part of visual perception—i.e. what cognitive scientists refer to as early vision—lies at the root of all recognition and categorization tasks. In particular, traditional epistemologists like to think of the content of visual PEs as consisting of the non-conceptual output of early vision and the conceptual output of late vision. Clearly, if the non-conceptual content of visual PEs is cognitively impenetrable, then PEs could serve as an epistemic foothold for the justification of vision-based beliefs with conceptual content. If it were independent of cognition, no epistemic circularity worries (Siegel 2012; Lyon 2011) would have to arise, and the non-conceptual content of visual PEs would be certain to be not theory-laden in any interesting sense of the word (Kuhn 1962; Popper 1935; Raftopoulos 2009).

This picture is also attractive from a cognitive science perspective, given that it allows us to understand visual processing as essentially encapsulated, which suggests a modular and, thus, relatively simple architecture of the brain. According to the most simple picture, there is one module responsible for early visual processing, and the output of this module is feedforwarded into brain areas processing higher-level cognitive tasks such as recognitional and categorizational tasks. The alternative would be to replace the standard picture of bottom-up processing from an encapsulated visual module upwards to higher-order cognitive states by a more dynamic interplay between higher-order cognitive states and lower-order perceptual states (the inputs to the senses being the lowest state), which suggests a more cross-linked architecture and a continuous functioning between perception and cognition.

Over the years this traditional picture has come under attack from various sides. On the one hand, various epistemologists have argued against the picture from a purely philosophical perspective. In particular, some follow Davidson in arguing that the non-conceptual, cognition-independent content of PE cannot justify PBs, non-conceptual, cognition-independent content of PEs can only cause them. Others hold that the content of PEs is purely conceptual and thus, is not completely independent of cognition. Brewer (1999) and McDowell (1994), for example, are very explicit about this. The main motivation for this view is the position that only the conceptual content of PEs can be relevant for the justification of PBs and that such content is the only content of philosophical relevance. According to these philosophers, all content of PEs can also be content of beliefs (see Siegel 2015, §6). They thereby reject both Davidson's (1986) idea that PEs only cause PBs as well as the claim that only beliefs can justify beliefs—since for them, PEs can too. (McDowell admits that PE might include non-conceptual content in the sense that it includes, for example, picture-like images. However, since they are only causally relevant but not justificatorily relevant, he considers that part of PE to have no content.)

Leaving behind the traditional picture, the best contemporary theories of the justification of PBs rely on the conceptual content of PE to justify PBs (Pryor 2000, Siegel 2017).[2] The basic idea is:

PBs. Here we ignore this strong variant of empiricism. We also only briefly address the empiricist claim that we learn our first concepts on the basis of the non-conceptual content of PEs.

[2]For the moment we concentrate on internalist theories of justification. Externalist theories of justification that ascribe

if an epistemic agent has the PE that p (e.g., *this object is red*), then her PB that p (i.e., *this object is red*) is justified, prima facie. Given this approach, the most pressing question is: what concepts can figure in PE? Can it only contain concepts such as RED, LOUD, SALTY, or can it also contain concepts such as DIAMOND, CAR, HIGGS-BOSON, BREAKABLE, HAPPY? This question corresponds to the question of whether the content of PE is thin or rich. That the content of PE is rich in the sense that it includes highly complex or abstract conceptual content has been defended for the perception of causal relations between events, of intentions of agents, of natural and artificial kinds, and for the perception of a variety of aspects in the context of social cognition. (For a discussion of the thin vs. rich content view of PE see, for example, Newen (2016), Siegel (2005, 2010), and Toribio (2015).)

On the other hand, a number of empirical findings since the 1940s suggest that PEs depend in multifarious and relatively direct ways on higher-order cognitive processes such as beliefs and desires, etc. First, psychological findings from the early "new look psychology" (as initiated in the works of Brunner & Goodman (1947), Bruner & Postman (1947, 1949), and others) to contemporary results reported in Levin & Banaji (2006), Balcetics & Dunning (2010), Vetter & Newen (2014), Lupyan (2012), Proffitt & Linkenauger (2013), and others, are taken to demonstrate the influence of cognition on the non-conceptual content of PEs. (See Firestone & Scholl (2015) for further references and critical discussion.)

Second, neuroscientific data shows that brain areas associated with higher-order cognitive processes feed back into brain areas associated with early vision. Thus, in the light of our present knowledge about the neurobiological structure of the brain, a more cross-linked architecture and a continuous functioning between perception and cognition might be the more plausible picture. For example Vetter & Newen 2014 refer to Salin & Bullier 1995 who show that "[t]he neurophysiological connections of the primary visual system with the rest of the brain are such that there are much more feedback connections to primary visual cortex from higher cognitive areas than feedforward connections to higher cognitive areas", and according to Vetter and Newen this strongly supports that cognitive penetration is possible and occurs frequently. In addition, Vetter and Newen argue that the "time course of visual processes in V1 and V2 is such that we cannot presuppose simple serial feedforward processing" from from early vision to late vision and thus "our visual experience is not the product of a bottom-up encapsulated modular process but the product of an embodied perception-expectation-action loop which is implemented for a cognitive system by a highly flexible multiple integration of bottom-up and top-down processes" (Vetter & Newen 2014: 66). Given these empirical findings, a modular structure of the brain no longer holds the promise of providing a simpler model of the perception–cognition relation. To account for early visual processing in terms of informationally encapsulated modules one may have to posit not one but many such modules, each with its own set of functional input–output descriptions and each with its own individual neurological implementation and complex and varying patterns of interaction with other encapsulated modules. Thus, modular accounts of early visual processing might appear simpler a priori, but given the empirical findings reported in Vetter & Newen 2014 it is doubtful that they can live up to that promise. Whether these psychological and neuroscientific findings already establish that there is an effect of cognitive processes on perception, and if so whether it must be understood as cases of cognitive penetration, is of course hotly debated in the literature (see, for example, Raftopoulus 2001 and Firestone & Scholl 2015).

Third, this traditional picture has also come under attack from certain theoretical approaches in cognitive science. In particular, in the light of one recent theoretical approach in cognitive science, i.e. the predictive processing or the predictive error minimization framework (Friston 2009, Hohwy 2013, Clark 2013), one would rather expect that cognition alters PEs frequently and sometimes comprehensively. The framework suggests that perception is highly dependent on cognition in general and beliefs and desires in particular. Accordingly, many cognitive scientists believe that

---

no relevant epistemic role to PE are briefly discussed later in the paper. Externalist theories often concentrate on the process that leads to the PB: to the extent that it is reliable, the PB is justified whatever the content of the PE.

perception is *cognitively penetrable*.

There are two lines of defense that advocates of the traditional picture are building up. First, they argue that not every causal influence reaching from cognition to perception counts as an instance of cognitive penetration, because many of these influences do not invoke worries about epistemic circularity or the charge that the content of PE is theory-laden. The idea is, if changes in the content of one's PE are due to factors such as attention or the state of the perceptual system, etc., then the mentioned circularity worries do not necessarily arise. They only arise if those changes have been induced by the belief itself.[3] Thus, many require a semantic or rational connection between the content of cognition and the content of perception for cognitive penetration to hold. That such a rational connection is necessary for cognitive penetration to hold has been convincingly argued by Pylyshyn (1999) and very forcefully by Macpherson (2012). Pylyshyn writes:

> [An] influence that . . . originates outside of the visual system and affects the content of visual perception (what is seen) in a certain meaning-dependent way . . . we call cognitive penetration. . . . [I]f a system is cognitively penetrable then the function it computes is sensitive, in a semantically coherent way, to the organism's goals and beliefs, that is, it can be altered in a way that bears some logical relation to what the person knows. (Pylyshyn 1999: 343)

Accordingly, if the content of PE is cognitively penetrable, then the content of PE is sensitive, in a semantically coherent way, to the subject's cognitive states such as desires and beliefs. The content of that subject's PE can be altered by altering the content of her beliefs and desires, and during all those alterations there remains some logical, semantical or rational relation between the content of the subject's cognitive state and the content of that subject's PEs.

A second line of defense of advocates of the traditional picture consists in doubting the suggested interpretation of the empirical data. Defenders of the traditional picture argue against the psychological data, claiming that these experiments and questionaries only demonstrate that there is some influence on the conceptual content of PEs, respectively on content of the PBs. They do not demonstrate that there is some influence on the non-conceptual content of PE. Against the neuroscientific data they argue that we do not know how the output of early vision is influenced by higher-order cognitive states. That there are many more feedback connections to the primary visual cortex from higher cognitive areas than feedforward connections to higher-order cognitive areas does not show that cognitive penetration takes place. The top-down feedback connections from higher cognitive areas to primary visual experience might only be relevant for distributing attention to certain perceptual states, and we have absolutely no idea whether and how they might change the content of PE in a semantically coherent way. In particular, from the neuroscientific data alone, we have no reason to expect that the content of PEs is theory-laden or that it is the result of circular processing within the brain. Thus, they maintain the most plausible hypothesis

---

[3]The relation between attention and cognitive penetration is disputed in the literature. Obviously, changes in attention or in the state of the perceptual system can lead to epistemic problems as highlighted by Stokes (2015). The question is whether these epistemic problems are of a distinctively different form. According to philosophers such as Fodor and Pylyshyn, they are. Changes in attention might lead to similar epistemic problems as directing your gaze or measuring instrument in a different direction do, one can miss evidence. Changes in the state of the perceptual system might lead to similar epistemic problems as using faulty or unclean measurement instruments, one can obtain false evidence. However, they maintain that these epistemic problems are of a different form than the epistemic problems associated with theory-ladenness of observation and cognitive penetration. Circularity worries seem to arise only in the latter two cases where your theory, your belief, or respectively your language influences the content of perception experiences and perceptual beliefs. However, Mole (2015) and Marchi (2017) argue that, depending on the theoretical assumptions about attention and the cognitive system, attentional processes can be a contributing cause of cognitive penetration. For example, according to Marchi (2017) in the predictive processing attention is associated with expectations of precision of the bottom-up information. The lower the expectation of precision, i.e. the lower the attention, the greater the top-down influence. Thus, it is low attention in the sense of the predictive processing literature, that enables higher-order beliefs to have a top-down influence on lower-order PEs. Arguably, however, it is the content-dependence of PEs from beliefs that invokes our worry of epistemic circularity. It is just the case that certain attentional processes are necessary for the worry of epistemic circularity to arise.

is still that the output of early vision is cognitively impenetrable.

Nevertheless, even if those two lines of defense hold, epistemologists and philosophers of cognition have navigated themselves into a tricky position. More specifically, even if the two lines of defense hold this will not suffice for saving the justificatory role that we attribute to PE: neither the non-conceptual content of PEs nor their conceptual content can be used for the justification of PBs. Given the available empirical data, everyone seems to admit that the conceptual content of PEs, i.e. the output of late vision-associated recognitional and categorization tasks, is probably cognitively penetrated and, thus, ill-suited as an epistemic fundament for the justification of PBs. Indeed the second line of defense admits that the available experiments and observations demonstrate that there is some top-down influence on the conceptual content of PEs, respectively on the content of the PBs and the respective judgments. What they aim at saving is the cognitive impenetrability of the output of early vision, i.e. the non-conceptual content. However, in opposition to the traditional picture, epistemologists have also argued that only the conceptual content of PE is relevant for justifying PB. The problem then is neither the conceptual nor the conceptual content of PEs can serve as an epistemic fundament for justifying PBs. Attempts to tackle this challenge are very rare, one noticeable exception being Siegel (2017).

In the following section, we lay the groundwork for a third line of defense. We outline a formal theory of the non-conceptual content of PE and the content of PB to show that PBs can stand in semantic or rational relations with the non-conceptual content of PEs. The outlined theory shows that the non-conceptual content of PEs can serve as an epistemic fundament for the justification of PBs if its content is sufficiently independent of the initial beliefs of the agent, i.e. if the first and second lines of defense hold. In addition, this theory will also be relevant for investigating whether the non-conceptual content of PE is sufficiently independent of cognition as suggested by the second line of defense. As already argued, this question can only be answered when we understand how far the non-conceptual content of PE can be influenced by the content of one's beliefs in virtue of the rational or semantic relations between them.

3 The Content of Color Experience and the Content of Color Concepts

3.1 The Content of Color Experience

We want to advocate a picture according to which the content of PE includes non-conceptual content. The three most prominent reasons that are typically taken to support that PEs must include non-conceptual content are introduced in the following.[4]

First, our *discriminatory abilities* seem to be more fine-grained than our ability to draw conceptual distinctions (e.g. following Evans 1982, see also Peacocke 2001). Consider, for example, our ability to discriminate various shades of color. This discriminatory power can be best understood by assuming that there is some non-conceptual content within our color experiences that allows us to discriminate between different shades of green. After all, it seems plausible that we do not have and also could not have a concept for each of these shades of green that we are able to distinguish. (Of course, once we admit the existence of non-conceptual content we can refer to the non-conceptual content of PEs with the help of demonstrative concepts like 'that shade of color', 'that form', etc.)

Second, the *content of PE is analog* and not digital (see especially Goodman 1976 and Dretske 1981) (respectively, continuous and non-discrete). Colors, for example, can be more or less sat-

---

[4]They are prominently discussed in Evans 1982, Heck 2000, and Peacocke 2001. Macpherson (2015: 335-337) even discusses six features of PEs that point towards non-conceptual content.

urated and there is always a third color that is less saturated than the first and more saturated than the second (if we presuppose that the two initial shades of color have a different degree of saturation). The same seems to be true for the conditions of accuracy of PEs. The contents of beliefs are true or false and, thus, they are not more or less accurate. In contrast to the content of beliefs, the content of our PEs can be more or less accurate. If one perceives an object as being two arms' length away whereas it is actually three arms' length away, then this PE is more accurate than if one would have perceived it just as being just one arms' length away. (Here, of course, I assume that the representations which accuracy conditions are in question are not represented conceptually, but non-conceptually.[5])

Third, *explaining the behavior of animals and pre-conceptual infants* seems to require that we assume that PEs (as well as other cognitive processes) include non-conceptual content. Bermudez & Cahen (2015, §4.3), for example, argue to this effect. When one witnesses how animals react in response to certain observational situations, then we have to admit that we cannot explain this behavior without assuming that there is some content that is represented in the animal's mind. Since this content is, presumably, not conceptual, we have to assume it is non-conceptual; so the argument goes. The ability of children to learn new perceptual concepts is another case in point. If we want to understand learning concepts of infants, the traditional idea is the following: there is some non-conceptual content given in PEs and children acquire their first concepts by learning that PEs with very similar non-conceptual content are labeled with identical words. Especially empiricists embrace the idea that all or most of our concepts are learned via PEs. However, to explain concept learning without presupposing prior acquaintance with concepts we have to assume that there is some other content in PEs than conceptual content.
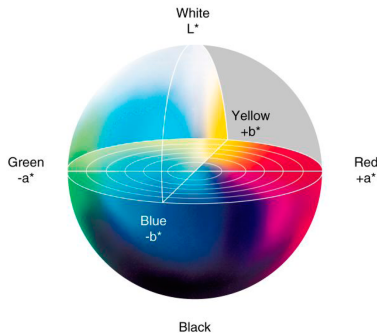
For the purpose of this paper, we presuppose that the above considerations demonstrate that PEs have non-conceptual content. However, merely stating that PE include non-conceptual content does not suffice for the purpose at hand: to show that we can reason rationally from PEs to PBs. We need to say more about the structure of the non-conceptual content and how this content can be used to provide a theory of rational reasoning from PEs to PB. It has been suggested that the non-conceptual content of PE can be understood in terms of sets of (centered) possible worlds (Stalnaker 1998) or as "scenario content" (Peacocke 1992). Peacocke's idea is that the non-conceptual content conveys geometrical information about the scenario around us. He insists that his account is an account of the content of PE itself in contrast to the content of PE as experienced by the agent. He writes: "I should also emphasize that the positioned scenario content is literally meant to be the content itself. It is to be distinguished from any mental representation of that content" (Peacocke 1992: 66). Thus, neither account helps us to describe the non-conceptual content of a specific PE as phenomenally experienced by the agent, nor describes how the agent processes that content. In both cases, agents do not have epistemic access to the non-conceptual content of PE (neither consciously nor unconsciously) for this we would need to know how that content is mentally represented. In the following we provide such an account.

Psychologists and cognitive scientists assume quite readily that there is non-conceptual content within PE and they provide formal models for this content. While philosophers commonly concentrate on the relationship between the PE and what it represents in the external world (in close analogy to philosophy of language, where extensional semantics dominate), psychologists and cognitive scientists focus on the content of PE as phenomenally experienced by the subject (in close analogy to intentional semantics). More specifically, when psychologists and cognitive scientists want to model the variety of possible phenomenal experiences a subject can experience, they often refer to phenomenal similarity spaces. For example, the color space represents the shades of colors a subject can experience: it represents shades of color with the help of a geometrical

---

[5]There also might be a conceptual representation of distance. Firestone and Scholl (2015), for example, defend the cognitively impenetrability thesis in the context of tasks involving the evaluation of distanc, by admitting that perceptual judgments of distance, respectively PBs in our terminology, can be top-down influenced. The non-conceptual content of PEs concerning distance, they maintain, are cognitively impenetrable.

space and the distance between points in that space is interpreted as representing the dissimilarity between two shades of color.

Example: the Phenomenal Color Space

The three quality dimensions of color perception:

- hue (red, yellow, green, blue, pink, and red again)
- brightness (white, gray, black)
- chromaticness/saturation (colored, gray)
- every perceivable color is a point in the three dimensional color space.

The idea is that within that color space we can represent shades of color as points in the space. Every point in that space specifies the hue of the color, its brightness and its saturation. Every other color shade differs from the first in at least one of these respects. According to this idea, the non-conceptual content of our PE of color can be represented by a point in the color space. Such a representation of the shade of color has the advantage that the content of color perception can indeed be said to be analog or continuous instead of being digital or discrete. The above color space, the CIE-LAB 1976 color space model, is informed by empirical investigations into how humans perceive color and how the cones on the retina are influenced by light of different wave lengths.[6]

Representing shades of color as points in some color space seems to be appropriate and in line with the state of the art in psychology and cognitive science. Indeed, the empirical research in this area of research does not question this; rather, it concentrates exclusively on the question of which space is the most suitable one for representing color appearance. By following this proposal we can also explain the above three observations that prompted philosophers to argue that there is non-conceptual content within PE. The continuous color space can account for the fine-grainedness of our discriminatory abilities of color if compared to our ability to draw conceptual distinctions based on color concepts. (It can even account for the vagueness associated with borderline cases of specific color concepts. Here the concept of a border is not just used in a metaphorical sense. The relevant literature is the following: Douven et al. 2013.) Of course, this representation of color can also account for why we characterize the content of PEs as analog or continuous, since it represents shades of color as points in a continuous space. Finally, we can use this representation of color experiences for explaining the behavior of animals and infants.[7] In principle, this approach also offers enough resources to explain how children learn color concepts in the first place (omitted for blind review).

One possible objection needs to be addressed. If we express the content of color perception by stating that an object o is presented to us in PE as having the color $\langle x, y, z \rangle$ in the color space (short: $c(o) = \langle x, y, z \rangle$), we seem to express a proposition which is true or false. Thus, one might object that the content of such an expression is digital or discrete and not analog or continuous as required for non-conceptual content of PE. Therefore, the content expressed is conceptual and not non-conceptual. This predicament is due to the fact that for expressing the non-conceptual content

---

[6]The best color appearance models nowadays use more than three dimensions to represent color as experienced by humans, and thus it is less than ideal to represent them pictorially. The CIE-CAM02 model uses the dimensions brightness (luminance), lightness, colorfulness, chroma, hue and saturation (where saturation could be understood as a dependent dimension that does not need to be added to the model). For the purpose of the paper, it is not important in which space we represent color experiences, it only matters that there is such a space and that we can represent a specific color experience as a point in that space.
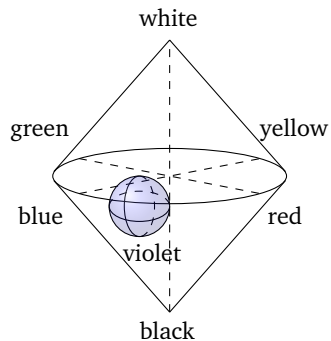
[7]In biology we need to use different phenomenal color spaces for different species to explain their behavior (sometimes they might even differ within species).

8

of color perception we need to rely on concepts. Using concepts to express the present account of non-conceptual content cannot be circumvented. We need another way to cash out the idea that $c(o) = \langle x, y, z \rangle$ expresses the non-conceptual content of color perception. We propose that we help ourselves by restricting how we can use these expressions in order to express the non-conceptual content of color perception. In particular, we suggest that the following content cannot be the outcome of a PE of color: $c(o) \neq \langle x, y, z \rangle$ and $c(o_1) = \langle x_1, y_1, z_1 \rangle \wedge$ *(or $\vee$)* $c(o_2) = \langle x_2, y_2, z_2 \rangle$. Given these restrictions it is reasonable to say that $c(o) = \langle x, y, z \rangle$ expresses the analog or continuous content of color perception. The discussion of to what extend this approach can be generalized for all aspects of visual experience such as the shape and location of objects, the view point of the observer and the illumination of the whole scene will be take up again in the final section of this paper.

3.2   The Content of Color Concepts

In this subsection we outline our account of the conceptual content of PBs concerning color. The conceptual content of PBs (and possibly of PEs if they possess conceptual content) should enable us to elaborate a rational relation between the non-conceptual content of PE and the conceptual content of PBs, in order to proceed rationally from PEs to PBs. In order to reach this goal we follow Gärdenfors and understand concepts as being associated with certain regions in a conceptual space. Conceptual spaces are geometrical spaces and their geometrical structure is used to capture semantical properties and relations. In this paper, we restrict ourselves to phenomenal similarity spaces instead of using the full-blown apparatus of Gärdenfors's Conceptual Spaces framework. The dimensions in the phenomenal similarity space represents respects in which objects can be experienced as similar or dissimilar and distances between points in the space represent the phenomenal similarity of objects within the space. With this restriction to phenomenal similarity spaces comes the confinement that we only aim at modelling the content of perceptual concepts, respectively the perceptual aspects of a given concept. We do not want to take any stance on the question of whether all semantic aspects of a given concept can be represented within the conceptual spaces framework.

In more detail, the idea is the following. Let us refer to those concepts that can figure in PBs (and possibly PEs) as *perceptual concepts*. One main purpose of perceptual concepts is to allow for the categorization of objects in a few linguistic categories on the basis of our manifold PEs of those objects. For example, the color concept VIOLET allows one to categorize objects according to their manifold shades of color as experienced by the agent. With the help of such concepts, we achieve two goals. First, these perceptual concepts can be understood as coming with what has been called 'language entry rules' (Putnam 1981: 11). They allow us to introduce a concept or a word for a given object on the basis of a non-conceptual fundament; thus, perceptual concepts enable us to climb the first step of rational belief acquisition. For example, the perceptual concept VIOLET enables us on basis of our PE to introduce the PB that the object is violet. Second, such concepts also allow us to group similar shades of color together by subsuming them under one concept or category. For example, the concept VIOLET subsumes various different shades of color under one label and delimits them from various other shades of color. Given this job description, it is only natural to understand the concept VIOLET as corresponding, at least in part, to a region in the phenomenal similarity space. This region then includes all those manifold points in the space that we want to group in one category under one label and it thereby delimits these points in the space from other points in the space that lie outside of the relevant region in the phenomenal similarity space. Schematically, we can depict this as follows:

white

green   yellow

blue   red

violet

black

The color term 'violet' (as used in the English language) is associated with the concept VIOLET which can be understood as corresponding to one specific region in the color space. (As already said, at least in part it should correspond to such a region. The whole concept might also be identified via its relations to other concepts.)

A first advantage of this approach to capturing the content of perceptual concepts is that it is simple and natural to integrate prototype semantics within that framework. In particular, various empirical findings suggest that objects instantiating a perceptual concept can be more or less representative of the category. For example, robins are more representative of the category BIRD than eagles, ravens and many other birds (Rosch 1975,1978; Mervis & Rosch 1981). Certain shades of violet are more prototypical than others. Similar prototypical effects have been shown to exist for various kinds of concepts. Within the present framework these effects can be modeled with the help of the geometric properties of the phenomenal similarity space. The main idea is the following: a given perceptual concept can be understood as a region in a phenomenal similarity space, and our prototype for that concept can be understood as a point or a prototypical region at the center of the region of that concept.

A second advantage is that we can understand how the non-conceptual content of PE can serve as a fundament for acquiring concepts in the first place. First, children have manifold PEs of the objects around them, and from competent speakers of the language they learn to draw the same boundaries (respectively sufficiently similar boundaries) between points in the phenomenal similarity space. Only when they learned distinguishing shades of violet from shades of blue, red, yellow, etc., for example, do we say that they have mastered the concept VIOLET. Of course, we still need to explain *how exactly* children learn their first concepts on the basis of the non-conceptual content of PE. Gärdenfors (2000) shows that we can use geometrical spaces and the idea of prototypes to explain concept acquisition. (omitted for blind review) offer an alternative account, according to which children use a quasi-Bayesian reasoning process to acquire concepts. For the purpose of the paper, the details of those accounts are not important. What is important is that the present approach allows us to frame various different theories about how children can acquire first concepts from the non-conceptual content of PE without presupposing that they already possess perceptual concepts. This provides additional support for the present proposals for understanding perceptual concepts and the content of PE. In particular, it shows that the present proposal is preferable to alternative proposals that do not display the potential to account for concept learning in the indicated way.

4 A Rational Relation Between PE and PB

In the following we outline how the non-conceptual content of PE can stand in rational relations with the conceptual content of PB; we do so by focusing on color perception and beliefs about color.

The challenge is to specify in precise formal terms under which conditions it is rational to form a belief about the color of an object on the basis of one's color experience of that object. In order to

meet this challenge, let us briefly discuss the standard theory of rational reasoning *within episte-mology and cognitive science*: Bayesianism.

Since the beginning of analytic philosophy, scientific philosophers and philosophical scientists such as de Finetti (1937), Keynes (1921), von Mises (1931), and Ramsey (1926) have argued that the foundations of rational reasoning are best captured in terms of probabilities. The reason for this is simple: the Bayesian approach to understanding the foundations of rational reasoning in terms of probability is the most intuitive and comprehensive approach to such an enormous challenge. Strevens (2008) summarizes the advantages of the Bayesian approach:

> The popularity of the Bayesian approach is due to its flexibility, its apparently effortless handling of various technical problems, the existence of various a priori arguments for its validity, and its injection of subjective and contextual elements into the process of [rational reasoning] in just the places where critics of earlier approaches had come to think that subjectivity and sensitivity to context were necessary. (Strevens 2008: 5)

As indicated by Strevens, there are many a priori arguments for Bayesianism, from the so-called Dutch book argument by Ramsey (1926) over the calibration argument by van Fraassen (1983) to the accuracy argument by Joyce (1999) and Leitgeb & Pettigrew (2010a,b). These arguments have lead to an almost universal acceptance of Bayesianism by those philosophers and psychologists (and economists and scientists) who employ formal tools to understand and model rational reasoning.

The application of the Bayesian framework has proven to be extremely fruitful in cognitive science, too. In particular, it has been applied to the problem of perception and categorisation and the problem of sensory coding and processing. The idea that perception should be understood as being inferential is often traced back to Helmholtz. Already Helmholtz used probabilities to model the nature of the unconscious inferential processes underlying perception and categorization. Contemporary approaches in psychology and neuroscience to model perception are discussed extensively in Knill and Richardson (1996) and Rao, Olshausen & Lewicki (2002). The more general Bayesian brain hypothesis, i.e. the hypothesis that in general (not just in perception) the brain codes and processes as if it were employing Bayesian methods or models, is discussed by philosophers of cognition and neuroscientists too, e.g. Colombo & Seriès (2012), Knill & Pouget (2004), Friston (2009) and Hohwy (2013).

According to the Bayesian norms of reasoning as they are understood in epistemology, an agent's degrees of belief should obey the axioms of probability and the agent should update these degrees of belief in the light of new evidence by using Bayes's Rule.[8] The axioms of probability theory state the following:[9]

Axioms of Probability. Let W be a set of possibilities (e.g. possible worlds) and let $\mathcal{A}$ be an algebra of subsets over the set W of possibilities. A function $\mathrm{Pr} : \mathcal{A} \to \mathbb{R}$ is a probability function on $\mathcal{A}$ if and only if for all $A, B \in \mathcal{A}$:

1. $\mathrm{Pr}(A) \geq 0$

2. If $A = W$, then $\mathrm{Pr}(A) = 1$

3. $\mathrm{Pr}(A \cup B) = \mathrm{Pr}(A) + \mathrm{Pr}(B)$, if $(A \cap B) = \varnothing$

---

[8]For the purpose of this paper, we ignore the update method of Jeffrey conditionalization.

[9]Especially for philosophical applications, the notions of conditional probabilities play an important role. Thus, the above definition of probabilities needs to be supplemented by a definition of conditional probabilities.

Definition. If $\mathrm{Pr}(B) > 0$, then $\mathrm{Pr}(A|B) = \mathrm{Pr}(A \cap B) / \mathrm{Pr}(B)$

Then Bayes's Rule states how an agent should update her degrees of belief in the light of new evidence.

Bayes's Rule. If Pr is the agent's probability function at time point t_0, E is the logically strongest proposition that the agent became absolutely certain of between time points t_0 and t_1, and $\mathrm{Pr}(E) > 0$, the agent's probability function after learning E should change to Pr_E, which is defined as follows:

$$\mathrm{Pr}_E(H) = \mathrm{Pr}(H|E)$$

for all $H \in \mathcal{A}$.

The problem for the present purpose is that in epistemology the elements of the algebra are typically interpreted as propositions that are either true or false. Thus, a proposition A is best understood as a binary propositional variable that can take on the values 1 or 0. The non-conceptual content of color experience, however, has been said to be continuous or analog and, thus, is best understood as a continuous random variable. As we recall, color experiences can be represented as points $\langle x, y, z \rangle$ in the color space. Color concepts, on the other hand, can be represented as regions in a space and they can be understood as binary random variables that take on the value 1 if the relevant object's shade of color falls under the given color concept and 0 otherwise.

It is possible to define probabilities for continuous random variables, and in the following we will introduce this machinery directly for the three dimensions of color perception.

Probabilities and Color Perception. *Let $\mathcal{X}, \mathcal{Y}, \mathcal{Z}$ be the continuous random variables associated with the dimensions of color perception and which can take on values x, y, and z in $\mathbb{R}$. Then $\mathrm{Pr} : \mathcal{C} \to \mathbb{R}$ is a probability function on $\mathcal{C}$ whose elements are the color concepts if and only if for all color concepts $C \subseteq \mathcal{X} \times \mathcal{Y} \times \mathcal{Z}$:*

$$\mathrm{Pr}\big(c(o) \in C\big) = \int_C f_{\mathcal{X}, \mathcal{Y}, \mathcal{Z}}(\langle x, y, z \rangle)\, \mathrm{d}x, \mathrm{d}y, \mathrm{d}z$$

where $f(\langle x, y, z \rangle)$ is the joint probability density function such that $f(\langle x, y, z \rangle) \geq 0$ and $\int_{\mathcal{X} \times \mathcal{Y} \times \mathcal{Z}} f(\langle x, y, z \rangle)\, \mathrm{d}x, \mathrm{d}y, \mathrm{d}z = 1$

This gives us the basis for understanding how one can reason rationally from the content of PE of color to the PB about the color of objects.[10] In particular, it shows how one can directly infer a belief about color from the content of perception. In our visual experiences, object o is represented as having the color $\langle x, y, z \rangle$ in color space. In cognition, the concept VIOLET corresponds to a region in color space that does or does not include the point $\langle x, y, z \rangle$. Based on the probability, we then can infer (from visual experience to belief) that object o is violet.

$$\mathrm{Pr}\big(c(o) \in \text{VIOLET} \mid c(o) = \langle x, y, z \rangle\big) = \begin{cases} 1 & \text{if } \langle x, y, z \rangle \in \text{VIOLET} \\ 0 & \text{if } \langle x, y, z \rangle \notin \text{VIOLET} \end{cases}$$

This formal theory accomplishes what we philosophers have sought after: it connects the non-conceptual content of PEs with the conceptual content of PBs in a rational way. Of course, this picture of the rational relation between the content of color perception and beliefs about color involves many idealizations, and a satisfactory theory of the rational relations between perception and belief has to overcome these idealizations. Indeed, overcoming these idealizations is the first topic we have to address when we discuss the directions of future research as regards the present approach.

---

[10]The relevant definition of conditional probabilities is a little bit more complicated.

Definition. $\mathrm{Pr}\big(A|c(o) = \langle x, y, z \rangle\big) = r$ iff $\mathrm{Pr}(A \wedge c(o) = \langle x, y, z \rangle) = \int_C f_{\mathcal{X}, \mathcal{Y}, \mathcal{Z}}(\langle x, y, z \rangle) \times r\, \mathrm{d}x, \mathrm{d}y, \mathrm{d}z$

5 Directions of Future Research

There are three important directions that future research concerning the approach presented here should explore. In the following we want to outline them and indicate the results that we can expect from such research endeavors. The first direction concerns the idealizations involved in the theory as it is presented so far. The aim must be to get rid of some of these idealizations. The second direction that future research should explore concerns the question of how one can generalize the present approach so that it is suitable for modeling the content of PE in general: the content of visual experiences as well the content of PE that originates from other sense modalities. The third direction should aim to put the present approach to work. The aim is to solve the important problems in epistemology and philosophy of cognition that motivated elaborating theories of non-conceptual content in the first place.

Let us begin with the idealizations within the approach as presented so far. First, the present approach assumes that the agent fully relies on her color perception and its rational relation with one's color concepts. Of course, this presupposition is often not satisfied. For example, suppose you are invited to a party and you see the suit of the party host only very briefly, perhaps for one second. Here you might not have gotten a clear and precise PE of the color of the suit and, thus, you might not be sure what PE you did have. Modeling this kind of uncertainty about the content of one's PE of color is possible. One can introduce second-order probabilities or rely on Jeffrey conditionalization properly formalized for continuous random variables. The relevant machinery is already introduced and discussed in detail in Skyrms (1980). Applying these formal tools to the presented theory of the non-conceptual content of PE seems to be appropriate and promising.

Another and slightly different example hints at a second kind of idealization: suppose you see the suit of the host in dim lighting conditions. It is hard to decide whether you should believe that the suit is dark blue or black, even though you have a clear and precise PE of the color of the suit. This example suggests that we should distinguish the concepts LOOKING BLACK and BEING BLACK. Given the PE you are justified to believe that the suit looks black in this example, but you are perhaps not justified to believe that it *is* black. (Indeed, given suitable background knowledge you might even be justified in believing that it is blue despite your PE of a black suit, e.g., if you have been told that the host will be wearing a dark blue suit.) Here the aim should be to provide a more realistic theory of the relation between the content of one's PEs and the content of PBs one is justified in entertaining in the light of the PE. For this purpose one must also model the content of one's PE of the lighting conditions, and only together with the content of your PE of the color of the suit are you justified in forming a PB about the color of the suit.

The third idealization involved in the examples is the assumption that the perceived shade of color is not a borderline case. Suppose you have a clear and precise PE of the color of the suit and you know that the lighting conditions are ideal for coming to a judgement about the color of the suit, but then you recognize that even given those ideal conditions you cannot form a PB about the color of the suit. This can happen if the perceived shade of color is a borderline case between blue and black. If we were dealing with a borderline color shade for the given color term, then assigning probability 1 or 0 to the relevant proposition would be inadequate. An approach for dealing with this problem is characterized in Decock & Douven (2014). In their paper they use conceptual spaces to model a graded membership relation, and the idea is to adjust the probabilities accordingly.

Now let us address the prospects for generalizing the presented approach to encompass all non-conceptual content of PE. The question is whether the content of PE in general can be represented by points in some geometric space. Previous work in the area of vision has been championed by David Marr (1982). According to Marr, all aspects of vision can be modeled geometrically. He writes:

> There are four main factors responsible for the intensity values in an image. They are (1) the geometry and (2) the reflectances of the visible surfaces, (3) the illumination of the scene, and (4) the viewpoint. In an image, all these factors are muddled up, some intensity changes being due to one cause, others to another, and some to a combination. The purpose of early visual processing is to sort out which changes are due to what factors and hence to create representations in which the four factors are separated. (Marr 1982: 41)
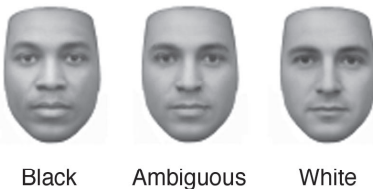
For all these four factors of a visual picture, Marr maintains that they can be represented as points in a complex geometric space. Churchland (1992, 1998), Gauker (2011) and especially Gärdenfors (2000, 2014) have made further significant advances to explain the various features of perception. They argue that we should explain these features by modeling the content of PEs in terms of similarity spaces. In addition, cognitive scientists have proposed phenomenal similarity spaces for sound, taste, and also the shape and movements of objects. (To my knowledge no one has yet proposed a geometric model for the non-conceptual content of PE of smell.) Thus, it is indeed a promising approach to model the non-conceptual content of PE in general with the help of geometric spaces. If this is so, then it is also promising to model one aspect of perceptual concepts as regions in the phenomenal similarity space. This way these concepts can function as language entry rules as envisioned by Putnam (1981). That these language entry rules can be understood in analogy to rational inference rules has been shown in the previous sections.

The final direction of future research that we want to address here is the application of the present theory to various philosophical problems in epistemology and the philosophy of cognition that motivated theories of non-conceptual content of PE in the first place. We have already discussed in some detail the traditional empiricist picture according to which one can justify PBs with the help of the non-conceptual content of PE. Another perhaps even more interesting application, which we already discussed in section three, is the possibility to provide a theory of concept learning from non-conceptual content of PE alone. This idea has been discussed in philosophy at least since the British empiricists and with the theory of the non-conceptual content of PE presented here it seems possible to overcome the challenges associated with it. If we associate a perceptual concept with the ability to categorize objects on the basis of our PE, we have to explain how pre-conceptual children acquire the ability to subsume different but similar objects under the same label. The key idea (pursued in Gärdenfors 2000, 2014; omitted for blind review) is to formulate the respective theory in terms of the similarity of the PE of those objects.

A further important and interesting application of the presented theory concerns the topic of cognitive penetration. As we recall, many epistemologists require that there is a semantic or rational relation between the content of the influencing belief and the influenced PE. For the sake of the presentation of the main idea of our approach we assume that the following is a case of cognitive penetration. (In our description of these empirical findings we follow Macpherson 2012.) Levin & Banaji (2006) studied the misperception of the lightness of faces due to the agent's beliefs. In particular, they report that

> "[w]hite faces were consistently judged to be relatively lighter than Black faces, even for racially ambiguous faces that were disambiguated by labels. Accordingly, relatively abstract expectations about the relative reflectance of objects can affect their perceived lightness." (Levin & Banaji 2006, abstract)

One of the experiments (Experiment 2) was designed as follows: First, three faces were constructed: one unambiguous Black face (B), one unambiguous White face (W), and one ambiguous face.

Black   Ambiguous   White

Second, "all subjects saw an instruction screen that included the ambiguous face next to one of the unambiguous faces underneath the labels Black and White. Subjects in the BW/W condition saw the ambiguous (BW) face paired with the White face, and therefore the BW face was labeled Black. In contrast, subjects in the B/BW condition saw the same ambiguous face labeled White and paired with the unambiguous Black face." Finally, on each trial, participants were presented with one face adjacent to an adjustable gray region and they were asked to match the lightness of the adjustable gray region to lightness of the displayed face. The interesting result was that the participants "chose darker samples for Black faces than for White faces and, in this case, chose a darker standard for the ambiguous face when it was labeled Black" than when it was labelled white. For Macpherson (2012), this is a clear case of cognitive penetration: our belief that the ambiguous face is White, respectively Black, influences our perception of the lightness of the face. If the label 'Black' induces in us the the belief that the ambiguous face is a Black person, then we perceive that face as darker than if the label 'White' induces in us the belief that the ambiguous face is a White person. Yet, if we follow traditional epistemologists then there is no semantic or rational relation between the content of the PE of the grayness of the face and one's belief that one sees a black person. There cannot be such a relation because the content of the belief is conceptual but the content of PE is not conceptual. However, the presented theory has the potential to account for the rational relation between the content of the belief and the content of the color experience. The concept BLACK FACE is associated with a region in the color space that includes darker shades if one compares it with the region in the color space that is associated with the concept WHITE FACE. Thus, in this case of top-down influence from beliefs to PE, the relevant criterion seems to be satisfied. The influence that "originates outside of the visual system and affects the content of visual perception (what is seen) in a certain meaning-dependent way" is a case of cognitive penetration. It is a case of cognitive penetration because the "function it computes is sensitive, in a semantically coherent way, to the [... agent's] beliefs, that is, it can be altered in a way that bears some logical relation to what the person knows (Pylyshyn 1999: 343).

References

[1] Bermúdez, J. & Cahen, A. 2015. Nonconceptual mental content. The Stanford Encyclopedia of Philosophy (Fall 2015 Edition), Edward N. Zalta (ed.).

[2] Bradley, D. 2015. A Critical Introduction to Formal Epistemology.Bloomsbury Publishing.

[3] Brewer, B. 1999. Perception and Reason. Oxford University Press.

[4] Carnap, R. 1950. [1965, second edition]. Logical Foundations of Probability. University of Chicago Press.

[5] Carnap, R. 1971a. A basic system of inductive logic, Part I. In R. Carnap, & R. C. Jeffrey (eds.), Studies in inductive logic and probability (Vol. I), University of California Press.

[6] Carnap, R. 1971b. Inductive logic and rational decisions. In R. Carnap, & R. C. Jeffrey (eds.), Studies in inductive logic and probability (Vol. I). University of California Press.

[7] Churchland, P. 1992. A Neurocomputational Perspective: The Nature of Mind and the Structure of Science. MIT Press.

[8] Churchland, P. 1998. Conceptual similarity across sensory and neural diversity: the Fodor/Lepore challenge answered. Journal of Philosophy 95: 5-32.

[9] Clark, A. 2013. Whatever Next? Predictive Brains, Situated Agents, and the Future of Cognitive Science. Behavioral and Brain Sciences 36:181-204.

[10] Colombo, M. & Seriès, P. 2012. Bayes in the brain. On Bayesian modelling in neuroscience. The British Journal for Philosophy of Science 63: 697–723.

[11] Davidson, D. 1986. A Coherence Theory of Knowledge and Truth. in Truth and Interpretation, E. LePore (ed.), Blackwell: 307–319.

[12] Decock, L. & Douven, I. 2014. What is graded membership? Noûs 48: 653-682.

[13] Douven, I. et al. 2013. Vagueness: a conceptual spaces approach. Journal of Philosophical Logic 42: 137-160.

[14] Dretske, F. 1981. Knowledge and the flow of information. MIT Press.

[15] Evans, G. 1982. The Varieties of Reference. Oxford University Press.

[16] Firestone, C., & Scholl, B. 2015. Cognition does not affect perception: Evaluating the evidence for 'top-down' effects. Behavioral and Brain Sciences.

[17] Friston, K. 2009. The free-energy principle: a rough guide to the brain? Trends in Cognitive Science 13: 293-301.

[18] Gauker, C. 2007. A critique of the similarity space theory of concepts," Mind and Language: 317-345.

[19] Gauker, C. 2011. Words and Images: An Essay on the Origin of Ideas. Oxford University Press.

[20] Gärdenfors, P. 2000. Conceptual Spaces—The Geometry of Thought. MIT Press.

[21] Gärdenfors, P. 2014. Geometry of Meaning: Semantics Based on Conceptual Spaces. MIT Press.

[22] Goodman, N. 1976. Languages of Art. Hackett.

[23] Hartmann, S. & Sprenger, J. 2010. In S. Bernecker & D. Pritchard (eds.): Routledge Companion to Epistemology: 609-620.

[24] Heck, R. 2000. Nonconceptual content and the space of reasons. Philosophical Review 109: 483–523.

[25] Hohwy, J. 2013. The predictive mind. Oxford University Press.

[26] Knill, D. & Pouget, A. 2004. The Bayesian brain: the role of uncertainty in neural coding and computation. Trends in Neurosciences 27: 712–719.

[27] Kuhn, T. S. 1962. The Structure of Scientific Revolution. University of Chicago Press.

[28] Leitgeb, H., Pettigrew R. 2010a. An objective justification of Bayesianism I: The consequences of minimizing inaccuracy. Philosophy of Science 77: 201–235.

[29] Leitgeb, H., Pettigrew R. 2010b. An objective justification of Bayesianism II: The consequences of minimizing inaccuracy. Philosophy of Science 77: 236–272.

[30] Levin, D. & Banaji, M. 2006. Distortions in the perceived lightness of faces: The role of race categories. Journal of Experimental Psychology: General 135, 501–512.

[31] Lyons, J. 2011. Circularity, reliability, and the cognitive penetrability of perception. Philosophical Issues 21: 289–311.

[32] Macpherson, F. 2012. Cognitive penetration of colour experience. Rethinking the issue in light of an indirect mechanism. Philosophy and Phenomenological Research 84: 24–62.

[33] Macpherson, F. 2015. Cognitive penetration and nonconceptual content. In J. Zeimbekis & A. Raftopoulos: 331-358.

[34] Marchi, F. 2017. Attention and cognitive penetrability: The epistemic consequences of attention as a form of metacognitive regulation. Consciousness and Cognition 47: 48–62.

[35] Marr, D. 1982. Vision: A Computational Investigation into the Human Representation and Processing of Visual Information. W. H. Freedman and Company.

[36] Mervis, C. & Rosch, E. 1981. Categorization of natural objects. Annual Review of Psychology 32: 89–115.

[37] McDowell, J. 1994. Mind and World. Harvard University Press.

[38] Mole, C. 2015. Attention and cognitive penetration. In A. Raftopoulos & J. Zeimbekis (eds.) *The cognitive penetrability of perception. New philosophical perspectives*, Oxford University Press,218–238.

[39] Peacocke, C. 1992. A Study of Concepts. MIT Press.

[40] Peacocke, C. 2001. Does perception have a nonconceptual content? Journal of Philosophy 98: 239-264.

[41] Popper, K. 1935. Die Logik der Forschung. Springer Verlag.

[42] Pryor, J. 2000. The skeptic and the dogmatist. Noûs 34: 517–549.

[43] Putnam, H. 1981. Reason, Truth, and History. Cambridge University Press.

[44] Pylyshyn, Z. 1999. Is vision continuous with cognition?: The case for cognitive impenetrability of visual perception. Behavioral and Brain Science 22: 341-365.

[45] Quine, W. V. O. 1960. Word and Object. MIT Press.

[46] Raftopoulos, A. 2009. Cognition and perception: How do psychology and neural science inform philosophy? MIT Press.

[47] Rosch, E. 1975. Cognitive representations of semantic categories. Journal of Experimental Psychology: General 104: 192–233.

[48] Rosch, E. 1978. Prototype classification and logical classification: the two systems. In Scholnik, E., (ed.) New Trends in Cognitive Representation: Challenges to Piaget's Theory.

[49] alin, P. & Bullier, J. 1995. Corticocortical connections in the visual system: Structure and function. Physiological Reviews, 75(1), 107–154.

[50] Siegel, S. 2010. The contents of visual experience. Oxford University Press.

[51] Siegel, S. 2012. Cognitive penetrability and perceptual justification. Noûs 46: 201–222.

[52] Siegel, S. 2013. The epistemic impact of the etiology of experience. Philosophical Studies, 162: 697–722.

[53] Siegel, S. 2015. The Contents of Perception. Stanford Encyclopedia of Philosophy, Edward N. Zalta (ed.).

[54] Siegel, S. 2017. The rationality of perception. Oxford University Press.

[55] Skyrms, B. 1980. Higher order degrees of belief. in D. H. Mellor (ed.), Prospects for Pragmatism. Cambridge University Press: 109–137.

[56] Stalnaker, R. 1998. What might nonconceptual content be? Philosophical Issues 9: 339–352.

[57] Stokes, D. 2015. Towards a consequentialist understanding of cognitive penetration. In J. Zeimbekis & A. Raftopoulos: 75–100.

[58] Vetter, P., & Newen, A. 2014. Varieties of cognitive penetration in visual perception. Consciousness and Cognition 27: 62–75.

[59] White, R. 2006. Problems for dogmatism. Philosophical Studies 131: 525–57.

[60] Zeimbekis, J., & Raftopoulos, A. (eds). 2015. The Cognitive Penetrability of Perception: New philosophical perspectives. Oxford University Press.