

Cognitive Penetration of Colour Experience: Rethinking the Issue in Light of an Indirect Mechanism

FIONA MACPHERSON

University of Glasgow

Can the phenomenal character of perceptual experience be altered by the states of one's cognitive system, for example, one's thoughts or beliefs? If one thinks that this can happen (at least in certain ways that are identified in the paper) then one thinks that there can be cognitive penetration of perceptual experience; otherwise, one thinks that perceptual experience is cognitively impenetrable. I claim that there is one alleged case of cognitive penetration that cannot be explained away by the standard strategies one can typically use to explain away alleged cases. The case is one in which it seems subjects' beliefs about the typical colour of objects affects their colour experience. I propose a two-step mechanism of indirect cognitive penetration that explains how cognitive penetration may occur. I show that there is independent evidence that each step in this process can occur. I suspect that people who are opposed to the idea that perceptual experience is cognitively penetrable will be less opposed to the idea when they come to consider this indirect mechanism and that those who are generally sympathetic to the idea of cognitive penetrability will welcome the elucidation of this plausible mechanism.

1. Introduction

Can the phenomenal character of perceptual experience be altered by the states of one's cognitive system, for example, one's thoughts or beliefs? If one thinks that this can happen (at least in certain ways that will be identified below) then one thinks that there can be cognitive penetration of perceptual experience; otherwise, one thinks that perceptual experience is cognitively impenetrable. Prior to encountering one particular counterexample, I found the cognitive impenetrability thesis rather plausible. Although there are myriad alleged counterexamples to the cognitive impenetrability thesis, there are a number of good strategies to deal with them. Most of the alleged counterexamples involve cases in which it is claimed there are two experiences with different phenomenal characters and the reason for this difference is best

explained by a difference in the propositional attitudes of two subjects who each have one of the experiences (or by a difference in the propositional attitudes of one subject who has one experience and one set of propositional attitudes at one time and the other experience and another set of propositional attitudes at a later time). One of the strategies to explain away these cases is to claim that there is no difference in the phenomenal character of the experiences in a given case – there is merely a difference (or change) in the judgments made on the basis of experiences with the same phenomenal character. Another strategy is to claim that, while the experiences differ in phenomenal character, this change hasn't come about because a state of the cognitive system has penetrated experience, but on account of a change in perceptual processing – such as a change in perceptual attention or in eye movement.

However, there is one potential counterexample to the cognitive impenetrability thesis that is not persuasively explained away by the usual strategies. The case is one in which it seems subjects' beliefs about the typical colour of objects affects how they experience the colour of objects (Delk and Fillenbaum (1965)). After explaining why I think this case is hard to explain away, I propose a mechanism of *indirect* cognitive penetration that explains how cognitive penetration may occur. The mechanism is indirect because the penetrating belief causes some non-perceptual state with phenomenal character to come into existence or causes some process to occur that typically would yield a non-perceptual state with phenomenal character, and then that non-perceptual state or process affects, interacts with, or modifies, the phenomenal character of the perceptual state. I show that there is independent evidence that each step in this process can occur. I suspect that people who are opposed to the idea that perceptual experience is cognitively penetrable will be less opposed to the idea when they come to consider this indirect mechanism. And I hope that those who are generally sympathetic to the idea of cognitive penetrability will welcome the elucidation of this plausible mechanism.

The structure of this paper is as follows. In section two, I will spell out in detail the thesis that perceptual experience is cognitively impenetrable and give reasons why people think that it is true. In section three, I will provide examples of challenges to the thesis and elucidate the ways in which one can try to reject them. In section four, I will outline the case concerning colour that I believe is hardest for the supporter of the cognitive impenetrability thesis to deal with. In section five, I outline the indirect penetrative mechanism that I believe may explain the cognitive penetration that seems to take place in the problematic case. I also give reasons to think that this mechanism could be realised in humans. The first stage of the mechanism happens frequently when

imagining, dreaming and hallucinating. We have evidence that second stage of the mechanism occurs in humans because we have evidence that the Perky effect occurs and evidence from the nature of some dreams and hallucinations. Putting together these two stages yields my proposal for the indirect mechanism. Finally, I speculate on other perceptual phenomena that may be explained by the same mechanism.

2. The Nature of the Cognitive Impenetrability Claim

The claim that perception is not cognitively penetrable is a thesis that has received much attention in philosophy and psychology. In fact, two different forms of the claim have been made. One form of the claim has been recently advocated at length by Pylyshyn (1999). He argues that the content of the early visual system cannot be altered by the higher-level, cognitive system of the brain. He defines the early visual system functionally, as a system that takes attentionally modulated signals from the eyes (and perhaps some information from other sensory modalities) as inputs, and produces shape, size and colour representations – representations of visual properties – as output. These representations are then categorised and identified by the cognitive system making use of memory, knowledge and judgment. Pylyshyn holds that the content of early vision is cognitively impenetrable, that is, cannot be altered in virtue of the *content* of the cognitive system (as opposed to other features of the cognitive system):

if 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 (1999: 343)

This condition is required to ensure that cases of the following kind do not count as cases of cognitive penetration: I believe that today is the day of an important exam. This belief causes stress and brings on a migraine. The migraine causes disturbances to my vision and I now experience flashing lights at the periphery of my visual field in addition to experiencing the scene in front of me. In this case, my visual experience has been altered on account of my belief but there is no intelligible connection between the *content* of the belief – my exam is today – and the *content* of my visual experience – the apparent flashing lights – even though we understand why the belief caused this experience (namely, stress can cause migraines and these can cause visual disturbances).

Pylyshyn also explicitly denies that the output of the early visual system should be thought of as being or determining one's visual

experience or, put differently, that the content of the output of the early visual system should be thought of as the content of visual experience. (The content of a system or of an experience is specified by saying what is represented by that system or experience.) Pylyshyn's thesis is, therefore, not about perceptual experience but about the subpersonal representational outputs of brain processing mechanisms.¹

However, in philosophy and in psychology, the claim that perception is cognitively impenetrable is *usually* taken to be one about perceptual *experience*. The claim is that the phenomenal character of perceptual experience, that is, what it is like to undergo a perceptual experience, cannot be altered in certain ways in virtue of the content of states of the cognitive system – states such as thoughts, beliefs, desires, and other propositional attitudes – that one may possess.^{2,3} Moreover, on the innocuous assumption that I will make for the purposes of this

¹ Nonetheless, paradoxically given his insistence that his claim is not about perceptual experience, in that article Pylyshyn offers evidence for his claim that relies on facts concerning the nature of the phenomenal character perceptual experience. For example, when he cites the fact that visual illusions persist, despite one's knowledge of the illusion, he is referring to the constant, unchanging nature of perceptual experience in the face of gaining relevant new beliefs at odds with the experience (1999: 344).

² The phrase "what it is like" was introduced into philosophy as a way to capture phenomenal character by Nagel (1974).

³ When thinking about what states of one's cognitive system are relevant to cognitive penetration one is naturally drawn towards thinking about the different beliefs and desires that one has. However, there might be other states of one's cognitive system that cognitively penetrate. One example comes from noting that the state of one's cognitive system is determined, in part, by which concepts one possesses. If the possession of a concept affected one's perceptual experience then that state of concept possession would be one that cognitively penetrated one's experience. (Of course some people think that which concepts one has is a matter closely tied to the issue of which beliefs and desires one has – indeed some even go so far as to claim that one determines the other.) This type of cognitive penetration would be a form of "category perception". We know that people with different concepts judge objects in the world differently. One person may judge that an object falls under a category that another person does not on account of their possession of different concepts. For example, it has been claimed that which colour concepts one possesses can affect one's judgements (and the speed of one's judgements) about colour (Winawer et al. 2007) and possession of concepts pertaining to types of animal based on the patterns of marking on their hide affects one's judgements (and the speed of one's judgments) about types of animal (Goldstein and Davidoff 2008). If the differing judgments are reflective of differing perceptual experiences then these would be examples of cognitive penetration. Another example of a state of the cognitive system that may affect one's experience is the cognitive system being primed so that certain concepts are likely to be triggered or activated. (In such a case one is in a cognitive state such that one is likely to form beliefs or desires involving those concepts, perhaps as opposed to other beliefs or desires.) If being in such a state affected one's perceptual experience then that state would be one that cognitively penetrated perceptual experience.

paper, namely, that there is no change in the content of a perceptual experience without a change in the phenomenal character of that experience and vice versa, then the claim will be that the phenomenal character of perceptual experience and the content of perceptual experience cannot be altered in certain ways in virtue of the content of states of the cognitive system.

This rough statement of the cognitive impenetrability claim needs to be made more precise. We all know that which beliefs or desires one has can alter what one chooses to perceive because they can motivate the turning of one's head, eyes and body – and the position of these will, of course, affect what perceptual experience one has. Similarly, if one is perceiving a stationary unchanging scene and the conditions for perceiving remain unchanged and one doesn't move one's body or eyes at all, then perhaps one can have different experiences just in virtue of attending to different parts of the scene. And, of course, the location one attends to in a scene can be influenced by one's beliefs and desires. This type of spatial attention is often thought of as a subtle form choosing what to perceive that, instead of involving any movements of the body, just involves a shift in spatial attention.⁴ Although both of these cases involve states of the cognitive system affecting which perceptual experiences one has, they are not typically taken to be cases of cognitive penetration. A case of cognitive penetration is a case in which that which is perceived, the viewing conditions, and the state of the sensory organ are held fixed. One way to capture this is to imagine the proximal stimulus that affects the sensory organ (light in the case of the eye, sound waves in the case of the ear, etc.) and the condition of the sensory organ remaining exactly fixed. In addition, the location of one's attentional focus is held fixed. If it is possible for two subjects in these conditions to have different perceptual experiences (different in respect of phenomenal character and content) on account of the differing states of their cognitive systems, or if it is possible for one subject at different times in these conditions to have different experiences on account of the difference between the states of their cognitive systems at those times, then cognitive penetration is possible. In addition, as we have already seen by looking at the migraine case, it needs to be added that there have to be some links between the content of the cognitive state and the content of the perceptual state that is affected of a nature such that the effect on the content of the perceptual experience is made intelligible, or in some very minimal sense rational, in light of the content of the cognitive state.

⁴ See Pylyshyn (1999) and Raftopoulos (2001).

Thus, perceptual experience is cognitively impenetrable if it is not possible for two subjects (or one subject at different times) to have two different experiences on account of a difference in their cognitive systems which makes this difference intelligible when certain facts about the case are held fixed, namely, the nature of the proximal stimulus on the sensory organ, the state of the sensory organ, and the location of attentional focus of the subject. It is this version of the cognitive thesis that is under investigation in this paper.

The cognitive impenetrability thesis about experience is often held *in conjunction* with a thesis like Pylyshyn's about early visual processing. The thought is that the phenomenology of perceptual experience is determined by the output of some specifiable early visual system. Thus, frequently, claims about early vision and claims about perceptual experience go hand-in hand.


For example, one way the cognitive impenetrability thesis is stated is by claiming that the visual modalities – vision, audition, touch, and the like – are modular. The idea is that the processing that goes on in the brain in, say, vision is specialised and isolated – that is, is *informationally encapsulated* – from the other perceptual modalities and from the more general cognitive system. Information from a sensory organ enters the brain and is processed discretely in the perceptual mechanism associated with that organ and the output of this mechanism serves as input to the cognitive system. There is no top-down processing, that is no feedback, from the cognitive system into the perceptual modules (although there may be top-down processing *within* the perceptual system).⁵ (It is further claimed that the perceptual modules process sensory input on the basis of certain assumptions or information about the world built into the modules. These assumptions are fixed and themselves cognitively impenetrable.) Unlike Pylyshyn's view, the thesis is not taken to be one merely about subpersonal brain processing and architecture. Rather, it is assumed that the output of the sensory modules determines, or is identical with, the content and phenomenal character of perceptual experience. This view is promoted at length by Fodor (1983). However, it is only the claim that perceptual experience is cognitively penetrated that is the one at issue in this paper.

One major motivation for the view that perception is cognitively impenetrable is the persistence of visual illusions.⁶ In such illusions,

⁵ This is a brief sketch of what is meant when it is said that a system is modular. See Fodor (1983) for much greater detail. Much debate has taken place in the literature about how one should spell out modularity. This debate is beyond the scope of this paper – but harmlessly so. The details of the nature of modularity are tangential to the arguments here.

⁶ See Fodor (1983).

one's visual experience presents the world as being one way, and continues to do so, even when one knows that the world is not that way. For example, in the Müller-Lyer illusion one's visual experience is such that the horizontal lines seem to be of different lengths, and this experience persists even when one knows that the lines are the same length.



Müller-Lyer Figure

Of course the existence of one example where some specific belief about the world does not impact upon perceptual experience does not prove the thesis that perceptual experiences are *never* cognitively penetrated. At best, it establishes that the content of perceptual experience is not responsive to the content of belief in the same way as, or to the same extent as, our other beliefs are. For example, at least on pain of irrationality but perhaps *tout court*, our beliefs must have consistent content – certainly this seems to be true in the case of conscious beliefs with explicit contradictory content. But in the case of persisting illusions, in which there is no plausible reason to believe that one is irrational, the content of one's experience and the content of one's conscious belief is not consistent. So the dispute about whether perceptual experience is cognitively impenetrable or not tends to proceed on the basis of cases. Those who are against cognitive impenetrability present examples in which it looks as if subjects' propositional attitudes affect the nature of the phenomenal character and content of their experience. Those who defend the thesis try to explain away these examples.

There are several disputes in current philosophy that are closely related to the question of whether perceptual experience is cognitively penetrable. One such dispute is whether observation is "theory-neutral". Let me explain.

The claim that observation is theory-neutral, as advocated by Fodor (1984), is the claim that there are a core group of beliefs – those formed directly on the basis of perceptual experience – and those beliefs are not

affected by one's other beliefs. Arguing for the opposing side, Churchland (1979 and 1988) claims that there are no perceptual beliefs unaffected by one's other, theoretical, beliefs. All beliefs are theory-laden.

A natural position for one such as Churchland to adopt would be that experience is thoroughly cognitively penetrated, for this would explain why perceptual beliefs are also affected by other beliefs. And indeed Churchland does adopt this position. But it would be open to Churchland not to do so. He could hold that perceptual experiences are cognitively impenetrable but all beliefs that we form on the basis of such experiences are dependent on our other beliefs and theoretical commitments. Indeed, Churchland seems amenable to thinking that there is a core element of perceptual experience, which he calls *perceptual sensation*, that is cognitively impenetrable, however, he thinks perceptual experience consists in more than this sensational nub. In short, if one thinks that there are no theory-neutral beliefs it is open to one to hold either that perceptual experience is cognitively penetrable or cognitively impenetrable.

However, a proponent, such as Fodor, of the theory-neutrality of observation beliefs *must* be committed to the cognitive impenetrability thesis. For if one's existing beliefs can affect which perceptual experiences one has then they will indirectly affect which perceptual beliefs one has, because different experiences will cause different perceptual beliefs. Thus, the cognitive impenetrability of perceptual experience is a necessary but not sufficient condition for the truth of the theory-neutrality of observation but, as we saw in the previous paragraph, it is neither a necessary nor a sufficient condition for the truth of the theory-ladenness of observation.

Another dispute in current philosophy that is closely related to the question of whether perceptual experience is cognitively penetrable is the debate concerning the admissible contents of experience. This dispute concerns which properties are represented in perceptual experience. At present, the debate in the literature focuses on the content of human visual experience. One view is that our visual experiences have only "low-level" content, such as content about shape, colour, position in space, and perhaps object-hood. On this view, your experience might represent an object of a certain size, shape and colour in front of you. Of course, if the occasion arose that the object in front of you was your brother, and you recognised that, it would be natural to think that that object also looked to you like a human being and your brother. But on this view, that thought is to be cashed out by saying that on the basis of your experience that only represented low-level properties, such as size, shape and colour, you *judged* that it was a human being and your brother. Forming this judgment might be

unconscious, spontaneous, unbidden, and take place incredibly quickly, but it is a judgment nonetheless. The representation of a human being and your brother is not part of the content of the visual experience; rather, it is the content of some cognitive state formed on the basis of the experience together with one's memory and background knowledge. Thus, if you had an experience that represented an object of a certain size shape and colour, and at first you didn't realise that it was your brother but then you came to realise that it was your brother, this view would say that your visual experience did not change in content or in phenomenal character. The change in your mental life is a change in non-perceptual cognitive states, such as beliefs or judgments.⁷

The opposing view is that visual experiences have "high-level" content, such as content about natural kinds (for example about human beings or tigers or pine trees), content about causation, about particular individuals (such as your brother, as opposed to something with the look of your brother) and other content that outstrips, roughly, properties pertaining to shape, size, position, colour and object-hood. On this view, if you had an experience that represented an object with certain low-level properties, and at first you didn't realise that what you was seeing was your brother but then you came to realise that it was your brother, then an acceptable explanation might be that your visual experience did change in content and in phenomenal character. Your experience came to represent that a human being that was your brother was present. Thus, on this view, the change in your mental life can be a change in the character or content of your perceptual experience.

People who subscribe to the existence of high-level content in visual experience are likely to reject cognitive impenetrability. This is because many of the arguments for high-level content proceed by arguing that *learning* can affect which visual experience one has. The idea is that the cognitively unadulterated visual system will have limited representational powers – no doubt powers to represent only that which is likely to be encountered by most humans in different locations and at different times – powers that are likely to be useful to all. These representational powers will extend to properties such as shape, colour and size. But, depending on the specifics of the environment and the skills that it will be useful for particular humans to have, the visual system can come to represent new objects and properties via repeated exposure to them and via top-down influences from the cognitive system.

⁷ Some people think that states of the cognitive system, such as beliefs and desires can have phenomenal character; others deny it. Those that deny it that wish to defend the idea that experiences only have low-level content would have to deny that there was any change in phenomenal character in such cases.

For example, plausibly, the visual system does not come ready-made to represent pine trees or the specific individual that is my brother. This is a type of tree and a particular human that not all humans will encounter. If all such specific representational abilities had to be built-in to the visual system it would be enormous and unwieldy. But, perhaps on repeated exposure to pine trees or to my brother, a subject can come to notice features that all and only such trees have and that all and only that person has. And perhaps the subject's knowledge of these features can feed into their visual system so that they come to have visual experiences that are sensitive to those features, at least sensitive in a way that they were not before. In this way, the subject would come to have visual experiences that they did not have before – visual experiences that represent pine trees or my brother. In other words, a very plausible mechanism for visual experiences coming to have high-level content is that the visual system is penetrated by the cognitive system.

Of course someone might be a high-level theorist and hold that the visual system comes ready-made to make pronouncements about all sorts of contents, including those involved in high-level content. Or they may think that although the visual system can be altered during one's lifetime, it is not altered by cognition. (Perhaps they think it is altered by some non-cognitive process, such as an associative process that takes place solely within the visual system.) Nonetheless, high-level content theorists are likely (though not compelled) to want it to be the case that cognitive penetration of perceptual experience can occur.

Those who believe that visual experience represents only low-level contents have one less reason than those who believe visual experiences can represent high-level content to think that experience is cognitively penetrable. For if high-level contents cannot feature in visual experience then the postulation of penetration is not required to explain alleged high-level contents in visual experience – instead, such content is to be denied. However, such a theorist needn't be committed to cognitive impenetrability. Perhaps they think that cognition can affect the representation of various low-level contents, such as shape. For example, there is some evidence that whether one is subject to the Müller-Lyer illusion depends on the nature of the environment that one has grown-up in. For example, it seems that if one has not grown-up in a “carpentered” environment (one that contains many right angles) then one is less likely to be subject to this illusion.⁸ This

⁸ See McCauley and Henrich (2006). The evidence is not to my mind conclusive but further discussion of this case takes me too far from my aims in this paper. Nonetheless, I will return to consider whether the mechanism for cognitive penetration that I postulate could account for this case *if* it is a case of cognitive penetration in section 5.

evidence has been taken to suggest that people who live in carpentered environments have implicitly learned certain facts about typical relationships between distances, lengths, and the appearance of angles where three lines or surfaces appear to converge, and thus beliefs with contents consisting in these facts have become part of their cognitive system. These facts, it is postulated, feedback down and influence the lengths that our experiences represent certain lines as having – sometimes inappropriately, as when the Müller-Lyer illusion is experienced.

Nonetheless, candidate cases of low-level cognitive penetration have, in the past, been thought difficult to find – far more difficult than high-level ones. Thus, there has often been little pressure for someone who believes only in low-level content to think that cognitive penetration occurs. But, as will become apparent, the case that I think is a persuasive case of cognitive penetration does involve low-level content – content about colour. Thus, I think low-level theorists who do not believe in cognitive penetration need to re-consider their position.

3. Responses to Alleged Cases of Cognitive Penetration


In this section, I outline some cases that have been thought to challenge the idea that perception is cognitively impenetrable. I do so in order to show typical ways that a defender of cognitive impenetrability can defend their thesis against such attack. I do not intend to cover all cases of alleged penetration. In omitting some cases I do not mean to suggest that they are not powerful challenges or ones that do not require addressing by a defender of cognitive impenetrability. My goal here is only to elucidate the common ways in which many sorts of potential counterexamples can be defused. These ways of defusing alleged counterexamples are ones that one might think one can use to respond to the case about colour that I will consider in more detail in the next section. But, as we will see in that section, these ways of replying are not good ways of replying to the colour case – which, together with the fact that colour is a low-level property, is what makes it such a powerful example.

A common form of alleged counterexample to the cognitive impenetrability thesis consists of two experiences that are had by a subject at different times that intuitively have different phenomenal characters. Then a case is made that the best explanation for this difference is that the cognitive system of the subject is different at those different times. For example, in the previous section, I discussed the view that there can be high-level content in perceptual experience. Some arguments for the existence of high-level content in experience can be elucidated in a way that makes them arguments that perceptual experience is cognitively penetrated.

Consider, for illustrative purposes, the following argument based on Siegel's (2005) argument concerning pine trees. Imagine a scenario in which someone gains a recognitional capacity – the capacity to tell that a tree is a pine tree by looking at it. Consider the experience of the subject before they have the recognitional capacity when they look at a pine tree E_1 , and call the experience they have after gaining the capacity when looking at the same tree from the same position, in the same lighting conditions, when their eyes are working similarly, and when the same parts of the scene are being attended to, E_2 . It is plausible that the phenomenology of the subject will be different overall when having E_1 and E_2 . After all, when having E_1 the subject will not know what kind of tree they are looking at. When having E_2 they will. One explanation of this difference is that the visual experiences E_1 and E_2 are different, and different only because the second experience represents that a pine tree is present, while the first does not. If this were the case then there would be reason to think that gaining the recognitional capacity to identify pine trees, which plausibly involves one's cognitive system gaining beliefs about how pine trees look and a recognitional capacity to identify pine trees, affects the experience one has. E_2 represents a pine tree, while E_1 does not and this difference in content is manifested in the different phenomenal characters of the experiences. Hence, cognitive penetration has taken place.

Consider another kind of case, which also conforms to the common form of potential counterexample. Certain figures are ambiguous. When one looks at these figures there are two distinct ways of seeing them. It is plausible to think that one has two different experiences corresponding to these two different ways of seeing. Examples include the Necker cube (in which the line AB in the diagram can be seen either as being on the front face of the cube or on the back face), the duck/rabbit picture and the dolphin/naked couple picture.


One might think that the two different experiences had when seeing the figure in the two ways possible come about because of influence from the states of one's cognitive system. For example, with practise, one can learn to have one of the two different experiences almost at will – one can voluntarily change how the figure looks. It could be argued that here one's thoughts about which experience one wants to have – or one's thoughts about what one wants to experience – changes one's experience. Another way in which one might think that one's experience is influenced by one's cognitive system concerns which experience one has when looking at the figures, or which experience one is most likely to have or has first. For example, it is sometimes claimed that adults typically have the experience of the naked




A

B

Necker Cube



Duck/Rabbit Picture



Dolphins/Naked Couple

couple when they first look at the dolphin/naked couple ambiguous figure, whereas children typically have an experience as of the dolphins. If this is true, one might think that it is explained well by the fact that children have very different thoughts, concerns, expectations and proclivities compared to adults.

There are two main ways to respond to these types of alleged counterexamples to the thesis that perceptual experience is cognitively impenetrable. The first way is to deny that the two perceptual experiences in question have different phenomenal characters. One might claim that instead there is only a difference in the judgments one makes about the world or the judgments one makes about one's perceptual experiences. (Then it is open to one to either claim or deny that these judgmental differences bring with them some form of phenomenal differences). For example, in the pine tree case one might claim that the difference before and after one gains the recognitional capacity is not that E_1 and E_2 are different but that one's judgments are different. One judges that there is a pine tree in one case and one doesn't judge that in the other case. Of course one might make additional judgments

in the case one has the recognitional capacity for identifying pine trees. For example, one might make subtle judgments about the outline shape of the tree and its leaves or their colour.

In the ambiguous figures cases, one might hold that there are not two experiences associated with seeing each of the figures. Rather, there are two different judgments that people have when looking at the figures – for example, that there is a picture of a duck present or that there is a picture of a rabbit present. On account of making these different judgments people misjudge that their experience is changing. Perhaps certain features of the experience become more salient to the people on one occasion compared to another and so the people judge that the picture is of different things. Perhaps their doing this helps to explain why people think that their experience itself is changing when in fact it remains unchanged.

The second way to respond is to agree that there is a difference between the phenomenal characters of the perceptual experiences but to hold that this is not best explained by cognitive penetration. For example, in the case of the pine trees, one might hold that in gaining a recognitional capacity for identifying pine trees one learns where to focus one's attention – perhaps on the outline shape of the tree or leaves, etc. – in order to identify the tree. As a result of this change in focus of spatial attention, the phenomenal character of one's experience is altered. As discussed in the previous section, however, changes of spatial attention, even if these are brought about by a change in one's cognitive system, are typically not counted as cases of cognitive penetration. Similarly, in the case of ambiguous figures, one might think that one's judgment about what the figure is a picture of affects one's eye movements or where one focuses attention on the figure. Again, this would result in a difference of experience – but it is a case where one's judgment may affect one's experience that is not a case of cognitive penetration. An alternative explanation in the case of ambiguous figures is that the early visual system can autonomously process some inputs in two distinct ways yielding two distinct experiences, but that these different ways of processing are unaffected by the cognitive system.

Which method of explaining away a potential counterexample is best will depend on the details of the case. I don't wish to defend particular responses to the examples above save to note, without argument, the fact that I find the first strategy more plausible in the pine tree case and the second strategy more plausible in the ambiguous figure case. However, there is one case that I believe is not well explained by these types of response. It is to this case that I turn in the next section.

4. The Colour Counterexample

In 1965, Delk and Fillenbaum published a study that is a candidate for being a case where subjects' beliefs about the characteristic colour of an object can affect the colour that they experience that object as having.⁹ Their experiment is, somewhat surprisingly, not discussed as often as it should be in the literature on cognitive penetration, given how persuasive an example of it I believe it to be. In this section, I relate the experiment and explain why the standard strategies outlined in the previous section, designed to explain away alleged counterexamples, do not seem to explain away this case.

The experiment was carried out as follows. The experimenters took a sheet of paper of a uniform orange colour. They cut out shapes of various objects from it. Some of these objects were characteristically red: a heart (a love-heart shape), a pair of lips, an apple. Some were not characteristically red: an oval, a circle, a square, an ellipse, a horse's head, a bell, a mushroom. One at a time, the cutout shapes were placed in front of a coloured background that could be altered. The background could be changed from a yellow colour through orange to a red colour by twisting a knob. The subjects were told to instruct the experimenter to make the background colour more yellow or more red until it was the same colour as the cutout shape in front of it, so that the cutout shape could no longer be distinguished from the background.¹⁰ When the cutout shape of an object that had a characteristically red colour was placed in front of the background, the subjects selected a background colour that was more red than the colour that they selected the background to be when the cutout shape was of an object that didn't have a characteristically red colour. When the cutout shape was of an object that didn't have a characteristically red colour the subjects' instructions yielded a background that was less red and more yellow. The experimenters claimed that:

One is led to the conclusion that past association of color and form does in some way influence perceived color, since that is the one

⁹ This study was based on similar but less conclusive studies carried out along the same lines. See Duncker (1930) and Bruner et al. (1951). Variants of this experiment have been carried out more recently achieving similar results. See Hansen et al. (2006).

¹⁰ In fact the subject had to look at the cutout against the coloured background through a sheet of wax paper that was almost, but not quite, transparent. The purpose of the wax paper was to slightly blur the boundary area of the cutout shape against the background. I believe this detail is unimportant for my discussion of the case. One can imagine the experiment being repeated in modern conditions on a computer screen where such paper would not be required.

respect in which the figures did clearly differ. (Delk and Fillenbaum 1965: 293)

Although Delk and Fillenbaum don't themselves suggest a mechanism that explains why this happens, one might think that what is happening is that the subjects' beliefs, that certain of the cutout shapes were shapes of objects that were characteristically red, penetrated their perceptual experience of those cutout shapes thereby altering the content and phenomenal character of those experiences. The subjects had experiences of those cutout shapes that represented them as being more red than they really were. This did not happen when the cutout shape was not of a characteristically red object. In that case, the subjects had accurate experiences of the orange colour of the cutout shape.

Should we accept that this is a case of cognitive penetration of colour experience? The case is rather persuasive for neither of the two strategies discussed in the section above that one might use to undermine the idea that cognitive penetration is occurring is a very attractive alternative explanation of what is going on in this case.

The first strategy one could employ is to claim that there is not a difference in the subjects' experiences of the characteristically red and not characteristically red objects; there is simply a change in their judgments. Thus, one could claim that the experiences of the subjects when they look at the cutout shape of a characteristically red object are not as of a more red colour than the cutout shape really is and are not different to their experiences of the colour of the cutout shape of objects that are not characteristically red. Instead, one might claim that the subjects' experiences represent all the cutout shapes accurately as the shade of orange that they are. The subjects merely judge that a cutout shape is more red than it really is when it is a cutout of a characteristically red shape.

However, this explanation requires the postulation of a misjudgment on the part of a subject. The subject misjudges that the colour of the orange cutout and the colour of the red background it is on are the same colour when, according to this strategy, the subject's experience doesn't represent them as being the same. This type of misjudgment is not one of the sort that one commonly gets in illusions – where one judges the world to be one way when it is not because one's experience misleads one by being inaccurate. Rather, in this case the misjudgment in question involves judging that the world is a certain way (the shape and its background are the same colour) when not only is it not that way but also one's experience tells one that it is not that way. One's experience tells one – represents – that

the colours are different, according to this strategy. Thus, subjects are not only misjudging the way the world is, they are doing so in the face of their experience telling them accurately how the world really is. Thus, one might say, they are misjudging the nature of their experience too.

Note that to employ this type of strategy in the pine tree case, discussed in the section above, one didn't have to attribute *any* misjudgment or error to the subject – simply an accurate judgment that a pine tree was present after the recognitional capacity was gained, and a lack of judgment beforehand. In addition, the type of misjudgment that this response to the colour case requires is not an explicable one – as was the case when this strategy was employed to explain the ambiguous figures example, also discussed in the section above. (In the ambiguous figures case, the subject's judgment that his or her experiences are different was explained by the fact that he or she judges that the picture is of different things on different occasions.) Rather, the type of misjudgment that the experimental set-up requires one to attribute to the subject is brute and inexplicable and of a very fundamental kind. To see this, consider the types of error a subject might make on the basis of having an accurate colour experience.

A subject might have an accurate experience of an orange colour and judge that the colour is orange, but due to mere incorrect word selection say that the colour is “red.” Such cases occur and do not involve an error of judgment only an error of expressing one's judgment. This is not the type of error that must be postulated to explain away the alleged case of cognitive penetration under consideration.

Second, a subject might have an accurate experience of an orange colour but miscategorise that shade as being a shade of red. Nonetheless, the subject might alter a background colour to be a shade that accurately matches the miscategorised shade of colour. This type of misjudgment seems perfectly possible. It would appear to be one in which, in one sense, the subject is aware of exactly the nature of their colour experience, which they demonstrate via the colour-matching task. Their error is simply one of which colour category to assign the shade of colour to – red and not orange. This is also not the type of misjudgment that must be posited by the person who wants to explain away the apparent cognitive penetration going on in the colour case using this strategy.

Rather, the type of misjudgment that must be going on is one in which a subject has an accurate experience of an orange colour (of the cutout) and an accurate experience of a shade of red (of the background), nonetheless, when looking at both the shades simultaneously,

they judge that the shade of red and the shade of orange are the same shade.¹¹ If the subject were misperceiving one of the colours, the misjudgment would be explicable, but this explanation is not available as the case is one in which, by hypothesis, the subject's experiences are completely accurate. What we are to imagine is that the subject's experiences of the colours are different, yet the subject judges the colours to be the same. While such cases do not seem impossible, they involve attributing a gross form of misjudgment to the subject. This is not a mere miscategorisation of which colour a shade is. The subject is making a large and brute error of judgment *despite* their experience. We have no explanation of why the subject makes such an error – or why they make such errors in the systematic way the experiment shows they do. One might think that an explanation lies in the fact that the subjects believe that the cutout is of a shape that is characteristically red and that this belief affects their judgment but, firstly, that doesn't explain why they should make this gross error when their experience is telling them otherwise – for frequently our judgments *are* corrected by our experiences. Secondly, this alleged explanation is the very hypothesis under consideration, namely that the subject misjudges despite their experience, so it is hardly a suitable candidate for explaining *why* the subject so misjudges that when, at the same time, their experience is telling them otherwise. That the subject does make this gross error when they believe the cutout's shape is of a characteristically red object is still a fact that stands in need of explanation. Because of this, and because there is no independent evidence that the subject's experiences of the cutout and the background it is on are different, the postulation of this error and the postulation of the difference in experience with respect to the shape and the background it is on looks to be ad hoc – a postulation only to save the theory that perceptual experiences are cognitively impenetrable.

Compare the above explanation to the alternative – that cognitive penetration is taking place. If cognitive penetration is taking place the subject merely misperceives one of the colours (the cutout of the shape of a characteristically red object) and so has experiences as of the same colour when looking at that cutout and the background. This case involves no postulation of error in the subject's judgment. Rather it

¹¹ The only alternative would be to suppose that the subject misperceives the background. But if that were the case then that experience would appear to be being cognitively penetrated by the belief about the characteristic colour of the heart-shaped cutout. Not only would that be odd, it would be another alleged case of cognitive penetration that would have to be explained away. Explaining it away would presumably be done by one or other of the strategies that I am outlining here, but these strategies would face exactly the same problems as the ones I am outlining here for them.

merely involves an inaccuracy that the subject can't be held responsible for – an inaccuracy in their experience. On this account, the subject accurately reports the way the world looks to them – even if the world is not that way.

Thus, explaining away the alleged cognitive penetration of experience by means of this strategy involves, in this case, the postulation of a gross, brute and inexplicable error on the part of the subject – an error on the personal level that we think subjects are responsible for. No doubt many people think errors of this kind are possible, but this view predicts that these errors occur in a systematic way – for example they will always occur in the conditions that the experiment specifies. And that thought is unpalatable, for what then ensures any of our judgments reflect our experience? The alternative explanation involves no such postulation of brute, inexplicable error on the part of the subject. For this reason, the explanation that tries to explain away cognitive penetrability is ad hoc and far less attractive than the alternative.

Furthermore, note that colour is a low-level property – it is a property that all people agree is represented by visual experience – as opposed to a high-level property, like a natural kind property. Therefore a common strategy that is employed by low-level theorists to maintain that two experiences are the same, and that they represent the same properties, cannot be employed here. The strategy is to claim that any evidence that the experiences are different is really only evidence that the contents of judgments formed on the basis of those experiences are different, on the grounds that experiences cannot represent the properties in question, as they are high-level properties. But this strategy can't be applied to this case for the properties at issue – colour properties – are low-level properties.

I have just been considering the first broad strategy that might be employed to explain away the case of colour perception that seems to show cognitive penetration takes place – postulating that the experiences that the subject has of the two different kinds of cutouts are the same and that it is their judgments about the colour that differ. The second broad strategy that one might use to undercut the idea that the case under discussion is a case of cognitive penetration is to accept that the experience of the cutout shapes of characteristically red objects is different from the experience of the cutout shapes of the objects that are not characteristically red, but to claim that this fact is not explained by cognitive penetration. There are two factors that one might appeal to in order to explain what other than cognitive penetration accounts for the difference between the experiences:

- (i) a difference in the cognitive states of the subject, but in a way consistent with the cognitive impenetrability of perceptual experience (by a shift in the location of one's attention) or

or

- (ii) some factor other than the differing cognitive states of the subject.

Let us look at each in turn.

As we saw in section three above, one might claim that the subject's beliefs alter the focus of their attention and this is responsible for the change in their experience. At the same time, we noted that such alterations in experience by cognition by means of attention are not counted as cases of cognitive penetration. However, I will argue that the kind of difference in attention that will have to be postulated to successfully explain the change in experience in the colour case is rather special and that it is not one that is clearly compatible with the cognitive impenetrability of perceptual experience.

What kind of change in attention that results in a change of experience is compatible with the cognitive impenetrability of perceptual experience? It is changes in spatial attention that people have in mind. Indeed, I know of no other discussion of changes in attention in this context except for changes in spatial attention. One type of case that is discussed in this regard is the case of ambiguous figures. Consider the duck/rabbit ambiguous figure. It is often held that the difference in experience that one has in seeing the figure as a duck and then as a rabbit is in *where* one focuses one's attention. For example, someone might hold that if one focuses attention on the left-hand side of the figure then one will see the figure as a duck, while if one focuses attention on the right-hand-side of the figure then one will see the figure as a rabbit. (No doubt that explanation is in fact too simplistic and patterns of spatial attention are required to fully explain the case, but it will do no harm to ignore such complications here.)

Would a change in *where* attention is focused be a good explanation of why experiences of uniform patches of colour were different on different occasions? No – for the colour is uniform all over the cutout, so a change in attention would not plausibly yield a change in the experience of colour. One might think, therefore, that this strategy is doomed to failure. However, this would be too quick. What we should conclude is that if some form of attention shift is going to explain this case it must be some form of attention other than spatial attention. Reflection on the case, I believe, will show that

the best explanation involving attention that one can give here is that when one sees the cutout shapes of objects that are characteristically red, one pays attention to the *redness* within the orange colour of the paper at the expense of the yellowness within the orange colour of the paper. When one sees the cutouts of the objects that are not characteristically red one doesn't pay special attention to the redness within the orange colour. (Orange is a binary colour, which is to say that it is seen as a *combination* or mixture of yellow and red.¹² Thus orange is experienced as having a red component and a yellow component. It is focusing attention on the red component, at the expense of the yellow component, that I have in mind here.) This change in attention towards the red component in the orange colour of the paper might alter the nature of one's experience and this would explain well why one chose a more red colour as a match for some of the cutout shapes compared to others.

However, while it is frequently held that if a cognitive state changes the *location* that one attends to and thereby changes one's experience then it needn't count as a case of cognitive penetration, it is not obvious that other changes in attention caused by a cognitive state should not count as cases of cognitive penetration. To see this, note that the reasons that explain why people think that changes of spatial attention directed by one's cognitive states should not count as cases of cognitive penetration are two-fold. The first is that such cases are akin to selecting what spatial location to attend to by means of changing the position of one's body or eyes. But the kind of attention that has to be postulated to explain the colour case is selection of one feature or property of an object (its redness) that is instantiated at the very same location as another feature or property of an object (its yellowness). Thus, it is not the sort of selection of one feature over another that could be brought about by a change in the position of one's body or eyes. The second reason people have for allowing a change in the focus of spatial attention caused by a cognitive state to be consistent with the cognitive impenetrability of experience is that they think that allocation of attention to a spatial location occurs before perceptual processing occurs. The thought is that the nature of the stimulus, the nature of the perceptual conditions, the state of the sensory organ and where attention is focused are all conditions that determine what will get perceptually processed – but the perceptual processing itself, that proceeds after these conditions are in place, is impenetrable. However, it is far from clear that attention to the redness within a patch of orange is an attentional effect that is a condition that fixes what will be

¹² See Hardin (1988: 39).

perceptually processed in advance of any such processing taking place. On the contrary, it is reasonable to think that attention to the redness within the orange only occurs once perceptual processing has got underway, when orange has been registered by the brain or, if one is thinking of attending as something that the subject does (as opposed to something done by a subpersonal mechanism) when the subject becomes aware of the orange colour. After all, how can there be attention to a feature of the orange colour if one's brain has not detected the orange colour or if one is not aware of – experiencing – the orange colour? In this respect, this form of attention is rather different from spatial attention, about which there is evidence that one's attention can be focused on a location prior to one's having a visual experience. For example, if one had one's eyes closed (and thus suppose one had no visual experience) a noise might lead one to attend to the apparent location of the noise and, on subsequently opening one's eyes, one's experience might in part be the way it is because one's attention is focused on that location. The thought is that spatial cues in other modalities, or spatial cues in the same modality prior to having the particular perceptual experience in question, or voluntary thought of a location alone, can focus attention on a location and thereby alter subsequent visual processing and, hence, subsequent experience.¹³

In short, if one adopts the strategy of claiming that cognitive penetration is not taking place in the colour case because, although one's experiences of the colours of the shapes are different, this can be explained by one's cognitive states affecting attention, then certain further explanatory tasks become incumbent upon one. First, one needs to explain what kind of change in attention could change the nature of one's experience. Second, one has to ensure that the kind of change of attention cited is one that is compatible with the cognitive impenetrability of experience. I have given some reasons in the above paragraph to think that the most plausible form of change of attention that would explain one's experience in the colour case is not a form of spatial attention. It is attention to one of the colour components within a binary colour – attention to the redness that exists within orange. And I have argued that this is not a form of attention that can be easily construed as not being a form of cognitive penetration.

Indeed, consideration of this case brings to the fore the more general difficulty of establishing in any case – even in the case of spatial attention – whether the contribution of attention to perception really is a

¹³ Evidence for this type of effect has been found, for example, in: Spence and Driver (2004), Santangelo et al. (2006), Santangelo et al. (2009) and Störmer et al. (2009).

pre-perceptual phenomenon. As this is a difficulty that plagues all forms of attention then so much the worse for any strategy of explaining away alleged cases of cognitive penetration that relies on the claim that changes in attentional focus caused by changes in cognitive state can occur without cognitive penetration.¹⁴

I turn now to consider (ii), that is, whether some factor other than the differing cognitive states of the subject could be responsible for the difference between the experiences. Delk and Fillenbaum explicitly say that they tried to ensure that the only relevant difference between the cases where the subjects were presented with cutout shapes of characteristically red objects and non-characteristically red objects was the subject's beliefs. For example, they tried to ensure that the size of all the cutouts was similar and that each group of characteristically red and non-characteristically red cutouts contained members that had both curved and straight lines. I think that it is reasonable to believe that they ensured that no factor external to the subject was responsible for the different behaviour of the subject towards the characteristically red and non-characteristically red cutouts, but did they ensure that there was no relevant difference within the subject apart from the subject's beliefs?

One thing undetermined is whether the subjects' visual systems reacted differently when they saw the cutouts of the characteristically red objects compared to the non-characteristically red objects but not because of influence from the cognitive system. It might be that the early visual system, autonomously from belief or other cognitive states, alters the colour experiences of characteristically red shapes, to make them appear more red than they really are. This might happen in accordance with associationist principles, so that it is past exposure to a certain shape having a certain colour that has altered the way the visual system processes that shape's colour.

One motivation for thinking that this might be the case comes from considering what would happen if you suddenly lost your belief that one of the characteristically red shapes, say the heart, was red (for example, a God appeared to you and told you that this was the case). My intuition says that it is plausible that you would still undergo the colour effect. You would still experience orange hearts as more red than orange non-characteristically red shapes.¹⁵ If this is right, then perhaps it is not the cognitive system that is responsible for the difference in the colour experience that subjects have.

¹⁴ The question of whether attention is a pre-perceptual process or not is discussed in Raftopoulos (2005) and Rowlands (2005).

¹⁵ Thanks to Tim Bayne for making this motivation for the view clear to me.

There are two ways in which one can reply to this suggestion for what other than cognitive penetration explains Delk and Fillenbaum's results. The first is to agree that perhaps beliefs with the content that hearts are characteristically red cannot penetrate your experience. However, one might think that other beliefs can penetrate it. Perhaps the beliefs that can penetrate are ones that it would be very difficult for a God to overturn (just by telling you, without tampering directly with your brain). Examples of such beliefs are the belief that most of the hearts I have experienced have seemed red and the belief that for some time I believed that most hearts were red. A very similar response would be to claim that it is not any specific belief that is responsible for the effect on the visual system, rather, because of your (at least apparent) exposure to red hearts, the concept "red" gets primed when you have an experience as of a heart. Your cognitive system, in this primed state, then affects the visual system. This does not require any particular beliefs to be present in order for cognitive penetration to occur. In fact, one way to spell out what it is for the cognitive system to be so primed is that one is disposed to bring to mind some beliefs or other about redness, but which ones is irrelevant.

The second reply that one can make to the suggestion that the visual system affects your experience of characteristically red shapes without the cognitive system penetrating it is to admit that perhaps in some cases this is the mechanism at work, but to claim that there are an important class of cases where this is not a plausible explanation. This reply is convincing because there are cases where the objects that elicit the effects potentially explicable by cognitive penetration are ones that have to be classified by the cognitive system in one way rather than another for the effect to occur. They are cases where it is implausible to think that low-level vision can do this categorisation alone and belief, desire or other cognitive, conceptual states play a crucial role in the occurrence of the effects.

Examples of such cases include the shapes of characteristically red objects used in the colour case, other than hearts, such as apples. While it may be plausible to think that the visual system is sensitive to heart shapes and autonomously responds to that shape making it seem more red than it is, it is far less plausible to think that the visual system can classify certain objects as being apples and autonomously alter the apparent colour of such objects. The reason is that classifying something as an apple doesn't depend just on simple features, such as shape, that the visual system is responsive to, for many things have the shape of apples, such as cherries and other fruits. And while the visual system can respond to shape, it surely can't classify an object as an apple, as it

is not sensitive to the features that are required to do so. It is surely only the cognitive system that is capable of doing such classificatory work.

An even better example of objects that elicit the effects potentially explicable by cognitive penetration, and are ones that have to be classified by the cognitive system in one way rather than another for the effect to occur, are certain faces. A series of experiments carried out by Levin & Banaji (2006) showed that greyscale pictures of faces that have stereotypical features of white people and similar pictures of faces that have stereotypical features of black people are matched by subjects for lightness to different samples of grey, even though the pictures of the faces have the same surface luminance. Subjects were instructed to change a patch of grey that could be altered from dark to light to match a picture of a face placed beside it. The faces that looked like faces of white people were matched to lighter shades of grey than those of black people. Remarkably, the effect appears to increase when one knows about the effect and stares hard at the faces, as readers will be able to verify for themselves by looking at the examples of faces used in the experiment re-published below. In a variant of the experiment, a racially ambiguous face was labelled either as the face of a white person or the face of a black person and this factor alone determined what shade of grey the subjects chose as a match for the lightness of the face. This second result is particularly significant, as the categorisation was clearly done at the cognitive level as it was the labelling of the face that was responsible for the effect. Levin and Banaji state:

White 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. (2006: 501)

These face cases provide examples involving colour matching, as did Delk and Fillenbaum's experiment, which we have already seen lend strong support for the thought that two different experiences of colour (in this case lightness) are being had by subjects. At the same time they provide an excellent example where the categorisation of the object, which is necessary for the different experiences of lightness to be had, takes place at the cognitive level.

Three faces used and published in Leven and Banaji (2006). The stereotypically black face on the left has the same luminance as the stereotypically white face in the middle. The face on the right is an example of a racially ambiguous face.



Thus, experiments of the type carried out by Delk and Fillenbaum, and Levin and Banaji, can provide very good reasons for thinking that colour experience is cognitively penetrable. The typical ways in which one might explain away these cases face clear difficulties. So, I believe that the simplest and best explanation of these cases is that cognitive penetration is occurring. A subject's experiences of the colour of different objects is not the same, despite the colour of the objects actually being the same. This is because their experiences are penetrated by either the subject's belief that some of the objects are of a type that have a characteristic colour, or by some state of the subject's cognitive system in which it is primed to make certain beliefs about that characteristic colour conscious.

5. A Mechanism for Penetration

If one is persuaded that the colour case outlined in the previous section is a case of cognitive penetration, then the question arises of how a state of the cognitive system can affect perceptual experience. It certainly doesn't seem that we can affect our perceptual experiences at will. If only we could not experience the sound of the fingernails scraping down the blackboard, or taste the horrible mushrooms hidden in the stew, or see the litter spoiling the otherwise beautiful view, just by wishing! Similarly, how we expect the world to be is often not how we perceive it to be, either to our delight or disappointment. Resistance to the thought that cognitive penetration can occur is, I think, in part a concern about the nature of the mechanism of penetration. In this section, therefore, I propose a plausible psychological mechanism to explain cognitive penetration. It is a two-step process and I will bring forward evidence, independent of the phenomenon of cognitive penetration, that shows that each step can occur. In virtue of the mechanism having these two steps, I call it *indirect*. I think that the mechanism may make the idea that cognitive penetration occurs more palatable to those who are (or were prior to exposure to the colour

case of section 4) apt to think that cognitive penetration does not occur. And for those people who think cognitive penetration is widespread, this mechanism may explain a large number (although I don't claim all) of the cases in which they think cognitive penetration occurs.

The first step of the mechanism involves our cognitive states causing some non-perceptual state with phenomenal character to come into existence or to alter the phenomenal character of some existing non-perceptual state that has phenomenal character.¹⁶ (By "non-perceptual state" I mean here a state not involved in perceiving the world. On this disambiguation of the term, hallucinatory experiences are non-perceptual states. Of course the term "non-perceptual state" often is used otherwise, namely, to refer to states other than the experiences involved in both seeing and hallucinating. But it is convenient for my purposes here to use the term in the more inclusive sense.) Such cases are commonplace. Here are some examples:

- You are imagining nothing. I then ask you to imagine looking at George Square in Glasgow from Queen Street station. You comply. Your cognitive state has caused a non-perceptual state with phenomenal character to come into existence. You deliberately imagined what you did.
- You are daydreaming about sailing down the river Clyde. You are imagining the sights and sounds and the feeling of smirr on your face. Suddenly you think of me and find yourself imaging sailing down the Clyde with me. Your thought about me has changed your daydream, and altered its phenomenal character. This whole imaginative experience might not have been deliberate. The imaginative experiences might have come unbidden.
- You dream of me because you met me today for the first time.
- You visually hallucinate me and part of the explanation is that you know me well.

It is widely agreed that imagining, dreaming and hallucinating always or frequently involve being in a state with phenomenal character. Philosophers are wont to imagine that some such states have the very same phenomenal character as some perceptual experiences. But often

¹⁶ As we will see later, it might be best to think that what happens in the first step is that our cognitive states cause the initiation of an imaginative process that would normally result in an imaginative experience with phenomenal character.

these states are very unlike the typically perceptual experiences that we have. For example, the phenomenal character of imaginative experience is said to be “weaker” or “less lively” than that of perceptual experience. And imagination typically does not have the force of presenting the world as being the way imagination represents it to be. In addition, imagination can be much less determinate than typical perceptual experience. Hallucinations are sometimes recognised as such for their bizarre content and, as J. L Austin pointed out, dreaming that one meets the Pope is rarely, if ever, like really meeting the Pope.¹⁷

Imagination, dreams and hallucination clearly frequently involve some phenomenal character, and there is no doubting that the content and phenomenal character of these states is often generated, affected by, and dependent on our beliefs and desires. In this sense, these states are cognitively penetrable. The content of such states is overwhelmingly content dependent on one’s cognitive states, although there may be some content of such states that is not so dependent. One oft cited example of the latter is when the content of a dream is affected by one’s alarm clock going off and the sound of the alarm may feature as an element of the dream.¹⁸ Another case might be one in which an evil scientist tampers with your brain and the content of your resulting hallucination is attributable to that tampering and not to your cognitive system. But such examples don’t detract from the thought that often the content of such states is generated by or dependent on one’s beliefs, thoughts, desires, expectations and previous experiences with the world. Indeed, I believe that the content of the majority of such states are, as a contingent matter of fact, affected by the states of one’s cognitive system.

The second step involves the phenomenal character of these non-perceptual states interacting with and affecting the phenomenal character and content of perceptual experiences. For example, in the colour case discussed in the section above, perhaps subjects imagine the cutout shape being red and the phenomenal character of this imaginative state interacts with the phenomenal character of their visual experience. In this case, it doesn’t seem plausible to suggest subjects are aware of two states or two phenomenal characters. So the best way to put the idea is that the contribution of the imagination and the contribution of vision combine producing one phenomenal state. If the imagination would have produced a state with the phenomenal character of an experience

¹⁷ See Austin (1962).

¹⁸ One might claim that the sound of the alarm isn’t part of the content of the dream and is part of the content of one’s auditory experience. But even if this example doesn’t provide a good example, there are others and it doesn’t detract from the main thrust of my argument which is simply that much of the content of dreams, hallucinations and imaginative states is determined by one’s cognitive system.

had when imagining a red object if no contribution from the visual state had been present, and if the visual system would have produced a state with the phenomenal character of a visual experience had when veridically seeing an orange object if no contribution from the imaginative state had been present, then the resulting phenomenal character of the combination of vision and imagination would plausibly be the phenomenal character of an experience had when veridically seeing a reddish-orange object. This would explain why the subjects matched a more red colour to cutouts of an orange shape with a characteristically red colour compared to those that weren't characteristically red.

The next question to address is why should one think that perception can interact with other mental processes – processes such as imagination, dreaming and hallucinating that are affected by states of the cognitive system – to yield a state whose phenomenal character must be explained by the contribution of both perception and those other process? After all, one shouldn't think that the other process provides, say, some phenomenal redness that mixes with, say, the phenomenal orangeness provided by vision in the way that paints can be mixed. This is true, but shows only that this suggestion should not be the model of such interaction. Nonetheless, we have reason to believe that this kind of interaction between such states occurs. There are examples, completely separate and independent of potential cases of cognitive penetration, where we believe that the phenomenal character of a state is brought about by contributions from both perception and either imagination, dreaming or hallucinating.

The first example is illustrated in the Perky effect. Perky conducted an experiment in which subjects were asked to fixate at a point on a blank white screen and to visually imagine various objects on it, such as a tomato, a banana, a book, and a leaf. While the subject did this, a faint image of the particular object was projected onto the screen. The image was above the typical threshold for visibility. The subjects didn't realise that this was going to happen. Nor did they report that they saw anything at the time or when explicitly asked afterwards whether they had seen anything or imagined everything. However, when they reported what they were imagining, it was clear that what was projected onto the screen affected their experience. For example, one subject reported their surprise when they tried to imagine a banana in a horizontal position and claimed that they had ended up imagining it in a vertical position. Another subject was surprised as she said that she had intended to imagine an elm leaf but ended up imagining a maple leaf.

Perky's experiments are typically reported as being ones in which perceptual experiences were mistaken for imaginative mental states and

hence that the phenomenal character of some visual experiences are the same as that of some imaginative experiences or at least that any differences are a difference of degree and not of kind.¹⁹ However, others have thought that this conclusion is too strong. Segal, who followed up Perky's work with subsequent similar experiments, held that what was happening was that imagining raised subjects' perceptual threshold so that they did not have a perceptual experience of the image, yet unconscious perception of it influenced what they did imagine.²⁰

What does seem clear is that the phenomenal character of the resulting state – be it an imaginative one, or a perceptual one mistaken for an imaginative one, or some combination of perceptual or imaginative state – has aspects contributed by perceptual and imaginative processes. For example, as we have seen, the image projected on the screen affected the reported phenomenology of the subsequent experience. But imagination did too. One subject reported visually imagining red vein-like markings on the leaf, when none were present on the image projected on the screen. Another subject claimed that the book they were imagining had text on it, when the projected image of the book had no writing on it. Colours from the projected image can affect the resulting experience. Segal (1972, as reported in Thomas (2008)) performed an experiment like Perky's in which she asked subjects to imagine the skyline of New York. An image of a tomato was projected onto the screen. Observers didn't report the tomato image but several reported imagining seeing New York at sunset. Thus, we have examples where perceptual elements and imaginative elements combine to produce the phenomenal character of what seems to the subject to be one state – be it really a perceptual experience, really an imaginative experience, or some amalgam of the two.

The second case has already been discussed briefly – that of what is often described as the incorporation of external stimuli into dreams. As with the Perky case, I am not interested in whether it is correct to say that a subject who has, say, the sound of their alarm clock incorporated into their dream, is dreaming the sound or hearing the sound, or some amalgam of the two. But, like the Perky case, it does not seem as if the subject is aware of two separate states – a dreaming one and a perceptual one. And the resulting phenomenology, of what is often reported to be the one dream state, is made up of elements, such as the noise of the alarm clock, that have been perceived and dream elements which are, as discussed above, typically highly dependent on cognitive states. For example, one can imagine that the noise of one's alarm

¹⁹ See, for example, Thomas (2008).

²⁰ See Segal (1971) and Segal and Fusella (1971) cited in Thomas (2008).

clock is experienced as the sound of a bin-lorry reversing in one's dream. Or one could imagine that a loud crash sound heard whilst asleep might be experienced as a slightly different sound – that of thunder – in one's dream.

The third case is the case of hallucinations affecting visual experience. Many hallucinations are such that the objects that the subjects hallucinate appear to be in the world, at least of part of which the subject really is perceiving. This is true in cases in which the subjects realise that they are hallucinating and in cases in which they don't. For example, in Lilliputian hallucinations, very small people are hallucinated and seem to be seen, for example, on the carpet or peeking out from behind the curtains or on the hallucinator's food.²¹ Some hallucinations in Parkinson's disease involve the apparent perception of objects such as cats moving around the room between chair legs and even sitting on the subject's lap.²² Real perceived objects can occlude and be occluded by hallucinated objects. If these hallucinations are in part caused by states of the subject's cognitive system then, together with perception of the environment, they seem to yield a mixed perceptual and hallucinatory state, the phenomenology of which is a mixture of both perceptual and hallucinatory elements.²³

These three examples, the Perky effect, incorporation of perceived elements into dreaming, and mixed perceptual and hallucinatory states, suggest that the phenomenology of perceptual states can be altered by non-perceptual states or processes or vice versa. This shows that it is not implausible to think that processes affected by cognition can affect the phenomenology and content of one's perceptual experience. This shows that the second step of the proposed indirect method for cognitive penetration could plausibly occur.

So far in this section, I have proposed a two-step, indirect, mechanism for how cognitive penetration may occur. First, one's cognitive states generate an imaginative phenomenal state or process that typically would produce such a phenomenal state. Second, the imaginative

²¹ See Chand and Murthy (2007).

²² See Manford and Anderman (1998).

²³ Indeed, it is not obvious that these cases that involve a mixture of hallucination and perception are not themselves cases of perceptual experiences being cognitively penetrated. To my knowledge these cases are never discussed in the cognitive penetration literature. One might resist the thought that these cases are instances of cognitive penetration by holding that the non-hallucinated elements of the experience are not cognitively penetrated. But it is not obvious to me that we should not just think of these cases as cases of cognitive penetration heretofore not classified or recognised as such.

state or process interacts with the perceptual state or process yielding one state that is an amalgam of the two states and has a phenomenology that has elements contributed from both perception and imagination. Thus, for example, in the colour case discussed above, the first step would be that a subject's knowledge that the cutout was of a shape that is of a characteristically red object would affect the subject causing them to imagine a red object, or generate the process that would typically produce an imaginative experience of a red object. The second step involves the phenomenal character of this imaginative state, or the imaginative process that typically would produce it, interacting with the phenomenal character of the visual experience of the orange cutout shape, or the perceptual process which would typically yield such a visual experience. The result is an experience as of a reddish-orange colour. In this case, it doesn't seem plausible to suggest subjects are at any stage aware of two states or two phenomenal characters. So, in this case and perhaps almost all, the relevant imaginative and perceptual processes simply produce one state with phenomenal character whose nature has contributions from both the imaginative and perceptual processes. If the imaginative process would have typically produced an imaginative experience of red if unaffected by the perceptual state, and if the perceptual system would have typically produced a visual experience as of orange if unaffected by the imaginative state, then it is unsurprising that the phenomenal character of the state produced by both processes would be the phenomenal character of an experience had when accurately seeing a reddish-orange object. Having this experience would explain why the subjects matched a more red colour to cutouts of shapes with a characteristically red colour compared to cutouts of shapes that weren't characteristically red.


I would like to stress at this point that imagining may be a process that is deliberate and initiated by an agent or it may not be. One may choose to imagine and what to imagine, or one may simply find oneself imagining, and one may have no choice about whether or what one imagines. I think that the sort of imaginative process that will typically take place in my two-step mechanism is the non-deliberate, unbidden kind. Indeed, typically, one will not even realise that one is undergoing an imaginative process or that one's imagination is influencing one's experience. (The precedents for the existence of such cases, where subjects are unaware of the combination of processes contributing to their experience is, as mentioned above, the Perky case, perceptual insertion into dreams, and cases of mixed perceptual and hallucinatory states, where people don't realise that they are hallucinating.)

This two-step mechanism may explain other cases of cognitive penetration if, indeed, there are others. (But perhaps it will not explain all cases of cognitive penetration that exist – I do not argue that it will in this paper.) For example, if we have two different experiences when looking at pictures of ambiguous figures and this is due to cognitive penetration then my two-step indirect mechanism may explain what happens. Consider the Necker cube pictured above. Perhaps one imagines the line AB either at the front or at the back of the cube, or at least the process by which this imagining would take place is set in motion. Then the imaginative experience or process interacts with one’s perceptual process to yield one or other of the experiences of the cube.

Another case that the two-stage model may explain is the cognitive penetration that some people think takes place in, and explains, the Müller-Lyer illusion. Recall from section 2 that there is evidence that if one has not grown-up in a “carpentered” environment (one that contains many right angles) then one is less likely to be subject to this illusion. Suppose, for the sake of argument, that cognitive penetration does take place when one undergoes the Müller-Lyer illusion – that is, suppose one’s beliefs and expectations based on living in a carpentered world, and which one would not have if one lived in a different sort of environment, cause one to experience the lines as having different lengths. What kind of explanation might the two-stage model give of this case?

It may be that people who suffer the illusion are not able to help themselves imagining lines of the upper horizontal sort in the Müller-Lyer figure above as being further away than the bottom horizontal sort of line because of their experiences with corners of rooms or buildings. And their imagination may combine with their perception of the two lines to yield an experience in which the top line looks longer (because it is imagined to be further away) than the bottom line.

Another example concerns music. It is sometimes claimed that a person who is knowledgeable about music – knows about chords, keys, and music theory – will have a different experience of music than someone who lacks this knowledge. Suppose that cognitive penetration occurs in such a case. (I don’t argue that it does.) It would be plausible to think that it occurs because the person imagines the structure of the notes and music corresponding to the musical theory they know and that their imagining affects their perception of the music. For example, musical theories are sometimes illustrated by three-dimensional models. One example is the pitch spiral, which is a model of Western music’s pitch structure. Ascending pitches are represented as positions on an ascending spiral and parts of the spiral directly



above or below one another always correspond to notes that are one octave apart.

If one is familiar with the model then perhaps when one hears a melody played in one octave and then in a higher octave one imagines the same melodic structure being repeated higher-up the spiral and perhaps this imaginative process affects one's auditory experience.

The final example I'll discuss comes from cricket. One might think that a cricket umpire who is a specialist in making fine judgments about the behaviour of the cricket ball has a different experience of a ball hitting a batsman's batting pad than that of a non-expert. The non-expert may see just the ball hit the pad and then have to judge whether the ball would have gone on to hit the wicket – something they need to determine in order to judge whether the batsman was leg before wicket. The expert umpire may see the ball hit the pad but also have an experience as of the ball travelling through the space occupied by the batsman's leg and hitting or missing the wicket – an experience that would be somewhat like looking at the "Hawkeye" animations that are produced by computer on a screen for the third umpire who is off the pitch and for the viewers on television. *If* umpires have such experiences (and it not clear that they do, nor have I argued for it here) then it would be reasonable to think that it was their imagination, fuelled by their experiences of the behaviour of balls in the context of the game and their experience of the Hawkeye pictures, in combination

with the contribution from their perception, that produced this experience.²⁴

Alleged cases of cognitive penetration that plausibly wouldn't be explained by my two-stage, indirect mechanism involve expert perception where the only difference between the novice and the expert is that the expert is accurately picking up on some objective properties that the novice isn't. In such cases, degrees of sensitivity to the way the world is, rather than imagination, may provide a better explanation. My mechanism will be better at explaining cases in which there is misperception, illusion, mere subjective differences in perception, or the addition of content to experience beyond that which is strictly speaking seen (such as in the cricket umpire case).

Of course, it may not be that all imaginative states can combine with perceptual ones to change the phenomenal character of the resulting state. I do not suppose that occurs, in part, because I can't change my perceptual experiences at will by imagining. Moreover, I don't suppose that one can come to see any two lines of equal length as being different in length by imagining that they unequal or by imagining that one is further away from another. Very specific conditions for interactions between perceptual and imaginary processes may exist. For example, the difference between voluntary and involuntary imagination may have some, as yet unknown, role to play, as may one's familiarity with what one is imagining, as may some relations between imagined and perceived properties. These are matters for further philosophical reflection and no doubt in part for psychologists to investigate.

Psychologists may, indeed, be able to test my theory of the two-stage mechanism of indirect cognitive penetration. One way of testing it would be to see if people in whom perception is preserved whilst visual imagery is destroyed are as susceptible to the colour illusion and the Müller-Lyer illusion as other subjects in whom imaginative capacities are preserved. There is a vigorous debate in the current literature as to whether visual perceptual and visual imaginative processes take place in the same part of the brain. Those that argue that they do hold that one cannot have visual perception without visual imagination and vice versa. However, evidence of a double dissociation between visual perception and visual imagination has emerged.²⁵ If this evidence stands

²⁴ Thanks to John Sutton for suggesting this example to me. (And thanks to Jonathan Bird for setting me straight on the cricketing terminology!) In fact, this would be a case of someone having an experience like the experiences of those with "heavenly vision". Heavenly vision was posited to be the vision that angels have. They were said to be able to see solid objects as solid and yet still see other solid objects directly behind them, without the use of mirrors, etc.

²⁵ See Moro et al. (2008), Bartolomeo (2008), and Bartolomeo et al. (1998).

up to scrutiny, there would be a way of empirically disconfirming my hypothesis: find a subject that cannot visually imagine but is still subject to experiences that we believe are cognitively penetrated and that are so by the mechanism I propose. If the evidence does not stand up to scrutiny, however, then the evidence that visual perception and visual imagination are, at least partly, processed in the same part of the brain would help to explain why interaction between the two processes occurs.

Another way to test my theory wouldn't depend on finding people who can't do or undergo *any* visual imagining, rather, it would try to determine whether people who score highly on tests that determine that they are good at visual imagining are more susceptible to the apparent effects of cognitive penetration, and people who score lowly are less susceptible. For example, one could test whether people who are good at visual imagery adjust a background in front of which have been placed orange cut-outs of characteristically red objects to be a more red colour than those who are bad at visual imagery. Similarly, one could test whether there is a correlation between susceptibility to the Müller-Lyer illusion (among people who live in carpentered environments) and visual imaginative abilities.

6. Conclusion

I have argued that there is one particularly difficult case to account for by those who don't believe that cognitive penetration of perceptual experience can take place. The case involving colour perception cannot be easily explained away by standard strategies for dealing with such cases.

I proposed a two-stage, indirect mechanism that would explain how cognitive penetration could occur. The mechanism is indirect because the penetrating belief causes some non-perceptual phenomenal character to come into existence, or causes some process that would typically produce that phenomenal character to initiate. Then that non-perceptual phenomenal character, or that process affects, or interacts, with or modifies, the perceptual phenomenal character that one has. I provided independent evidence that each step in this process can occur.

I suspect that people who are opposed to the idea that perceptual experience is cognitively penetrable will be less opposed to the idea when they come to consider this mechanism because it is indirect and proceeds by one state with phenomenal character interacting with another or processes that each produce phenomenal character interacting with one another. I also suspect that those who are generally

sympathetic to the idea of cognitive penetrability will welcome the elucidation of this plausible mechanism. The existence of this mechanism is open to empirical testing, as are the conditions under which it operates, as I have outlined above.²⁶

References

- Austin, J. (1962) *Sense and Sensibilia*, edited by G. J. Warnock, Oxford: Oxford University Press.
- Bartolomeo, P. (2008) "The Neural Correlates of Visual Mental Imagery: An Ongoing Debate," *Cortex*, 44: 107–108.
- Bartolomeo, P., Bachoud-Lévi, A. C., de Gelder, B., Denes, G., Dalla Barba, G., Brugières, P., et al. (1998) "Multiple-Domain Dissociation Between Impaired Visual Perception and Preserved Mental Imagery in a Patient With Bilateral Extrastriate Lesions," *Neuropsychologia*, 36(3): 239–49.
- Bruner, J. S., Postman, L. and Rodrigues, J. (1951) "Expectation and the Perception of Color," *The American Journal of Psychology*, 64(2): 216–227.
- Chand, P. K. and Murthy, P. (2007) "Understanding a Strange Phenomenon: Lilliputian Hallucinations," *German Journal of Psychiatry*, 10: 21–24.
- Churchland, P. M. (1979) *Scientific Realism and The Plasticity of Mind*, Cambridge: Cambridge University Press.
- (1988) "Perceptual Plasticity and Theoretical Neutrality: A Reply to Jerry Fodor," *Philosophy of Science*, 55(2): 167–187.
- Delk, J. L. and Fillenbaum, S. (1965) "Differences in Perceived Colour as a Function of Characteristic Color," *The American Journal of Psychology*, 78(2): 290–93.
- Duncker, K. (1939) "The Influence of Past Experience upon Perceptual Properties," *The American Journal of Psychology*, 52(2): 255–265.
- Fodor, J. A. (1983) *The Modularity of Mind*, Cambridge MA: MIT Press.
- (1984) "Observation Reconsidered," *Philosophy of Science*, 51(1): 23–43.
- Goldstein, J. and Davidoff, J. (2008) "Categorical Perception of Animal Patterns," *British Journal of Psychology*, 99: 229–43.

²⁶ Thanks to David Chalmers, Parker Crutchfield and Stuart Crutchfield and audiences at the University of Toronto (particularly Stephen Biggs, Mohan Matthen and Dustin Stokes), the University of Oxford (particularly Tim Bayne), the University of Cornell (particularly Susanna Siegel and Nico Sillins) and at CenSes: Centre for the Study of the Senses at the Institute of Philosophy, University of London (particularly Ophelia Deroy) for discussing these issues with me.

- Hansen, T., Olkkonen, M., Walter, S. and Gegenfurtner, K. R. (2006) "Memory Modulates Color Appearance," *Nature Neuroscience*, 9(11): 1367–8.
- Hardin, C. L. (1988) *Color for Philosophers: Unweaving the Rainbow*, Indianapolis, Indiana: Hackett Publishing.
- Levin, D. T. and Banaji, M. R. (2006) "Distortions in the Perceived Lightness of Faces: The Role of Race Categories," *Journal of Experimental Psychology: General*, 135(4): 501–12.
- Manford, M. and Anderman, F. (1998) "Complex Visual Hallucinations: Clinical and Neurobiological Insights," *Brain*, 121: 1819–40.
- McCauley, R. N. and Henrich, J. (2006) "Susceptibility to the Müller-Lyer Illusion, Theory Neutral Observation, and the Diachronic Cognitive Penetrability of the Visual Input System," *Philosophical Psychology*, 19: 79–101.
- Moro, V., Berlucchi, G., Lerch, J., Tomaiuolo, F. and Aglioti, S. M. (2008) "Selective Deficit of Mental Visual Imagery With Intact Primary Visual Cortex and Visual Perception," *Cortex*, 44: 109–118.
- Nagel (1974) "What is it Like to Be a Bat?," *Philosophical Review*, 83: 435–50.
- Perky, C.W. (1910) "An Experimental Study of Imagination," *American Journal of Psychology*, 21: 422–452.
- Pylyshyn, Z. W. (1999) "Is Vision Continuous with Cognition? The Case for Cognitive Impenetrability of Visual Perception," *Behavioral and Brain Sciences*, 22: 341–423.
- Raftopoulos, A. (2001) "Is Perception Informationally Encapsulated? The Issue of the Theory-Ladenness of Perception," *Cognitive Science*, 25: 423–51.
- (2005) "Cognitive Penetrability of Perception: A New Perspective," in A. Raftopoulos (ed.) *Cognitive Penetrability of Perception: Attention, Action, Strategies, and Bottom-Up Constraints*, Hauppauge, NY: Nova Science.
- Rowlands, M. (2005) "The Cognitive Penetrability of Perception," in A. Raftopoulos (ed.) *Cognitive Penetrability of Perception: Attention, Action, Strategies, and Bottom-Up Constraints*, Hauppauge, NY: Nova Science.
- Santangelo, V., Van der Lubbe, R. H., Belardinelli, M. O. and Postma, A. (2006) "Spatial Attention Triggered by Unimodal, Crossmodal, and Bimodal Exogenous Cues: A Comparison of Reflexive Orienting Mechanisms," *Experimental Brain Research*, 173(1): 40–8.
- Santangelo, V., Olivetti Belardinelli, M., Spence, C. and Macaluso, E. (2009) "Multisensory Interactions Between Voluntary and

- Stimulus-Driven Spatial Attention Mechanisms Across Sensory Modalities,” *Journal of Cognitive Neuroscience*, 21: 2384–97.
- Segal, S. J. (1971) “Processing of the Stimulus in Imagery and Perception”, in S. J. Segal (ed.) *Imagery: Current Cognitive Approaches*, New York: Academic Press.
- (1972) “Assimilation of a Stimulus in the Construction of an Image: The Perky Effect Revisited,” in P. W. Sheehan (ed.), *The Function and Nature of Imagery*, New York and London: Academic Press.
- Segal, S. J. and Fusella, V. (1971) “Effects of Images in Six Sense Modalities on Detection (d’) of Visual Signal from Noise,” *Psychonomic Science*, 24: 55–56.
- Sekuler, R., Sekuler, A. B. and Lau, R. (1997) “Sound Alters Visual Motion Perception,” *Nature*, 385: 308.
- Siegel, S. (2005) “Which Properties Are Represented in Perception?,” in T. Szabo Gendler and J. Hawthorne (eds.) *Perceptual Experience*, Oxford: Oxford University Press.
- Spence, C. and Driver, J. (2004) (eds.) *Crossmodal Space and Crossmodal Attention*, Oxford: Oxford University Press.
- Störmer, V. S., McDonald, J. J. and Hillyard, S. A. (2009) “Crossmodal Cueing of Attention Alters Appearance and Early Cortical Processing of Visual Stimuli,” *Proceedings of the National Academy of Sciences of the United States of America (PNAS)*, 106(52): 22456–61.
- Thomas, N. J.T. (2008) “Mental Imagery”, in E. N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy (Winter 2008 Edition)*, URL = <<http://plato.stanford.edu/archives/win2008/entries/mental-imagery/>> .
- Winawer, J., Witthoft, N., Frank, M. C., Wu, L., Wade, A. R. and Boroditsky, L. (2007) “Russian Blues Reveal the Effects of Language on Color Discrimination,” *Proceedings of the National Academy of Sciences of the United States of America (PNAS)*, 104(19): 7780–5.