

The Cognitive Impenetrability of Recalcitrant Emotions

Raamy Majeed

University of Cambridge

Draft: 2017

ABSTRACT: Recalcitrant emotions are emotions that conflict with your judgements, e.g. fearing flying despite judging that it is safe. Much of the present controversy concerning these emotions has to do with spelling out the precise nature of this conflict, and determining what this, in turn, tells us about a theory of emotions. This debate, however, leaves unexamined a crucial feature of these emotions, *viz.* their recalcitrant, mulish nature. This paper aims to make up for this neglect. In particular, I argue the recalcitrant nature of recalcitrant emotions can be accounted for by their cognitive impenetrability. I also argue one advantage of doing so is it helps link the literature on emotional recalcitrance with empirical findings on emotion generation and regulation; especially in a way that suggests the recalcitrant nature of such emotions, at least in some instances, is a product of how they are generated.

KEYWORDS: recalcitrant emotions; cognitive penetration; information encapsulation; affect-programs; emotion generation; emotion regulation

1. Introduction

Consider the following emotional episodes. You fear Fido, your neighbour's dog you judge to be harmless. You are angry with your colleague, even though you know his remark wasn't really offensive. You are jealous of your partner's friend, despite be-

believing she doesn't fancy him. D'Arms and Jacobson (2003) call these *recalcitrant emotions*: emotions that exist “despite the agent’s making a judgment that is in tension with it” (pg. 129). This phenomenon has been of interest to philosophers of emotion on account of the challenge they pose for a theory of emotion. Drawing on the work of Greenspan (1981) and Helm (2001), Brady argues this is “to explain the sense in which recalcitrant emotions involve *rational* conflict or tension” (2009: 413).

Whether we require *rational* conflict to account for emotional recalcitrance is debatable.¹ Indeed, much of the present controversy involves spelling out the precise nature of this conflict. But conflict, rational or otherwise, isn't the only feature that is pertinent to the phenomenon. What is unaccounted for is precisely what gives these emotions their name, *viz.* their recalcitrance. That is, it is in the very nature of recalcitrant emotions that they are mulish, that they remain resistant, and don't succumb to our judgements.

Here is an example. Suppose I judge that flying is safe, but feel instantly afraid as soon as my plane starts to take off. But suppose, also, once I realize that my fear is irrational, or at least, that it is in tension with my judgement, my fear dissipates. This won't count as an instance of emotional recalcitrance. By contrast, say I remain fearful despite my judgement. I keep thinking to myself, ‘I know this is safe’, and yet I continue to feel afraid. This, I venture, better captures what we mean by emotional recalcitrance. *Mutatis mutandis* for being afraid of Fido, being jealous of your partner's friend etc. All familiar cases of emotional recalcitrance seem to share this feature. It is certainly present in examples used to introduce the notion of emo-

¹ E.g. Döring (2014) defends a perceptual theory emotions, which seeks to explain away any irrationality concerning the conflict.

tional recalcitrance. The question is, What accounts for it? This question has, surprisingly, not been taken up.

The aim of this paper is to address this question. It does so by borrowing a notion from cognitive science, i.e. ‘cognitive penetration’. In broad terms, a function a system performs is cognitively penetrable if it can be affected by purely cognitive factors, such as beliefs, desires, inferences, tacit knowledge, and so on.² Applying this to emotions, Goldie suggests, “someone’s emotion or emotional experience is cognitively penetrable only if it can be affected by his relevant beliefs” (2000: 76). For present purposes, assume an emotion is cognitively impenetrable if it fails to be penetrable in this sense. The hypothesis I want to explore in this paper is the thesis that emotions are recalcitrant to the extent that they are cognitively impenetrable.

This hypothesis will strike some as a reformulation of definitions of emotional recalcitrant already on offer, e.g. by D’Arms and Jacobson. This paper seeks not to (re)define ‘recalcitrant emotions’, nor provide a conceptual analysis of the phenomenon to be picked out by these terms. Rather, the aim is to provide an *a posteriori* hypothesis that seeks to *explain* a specific feature of the phenomenon of emotional recalcitrance, *viz.* their recalcitrant nature. Moreover, it does so in such a way that this feature can be operationalised. In other words, our hypothesis is formulated as such that the claim it makes about the recalcitrant nature of recalcitrant emotions is subject to empirical scrutiny.

The merit I am claiming for my hypothesis, besides making the phenomenon easy to understand, is that it is empirically fruitful in the following two respects. First,

² See Nicholas *et. al.* (1996). ‘Cognitive penetration’ gets its name from a debate in philosophy of perception. See Stokes (2013) for an introduction.

we know how to operationalise cognitive impenetrability, but we don't yet know how to operationalise the recalcitrant nature of recalcitrant emotions. By demonstrating that the one leads to the other, we then see how we can have an operationalisation for both. In this way, one way the present hypothesis makes the phenomenon clear is that it is now amenable to empirical research. Second, and relatedly, understanding the phenomenon in a way that can be operationalised gives us another purchase on the empirical literature. That is, we stand to learn the nature of the causal mechanisms responsible for emotional recalcitrance.

In what follows, I provide an exposition of the impenetrability hypothesis (§2), and then go on to explain how this hypothesis helps link the debate surrounding emotional recalcitrance with recent empirical findings on emotion generation and emotion regulation (§3). One interesting result of forging such links is that it suggests the recalcitrant nature of recalcitrant episodes, at least on some occasions, is a product of how they are generated.

2. Exposition of the Impenetrability Hypothesis

The first, and as far as I can tell the last, major discussion of cognitive impenetrability, as it pertains to “irrational emotions”, is owing to Griffiths (1990, 1997). Though (some of) what Griffiths takes to be irrational emotions overlap with those presently picked out as recalcitrant emotions, his account is neglected by much of the current literature. This is warranted in one respect. Griffiths doesn't explain the precise sense in which recalcitrant emotions involve rational conflict, which, as mentioned earlier, is one of the central controversies. Rather, what he ends up explaining is how some

emotions can occur without any evaluative judgements; what Greenspan (1981) identifies as “groundless emotions”.

The neglect is, nonetheless, unfortunate in another respect, as Griffiths’s account proves informative when we try to understand how recalcitrant emotions might come about. Very briefly, groundless emotions occur when emotions are generated without the cognitive process of belief-fixation that gives rise to judgement. (This happens when emotions are generating via pathways that are cognitive impenetrable)³. Since some of these emotions may conflict with our preexisting judgements, this will also explain the possibility of emotional recalcitrance, though not the very nature of the conflict itself.

What this brings out is that there are (at least) three explanatory phenomena relevant to recalcitrant emotions: (1) that there is a tension between emotions and judgements, (2) the exact nature of this tension, and (3) their recalcitrant nature. Griffiths’s account can help us explain (1), whereas much of the contemporary debate on emotional recalcitrance concerns itself with (2).⁴ This paper, by contrast, aims to explain (3). It is precisely here that the cognitive impenetrability of emotions is invoked. In what follows, I provide an exposition of what it means to say that an emotional episode is cognitively impenetrable with respect to explaining (3).

At first estimates, an emotion is cognitively penetrable if it can be influenced by cognition, and impenetrable if not. There are multiple ways to specify this notion, depending on what kinds of emotion are at issue, what counts as cognitive influence,

³ A more detailed account of this process will be presented in §3.

⁴ E.g. see Helm (2001), Brady (2009) and Döring (2014).

and what is the relationship that characterises this influence.⁵ What is penetrated, for our purposes, is best thought of as emotional episodes, not emotions *per se*. This is because instances of emotional recalcitrance cut across emotion types. For example, some episodes of jealousy are recalcitrant, whereas others are not. Likewise for fear, anger etc.

How to characterise cognitive influence is a thorny issue, not just because where to draw the line between processes that are cognitive and ones that aren't is evasive, but because on certain characterisations of the mind, all its operations, including all patterns of influence, count as cognitive.⁶ We can side-step these worries since the cognitive process involved in paradigmatic cases of emotional recalcitrance are higher-order cognitive processes, roughly "processes in which people use the information of the sort they verbally assent to (traditional beliefs) and the goals they can be brought to recognize (traditional desires) to guide relatively long-term action and to solve theoretical problems" (Griffiths 1997: 92). Thoughts, beliefs, and judgments, then, all count as potential cognitive penetrators in this sense.

To this let us add one more qualification. It is unclear whether cases of emotional recalcitrance concern emotions that are in tension with a subject's overall system of beliefs or a belief (or a set of beliefs) she happens to entertain at a given moment. There is a tension for the phobic between her judgement that flying is safe and her recalcitrant bout of fear. But this leaves open whether she also has beliefs, perhaps unconscious, pertaining to flying being dangerous. It is precisely the need to commit

⁵ Here I follow Siegel's (2015) exposition of the cognitive penetrability of perception.

⁶ E.g. see predictive processing accounts of the mind, such as those endorsed by Hohwy (2013) and Clark (2016).

to such beliefs that gets cognitivists - i.e. those who take judgements to be essential to all emotions - in a bind. That is, they have to claim that the subject undergoing a recalcitrant bout of fear judges that flying is both safe and unsafe, and is thereby (allegedly) irrational.⁷

The cognitivist position, though perhaps implausible, is certainly conceptually coherent, and we can allow for it by qualifying that, in a given recalcitrant case, the cognitive process that fails to penetrate the emotional episode is not just any higher-order cognitive process, but some relevant higher-order cognitive process: very roughly, a process which would normally penetrate the emotional episode if it weren't recalcitrant. For the phobic who fears flying, it would be the judgement that flying is safe, for the jealous lover, the judgement that his partner's friendships are innocent, and so on. This will leave room for the possibility that while an emotional episode, e.g. fearing flying, is not cognitively influenced by a relevant process, e.g. my judgement that flying is safe, it might still be influenced by other cognitive processes, some of which may themselves be relevant, e.g. my (contradictory) judgement that flying is dangerous, or random, e.g. my judgement that today is Tuesday. Further exposition concerning the notion of relevance will depend on how we understand the penetrability relation itself.

This relation, at minimum, is to be understood as a causal relation. Less minimally, it is also a pattern of (causal) influence that generates responses 'top-down'. The bottom-up vs. top-down distinction presupposes a hierarchal view of information

⁷ Cognitivists include Solomon (1980) and Nussbaum (2001). See Greenspan (1981) for the irrationality objection. Also see Helm (2001), Roberts (2003) and Brady (2009) for neo-cognitivist positions that attempt to avoid the objection.

processing, where the ‘core’ idea, as Shea (2015) notes, is the distinction between incoming and pre-existing information. For example, a typical way of drawing the distinction has it that “Bottom-up processing is processing which depends directly on external stimuli, whereas top-down processing is processing which is influenced by expectations, stored knowledge, context and so on” (Eysenk 1998: 152).⁸

Another way of making the same point is in terms of *information encapsulation*. A system is informationally encapsulated if the function it computes is insensitive to information stored in other (cognitive) systems. Top-down processing involves processes that are informationally unencapsulated (i.e. has access to other systems), whereas bottom-up processing involves ones that are encapsulated (i.e. lacks access). The penetrability relation being a top-down causal process, then, is tantamount to it being a process that is informationally unencapsulated.⁹

Whether information unencapsulation suffices to capture the penetrability relation is a point of controversy.¹⁰ For our purposes of understanding emotional recalcitrance via cognitive impenetrability, we need the penetrability relation that fails to obtain in recalcitrant cases to be due to not just processes being informationally encapsulated, but encapsulated with regards to some *relevant* higher-order cognitive

⁸ See Shea for a critique of such ways of drawing the top-down/bottom-up distinction.

⁹ Here I follow Fodor (1983: 69) and Griffiths (1997: 93). By contrast, information encapsulation, sometimes, is understood as a system lacking access to information in distinct systems, cognitive or otherwise. Thus understood, information encapsulation won’t capture the hierarchical view of information processing. But more importantly, the cognitive impenetrability of a system won’t entail its information encapsulation, even though cognitive penetrability would still entail information unencapsulation.

¹⁰ Some, e.g. Fodor (1983: 77) and Nichols *et al* (1996: 47), assume that it does, whereas Pylyshyn (1980, 1999) argues that it doesn’t.

process. Pylyshyn (1980, 1999), who coined the term ‘cognitive penetration’, provides an additional constraint on the penetrability relation, which helps us address the question, What counts as relevant in recalcitrant instances?

According to Pylyshyn, cognitive penetration requires a ‘semantically coherent’ relation. What this exactly amounts to is unclear, but the idea behind it is easy enough to understand. We want the penetrability relation to be intelligible. Stokes (2013) elucidates this by claiming the semantic relation has to be an inference-supporting one. This isn't to say that the relation itself is inferential. We can be neutral on this. What is important is we can explain, or give reasons for, why an emotional episode was affected as it was by a cognitive process in virtue of the nature of this process. For instance, the penetrated emotional process of jealousy is explicable via its penetrating cognitive process, *viz.* my belief that my partner is flirting with her friend. Jealousy, in this instance, would be a contender for a cognitively penetrated emotional episode. By contrast, a feeling of joy caused by this very belief isn't *ceteris paribus* intelligible, and thereby won't be an instance of cognitive penetration.

This constraint is crucial for explaining recalcitrant emotions, as examples of emotional recalcitrance concern emotions being unaffected by some higher-order cognitive process that would normally affect them, and in a way that would support inferences from one to the other. The judgement that flying is safe explains why I don't fear flying. In the recalcitrant case, it is precisely this judgement, and ones like it, that fail to influence the bout of abject fear one feels. It is for this reason that we, at least partly, feel a tension when undergoing a recalcitrant episode. Building in semantic coherence as a constraint on the penetrability relation ensures that we track emotional episodes where this tension is manifest.

Thus far we have seen some ways to specify the penetrability relation and its relata, *sans* a precise definition of the phenomenon. It is not my intention in this section to provide a definition, as providing counterexample-proof definitions in philosophy prove to be notoriously difficult. In any case, a rough working characterisation will suffice to get a handle on our hypothesis. The characterisation on offer is that an emotional episode is cognitively penetrable if it is susceptible to top-down (causal) influence from higher-order cognitive processes, which would support inferences from them to the emotional episode itself.

Now, cognitive penetrability, as Fodor notes, admits of degrees. Thus emotional episodes will be susceptible to cognitive influence in varying degrees. Our hypothesis, recall, is that emotions are recalcitrant to the extent that they are cognitively impenetrable. In light of our characterisation, this means recalcitrant emotions are recalcitrant to the degree to which these emotional episodes are resistant to top-down influence from the relevant higher-order cognitive processes.

3. Effects of Emotion Generation on Emotion Regulation

One of the advantages of understanding the phenomenon in this way is that we stand to learn the nature of the causal mechanisms which underlie it. The recalcitrant nature of recalcitrant emotions concerns emotion regulation. That is, it is a feature of emotional recalcitrance that the relevant emotional episodes are hard to regulate: they remain resistant despite our best efforts. In this section, we shall see how emotion generation bears on emotion regulation, and can thereby provide a (causal) explanation of emotional recalcitrance.

This endeavour consists of assessing two separate, though related, claims. First, emotions can be generated in more than one way. Second, the way emotions are generated impacts emotion regulation. In particular, emotions aroused in one of these ways is less susceptible to influence from cognition. Evidence for the first claim is more familiar, so here I will mainly focus on the second.

The affect-program theory, as outlined by Griffiths, is one way of modelling the first claim.¹¹ This theory comprises of an account of what affect-programs are in conjunction with a thesis about the appraisal mechanisms that generate them. Affect-programs are “the coordinated set of changes that constitute the emotional response” (Griffiths 1997: 77). These include the physiological changes in our body, and the ‘emotional feelings’ that accompany them.¹² Sometimes these responses are automatic, akin to reflexes, whereas sometimes they aren’t. In instances where they are automatic, the appraisal mechanisms that give rise to them are ‘more or less’ cognitively impenetrable.

All the familiar empirical data that are employed to support this theory also support the ‘twin-pathway’ model of emotion generation more generally. For example, the twin-pathway neurobiological account of emotion generation for fear is regarded as being confirmed by LeDoux (1996). On this account, two distinct ‘emotion circuits’ are involved in fear responses: roughly, (i) a thalamus-to-amygdala circuit, which bypasses the cortex, is ‘quick and dirty’, and occurs without the conscious experience of the stimulus, and (ii) a thalamus-to-cortex-to-amygdala circuit, which is slow, and occurs with the conscious experience of the stimulus.

Circuit (i) allows for the rapid response to threatening situations, which is needed for survival, and (ii) ensures a detailed evaluation of the emotional significance of the

¹¹ See Teasdale (1999) for a review of other ways of modelling dual, or multi-level, theories of emotion generation.

¹² He remains neutral as to whether ‘cognitive phenomena’, e.g. beliefs, should be added to the list.

situation such that we can respond in the most appropriate way. The existence of these two separate circuits is supported by several findings.¹³ Morris *et. al.*, (2001) for instance, confirm significant activation of the amygdala in response to fearful faces in a patient with damage to his primary visual cortex, and had no conscious visual perception of the stimuli. There is empirical evidence to suggest that other emotions too can be aroused by distinct pathways.¹⁴

As for the second claim, McRae *et. al.* (2012) sought to test “whether the way that an emotion is generated influences the impact of subsequent emotion regulatory efforts” (pg. 253). Emotion generation was measured based on whether emotions were generated either bottom-up or top-down. The core idea behind the distinction, recall, is a distinction between incoming and pre-existing information. For McRae *et. al.*, emotion generated bottom-up involve responses to “inherently emotional properties of stimuli” (pg. 253), whereas those generated top-down involve cognitive appraisals (more on this shortly).

Cognitive penetration involves top-down influences on processing, but not all top-down influences need be instances of cognitive penetration, at least not in the sense relevant for emotional recalcitrance, as the penetrators at the higher-levels of processing mightn’t themselves be higher-order cognitive processes. This distinction is, however, negligible in the present context, as the relevant top-down influences tested are paradigmatic instances of higher-order cognitive processing. Top-down emotion generation was understood as “the elicitation of emotion by the activation of appraisals that a situation is relevant to an individual’s goals”, which in the lab involves “the use of language, in the form of tailored autobiographical scripts, or narrations of events that might elicit an emotion-inducing appraisal” (McRae *et. al.* 2012: 254). Rather germane for our purpose, what McRae *et. al.* are testing when they seek to confirm whether emotion regulation is

¹³ See Öhman (2005) for an overview.

¹⁴ E.g. see Phelps *et. al.* (2001), Wagner *et. al.* (2008) and Oshner *et. al.* (2009).

influenced by the way an emotion is generated, then, is whether emotion regulation is dependent on the particular appraisal mechanism by which the emotion is generated, i.e. whether the mechanism is cognitively penetrated.

The notion of cognitive penetration proves significant at a further level of the experiment as well. Emotion regulation, following Gross (2007), is broadly understood as “any process an individual uses to influence the onset, offset, magnitude, duration, intensity or quality of one or more aspects of an emotional response” (McRae *et. al.* 2012: 254). The particular type of emotion regulation measured in the present study, however, is cognitive *re*-appraisal, which typically involves “the re-consideration or re-framing of an emotional event in less emotional terms” (pg. 253). Emotion regulation, in the present context, turns out to be dependent on whether emotional responses themselves are cognitively penetrable. In this way, a different description of what McRae *et. al.* are seeking to test is whether the cognitive penetrability of emotion generation mechanisms influence the cognitive penetrability of the emotional episodes themselves.

In the experiment, participants were presented with either fearful faces (bottom-up generation) or negative sentences, e.g. “her son is in the burning building” (top-down generation). This was followed by a neutral face with the instruction to either look and respond naturally (emotion generation) or decrease the negative effect using reappraisal (emotion regulation) — both whilst taking the background information into consideration (Figure 1.). Note: subjects were trained prior to come up with relevant reappraisals, e.g. “the situation is not as bad as it first seemed”. Moreover, a neutral scrambled face or sentence (e.g. “zzzzz”) was used as a control for emotion generated top-down or bottom-up, respectively.

These findings confirm the second broad claim, i.e. the way emotions are generated impacts emotion regulation. In particular, we see that emotions generated top-down are more susceptible to emotion regulation than those generated bottom-up.¹⁵ This in turn confirms our re-reading of the experiment hypothesis: the cognitive penetrability of emotion generation mechanisms influence the cognitive penetrability of the emotional episodes themselves. More specifically, emotional responses generated by cognitively penetrated appraisal mechanisms are more susceptible to cognitive influence than those generated by mechanisms that aren't thus penetrated.

Where does this leave us? The first part of the paper, which presented an exposition of the impenetrability hypothesis, can be seen as spadework. It offered an account of the mulish nature of recalcitrant emotions in a way that made this aspect of the phenomenon of interest empirically tractable. In this section, we begin to see what we could possibly learn from this picture.

The evidence, of course, is tentative. The present study only looked at a specific type of emotion, *viz.* fear. Moreover, it sought to study emotion regulation in general, as opposed to anything specific to do with the regulation of emotional episodes involving conflict; rational or otherwise. But suppose these findings generalise to all sorts of emotion types (at least the types of which we can find examples of emotional recalcitrance, e.g. anger, jealousy etc.), and to instances of emotion regulation where the episodes under consideration involve the relevant kind of conflict. If these two (quite plausible) suppositions hold true, we stand to learn at least two new things about recalcitrant emotions from the empirical data. First, the recalcitrant nature of recalcitrant emotional episodes concern their cognitive impenetrability. Second, the cognitive impenetrability of these emotional episodes, at least on some occasions, is a product of how they are generated.

¹⁵ This finding is developed further by Otto *et. al* (2014), who show that overlapping neural processes are involved in cognitive reappraisal and top-down emotion generation.

We also stand to learn something about how we ought to treat debilitating cases of emotional recalcitrance — though this is even more tentative. One of the findings by McRae *et. al.* was that cognitive re-appraisal in cases where emotions were generated bottom-up had the reverse effect: they heightened negative emotional responses rather than reduced them. This suggests harmful patterns of emotional recalcitrance might be best treated by means other than training patients to try to cognitively down-regulate them.

According to Griffiths, one of the advantages of his theory of emotions *vis-a-vis* cognitive impenetrability is that emotional recalcitrance, or what he calls ‘irrational emotions’, is “converted from a philosophers' paradox into a practical subject for psychological investigation” (1990: 192). Griffiths didn’t fully accomplish this task. But we haven’t either. In both instances, the exact sense in which recalcitrant emotions involve conflict, rational or otherwise, is left untouched. This paper draws inspiration from Griffiths in supposing that issues to do with cognitive penetrability are relevant to emotional recalcitrance. Nonetheless, we have shifted the explanatory target. What they explain is not emotional recalcitrance *per se*, but a specific feature of it, *viz.* their recalcitrant (mulish) nature. Moreover, by explaining this feature as we do, we retain a similar explanatory advantage: insofar as the hypothesis argued for in this paper is correct, we open up the recalcitrant nature of these emotions for future scientific investigation.

References

- Brady, M. S. (2009). 'The Irrationality of Recalcitrant Emotions'. *Philosophical Studies* 145: 413–30.
- Clark, A. (2016). *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. Oxford and New York: Oxford University Press.
- D'Arms, J., & Jacobson, D. (2003). 'The Significance of Recalcitrant Emotion'. In *Philosophy and the Emotions*, edited by A. Hatzimoysis. Cambridge: Cambridge University Press.
- Döring, S. A. (2014). 'Why Recalcitrant Emotions Are Not Irrational'. In *Emotion and Value*, edited by R. Roeser and C. Todd. Oxford: Oxford University Press.
- Eysenk, M. W. (1998). *Psychology: An Integrated Approach*. Harlow: Addison Wesley Longman.
- Fodor, J. A. (1983). *The Modularity of Mind*. Cambridge and London: MIT Press.
- Goldie, P. (2000). *The Emotions: A Philosophical Exploration*. Oxford: Oxford University Press.
- Greenspan, P. S. (1981). 'Emotions as Evaluations'. *Pacific Philosophical Quarterly* 62: 158–69.
- (1988). *Emotions and Reasons*. London: Routledge.
- Griffiths, P. E. (1990). 'Modularity, and the Psychoevolutionary Theory of Emotion'. *Biology and Philosophy* 5: 175-96.
- (1997). *What Emotions Really Are*. Chicago and London: Chicago University Press.
- Gross, J. J., editor. (2007). *Handbook of Emotion Regulation*. New York: Guilford Press.
- Helm, B. (2001). *Emotional Reason*. Cambridge: Cambridge University Press.
- Hohwy, J. (2013). *The Predictive Mind*. Oxford: Oxford University Press.
- LeDoux, J. E. (1996). *The Emotional Brain: The Mysterious Underpinnings of Emotional Life*. New York: Simon & Schuster.

- McRae, K., Misra, S., Prasad, A. K., Pereira, S. C., Gross, J. J. (2012). 'Bottom-up and Top-down Emotion Generation: Implications for Emotion Regulation'. *Social Cognitive and Affective Neuroscience* 7: 253-62.
- Morris, J. S., DeGelder, B., Weiskrantz, L., & Dolan, R. J. (2001). 'Differential extrageniculostriate and amygdala responses to presentation of emotion faces in a cortically blind field'. *Brain* 124: 1241-1251.
- Nicholas, S., Stich, S., Leslie, A., and Klein., D (1996). 'Varieties of Off-line Simulation'. In *Theories of Theories of Mind*, edited by P. Carruthers and P. Smith, 39-74. Cambridge: Cambridge University Press.
- Nussbaum, M. (2001). *Upheavals of Thought: The Intelligence of Emotions*. Cambridge: Cambridge University Press.
- Ochsner, K.N., Ray, R.R., Hughes, B., *et al.* (2009). 'Bottom-up and top-down processes in emotion generation: common and distinct neural mechanisms'. *Psychological Science* 20: 1322–31.
- Öhman, A. (2005). The role of the amygdala in human fear: Automatic detection of threat. *Psychoneuroendocrinology* 30: 953–958.
- Otto, B., Misra, S., Prasad, A., and McRae, K. (2014). 'Functional overlap of top-down emotion regulation and generation: An fMRI study identifying common neural substrates between cognitive reappraisal and cognitively generated emotions'. *Cognitive, Affective, & Behavioral Neuroscience* 14 (3): 923-938.
- Phelps, E. A., O'Connor, K. J., Gatenby, J. C., *et al.* (2001). 'Activation of the left amygdala to a cognitive representation of fear'. *Nature Neuroscience* 4: 437–41.
- Pylyshyn, Z. (1980). 'Computation and Cognition: Issues in the Foundations of Cognitive Science'. *Behavioral and Brain Sciences* 3: 111–32.
- (1999). 'Is Vision Continuous With Cognition?' *Behavioral and Brain Sciences* 22: 341-65.

- Roberts, R. (2003). *Emotions: An essay in aid of moral psychology*. Cambridge: Cambridge University
- Shea, N. (2015). 'Distinguishing Top-down from Bottom-up Effects'. In *Perception and its Modalities*, edited by D. Stokes, M. Matthew, and S. Biggs, 73-91. Oxford: Oxford University Press.
- Siegel, S. (2015). 'Afterword: Epistemic Evaluability and Perceptual Farce'. In *The Cognitive Penetrability of Perception*, edited by J. Zeimbekis and A. Raftopoulos, 405-424. Oxford: Oxford University Press.
- Stokes, D. (2013) 'Cognitive Penetrability of Perception'. *Philosophy Compass* 8 (7): 646–663.
- Teasdale, J. (1999). 'Multi-level Theories of Cognition-Emotion Relations'. In *Handbook of Cognition and Emotion*, edited by T. Dagleish and M. J. Power: 665-681. John Wiley & Sons.