

IGNORING FACES AND MAKING FRIENDS

Matthew Parrott
University of Birmingham

(DRAFT: Forthcoming in *Rethinking the Biopsychosocial Model*)

Abstract: Experimental evidence shows that experiencing maltreatment in early childhood affects the way a person responds to angry faces. This is viewed by many developmental psychologists as an adaptive response to an abusive or hostile environment. However, it is also widely thought that, in a non-abusive environment, these alterations to the way a person responds to angry faces also puts an individual at greater risk than average for developing psychiatric disorders. Why is this the case? In this chapter, I shall develop a potential explanation suggested by McCrory and Viding (2015) for why maltreatment-induced alterations to neurocognitive systems increase a person's vulnerability to developing psychiatric disorders. According to their general approach, spectators visually code or classify angry faces as 'threat cues'. This type of response makes sense in a hostile or abusive environment, because, in such an environment, angry faces are reliable indicators of an environmental threat. However, I shall argue that in ordinary non-abusive environments, angry faces do not reliably indicate the presence of an environmental threat and that, for this reason, a person cannot visually identify another person as a threat on the basis of her facial expression. In contrast to what McCrory and Viding suggest, I shall argue that what increases one's vulnerability to psychiatric disorder is not misidentifying people as angry, but rather misclassifying angry people as threatening. More specifically, by treating angry people as if they were physical or emotional threats, one exemplifies a pattern of sustained behaviour that shapes one's environment in a way that puts one at greater risk of developing a psychiatric disorder.

A staggering amount of scientific evidence shows that individuals who are physically, sexually, or emotionally abused, or neglected when they are infants or very young children are at a much greater risk of developing psychiatric disorders later in life, including anxiety, depression, and schizophrenia (for an overview see Gilbert et al. 2009, cf. McCrory and Viding 2015). As one would expect, not everyone who suffers from childhood maltreatment goes on to develop psychiatric symptoms. Many adult victims of child abuse or neglect lead well-adjusted, happy, and healthy lives. However, the fact that childhood maltreatment is highly correlated with the manifestation of psychiatric symptoms later in life does suggest that something about being abused in early life significantly alters one or more aspects of a person's psychological or neurocognitive functioning, in ways that

make the person more vulnerable to other factors that might contribute to the eventual emergence of psychiatric symptoms.¹

If we wish to understand precisely how childhood maltreatment might make a person susceptible to developing psychiatric symptoms, we need to clarify two things. First, we need to understand precisely how prolonged exposure to an abusive environment in early life changes an individual's psychological or neurocognitive systems. How exactly are information-processing pathways within the mind adjusted in response to the extreme trauma of child abuse or neglect? Second, we need a much clearer picture of how the identified maltreatment-induced alterations in these systems could put someone at higher risk than average for developing a psychiatric disorder later in life. How exactly might certain functional changes to psychological or neurocognitive systems increase the probability that a person will develop anxiety, depression, or some other kind of psychiatric condition? Answering these questions will not only give us a substantive theory of how childhood maltreatment makes a person vulnerable to various psychiatric disorders, it would also go a long way toward helping us design effective medical interventions to target specific parameters implicated in the onset of those disorders.²

Childhood maltreatment is correlated with an astonishing number of psychiatric symptoms, not to mention other negative outcomes like poverty and poor physical health. It is therefore very likely

¹ Some might be reluctant to go so far as to say that childhood maltreatment is a *cause* of the psychiatric disorders that emerge in adulthood. But if we think of a cause as 'something that makes a difference', it would be hard to deny, given the data, that childhood maltreatment makes a difference to whether or not one becomes depressed, anxious, or unusually aggressive. Similarly, it is pretty clear that if we were interested in significantly decreasing the incidence of psychiatric disorders in the adult population, we could do this by intervening upon the rates of childhood maltreatment. The promise of such a social intervention suggests that there is a type of causal relationship between these variables (cf. Pearl, 2009, Woodward 2008, 2005).

² McCrory and Viding (2015) term a theory of how maltreatment makes one susceptible to developing psychiatric disorders a 'theory of latent vulnerability' (cf. McCrory, et al. 2017).

that more than one psychological or neurocognitive system is transformed by experiences of maltreatment in early childhood. In this essay, I shall focus on only one system that looks to be rather significantly affected by sustained childhood maltreatment.

A number of studies have shown that individuals who experience childhood maltreatment exhibit significant differences in the way they process information associated with threat. More specifically, several studies have shown that individuals who are maltreated as young children respond differently to angry facial expressions, which are commonly thought of as 'threat-related cues' or 'threat-related stimuli' (McCrary and Viding 2015). The experimental data seem to indicate that experiences of abuse makes a person more adept at identifying or recognizing visually-presented stimuli of angry facial expressions.

Since sensitivity to actual threats in the environment is crucial to an organism's survival, one very plausible hypothesis is that an unusually hostile or abusive environment changes the neurocognitive systems responsible for processing information concerning angry faces precisely because, in a hostile environment, angry faces are extremely dangerous and threatening (McCrary, et al. 2013, McCrary, et al. 2017, McCrary and Viding 2015, Pollak, et al. 2009, Pollak and Sinha 2002). In an environment in which angry faces reliably indicate the presence of an immediate and significant physical or emotional threat to an organism, it would be important for the organism to be highly attuned to the presence of angry faces. Thus, heightened responsiveness to angry facial expressions can be seen as a kind of information-processing adaptation to an unusually hostile environment.

This naturally suggests a corollary hypothesis about how the same heightened responsiveness to angry faces may contribute to the eventual onset of psychiatric symptoms. The basic idea is that a specific calibration of an information-processing system can be adaptive or beneficial to an organism in one kind of environment but detrimental in another. As McCrary and colleagues claim, 'such adaptations are equally thought to incur a longer term cost as they may mean that the individual is poorly optimized to negotiate the demands of other, more normative environments, thus increasing

vulnerability to future stressors.’ (McCrorry et al. 2017, p. 339, cf. McCrorry et al. 2012) So, whether or not amplified responsiveness to angry faces is beneficial to an organism depends crucially on the type of environment in which the organism is embedded. Generally, this is a plausible idea, but we still need to understand precisely why heightened responsiveness to angry facial expressions becomes ‘maladaptive’ in a non-hostile environment. How exactly does an enhanced capacity to visually recognize or identify angry facial expressions increase a person’s vulnerability to developing psychiatric symptoms?

In this chapter, I shall focus on this last question. After reviewing some of the empirical evidence that shows how childhood maltreatment affects the ways in which a person responds to angry facial expressions, I shall briefly present three hypotheses suggested by McCrorry and Viding (2015) to account for why an individual who manifests heightened responsiveness to angry facial expressions might thereby be more vulnerable than average to developing psychiatric symptoms. The aim of this chapter is to develop the second of these hypotheses, which claims that maltreatment-induced alterations to neurocognitive processing ‘leads to behaviours that shape the child’s environment over time in ways that increases the likelihood of stressor experiences and decreases the likelihood of protective experiences.’ (McCrorry and Viding 2015, p. 500) Although I shall agree with McCrorry and Viding that maltreatment-induced alterations in neurocognitive systems ultimately contribute to a more aversive and distressing social environment, one in which a person is exposed to factors that contribute to developing psychiatric symptoms, I shall develop this proposal in a slightly different direction than the one they suggest.³

³ I think the proposal presented in this chapter is a plausible way of expanding upon McCrorry and Viding’s brief presentation of the second hypothesis (2015. Cf. McCrorry et al. 2017). However, It is worth noting that since it is not something that they explicitly discuss, it is not clear whether or not it will be something which they find congenial.

McCrorry and Viding think that angry facial expressions are classified or coded by spectators as 'threat-related cues'. There are two senses in which a presented stimulus might be thought of or categorized as a 'threat-cue'. First, certain stimuli tend to elicit a range of proprietary responses in various physiological or neurocognitive systems that seem to prepare an organism for an encounter with something harmful or dangerous. For instance, they tend to generate defensive responses in the autonomic nervous system, such as an increase in heart rate and blood pressure, and the release of hormones like ACTH and epinephrine. In this sense, a stimulus qualifies as a 'threat-cue' by virtue of the effects that it has on an organism. But, it is very natural to think that stimuli have these internal consequences precisely because they reliably indicate the proximity of actual environmental threats. Thus, a stimulus may also be a 'threat-cue' in the sense of being a reliable indicator of an environmental threat to an organism. These two senses of 'threat-cue' are not necessarily independent, since some types of autonomic responses, such as high blood pressure, are, if sustained, dangerous for an organism. So, any external stimulus which caused such a response would be a kind of environmental threat.

It seems clear that angry faces generally elicit defensive physiological and behavioural responses in a subject's threat-processing systems and therefore are 'threat-cues' in the first sense (McCrorry et al. 2017, cf. LeDoux 2000, Ohman 2002, Whalen et al. 2001). But it is plausible to think that this is partly because angry faces reliably indicate the proximity of environmental threats and so are 'threat-cues' in the second sense (cf. Ohman, 2009; Ohman et al. 2001; Fox et al., 2000). The type of environmental threat one faces may vary; for example, an angry face may indicate the presence of a severe physical danger in a hostile environment, but only the presence of some social threat in an ordinary environment. Nevertheless, one might naturally think that the visual system categorizes or codes angry facial expressions as threatening in the second sense precisely because angry faces reliably indicate the presence of some kind of environmental threat. This would explain why other neurocognitive threat-processing systems respond in ways which prepare the organism for aversive consequences.

In this chapter, I shall argue that this last idea is mistaken and propose instead that the reason heightened responsiveness to angry facial expressions can be detrimental to an organism in a non-hostile environment is because, in that kind of environment, angry faces do not reliably indicate the presence of threat, neither physical nor social. I shall also claim that for this reason engaging in friendly, mutually supportive, pro-social behaviour with others, including those who are most likely to actively contribute to our personal welfare and further our interests, requires a person to learn to attenuate behavioural responses to angry faces.

1) Experimental Evidence

A number of studies show that individuals who have experienced childhood maltreatment respond differently to angry faces. In several experiments, Seth Pollak and colleagues have shown that victims of abuse perform better than controls in visual recognition tasks that involve stimuli of angry facial expressions. For example, in one study (Pollak, et al. 2009), children were shown sequences of stimuli which correspond to what Pollak and colleagues call a 'naturalistic unfolding of emotional expressions' (2009, p. 245). They were asked to determine, for each image in a sequence, what emotion was being felt by the person in that image. In another study (Pollak and Sinha 2002), children were presented with a sequence of images of a single emotional expression, for instance anger, where the image quality gradually becomes less degraded as one progresses through the members of the sequence. Once again, for each image in the sequence, children were asked to identify the emotion depicted. In both experiments, children who have experienced child abuse were able to accurately identify angry facial expressions on the basis of less perceptual information (i.e., at an earlier stage in the sequence). In a related study (Pollak and Kistler 2002), children were presented with a continuum of facial expressions ranging, for example, from definitely angry to definitely sad.⁴ Maltreated children categorized a wider range of expressions on this continuum as

⁴ This is an interesting result, but it is difficult to say what it indicates. The continua of emotions that Pollak and Kistler present have something like an artificial borderline between two different types of

'angry' than controls. Furthermore, in all of these studies, there was no observed difference in the way maltreated children responded to facial expressions of emotions other than anger.

These experimental results seem to be confirmed by neurophysiological studies. EEG studies have shown that, when they are asked to search for or actively pay attention to angry faces, maltreated children manifest a comparative increase in brain activity, specifically a higher P3 amplitude, which is correlated with high levels of anxiety (McCrary and Viding 2015, Pollak, et al. 1997). In addition, a number of recent fMRI studies show neuro-functional differences in the way maltreated children respond to angry faces. For example, McCrary and colleagues (McCrary, et al. 2013) used fMRI to measure pre-attentive neural responses to facial expressions of emotion and found that maltreated children exhibited greater activation in the right amygdala when processing information about angry faces.⁵ A related fMRI study (McCrary, et al. 2011a, cf. McCrary, et al. 2017) also found that maltreated children exhibited greater activation in the right amygdala as well as in the anterior insula when processing angry facial expressions.

Among many developmental psychologists, these studies are taken to demonstrate that children who have been maltreated at a young age respond differently to angry faces. In what follows, I shall assume that this is the correct interpretation of the experimental results.

expressions, which is simply the middle image on the continuum, an artificially created 50% blend of each emotion being compared. It isn't clear to me that this artificial image corresponds to anything like a natural border between facial expressions of, for instance, sadness and anger. So, rather than thinking that the experimental results show that abused children categorize too many facial expressions as angry, one may interpret the results as showing that controls categorize too few facial expressions as angry. It isn't clear that either interpretation is preferable. It is also notable that neither group set their subjective border near the artificial midpoint.

⁵ A second interesting result of this study is that these children also showed greater right amygdala activation when processing happy faces. This is slightly at odds with the experimental results of Pollak and colleagues (Pollak and Kistler 2002, Pollak and Sinha 2002, Pollak, et al. 2009) which showed no information-processing differences with respect to expressions of happiness.

2) Explaining Neurocognitive Alterations

The notion that a neurocognitive system might develop in a manner that displays heightened responsiveness to angry faces is not difficult to grasp. As many theorists have remarked, it is quite plausible that abused children are exposed to an unusually high number of angry faces in early childhood (McCrory, et al. 2013, McCrory, et al. 2017, McCrory and Viding 2015, Pollak, et al. 2009, Pollak and Sinha 2002). More importantly, in an extremely hostile environment, the presence of an angry face is plausibly associated with immediate harmful outcomes. Since children in these environments need to protect themselves from various physical and emotional harms, angry facial expressions naturally acquire heightened salience. Therefore, it is plausible that learning to rapidly detect an angry face may be beneficial in such an environment. As Pollak and Sinha claim, ‘physically abused children learn to make decisions about the signalling of anger using minimal visual information. The development of increased perceptual sensitivity for the fine-grained details of variation in affective expressions may provide a behavioral advantage for living in threatening contexts, allowing earlier identification of emotion.’ (2002, p. 786, cf. McCrory, et al. 2017)

Indeed, mere exposure to a greater frequency of angry faces may itself be sufficient for explaining why a subject acquires the ability to accurately identify or recognize angry facial expressions quickly, or on the basis of less perceptual information. According to standard theories of perceptual learning, perceptual systems can learn to identify and discriminate a type of visual stimulus more quickly after repeated exposure to that stimulus (Goldstone 1998, Hagemann, et al. 2010). We therefore have some basic grasp of why someone exposed to an unusually high number of angry faces, especially in early life, might become more proficient at visually recognizing them.

Let’s assume that an explanation along these lines can account for why maltreated children exhibited augmented visual responsiveness to angry faces. It is still not obvious why this type of information-processing bias would increase the probability of developing a psychiatric disorder. A plausible idea is that alterations to the way a subject processes information about angry faces is only

beneficial in a hostile environment. When subjects are removed from this type of environment, it is no longer adaptive, perhaps because the frequency of exposure to angry faces diminishes. But, although it may be correct to think that a superior capacity to visually recognize angry faces on the basis of less information is less beneficial in a non-hostile environment, it does not follow from this that it would thereby be maladaptive or somehow put one at risk for developing a psychiatric disorder.

Suppose that I am born and raised in an environment where the primary source of food is a certain type of berry, but that 30% of these berries are poisonous. It would not be an evolutionary surprise if I learned to become very good at picking out or identifying these poisonous berries. Maybe I could learn to do this on the basis of the way they look or smell. If so, I might also become more attentive to the way these berries look or smell, especially when compared to other visible or olfactible properties. My capacity to process information relevant to identifying and recognizing these berries would be a kind of adaptive response to my environment. But now suppose that I move to an environment where only 0.05% of the berries are poisonous and there is also a much wider variety of nourishing food. My capacity to visually recognize poisonous berries may no longer be a clear evolutionary advantage, but it does not thereby become a deficit or something that puts me at risk of developing psychiatric or social problems. By analogy, it is difficult to see why being especially adept at visually identifying angry faces would itself to constitute a disadvantage to an organism.

So even though we have evidence linking childhood maltreatment to altered neurocognitive processing of angry facial expressions, we still need to understand how these processing alterations contribute to the onset of a psychiatric disorder. McCrory and Viding (2015) propose the following three candidate hypotheses. First, they suggest that having an enhanced capacity to visually recognize or identify angry faces is best accounted for in terms of a subject being hypervigilant to angry faces. The thought is that individuals who have experienced abuse are able to visually recognize angry faces on the basis of less perceptual information because they are more vigilant to

the presence of angry faces. As a result of this hypervigilance, their attentional resources are diverted from other domains. Specifically, McCrory and Viding propose that less attention will be allocated to domains relevant to 'optimal development in nonthreat environments.' (p. 500; cf. McCrory, et al. 2017).

The second hypothesis presented by McCrory and Viding is that heightened responsiveness to angry faces leads to propensities for certain types of behaviour that 'shape the child's environment over time in a way that increases the likelihood of stressor experiences and decreases the likelihood of protective experiences.' (2015, p. 500) How might this happen? McCrory and Viding again appeal to hypervigilance. They think it is plausible that children who are hypervigilant to angry faces are more likely to *misattribute* threat to peers and then behave toward those peers as if they were threatening (cf. Pollak and Kistler 2002). As a result, their peers are more likely to become reactively aggressive, which in turn gives rise to a much more stressful social environment, one in which pro-social relationships are difficult to create and sustain. It is not easy to form a mutually supportive relationship with someone whom one treats as a threat, and, without such relationships, one is at greater risk of developing psychiatric symptoms.

McCrory and Viding's final hypothesis is that greater sensitivity to angry faces recalibrates a person's affective systems such that they manifest amplified responses to future environmental stressors. For example, they may have generally higher levels of anxiety. Experiencing highly aversive affective responses may also lead an individual to develop negative coping strategies such as avoidance or repression, which would also increase one's chances of developing psychiatric symptoms.

For all three of these hypotheses, it is not the case that being adept at visually recognizing angry faces is itself what puts someone at risk of developing a psychiatric disorder. It is rather that changes to the neurocognitive systems underlying this discriminatory capacity have additional negative consequences, such as misattributing anger to non-angry people, or having generally high levels of anxiety. The three hypotheses are not mutually exclusive. It is possible that any or all of them could

be true.⁶ At this stage, they are highly speculative and in need of empirical confirmation (or disconfirmation). Nevertheless, in the following section, I would like to explore the second hypothesis more fully.

3) Anger and Threat

McCrorry and Viding's second hypothesis is that, by virtue of being hypervigilant to angry faces, individuals who have experienced childhood maltreatment develop a tendency to treat their peers as threats. *Prima facie* this sounds plausible. It does seem natural to think that an individual who has been exposed to a large number of angry faces would frequently be on the look-out for them. In that case, it would make sense if the individual had a greater propensity to miscategorise facial expressions as angry. This would explain the results of the experiment conducted by Pollak and Kistler (2002) in which subjects of childhood maltreatment categorized a greater range of facial expressions as angry. Hypervigilance to angry faces would also plausibly explain why individuals exposed to maltreatment were more adept at recognizing actual angry faces on the basis of less perceptual information.

McCrorry and Viding's presentation of this second hypothesis is extremely brief, so it will be useful to highlight its central claims. First, McCrorry and Viding are in a position to claim that hypervigilance to

⁶ These three hypotheses account for vulnerability to psychiatric disorder in terms of neurocognitive functioning. But they are not intended to exclude other explanations in terms of neurophysiology, genetics, or epigenetics. Indeed, there is evidence that many psychiatric symptoms associated with childhood maltreatment are heritable. For example, Caspi and colleagues have demonstrated that individuals who carry the MAOA-1 allele are at greater risk for developing patterns of anti-social behaviour (Caspi et al. 2002, cf. McCrorry et al. 2011b). Genotypes that predict a high risk of developing psychiatric disorder might also help to explain why certain subjects select high-stress or aggressive social environments, in which case a genetic account would also underlie a functional neurocognitive explanation (for further discussion see McCrorry et al. 2017).

angry faces puts one at greater risk of developing a psychiatric disorder because they assume that misattributing anger to someone is closely connected, or perhaps equivalent, to misattributing threat (cf. Ohman et al. 2001). Yet, it is really the latter that is more crucial to their explanation. The reason why they think individuals who have experienced childhood maltreatment become more vulnerable to psychiatric disorder is because they think these individuals are far more likely to treat non-threatening peers as if they were serious threats. The sort of treatment they have in mind consists in robustly exhibiting avoidance behaviour or anti-social behaviour such as acts of aggression. To exacerbate matters, in response to this type of behaviour, non-threatening peers are more likely to engage in reactive aggression.

This proposal makes sense but it is worth noting that the last step of McCrory and Viding's hypothesis has two stages. First, non-threatening peers are mistakenly categorized or identified as threats. For this to be at all relevant to experimental work on how individuals respond to facial expressions, there must be some connection between visually identifying a facial expression as angry and visually coding or classifying that same expression as threatening. As we have seen, McCrory and Viding follow a tradition of thinking of angry faces as 'threat-cues'. On the assumption that all angry faces are 'cues' or 'indicators' of some kind of environmental threat, identifying someone as expressing anger implies classifying them as a threat.

The second stage of McCrory and Viding's explanation claims that after someone identifies a person's facial expression as threatening, the subject will subsequently come to treat that person as a serious threat. According to McCrory and Viding, behaving toward one's peers in these ways directly contributes to a more stressful and detrimental social environment, which in turn increases one's chances of developing a psychiatric disorder. Notice that 'treating' one's peer as if they were threatening is not merely displaying a pattern of autonomic responses. As we have already seen, there is a sense in which an organism 'treats' a stimulus as a threat simply insofar as the stimulus activates certain processes in one's autonomic nervous system. But those responses are not

sufficient to generate a detrimental social environment; rather, they need to be connected to sustained patterns of anti-social or avoidance behaviour. So even if a wider range of facial expressions cause autonomic arousal in children who have experienced childhood maltreatment, we still need to understand how this leads to the sort of anti-social or avoidance behaviour that put individuals at risk of developing psychiatric disorders, or at least understand why these sorts of responses are not regulated or attenuated (cf. Lyons et al. 2009; Sell et al. 2009)

With this in mind, there seem to be four steps to McCrory and Viding's proposed explanation:

- 1) People who have experienced childhood maltreatment become highly sensitive (hyper-vigilant) to angry facial expressions. For this reason, they tend to miscategorise facial expressions as angry.
- 2) Assuming that angry facial expressions are 'threat-cues', categorising someone's face as angry means categorizing them as a threat.
- 3) Individuals who have been maltreated as children respond to the people that they have classified as threatening with anti-social behaviour like aggression, or with avoidance.
- 4) Anti-social and avoidance behaviours undermine pro-social and potentially supportive relationships (in part by provoking reactive aggressive responses from one's peers).

According to McCrory and Viding's proposal, what primarily contributes to the greater risk of psychiatric disorder is the way an individual treats people who are not really expressing anger, rather than the way they react to angry peers (although it is consistent with their view that the latter may also be a contributing factor). This an attractive hypothesis, but it is worth reflecting on whether angry faces really are 'threat-cues' in the sense of reliably indicating the presence of an environmental threat. As we have already seen, in a hostile environment, angry faces reliably precede harmful outcomes, and so they are plausibly reliable indicators of immediate and serious physical or emotional threat. This partially explains why a high frequency of angry faces in such an environment recalibrates a child's information-processing systems such that the child becomes more

responsive to them. It is natural to think that angry faces also reliably indicate the presence of a threat in non-hostile environments, even though the degree and nature of that threat may be different. Indeed, it is natural to think that basic emotional expressions generally have a function of reliably indicating properties in the environment relevant to an organism's interests (cf. Frijda and Mesquita, 1994). However, although this is a common way to think of emotional expressions like anger, I would like to present an alternative conception of expressions of anger in the following two sections.

4) Looking Angry

In the experiments presented earlier, we saw that subjects were presented with visual stimuli of angry faces for an extremely brief interval. So classifying or coding a facial expression as angry is a task carried out by the visual system, at least in the first instance. Let's assume this is true. The visual system responds to the way angry facial expressions look or appear, and it categorizes or identifies a person's face as angry because of how it looks or appears. The question is whether the visual system can also classify the person's face as threatening.

In what follows, I shall define a *distinctive look* as follows: the way something looks is a *distinctive look* of some property F if and only if most of the things in a spectator's environment that manifest that look really are F (Martin 2010, Millar 2000). So, for example, being round and yellow in a specific way is the distinctive look of *lemons*. Why? Because most of the things in our environment that look that way, most of the things that *look lemony*, really are lemons.

Of course, not all lemons look lemony; some of them have begun to rot and look rather different. An object can *be* F (it can be a lemon) without manifesting the distinctive look of F (without looking lemony). Moreover, not everything that looks lemony is a genuine lemon. My daughter has a plastic toy replica lemon that is a dead ringer for the real thing. But it isn't. So an object can manifest the distinctive look of F (it can look lemony) without in fact being F (it can be a plastic toy). Nevertheless,

since most of the things in our environment that look lemony are lemons, a lemony look is the distinctive look of lemons.

This shows that whether or not a look or appearance is distinctive of some property depends on features of a spectator's environment. In my environment, the majority of the things that look lemony really are lemons. But if I took a lemon from my fruit bowl to a non-standard environment, such as one where there are a number of factories manufacturing plastic toy replicas, then the way my lemon looked would no longer be distinctive of *lemons*. That is because it would not be true in that environment that most of the things that manifest a lemony look are lemons; rather most would be plastic toys. That does not mean that my fruit's appearance has changed. It looks the same no matter where you take it. What changes is whether the way it looks is *distinctive* of some property F (being a lemon or being a certain type of plastic toy).

All visible things have looks or appearances. Yet subjects do not always have the appropriate visual capacities to be sensitive to every way that an object looks. For instance, my neighbour is an avid aviculturist and has the capacity to visually discriminate a female house finch from a female Cassin's finch. I don't. But even though the two of us have different visual recognitional capacities, the two birds manifest the same appearance to both of us. It is just that one of us has learned to be visually sensitive to the way a female house finch looks and the other one of us has not.

Thus, whether or not a spectator is able to visually recognize or identify objects to be lemons on the basis of how they look depends on two things. First, a spectator must have the appropriate type of visual capacities; she must have the ability to visually recognize lemony looks. Second, the way the object looks must be distinctive of lemons. In the imagined scenario with toy factories, I would not be able to visually recognize or identify a piece of fruit as a lemon on the basis of how it looks. This is not because I am not sufficiently visually sensitive to that look, but because, in that environment, the way it looks is distinctive of plastic toys, not lemons.

How does this relate to our discussion of facial expressions? Facial expressions are standardly identified or classified by the way they look. This makes it plausible that angry facial expressions have a distinctive look.⁷ Given our definition, the way a face looks is distinctive of anger if and only if most of the faces in a spectator's environment that manifest that look really are angry. In such an environment, someone with the appropriate recognitional capacities will be able to visually recognize or identify faces as angry on the basis of how they look. It seems very clear that human beings have this ability from a very young age (Kahana-Kalman and Walker-Andrews 2001, Kestenbaum and Nelson 1990, Pollak, et al. 2009).

Consider now a hostile environment in which a child is subjected to severe physical and emotional maltreatment. We have seen that in that type of environment angry faces immediately precede harmful outcomes and are therefore a reliable indicator of threat. It follows from this that, in a hostile environment, the way an angry face looks is also a distinctive look of threat. If most of the faces that manifest a particular type of facial expression (i.e., the ones that look angry) are in fact threatening, it follows that most of those faces *look* threatening. Generally, for any property G that is reliably correlated with F, if the way something looks is distinctive of F (if most of the things in the environment which manifest the look are F), then it will also be distinctive of G. So, in a hostile environment, angry facial expressions present a distinctive look of *both* anger and threat, and a subject with the appropriate recognitional capacities will be able to visually recognize or identify threats on the basis of how they look.

But what about a non-hostile environment? Is it the case that in this type of environment angry facial expressions present a distinctive look of threat? I would like to propose that they do not

⁷ There is a further question of whether the ability to recognize angry facial expressions on the basis of how they look is sufficient to secure perceptual knowledge that a person is angry. This question is outside the topic of this essay but is discussed at some length in Parrott, 2017 (cf. Gomes forthcoming, McNeill forthcoming).

because, in this type of environment, it does not seem that angry faces are reliably correlated with the presence of threats. Most of the people that manifest anger toward me do not pose a threat to me and certainly do not harm me physically or emotionally. One group of people that I encounter with angry facial expressions are complete strangers – angry drivers, angry bystanders, angry commuters on the London Underground. Since these people are expressing anger, I'm able to visually recognize them as angry. But after our brief encounter, I never see them again. Partly for that reason, they are not threatening to me and, unlike in an abusive environment, their angry facial expression is not followed by harmful consequences. The other class of people that look angry are my family members and friends, people that tend to share my interests and wish to promote my well-being. These people are also not threatening. Rather than causing me serious harm, they tend to foster my health and happiness.⁸ So, in an ordinary environment, it seems that most of the faces that look angry are not threatening.⁹

This last claim gains further support from the view that the emotion of anger evolved because of its social function. Many evolutionary psychologists believe that anger evolved in order to organize complex human behaviour into cooperative social relationships (Haidt 2003, Keltner and Haidt 1999,

⁸ One might think that these people are threatening, it is just that the type of threat they present is different from the sort of physical or emotional threat once faces in an abusive environment. Thus, one might feel emotionally shaken by angry commuters or, especially if your loved ones are angry at you, their angry facial expressions might be thought to indicate a type of social threat, such as displeasure or the withdrawal of affection. This is an important line of objection to which we shall return in the following section.

⁹ Angry faces may nevertheless be threatening in the sense of causing defensive autonomic responses. But if the argument in this section is right, then this causal relation is brute—it is not a response to the visual system having recognized or identified an environmental threat.

Frijda and Mesquita 1994). As such, the primary function of anger is not to indicate the proximity of physical or social harm, but rather to signal to one's peers that there has been a perceived violation of a norm or expectation governing social cooperation. In this sense, expressions of anger function to facilitate cooperative behaviour. For instance, according to the model developed by Sell and colleagues, the function of anger is to 'recalibrate the target of the anger by showing the target that it will be worse off by continuing to behave in ways that place too little weight on the actor's interests.' (Sell et al. 2009, p. 15074). Moreover, since the function of anger is to promote cooperation, it is normally diminished once the event which initially triggered the anger is appropriately redressed (Keltner and Haidt, 1999; Lerner et al. 1998). If this is right, then adults in non-hostile environments will normally have learned that expressions of anger pose no threat of physical or social harm to those who are appropriately engaging in cooperative behaviour. By contrast, children raised in a hostile or abusive environment experience unprovoked, random, or dysfunctional anger, which does not diminish in response to clear signals of cooperative behaviour. Thus, these children have not had an opportunity to learn how expressions of anger function to foster social bonds.

By contrast, in non-hostile environments, a spectator will likely encounter a fair number of faces that manifest the distinctive look of anger, but the majority of these will not be threatening. It follows from this that, in a non-hostile environment, angry facial expressions do not present a distinctive look of threat. So a spectator in this environment will not be able to visually recognize or identify *threats* on the basis of how someone's face looks, just as a spectator is not able to visually recognize lemons on the basis of a lemony look in the imagined toy-factory scenario. Of course, it is possible that threatening individuals manifest some kind of distinctive look, which is to say that there may be some way that people can look such that most of the people that look that way are actually threatening. But it is not the look of an angry face. For that reason, it is not clear whether this would be a look that an ordinary spectator has the ability to recognize, or whether it is more like the look of a female house finch, which can only be recognized after significant training.

This suggests an alternative way of developing the second hypothesis presented by McCrory and Viding. Recall that according to their proposed explanation the proximate factor contributing to the higher risk of developing a psychiatric disorder is a more stressful environment shaped the fact that an individual treats non-threatening people as if they were threats. As we saw in the previous section, McCrory and Viding claim that this results from misidentifying the facial expressions of people who are not actually angry. However, we can now see that rather than misattributing anger to non-angry facial expressions, it is possible that individuals mistakenly attribute *threat* to people which their visual system correctly identifies as angry. If, as I have claimed, angry facial expressions do not present a distinctive look of threat, then it is possible for a spectator to correctly recognize an angry face on the basis of how it looks, but lack an adequate basis for visually identifying that face as threatening. Nevertheless, a spectator may take themselves to be recognizing a threat because they mistakenly think angry facial expressions present a distinctive look of threat.¹⁰

We actually should expect this sort of mistake whenever a spectator's environment is significantly changed. Suppose that I go to the environment dominated by lemon factories after having spent some time growing up on a lemon farm. I would naturally think that I am able to visually recognize *lemons* on the basis of how they look and I would take myself to be exploiting this capacity in the new environment. However, in that environment I am not in a position to visually discriminate lemons on the basis of how they look, only plastic toys. Nonetheless, since my behavioural responses are attuned to the lemon farm, I would naturally treat the plastic toys as if they were lemons, at least initially. Especially if I am not aware of the extent of environmental changes, I might not be in a position to know that I cannot visually identify genuine lemons on the basis of their

¹⁰ Notice that this suggestion is compatible with McCrory and Viding's idea that the individuals are also misidentifying people as angry. Of course, if these people are also misidentifying people as angry then the problem will be amplified.

lemony look. So, it would not be surprising if I responded to the plastic toys as if they were lemons because I take their look to be the distinctive look of lemons.

As we have seen, children who are abused in early life learn to visually discriminate facial expressions in a hostile environment wherein angry faces do present a distinctive look of immediate threat. This explains why they are predisposed to treat people with angry faces as threats. This predisposition would naturally continue after one is removed from a hostile environment. So, someone who has learned to associate angry faces with threatening or harmful people will naturally respond to angry faces as threats, even after she has moved to a non-hostile environment. Moreover, she will not have had an opportunity to learn the appropriate responses to expressions of anger in a non-hostile environment in which the emotion functions to promote cooperative social relationships.

Let me note how the explanation proposed in this section differs from the one presented by McCrory and Viding. First, both explanations hold that people who have been severely maltreated as children develop heightened visual responsiveness to angry facial expressions. For this reason, they are especially adept at visually recognizing angry facial expressions. But whereas McCrory and Viding appeal to hypervigilance to account for this heightened responsiveness, the explanation proposed in this section does not. Instead, it claims that individuals are able to acquire a set of heightened recognitional capacities in a hostile environment where people express anger are threatening, even though, in a non-hostile environment, they are not in a position to visually discriminate threats. This seems anodyne on McCrory and Viding's view because they assume that angry faces are correlated with some type of environmental threat.

This brings us to the second difference between the two proposals, namely that the proposal in this section rejects the notion that angry faces are reliable indicators of threat in all environments. Instead, it proposes that angry faces present a distinctive look of threat only in a hostile environment. In a non-hostile environment, it is more common for non-threatening people to look

angry. This is partly why behaviourally responding to angry people as if they were threatening can be detrimental to an individual in a non-hostile environment. Thus, whereas McCrory and Viding's explanation focuses on problematic responses to peers who are not really angry, the explanation proposed in this section focuses on how certain kinds of responses to angry people might contribute to a higher risk of developing psychiatric disorders.

Despite these differences, the explanation I have sketched in this section is compatible with McCrory and Viding's suggestion that individuals who are exposed to childhood maltreatment misattribute anger to non-angry people. So it is not really a competing hypothesis. The explanation also accepts their claim that the most proximate factor contributing to an increased risk of developing a psychiatric disorder is the more stressful environment generated by treating one's peers as if they were threatening. As McCrory and Viding point out, responding to non-threatening peers in this manner makes it much more difficult to create and sustain supportive and mutually beneficial relationships, which plausibly increases one's chances of developing psychiatric problems.

If your closest friends can be angry without being a threat to you, then so can your potential friends. Learning to dissociate expressions of anger from threat is one of the keys to forming lasting, mutually supportive social relationships. By taking angry facial expressions to be presenting a distinctive look of threat, subjects who have experienced childhood maltreatment risk missing out on these sorts of positive relationships.

5) Two Types of Emotion

The explanation I have just proposed relies on the notion that, in a non-hostile environment, expressions of anger are not reliably correlated with any kind of environmental threat. This might strike people as counterintuitive. The term 'expression' is factive--someone can *express* anger only if they are actually angry (cf. Gomes forthcoming, Green 2007). This means it is trivial that angry facial expressions are reliably correlated with actual anger. But we might naturally think that the emotion

of anger is itself a kind of threat. Especially in cases where people are angry at *you*, another person's anger might directly cause strong feelings of displeasure, or other kinds of negative arousal. Along these lines, some theorists have claimed that part of what it is to *be angry* with someone is to be a potential source of physical, emotional, or psychological harm to that individual. For example, Martha Nussbaum claims that it is part of the very concept of anger that it consists of 'a wish for things to go badly somehow, for the offender.' (2015, p. 46) Part of the reason Nussbaum thinks this is that she believes anger is a characteristic response to a perceived offense; it is a feeling that one has in response to being wronged by another. In that sense, even in cases where someone harms me accidentally, I feel angry in the sense of desiring some kind of retribution for the wrong done to me.¹¹ As we have already seen, in normal contexts, this anger will diminish so long as the offending party apologizes or redresses the wrongdoing. But even so, if part of what it is to be angry involves wishing for the object of one's anger to be harmed, then it might seem that all angry people are threatening, at least to the extent that they wish for some kind of negative outcome for the object of their anger. If this is right, then emotion of anger may be reliably correlated with threat, even in non-hostile environments. It is just that the nature and degree of threat varies significantly from the sort that one faces in a heavily abusive environment. But if angry faces express anger and anger is reliably correlated with threat, then the explanation I have proposed in the previous section cannot work.

Before responding to this objection, it is important to note a distinction between the idea that angry facial expressions reliably indicate the presence of some type of environmental threat and the idea that angry facial expressions are harmful. Many of us might find that expressions of anger, even by strangers, cause unpleasant feelings. We want those that we love to have positive feelings toward us, and the manifestation of anger can easily hurt, or cause anxiety, or other forms of distress. We

¹¹ My anger might very well vanish once the offender apologizes. But, on Nussbaum's proposal, even in this case, my short-lived anger will constitutively involve a 'wish for things to go badly somehow'.

have also seen that angry facial expressions directly stimulate defensive responses in the autonomic nervous system, at least some of which might reasonably be considered to be harmful. Nothing that I have proposed in this chapter is inconsistent with the idea that angry facial expressions can themselves be emotionally, psychologically, or physiologically harmful. But this is not equivalent to them indicating the presence of some kind of environmental threat. Whether or not someone is a threat depends on what they are going to do, either in the immediate future or at some other future time. So, for example, seeing a family member express anger might immediately cause negative arousal, but it might also indicate to me that the person is going to withhold affection from me for some length of time. Similarly, seeing a potential friend express anger might cause a detrimental emotional response, but it might also be taken to indicate that the individual is someone who is likely to cause negative effects in the future. Harm is a matter of degree and it is the fact that individuals who have experienced significant childhood maltreatment expect their peers to cause them serious harm in the future that undermines positive social relationships. This is the sense in which they treat their peers as threats.

Let's return to the objection, which claims that angry faces are reliable indicators of some type of threat. There are two ways to avoid this objection. The first would deny that the emotion of anger reliably manifests itself in harmful behaviours. Like any emotion, being angry involves dispositions to act and react in certain ways. But, there appear to be a wide range of different behaviours associated with anger (Kuppens, et al. 2004). So, although anger does give rise to aggressive or assertive behaviours, or to inhibition of them, there is evidence that anger also leads to pro-social behaviours (Kuppens, et al. 2004, Rime, et al. 1991). This should be expected if, as we saw previously, anger has a positive social function. So, it seems that anger does not invariably lead to harmful behaviours. Nevertheless, it may seem that the majority of cases of anger involve at least a *disposition* to treat the object of one's anger with some sort of aggression.

The second way to avoid the objection denies that expressions of anger are reliably connected to a disposition toward harmful or aggressive behaviour. To see how this could be true, it will be helpful to appeal to a distinction introduced by Richard Wollheim. According to Wollheim, our everyday vocabulary for speaking about the emotions is ambiguous. Sometimes we use words like ‘anger’ or ‘happiness’ to refer to psychological attitudes.¹² These are persisting, dispositional features of a person’s psychology that manifest themselves in the person’s behaviour. So, as we have seen, part of what it is to have an emotional attitude is to have a corresponding disposition or dispositions to behave in certain ways. In Wollheim’s terms, emotional attitudes provide a creature with ‘an *orientation*’ toward the world (1999, p. 15).

However, Wollheim also points out that we sometimes use the very same emotion-words to refer not to attitudes or dispositions, but to transitory mental *events* ‘in which the emotions manifest themselves.’¹³ (1999, p. 9) That is, we sometimes use words like ‘anger’ or ‘happiness’ to refer not to a persisting attitude, but to refer to an episodic conscious event that momentarily occupy our thought (cf. Frijda and Mesquita 1994).

We can avoid this ambiguity by describing conscious events as ‘feelings’. For instance, when I am abruptly cut off by another driver on the motorway, I might naturally feel angry without also *being* angry in a more robust sense. That is, my conscious mental life might be occupied with a particular sort of conscious episode, which we might associate with anger. I might vent this feeling by yelling or cursing, but it would be natural for the conscious episode to be short-lived. Unless there is some

¹² Wollheim classifies these as psychological dispositions.

¹³ Unlike many philosophers, Wollheim calls the transient events within a person’s mind ‘mental states’. But it is clear that he defines ‘mental states’ as events: ‘mental states are those transient events which make up the lived part of the life of the min, or to use Williams James’s great phrase, “the stream of consciousness”. (1999; pg. 1)

underlying pathology, a conscious mental event is not something that tends to persist in my mental life.

Crucially, the fact that I have a conscious feeling of anger does not entail that I *am* angry, in the sense of having an underlying attitude which disposes me toward certain kinds of behaviour. Of course, someone who is angry in the attitude sense also tends to have corresponding conscious feelings of anger. That is part of what it is to have the emotional attitude of anger. But the converse is not true. One can simply experience a conscious feeling of anger without that meaning they are behaviourally oriented toward the world in any specific way.

If this is right, then conscious feelings need not be reliably connected with emotional attitudes. Even though attitudes like anger typically cause conscious feelings of anger, in ordinary environments those feelings are frequently caused by environmental conditions. It is because I was cut off by a driver on the motorway or because my favourite café is closed that I feel angry. Features of my environment directly cause the feeling; they cause a particular conscious event without having to also cause any persisting dispositional attitudes. If cases where conscious feelings like anger are caused by environmental conditions, they do not reliably indicate the presence of any type of emotional attitude.

Let's assume that attitude of anger is some kind of environmental threat. The account I proposed in the previous section is committed to the existence of an ontological gap between facial expressions of anger and the emotional attitude of anger. We can now see that there could be such a gap in cases where angry facial expressions are caused by conscious feelings, rather than by emotional attitudes. In such cases, a spectator would be in a position to visually discriminate people who *feel* angry on the basis of their facial expressions, but they would not thereby be able to visually discriminate people with a persisting emotional attitude of anger. Yet it is only the latter that is plausibly a threat.

Facial expressions are transitory. One briefly vents a feeling of anger by grimacing or quickly expresses one's feeling of happiness by jumping for joy. Neither the grimace nor the jumping last for long. It seems to me that many facial expressions are directly caused by momentary conscious feelings. As such, the expressions are not a reliable indicator of any persisting attitudes or behavioural dispositions. For this reason, they have little evidential value in predicting how a person will be disposed to behave toward me, or whether the person's future actions will be helpful or harmful. Thus, even though certain kinds of facial expressions like anger may themselves be harmful in some sense, for instance, by causing momentary physiological distress, they indicate only what a person feels for a moment.

Of course, none of this is true in hostile environments. In that type of environment, conscious feelings are typically caused by underlying attitudes. The reason that an abusive parent tends to frequently experience conscious feelings of anger is that deep down he or she *is* extremely angry. The abusive parent has an underlying emotional attitude that orients her toward the world in a specific way and that attitude, rather than her surrounding environment, is directly responsible for her conscious feelings of anger. So, in a hostile environment, angry facial expressions are reliably connected to the emotional attitude of anger and since, especially in this type of environment, this attitude leads to extremely aggressive behaviour, it can reasonably be seen as a threat.

To make things worse, in an abusive environment, the emotional attitude of anger is itself dysfunctional. It occurs randomly or spontaneously and is therefore disconnected from the social function anger plays in non-abusive environments. Because the abusive parent's anger is clearly not a response to any offense or wrongdoing on the part of the child, and because it does not diminish in light of clear signals of cooperation, the maltreated child is unable to learn how to appropriately respond to anger, but responds instead with avoidance or aggression. There is simply no opportunity for the maltreated child to learn how anger serves to establish and regulate positive social relationships in non-abusive environments.

6) Conclusion

Correctly interpreting facial expressions is a complex problem. One needs not only to acquire the ability to visually discriminate the different ways a face can look, but one also needs to learn what these different expressive looks mean. For this reason, the environment in which one learns to visually recognize faces is crucial for determining precisely how one responds to different types of facial expressions.

As many psychologists have suggested, it is plausible that, in an abusive environment, angry facial expressions are indications of very serious threat. That is why it is beneficial for an individual's neurocognitive systems to develop to become highly sensitive to angry faces. When you are in a hostile environment, it is absolutely vital that you can recognize angry faces as quickly as possible.

Nevertheless, individuals whose neurocognitive systems manifest heightened responsiveness to angry faces seem to be at a much greater risk than average of developing psychiatric disorders later in life, once they are removed from an abusive environment. The main proposal of this chapter is that this risk can be explained as the direct result of the individual's behaviour shaping her environment to be more distressing and anti-social. Specifically, it is plausible that a person who has experienced childhood abuse will tend to treat non-threatening peers as if they were threats, which, significantly compromises the person's ability to develop supportive social relationships. This is the basic explanatory approach suggested by McCrory and Viding (2015).

However, I have also suggested that the way an individual who has experienced abuse responds to angry people may be inappropriate. This suggestion rests on the idea that in ordinary environments angry faces are not reliable indicators of threat, and so they are not a basis upon which a spectator can visually discriminate environmental threats. Indeed, in the previous section it was suggested that, in an ordinary environment, angry facial expressions are often not even caused by an

underlying attitude of anger. Rather, in those environments, many of the people one encounters express anger because they have a momentary conscious feeling of anger.

In an ordinary environment, emotional expressions are often fleeting outbursts that serve to vent conscious feelings. They do not indicate any long-term disposition, or intention, or plans on the part of the person who looks angry, or sad, or happy. So, we do not need to fear, or take precautions against people that express anger. The majority of these people either have our best interests in mind, or are completely innocuous. This is why in order to maintain positive relationships with friends and family members, and to craft new positive social relationships, one must learn to respond appropriately to angry facial expressions. One can form a lasting friendship with someone who expresses her feeling of anger, but only if one does not respond to the person as if she is some kind of threat. In a sense, making new friends, and keeping old ones, means learning to overlook angry faces.¹⁴

References

- Caspi, A., McClay, J., Moffitt, T.E., Mill, J., Martin, J., Craig, I.W., Taylor, A. and Poulton, R., 2002. Role of genotype in the cycle of violence in maltreated children. *Science*, 297(5582), pp.851-854.
- Fox, E., Lester, V., Russo, R., Bowles, R.J., Pichler, A. and Dutton, K., 2000. Facial expressions of emotion: Are angry faces detected more efficiently? *Cognition & Emotion*, 14(1), pp.61-92.
- Frick, P. J., & Viding, E. 2009. Antisocial behavior from a developmental psychopathology perspective. *Development and psychopathology*, 21(4), pp. 1111-1131.
- Frijda, N. H., & Mesquita, B. 1994. The social roles and functions of emotions. In S. Kitayama & H. R. Markus (eds.), *Emotion and culture: Empirical studies of mutual influence*. Washington, DC: American Psychological Association, pp. 51-87.

¹⁴ I would like to express my gratitude to Eamon McCrory and Essi Viding for offering many insightful comments on an earlier version of this chapter. I'm sure there are parts of this chapter with which they will disagree but I believe that my proposed explanation is in accordance with the spirit of their second hypothesis. Thanks also to Anil Gomes and Josh Shepherd for several helpful conversations on the topic of this chapter and to the participants of the London Mind Group.

- Gilbert, R., Widom, C. S., Browne, K., Fergusson, D., Webb, E., & Janson, S. 2009. Burden and consequences of child maltreatment in high-income countries. *The Lancet*, 373(9657), pp. 68-81.
- Goldstone, R. L. 1998. Perceptual learning. *Annual review of psychology*, 49(1), pp. 585-612.
- Gomes, A. forthcoming. Perception, evidence, and our expressive knowledge of others' minds. In A. Avramides and M. Parrott (eds.), *Knowing Other Minds*. Oxford: Oxford University Press.
- Green, M. 2007. *Self-Expression*. Oxford: Oxford University Press.
- Haidt, J., 2003. The moral emotions. *Handbook of affective sciences*, 11, pp.852-870.
- Hagemann, N., Schorer, J., Cañal-Bruland, R., Lotz, S., & Strauss, B. 2010. Visual perception in fencing: Do the eye movements of fencers represent their information pickup? *Attention, Perception, & Psychophysics*, 72 (8), pp. 2204-2214.
- Kahana-Kalman, R., & Walker-Andrews, A. S. 2001. The role of person familiarity in young infants' perception of emotional expressions. *Child development*, 72 (2), pp. 352-369.
- Keltner, D. & Haidt, J., 1999. Social functions of emotions at four levels of analysis. *Cognition & Emotion*, 13 (5), pp.505-521.
- Kestenbaum, R., & Nelson, C. A. 1990. The recognition and categorization of upright and inverted emotional expressions by 7-month-old infants. *Infant Behavior and Development*, 13 (4), pp. 497-511.
- Kuppens, P., Van Mechelen, I., & Meulders, M. 2004. Every cloud has a silver lining: Interpersonal and individual differences determinants of anger-related behaviors. *Personality and social psychology bulletin*, 30 (12), pp. 1550-1564.
- LeDoux, J.E., 2000. Emotion circuits in the brain. *Annual review of neuroscience*, 23 (1), pp.155-184.
- Lerner, J.S., Goldberg, J.H. and Tetlock, P.E., 1998. Sober second thought: The effects of accountability, anger, and authoritarianism on attributions of responsibility. *Personality and Social Psychology Bulletin*, 24 (6), pp.563-574.
- Lyons, D.M., Parker, K.J., Katz, M. and Schatzberg, A.F., 2009. Developmental cascades linking stress inoculation, arousal regulation, and resilience. *Frontiers in Behavioral Neuroscience*, 3, p.32.
- Martin, M. G. F. 2010. What's in a look. In B. Nanay (ed.), *Perceiving the World*. Oxford: Oxford University Press. pp. 160–225.
- McCrory, Eamon J., Stéphane A. De Brito, Philip A. Kelly, Geoffrey Bird, Catherine L. Sebastian, Andrea Mechelli, Sophie Samuel, and Essi Viding. 2013. Amygdala activation in maltreated children during pre-attentive emotional processing. *The British Journal of Psychiatry* 202 (4), pp. 269-276.
- McCrory, E. J., De Brito, S. A., Sebastian, C. L., Mechelli, A., Bird, G., Kelly, P. A., & Viding, E. 2011a. Heightened neural reactivity to threat in child victims of family violence. *Current Biology*, 21 (23), pp. R947-R948.

- McCrory, E. J. De Brito, S. A., and Viding E. 2011b. The impact of childhood maltreatment: a review of neurobiological and genetic factors. *Frontiers in Psychiatry*, 2, p.48.
- McCrory, E., De Brito, S. A., & Viding, E. 2012. The link between child abuse and psychopathology: a review of neurobiological and genetic research. *Journal of the Royal Society of Medicine*, 105 (4), pp. 151-156.
- McCrory, E. J., & Viding, E. 2015. The theory of latent vulnerability: Reconceptualizing the link between childhood maltreatment and psychiatric disorder. *Development and psychopathology*, 27 (2), pp. 493-505.
- McCrory, E. J., Gerin, M. I. & Viding, E. 2017. Childhood maltreatment, latent vulnerability and the shift to preventative psychiatry – the contribution of functional brain imaging. *Journal of Child Psychology and Psychiatry*, 58 (4), pp. 338-357.
- McNeill, W. E. S. forthcoming. Expressions, looks, and others' minds. In A. Avramides and M. Parrott (eds.), *Knowing Other Minds*. Oxford: Oxford University Press.
- Millar, A. 2000. The scope of perceptual knowledge. *Philosophy*, 75, pp. 73–88.
- Nussbaum, M. C. 2015. Transitional anger. *Journal of the American Philosophical Association*, 1(01), pp. 41-56.
- Öhman, A., 2002. Automaticity and the amygdala: Nonconscious responses to emotional faces. *Current Directions in Psychological Science*, 11(2), pp.62-66.
- Öhman, A., 2009. Of snakes and faces: An evolutionary perspective on the psychology of fear. *Scandinavian Journal of Psychology*, 50(6), pp.543-552.
- Öhman, A., Lundqvist, D. and Esteves, F., 2001. The face in the crowd revisited: a threat advantage with schematic stimuli. *Journal of Personality and Social Psychology*, 80(3), p.381.
- Parrott, M. 2017. The Look of Another Mind. *Mind*, 126(504), pp. 1023-1061.
- Pearl, J. 2009. *Causality*. Cambridge University Press.
- Pollak, S. D., Cicchetti, D., Klorman, R., & Brumaghim, J. T. 1997. Cognitive brain event-related potentials and emotion processing in maltreated children. *Child Development*, 68, pp. 773–787.
- Pollak, S. D., & Kistler, D. J. 2002. Early experience is associated with the development of categorical representations for facial expressions of emotion. *Proceedings of the National Academy of Sciences*, 99(13), pp. 9072-9076.
- Pollak, S. D., Messner, M., Kistler, D. J., & Cohn, J. F. 2009. Development of perceptual expertise in emotion recognition. *Cognition*, 110(2), pp. 242-247.
- Pollak, S. D., & Sinha, P. 2002. Effects of early experience on children's recognition of facial displays of emotion. *Developmental psychology*, 38(5), p. 784.
- Rime, B., Mesquita, B., Boca, S., & Philippot, P. 1991. Beyond the emotional event: Six studies on the social sharing of emotion. *Cognition & Emotion*, 5(5-6), pp. 435-465.

Sell, A., Tooby, J. and Cosmides, L., 2009. Formidability and the logic of human anger. *Proceedings of the National Academy of Sciences*, 106(35), pp.15073-15078.

Stein, M., Simmons, A., Feinstein, J., & Paulus, M. 2007. Increased amygdala and insula activation during emotion processing in anxiety-prone subjects. *American Journal of Psychiatry*, 164, pp. 318–327.

Whalen, P.J., Shin, L.M., McInerney, S.C., Fischer, H., Wright, C.I. and Rauch, S.L., 2001. A functional MRI study of human amygdala responses to facial expressions of fear versus anger. *Emotion*, 1(1), p.70.

Wollheim, R. 1999. *On the Emotions*. New Haven: Yale University Press.

Woodward, J. 2008. Mental causation and neural mechanisms. In J. Hohwy and J. Kallestrup (eds.), *Being reduced: New essays on reduction, explanation, and causation*. Oxford: Oxford University Press, pp. 218-262.

Woodward, J. 2005. *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.