LSE Research Online

THE LONDON SCHOOL OF ECONOMICS AND POLITICAL SCIENCE

Eileen Munro and Kristina Musholt

Neuroscience and the risks of maltreatment

Article (Accepted version)
(Refereed)

Original citation:
Munro, Eileen and Musholt, Kristina (2014) *Neuroscience and the risks of maltreatment*. Children and Youth Services Review, 47 (1). pp. 18-26. ISSN 0190-7409
DOI: 10.1016/j.childyouth.2013.11.002

© 2013 Elsevier Ltd

This version available at: http://eprints.lse.ac.uk/55179/
Available in LSE Research Online: November 2014

LSE has developed LSE Research Online so that users may access research output of the School. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. Users may download and/or print one copy of any article(s) in LSE Research Online to facilitate their private study or for non-commercial research. You may not engage in further distribution of the material or use it for any profit-making activities or any commercial gain. You may freely distribute the URL (http://eprints.lse.ac.uk) of the LSE Research Online website.

This document is the author's final accepted version of the journal article. There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

http://eprints.lse.ac.uk

Neuroscience and the risks of maltreatment

Eileen Munro and Kristina Musholt

Abstract

Findings from neuroimaging are increasingly being cited in policy debates to strengthen the case for early identification of, and intervention with, children at risk of maltreatment and poor outcomes. While agreeing that neuroscientific research into the risks of maltreatment is a very valuable and exciting area of study, this article challenges the confidence with which these findings are used in policy discussions. It critically discusses the reliability and validity of the relevant findings and the contribution they can currently make to our understanding of the causes and consequences of maltreatment. In addition, it is argued that this type of evidence, which is new in policy debates, is often being used in ways that are problematic. Many participants in the relevant policy debates seem to subscribe either to an implicit version of dualism about the relationship between the mind and the body, or to reductionism – the view that the mental can be reduced to the physical. Such assumptions threaten the way we think about human agency and moral responsibility but it is argued that they are misguided for conceptual reasons. It is concluded that neuroscience has the potential to contribute to our understanding of the causes and effects of maltreatment but cannot do so in isolation from the social sciences.

1. Introduction

Societies have long been concerned with reducing the number of children who suffer maltreatment in the family home. In the past, services have tended to become involved only when maltreatment has occurred, but there is increasing interest in intervening at an earlier stage, both to prevent maltreatment occurring and to prevent minor problems escalating to a serious level. This interest is strengthened by research findings that show that children who suffer from neglect, the largest sub-category of maltreatment, are at increased risk of poor outcomes in education, health, and behaviour, especially when combined with socio-economic disadvantage (HM Treasury 2003; Lindsey 2004; Sabates and Dex 2012) . There is clearly a moral argument for preventing harm and the possibility of doing so is now greater, as research finds evidence that a number of service interventions are producing positive results for children (Barlow and Schrader-MacMillan 2009; MacMillan, C, Barlow, Fergusson, Leventhal, and Tausig 2008).

Nonetheless, making the political and economic case for early intervention at a time of economic austerity is proving difficult. Inevitably, early intervention services are more costly than reactive services, at least in the short-term, since a larger number of families have to be offered a service. While research gives us some indicators of increased risk, the complexity and individuality of human development limits our ability to predict precisely which children will suffer harm. It is argued that prevention will lead to savings later on, so that such services are likely to be very cost-effective in the longer term. However, while there is considerable political interest in the policy of early intervention there is limited funding being made available. At the same time, poor economic growth and policies such as welfare

reform are potentially increasing the number of children at high risk of poor outcomes.

In this context, the findings from neuroscience research on the neurological impact of maltreatment are being seized on with enthusiasm as showing hard and compelling evidence of the damage children suffer. Interest in neuroscientific research in child welfare is linked both to questions relating to ways of reducing or preventing maltreatment, and to questions about the effective treatment of abused or neglected children. For example, Graham Allen's report for the UK Government (Allen 2011) has a copy of a CT scan on the front cover to emphasise the message of how seriously children are harmed by poor experiences in their early years and hence the importance of early intervention. The Harvard University Center on the Developing Child claims that "for children at unusually high risk neuroscience provides a compelling argument for beginning programs at birth if not prenatally"(Center on the Developing Child 2007 p.3). McCrory et al (2010) suggest that advances in neuroscience and genetics are "rapidly changing how we view early adversity, creating a neuro-biologically informed developmental narrative that has the potential to change social policy, societal perceptions of harm and the conceptual framework within which we think about clinical intervention and prevention". As O'Connor et al. (2012, p.225) suggest: "The brain has been instantiated as a benchmark in public dialogue and reference to brain research is now a powerful rhetorical tool". The case for providing early help to families to improve children's safety and development is compelling for a number of reasons, but this article aims to critically assess the contribution to the relevant public policy debates that neuroscientific evidence, and in particular evidence from neuroimaging research, is making at present. We are focusing on neuroimaging in

particular because this kind of research has a tendency to elicit particularly strong reactions due to the images it is producing and has recently received a lot of attention. Of course, much could be said about other neuroscientific approaches as well.[1]

Given the extensive social-psychological knowledge that already exists about the detrimental effects of abuse, the question arises as to whether neuroscience is adding something genuinely new to the discussion. We will argue that, at present, rather than providing us with new insights into *causal* mechanisms, neuroimaging studies present us with information about the neural *correlates* of the effects of child maltreatment. As we will explain in more detail, this can be an important first step towards the identification of the brain mechanisms involved in mediating the effects of maltreatment and hence towards the development of corrective or reparative interventions that target these mechanisms. However, the latter requires an interdisciplinary approach, and the finding of correlates as such is neither sufficient to provide new insights into causal mechanisms, nor should it be surprising – unless one is harboring dualist intuitions. Nonetheless, and despite the fact that at present results from neuroimaging research at best confirm what is already known from the social sciences, as pointed out above, much weight is currently put on these studies. We suggest that this is due to an implicit bias towards so-called 'hard sciences', which are somehow seen as providing better, more reliable evidence than 'soft sciences', such as the social sciences. However, this is problematic for various reasons including methodological problems, questions of validity, and the (implicit or explicit) reliance of neuroscience on social science research.

[1] For a critical discussion of the – often problematic – use of neuroscience in social policy discussions in more general terms, see Wastell and White (2012).

Moreover, neuroscience provides a type of evidence which is new to policy makers and which, due to the underlying assumptions that tend to be associated with this kind of evidence, has the potential to impact discussions about human agency and determinism in policy and practice debates in fundamental ways. This potential will be critically examined in the second half of this article. In particular, there is a tendency to ascribe a diminished level of responsibility to persons with altered brain structure or functionality. However, this threatens to undermine the persons' agency even though, as we will argue, the evidence to date does not warrant such a judgment.

We begin by examining the reliability and validity of structural Magnetic Resonance Imaging (MRI) and functional MRI (fMRI) scans (Section 2.1). There is a danger that people will overvalue findings from an unfamiliar but high status discipline such as neuroscience, so we go into some detail on how evidence is produced by neuroscientists. This draws attention to the number of hypotheses (many contested) and statistical analyses involved in producing the final MRI/fMRI picture. Note that this section is directed towards readers who are not familiar with neuroimaging techniques; readers with knowledge of this area of research may want to skip this section. We then (Section 2.2) critically appraise the neuroimaging studies that have looked at children who have been maltreated, examining what they currently tell us about the impact of childhood maltreatment on brain development. In so doing, we will also highlight some more specific methodological concerns regarding the existing imaging studies concerning children who have suffered from maltreatment.

Based on this, in Section 3, we critically discuss how this type of evidence influences debates about how to improve children's development and help the victims of

maltreatment. We analyse the assumptions implicit in current discussions about the relationship between the mind and the body, the disputed views on agency and moral responsibility, and the implicit assumptions in discussions about the relative status of the findings of the natural and the social sciences. In particular, we will question the claim that neuroscience provides independent evidence of the harm of maltreatment by showing how interpretations of the findings are reliant on existing social science research.

Nothing in this article should be taken as an attack on neuroscience per se. Neuroscience can add important and valuable information. A better understanding of the neurological processes underlying certain behaviours is both valuable in its own right and is beginning to produce some insights that contribute to better identification and response to problems. Nonetheless, it is important to be aware of the current technical as well as principled limitations of MRI/fMRI studies in order to be able to properly assess the data provided by this research. Moreover, as we will argue, in the context of a better understanding of the effects of maltreatment on children, insights from neuroscience are at their most valuable and useful when combined with insights from the social sciences (see Section 4). Thus, this article should be read as a plea for a more integrated and interdisciplinary approach to the study of maltreatment of children.

2. How reliable are the results from neuroimaging research ?

To a non-specialist it might seem as though brain scans provide us with something akin to a photograph of the brain, allowing direct insights into the brain's structure and

inner going-ons. Photographs can provide vivid and reliable evidence. A photograph of a bruise on a child's skin is treated with respect in a legal hearing. Even an x-ray of a fractured bone has some of the qualities of a photograph and conveys some of the same reliability. In the 1960s, x-ray evidence of fractured bones played a significant role in raising public concern about the prevalence of physical abuse. For some it seems that MRI and fMRI scans can play a similarly powerful role in current policy and practice debates. However, the images we find in the neuroscience literature are in fact far from being such simple records of damage or healthy development.[2] They are the results of a complex computational and statistical process and their interpretation requires considerable skills; they do not provide direct measures of neural activity. In being indirect measures, they do not differ from many other areas of study in child welfare: measures of individual traits, social constructs and emotional states are all of this kind. However, the apparent tendency in the literature to think of them as akin to photographs leads to a tendency to over-estimate their reliability. Moreover, neuroscience, and in particular neuroimaging research, is a relatively young and still developing field. While it generates a lot of excitement, it also shows teething problems with regard to the reliability and interpretation of its results. These need to be understood if the true potential of neuroimaging studies to inform our understanding of complex psychological and social phenomena is to be appreciated.

In the following, we will provide an overview of the underlying principles of neuroimaging research and its limits. The aim is to contribute to an informed discussion among researchers and policy makers who are not familiar with this

[2] Though note that there can be problems with reliability in reading x-rays, photos, and graphs, too. For example, for x-rays, trained and untrained observers may draw different conclusions.

technique as to how to properly assess neuroimaging results in the context of public policy debates regarding the consequences of the maltreatment of children. We shall begin by explaining the difference between a structural and a functional MRI (fMRI) and by presenting a very brief and strongly simplified description of the basic technical principles underlying structural and functional MRI research. A more detailed description can be found in Logothetis (2008, supplementary material), as well as in numerous textbooks on the topic (e.g. Buxton, 2002; Wood & Wehrli, 1998). We will then highlight some of the methodological problems associated with neuroimaging research; first in terms of the general challenges and limits of its methodology and then with a more specific focus on existing neuroimaging studies involving children who have suffered from abuse.

2.1 On the general reliability and validity of MRI studies

Magnetic Resonance Imaging (MRI) is based on the principle that atomic nuclei in the body (such as protons in fat and water) possess magnetic properties called spin. These spins can be aligned with the help of a magnetic field. The alignment can then be systematically altered through the application of radio frequency fields. When these are turned off, the spins of the protons return to their thermodynamic equilibrium. This process is called relaxation. The energy that is being released during relaxation can be measured. Protons in different tissues return to their equilibrium state at different relaxation rates. This allows for the generation of a contrast between different types of body tissue, which can be utilized for structural brain scans .

In the case of structural MRI, relaxation-rate differences at different sites can be used to generate anatomical contrast, as nuclei located in different parts of the brain will emit different signals during relaxation. Thus, MRI studies are able to demonstrate gross structural changes, such as shrinkage or growth of certain areas.

Similarly, in the case of functional MRI, changes over time in one or more relaxation rates in a single brain area can be used to image changes in the state of this area over time. This allows for the correlation of neural activity over time with behavioral measures, such as performance in a cognitive task. Put differently, the development of functional MRI scans allows for an observation of the brain 'in action'. The most common fMRI measurements rely on changes in relaxation times brought about by differences in blood oxygenation between test and control time intervals. Changes in blood flow, in turn, are thought to be related to neural activity. (Logothetis 2008) Thus, fMRI experiments rely on indirect measures of properties (such as changes in blood oxygenation levels and the altered magnetic properties related to these) that are thought to be correlated with mass neural activity. They do not provide direct measures of neural activity. Moreover, the relevant signal must be detected from a sea of noise, and various computational steps must then be performed to create the final image based on the signal.

This raises certain questions concerning the physiological assumptions underlying fMRI research, as well as questions with regard to the reliability of fMRI research. With regard to underlying physiological assumptions, as we have just explained, fMRI signals typically reflect changes in blood flow in a particular region of the brain, which, in turn, are thought to be associated with neural activity. However, even

the smallest measurable fMRI unit (or voxel) represents hundreds of thousands of different types of cells, not all of which are yet understood (Cacioppo & Decety, 2011; Logothetis, 2008). Moreover, in the past it was assumed that increased blood flow into a region meant that this specific region had just been highly active and had signaled other connected regions. This would imply that increased blood flow points to the *output* of a region. However, more recent studies of blood oxygen-dependent (BOLD) activity suggest that it reflects an increase in the *inputs* to a particular region from other areas of the brain (Amodio, 2010). Further, it has been pointed out that:

> "The fMRI signal cannot easily differentiate between function-specific processing and neuromodulation, between bottom-up and top-down signals, and it may potentially confuse excitation and inhibition. The magnitude of the fMRI signal cannot be quantified to reflect accurately differences between brain regions, or between tasks within the same region." (Logothetis, 2008, p. 877).

Thus, neuroimaging studies can only be said to study neural activity in a gross sense, and it is far from clear how exactly the signals generated by fMRI scanners are to be interpreted.

With regard to the reliability of fMRI studies, the basic problem confronted by researchers using fMRI consists in creating an acceptable signal-to-noise ratio. There are various sources of noise that make the acquisition of a good signal problematic, including thermal noise, system noise in the scanner, physiological/movement noise from the subject, non-task related cognitive processes, or changes in cognitive strategy over time (Bennett & Miller 2010). Additionally, the images produced by

brain scanners are influenced by a number of hardware variables, such as the strength of the primary magnetic field (which, in turn, can be correlated with artifacts), and acquisition parameters, such as voxel size, repetition time, and others. Thus, different hardware conditions as well as decisions made by researchers regarding data analysis will impact results and can differ substantially between different labs, causing difficulties for the replicability of studies as well as for estimates of reliability. To make matters worse, while there are a variety of methods to estimate the reliability of fMRI results, these focus on very different aspects of reliability and there is no consensus within the scientific community as to what constitutes an acceptable level of reliability for fMRI studies (Bennett & Miller, 2010) .

In addition, there have been problems with regard to the statistical analyses involved in the interpretation of imaging studies. For example, many studies involve the identification of regions of interest (ROI), that is, regions that appear to be most active during certain conditions. Problems can occur when the same data set is used both for selection of ROIs and for the analysis of changes during the experimental condition in question, as in this case the results statistics are not independent from the selection criteria (Kriegeskorte, Simmons, Bellgowan, & Baker, 2009). Moreover, neuroimaging studies generate huge amounts of data, and researchers often perform multiple statistical tests over the same range of data. This can lead to 'false positives', that is, results that appear statistically significant purely by chance (Bennett, Wolford, & Miller, 2009). While there are ways to correct for both types of problem, this is currently not routinely done (Bennett, Wolford, & Miller, 2009; Vul, Harris, Winkielman, & Pashler, 2009).[3]

[3] Many thanks to Julia H. Littell for alerting us to these statistical problems.

While it is reasonable to assume that many of these problems will be resolved in the future, as the technology develops and with the help of a combination of fMRI results with other techniques, it is important to understand that at present the technology is still very much in a stage of development.

2.2 On the reliability and validity of maltreatment studies

While the concerns mentioned above refer to imaging studies (in particular functional imaging studies) in general, there are more specific concerns regarding existing MRI and fMRI studies of children suffering from abuse and neglect.[4] For instance, a recent critical review of neuroimaging studies addressing the effects of child abuse (Hart & Rubia, 2012) found that while several structural imaging studies pointed to brain regions affected by childhood maltreatment, in particular the prefrontal cortex (which is thought to be involved in many executive control tasks), the hippocampus (which plays an important part in learning and memory), the amygdala (which is thought to be crucial for the processing of emotions, threat assessment, fear conditioning and behavioral regulation), the corpus callosum (which connects the two hemispheres and facilitates interhemispheric communication), and the cerebellum (which is part of the neural network that fine-tunes behaviour), there are major problems in interpreting the results of these studies. Most of the studies imaged subjects with associated psychiatric conditions, including Post Traumatic Stress Disorder, Major Depressive Disorder, Attention Deficit Hyperactivity Disorder, Oppositional Defiance Disorder,

[4] The following discussion is based largely on a recent critical review by Hart & Rubia (2012). As this review is not commonly referenced in the child welfare literature, it seems worthwile to us to go into quite a bit of detail here with regard to the content of this review.

Borderline Personality Disorder, Obsessive Compulsive Disorder, generalized anxiety disorder, phobias, eating disorders and substance abuse. This is problematic, because it makes it impossible to determine whether the changes in brain volume that were found are a result of the abuse, the psychiatric conditions, or an interaction between the two (Hart & Rubia, 2012). Moreover, the studies in question do not control for socio-economic variables, such as poverty. This is problematic, because poverty has similar impacts on child development and is overrepresented in individuals that are identified as maltreated (Cawson, Wattam, Brooker, & Kelly, 2000; Lindsey, 2004). In other words, the results of these studies don't give us clear information regarding the neural correlations of child abuse *per se*.

In addition, many of the structural MRI studies worked with small sample sizes (often less than 20 subjects per group). This is problematic because there can be large differences between individuals, so that the reliability of studies with small sample sizes is questionable. Moreover, even for studies with large sample sizes, inter-individual differences undermine any simple correlation between experience and an individual's brain development, in line with social and psychological knowledge of how varied the life course can be. While it might be possible to generalize that groups of children with a history of abuse may contain more subjects with an adverse factor than the average for the population as a whole, within those broad generalisations, there is considerable variation (Sabates & Dex, 2012). This degree of variation holds not only for structural MRI studies, but also for functional MRI studies. In fact, as pointed out by Bennett and Miller (2010), it is known that there are large and stable differences between individuals on almost any cognitive task, thus making this particular concern even stronger in the case of functional MRI studies.

Turning to fMRI studies in the area, again, changes in the activation of several brain regions have been reported for subjects who have suffered from maltreatment (surveyed by Hart and Rubia 2012). In particular, again, studies found altered activation of prefrontal regions; including those implicated in inhibition and working memory tasks, as well as those implicated in non-executive function tasks. Other studies found weaker responses to reward cues in the basal ganglia, and altered activation in the insula and cerebellum. The former is associated with diverse functions, such as perception, motor control, self-awareness, cognitive function and emotion processing. The latter, as mentioned above, is implicated in the fine tuning of behavior. However, again, functional neuroimaging studies of childhood maltreatment suffer from insufficient control of confounding co-morbid psychiatric conditions, such as PTSD (Hart & Rubia, 2012). Moreover, again, sample sizes in the studies that were surveyed were often small, and most studies tested adults with a history of childhood maltreatment, rather than abused children. In fact, according to the review by Hart & Rubia (2012) to date there have only been three fMRI studies with abused children. However, “it is important that both adults and children with histories of maltreatment are studied as brain changes can develop or normalise over time and the potential effects are likely to be very different in adults compared to adolescents [...]” (Hart & Rubia, 2012, p. 17). Further, to date there don't seem to be any longitudinal neuroscientific studies of subjects suffering from maltreatment to see if they produce similar findings to the social science studies that show that, for a large sub-group, the damage is enduring. While studies with adults who have a history of childhood maltreatment do point towards the enduring nature of the effects of maltreatment, due to the fact that many of the adults who were studied also suffered from co-morbid

disorders, it is impossible to pin down the precise effects of maltreatment *per se* on the basis of these studies. Nor do there seem to be studies of the impact of specific interventions. These gaps in the research make it very difficult to assess present results in terms of their prognostic potential and/or to use these results in order to draw conclusions about the efficacy of different types of intervention. While neuroscience has the potential to make a valuable contribution in this area, it is still only hypothetical.

Another concern regards the fact that many studies have used subject samples with very different maltreatment histories, including physical abuse, sexual abuse, emotional abuse, neglect and the witnessing of domestic violence. This is problematic because it is known that different forms of maltreatment present differently clinically, which suggests that they also have different effects on behavior and neurobiology (Hart & Rubia, 2012, p.18).

Finally, it should be mentioned that most of the structural MRI, but also fMRI studies have not tested for whole brain effects, but have rather focused on the regions of the brain that neuroscientists currently think are associated with the variables of interest, that is on a priori regions of interest (ROI). As Hart & Rubia (2012) point out, "this is a limited approach as we do not yet have extensive knowledge of all brain areas affected by childhood maltreatment and it is possible that ROI studies are not focusing on the most affected brain regions. This is even more problematic since most ROI studies were based on prior studies of relatively small sample sizes that furthermore have included high rates of comorbidities" (Hart & Rubia, 2012, p.18). We will discuss this problem in more detail in the next section.

Accordingly, although there is a growing body of evidence – both from neuroimaging studies and from other research in developmental neuroscience, including research into genetics -- showing that adverse life experiences in early childhood are likely to have damaging and potentially long-lasting effects on brain development, it has to be said that at present both MRI and fMRI studies on childhood maltreatment need to be treated with great care and are best seen as preliminary. While, as we pointed out above (and as we will explain in more detail in the next section), neuroimaging studies can provide an important first step towards the development of better interventions by narrowing down the brain areas and subsequently the mechanisms associated with the mediation of the effects of maltreatment, the results achieved to-date are still too imprecise and unreliable to be able to contribute to the development of new interventions in a significant way. Future studies that address the problems mentioned above are needed in order to gain more specific and reliable results.

3. The relationship between neuroscience and social science

3.1 The mind-body relationship

Having raised some concerns regarding the reliability of MRI studies, both in general and with a more specific focus on existing imaging studies of the effects of childhood maltreatment, we are now going to expose and challenge some of the assumptions that are implicit in the current literature as this type of evidence is brought into policy and practice discussions.

It seems that despite the limitations and concerns regarding neuroimaging studies mentioned in the previous section, many parties in the discourse show an (implicit or explicit) tendency to put stronger weight on neuroscientific evidence compared to data from the social sciences. Fonagy (2012) for example, praises an English charity because they have "nobly called for empirical researchers [neuroscientists] to take up and test its clinical insights in an attempt to offer them greater scientific validity" (Fonagy, 2012, p.97), suggesting that he takes other kinds of research, such as social scientific research, to be less valid. Weisberg et al (2008) also illustrate this tendency to be impressed by neuroscience in an experiment in which subjects without expertise in neuroscience were shown to take explanations that are accompanied by irrelevant neuroscientific information to be more satisfying than the same explanations in the absence of the irrelevant information, even when the explanations in question were manifestly bad.[5]

The fact that data from neuroscience research is seen as 'hard' evidence, compared to 'soft' social data, suggests an implicit dualism and a cultural bias towards ascribing higher status to the natural over the social sciences. As we will show in what follows, this is not only problematic for conceptual reasons, or because neuroscientific data has the potential to 'mask' bad explanations (as demonstrated by Weisberg et al. 2008), but also because of the implications for the way in which we think about agency, free will and determinism.

Looking at some of the recent public policy debates about the contribution of neuroscientific research one can get the impression that it is often taken to be an

[5] Many thanks to Julia H. Littell for alerting us to this evidence.

important, if not surprising, insight that psychological processes are indeed correlated with processes in the brain. This implies that the relevant parties to the debate are assuming some version of Cartesian dualism, that is, the view that the mental and the physical are fundamentally different types of substance. For example, Teicher (2000) criticises the view that the long-term difficulties arising from childhood maltreatment are brought about through psychological processes, offering as an alternative explanation the findings of neuroscience on how early maltreatment, "has enduring negative effects on brain development…we are beginning to understand how these abnormalities may account directly for the personality traits and other symptoms that patients manifest" (Teicher, 2000, p.50). This seems to suggest that he sees psychological and neuroscientific explanations as being in opposition to each other.

However, while this sort of dualism seems to be a common sense view, it raises serious problems in explaining how the two types of substance interact. How can an immaterial mind interact with a material body? Descartes posited that they interacted via the pituitary gland (an organ with no known function when he was writing) but this now seems implausible and does not, in any case, add much to our understanding of *how* the interaction occurs. The difficulties posed by dualism lead most philosophers and scientists to reject this position and instead accept the position that every psychological process is physically instantiated. This is not to say that psychological processes can be fully explained in terms of neuroscience. To the contrary, it is far from obvious how precisely descriptions at the psychological level relate to those at the physical level (we will return to this point below)[6]. But based on

[6] This is why many philosophers have adopted the position of "nonreductive materialism". On this view, there are no non-physical substances in the world, but this is not say that every phenomenon can be explained in terms of physics. This is partly because the concepts involved in, say, philosophy

a rejection of dualism, the correlations now being found between adverse experiences and processes in the brain are precisely what we should expect. So given that we already know, based on decades of social science and clinical research, that early maltreatment can have enduring and developmentally course-altering negative effects on people, it is only to be expected that these effects are reflected in the brain. In this sense, results from neuroscience research into the effects of maltreatment to date are simply confirming what is already known. Thus, to simply show that the brain is somehow involved in a psychological process does not yet constitute a major theoretical advance.

This is not to say that there is nothing to be learned from neuroscientific studies. Indeed, as we will see in a moment, a better understanding of the neurological processes correlated with known psychological processes may lead to innovative ideas and to the development of better interventions. But the insights gained from neuroscience research cannot simply consist in the fact that psychological processes have a physical counterpart in the brain. Nor can insights from the neurosciences be gained in isolation from findings in the social sciences. Rather, neuroscience can only contribute to a better understanding of the effects of maltreatment when it is seen in relation to the knowledge that we have from the social sciences.

In order to see this, we need to consider in more detail the different ways in which neuroimaging can be used. According to Amodio (2010), there are two main approaches in the context of social cognitive neuroscience. On the one hand, neuroimaging can be used for brain mapping, that is, in order to aid a better

---

or the social sciences on the one hand, and the natural sciences on the other have been devised for very different explanatory purposes and thus cannot be easily mapped onto or integrated with each other.

understanding of how the brain works. This kind of research is based on the correlation between the data obtained with the help of fMRI and behavioural and psychological data. It is used to help us understand the function of a given brain region.[7] On the other hand, neuroimaging can be used for psychological hypothesis testing, that is, to test hypotheses about the mechanisms underlying certain psychological processes, or the relations between different psychological processes. This approach starts from an assumption that a particular psychological process is associated with the activation of a particular brain area. To take an example from Amodio (2010), it is relatively well established that the amygdala is involved in fear conditioning (i.e. a form of learning in which an adverse stimulus becomes associated with a previously neutral context or stimulus). Now, social psychologists might develop the hypothesis that implicit racial bias is based on mechanisms of classical fear conditioning. They could then test this hypothesis by measuring activity in the amygdala while participants engage in a behavioral task that measures implicit racial bias. [8] If the amygdala were found to be active, researchers could link their knowledge about implicit racial bias with the knowledge about the mechanisms involved in fear conditioning, and develop new hypotheses for how implicit racial biases are acquired and potentially reduced. In this approach, the function of the brain region in question is already understood, and the research question focuses on experimental effects on psychological variables. This way, neuroimaging can help to inform psychological

[7] Notice though, that this "brain mapping approach" is problematic in itself. (Thanks to Julia H. Littell pointing this out to us.) For instance, the fact that a particular area of the brain is active during a particular task does not imply that the functions involved in solving this task are therefore located in that area (cf. Miller 2010). Rather, brain activity tends to be broadly distributed across the entire brain, and any one particular area can be involved in many kinds of different functions. One must also be careful to avoid ascribing properties or functions to parts of the person (such as parts of her brain) that can only be meaningfully ascribed to the whole person as such. Otherwise one commits a category error or "mereological fallacy" (Bennett and Hacker 2003).

[8] But notice that many tasks that are employed in measuring implicit bias do not appear to tap the affective dimensions of implicit bias, and so might not relate to amygdala activity.

theories, by adding knowledge about the mechanisms involved in certain psychological processes.

Based on such knowledge it might be possible to both develop better treatments or preventions. Moreover, this kind of knowledge might even lead to conceptual changes with respect to the way we think about certain psychological abilities or disorders. For example, we now have a reasonably detailed understanding of the neural mechanisms underlying addiction. This has lead both to a better understanding of what constitutes addiction, as well as to improved treatments (e.g., Colombo 2012, Shea 2013). Similarly, as we will see below, a better understanding of the brain mechanisms underlying our individual abilities for cognitive control might contribute to a more nuanced understanding of such notions as agency and the ways in which our abilities to exercise agency can be impaired.

In the study of maltreatment, as our knowledge of the cognitive functions of the brain grows, it may well be that in the future neuroimaging studies could, for example, provide insights into the involvement of brain areas (and the cognitive functions fulfilled by these brain areas) in behavioral changes associated with maltreatment that were hitherto unsuspected, or that precede changes in behavior. In this way neuroscience could contribute to a more detailed understanding of the phenomena in question and support the development of new psychological hypotheses and better diagnostic tools and treatment options.

However, it is important to see that both neuroimaging approaches in social cognitive neuroscience are built on existing social and psychological research and can only be

developed further if they go hand-in-hand with continued social-scientific (and philosophical) theory building. As mentioned, the brain mapping approach relies on the correlation of behavioral and psychological data with brain activation measurements. However, this requires a detailed analysis of the psychological processes under investigation, as one cannot simply assume a one-to-one mapping between a psychological concept or process and a pattern of brain activation. This is especially so in the case of highly complex psychological constructs, such as self, emotions, agency, etc. Such high-level concepts are difficult to define, yet without a clear understanding of one's psychological constructs, one cannot draw valid inferences about whether or how they map onto neurological structures or processes (Gillihan & Farah, 2005). Therefore, what is required for this kind of research is the careful conceptual analysis of the psychological phenomenon or process one wants to investigate, and the development of behavioral measures that provide valid manipulations of the psychological process in question and produce reliable and interpretable behavioral evidence (Amodio, 2010; Caicippo & Decety, 2011). In other words, it will require careful psychological (and philosophical) theorizing. However, note that it is not always possible to break down high-level philosophical, psychological or sociological concepts into units or processes that can be meaningfully assessed with the help of neuroimaing studies, nor is this always necessary or desirable. Often, social scientific explanations will be better suited for the explanatory purpose at hand, and given the difficulties in interpreting fMRI results that we have discussed above, behavioral measures will often provide more meaningful and more reliable results (Amodio, 2010).

The psychological hypotheses testing approach in neuroscience on the other hand relies on existing knowledge about the functional profile of different brain regions (obtained with the help of the 'brain mapping' approach). However, in addition to the general problems associated with this 'brain mapping' approach, for many brain regions this knowledge to-date remains limited, in particular with regard to children and with regard to anatomical and functional abnormalities in the brain. Moreover, brain areas can be involved in many different functions, so just because an area that is known to be involved in a particular function (such as fear conditioning, empathy or executive control) or is found to be active in a particular task (or structurally altered in a particular condition), this does not necessarily mean that it fulfills the same function here. Finally, as we have seen in the examples above, this approach will only lead to new insights if it is based on, and combined with, sound hypotheses based on psychological theorizing.

Thus, while there certainly is a potential for important contributions from neuroimaging towards a better understanding of the effects of maltreatment, in particular if it can help us to better understand the mechanisms involved in mediating the adverse effects of maltreatment, it is wrong to think that the contributions made by neuroscience are independent from or superior to insights from the social sciences. To the contrary, it is only through a careful combination of social science research and neuroscience that we can improve our understanding. This also means that, since we still know relatively little regarding the functions of specific brain areas, and given the principled limitations of imaging studies discussed above, the extensive social scientific knowledge we already have is often better suited for an understanding and explanation of the complex conditions associated with childhood abuse. This is

especially so in light of the fact that structural MRI can only provide us with a snapshot of the anatomy of the brain at a given moment in time, and that functional MRI relies on the investigation of brain activity during a limited range of cognitive tasks which can be performed in a scanner, such as the processing of the emotional pictures, or tasks related to memory or simple motor control. It is not obvious to what extent performance on these tasks, or a snapshot of someone's brain anatomy, is representative for the often complex psychological conditions associated with childhood abuse, and the various factors that contribute to or are relevant to these conditions.

3.2. Neuroscience and agency

Emphasising the close relationship between psychological and physical processes – and the corresponding need for a combination of neuroscientific research with social science research – is also important because many people seem to implicitly assume that behavior is either caused by the brain (in which case the person is not responsible for it), or intended by the person (in which case there is personal responsibility). The underlying assumption is, again, based on an implicit dualism. Thus, when confronted with findings that suggest altered structure or functionality in the brain, there is a tendency to assume that the person in question doesn't really possess agency over their behavior as a result of these changes. Accordingly, these findings are sometimes used to argue in favor of diminished moral responsibility in such cases. Put differently, neuroscientific findings about changes in the brain associated with maltreatment have the potential to change the way we think about agency and

responsibility, leading us to deny true agency (and thus moral responsibility) for victims of maltreatment.

Two recent studies provide evidence for this way of thinking about agency and moral responsibility (Aspinwall, Brown, & Tabery, 2012; Monterosso, Royzman, & Schwartz, 2005). In one study, Monterosso, Royzman & Schwartz (2005) asked participants to think about various situations in which an individual engaged in destructive behavior and then to decide whether the individual in question was to be blamed for the harm they caused. In addition, the researchers provided participants with information that either concerned the perpetrator's psychological history (for example, having suffered from abuse as a child), or their brain physiology (for instance an imbalance in neurotransmitters in the brain). They also varied the strength of connection between those factors and the behavior. The results showed significant differences in the responses to the two types of information. When given information about physiology, participants were much more willing to absolve the person in question from moral responsibility, to view their behavior as 'automatic' (rather than intentional) and to judge the behavior to be unrelated to the person's character, compared to when participants were given pyschological information about the person's history. In other words, they judged the behavior to be 'caused by the brain', rather than being intended by the agent. In contrast, while information about psychologically traumatic experiences tended to elicit sympathy, the behavior in question was still seen as intentional and the person as responsible for it. (Monterosso et al, 2005)

The fact that this kind of thinking is common is also supported by another recent study (Aspinwall et al, 2012), in which researchers surveyed 181 state trial court judges in the US and asked them to provide a verdict on a fictional case (loosely based on real events) in which a robber repeatedly strikes a restaurant manager who refused to hand over any money with a gun, resulting in brain damage for the manager. Participants were also informed that the perpetrator had been diagnosed with psychopathy. Half of them were presented with additional information from neuroscientists about the neurobiological mechanisms contributing to the development of psychopathy. The results showed that while the information regarding the psychopathy on its own was considered to be an aggravating factor overall (for instance due to the fact that the perpetrator would pose a continued threat to society) and led to longer sentences compared to normal sentences for assault, the additional information about the neurobiological mechanisms involved in psychopathology significantly reduced the sentence and significantly reduced the degree to which psychopathy was rated as aggravating. Although the judges did not agree with the explicit statement that the perpetrator had diminished moral responsibility, the introduction of neurobiological factors significantly increased the number of judges who invoked mitigating factors in their reasoning and balanced them with aggravating factors. Moreover, the mitigating factors that were mentioned often suggested that the judges did in fact consider the perpetrator to be "morally disabled" (Aspinwall et al, 2012, p.847). This was despite the fact that the neurobiological information did not change anything about the person's situation and characteristics (in particular the fact that they suffered from psychopathy), but rather simply constituted an additional piece of information regarding the neurobiological factors that are known to be correlated with psychopathy. This again suggests that people tend to rely on an implicit dualism

when reasoning about such cases – they see behavior as either caused by the brain, or as intended by the person, and consider these two processes as distinct (Monterosso et al, 2005).

However, this kind of thinking is misguided, because, as mentioned above, every psychological process is underwritten by a brain process, even though, as we have seen in the previous section, the relationship between psychological and neuroscientific levels of description and explanation is not straightforward and requires careful conceptual analysis. Thus, as Monterosso, Royzman & Schwartz (2005) point out, behavior cannot be *either* caused by the brain *or* intended by the person, as every behavior has a neurological basis. This means that when drawn to its logical consequence, the kind of thinking exemplified in the two studies above should lead us to the conclusion that all behavior is caused (as opposed to intended), with the implication that there is no free will. Indeed, some claims to this effect have been made by some neuroscientists as well as philosophers. However, while the question as to how to establish agency and responsibility in a deterministic world is a vexed and complicated philosophical issue, and while it would be beyond the scope of this paper to argue in detail for this position, we think that such a conclusion would likewise be misguided, for it would eliminate any meaningful notion of agency and moral responsibility for *all* human beings.

Rather, we should aim to distinguish between those types of behavior for which a person cannot be held responsible and those for which a person can reasonably be held responsible. Here, neuroscience can potentially be of use. For instance, if a person suffers from a traumatic brain injury or a brain tumor and subsequently

displays completely uncharacteristic behavior, this might give us reason to assume that the behavior in question is a result of the injury or tumor and that the person should not be held responsible. Likewise, a better understanding of the brain mechanisms underlying our ability to guide and control our behavior and to rationally respond to reasons, and the ways in which these can be impaired, might lead us to a better understanding of the underlying causes of an experienced inability to exert such control, and might lead to a better appreciation of the individual differences with respect to our ability to exercise agency. On the other hand, the simple fact that certain brain processes are involved in a certain type of behavior cannot in and of itself serve to absolve a person from agency and moral responsibility, for this would amount to the claim that nobody is ever responsible for anything because all our behaviors are underwritten by brain processes.

Naturally, such a distinction is not always easy to be drawn. Most cases are extremely complex, involving many different individual factors and circumstances. This is especially true for victims of abuse, which is why there is such a large variation in individual life histories. It is well known that not every victim of abuse will develop a psychiatric disorder or display problematic behaviors, yet it is also true that abuse does often have very adverse effects. Similarly, changes in brain structure or brain activation will not always lead to a specific change in behavior. Thus, what we need is a better understanding of the individual factors contributing to different outcomes, and a better appreciation of individual (and context-dependent) differences with the respect to the ability to exercise our capacities for agency. These will include genetic predispositions, environmental factors (including socio-economic and cultural factors), intervening variables (such as substance abuse or co-morbid psychiatric

disorders), specific types of abuse, the existence or absence of support networks, and many more. Neurobiological evidence can provide one important piece in this enormously complex puzzle, but we should not make the mistake of overestimating the significance of neuroscientific results in this context. In particular, the potential to generalize from these studies should not be overestimated and these results should not be interpreted in isolation from insights from social-scientific research.

4. A theoretical framework for integrating neuroscience with other kinds of research

Given that we are advocating a more integrated and interdisciplinary approach to the study of the effects of maltreatment, how can we best conceptualise the relationship between neuroscientific findings and other explanations of the impact of maltreatment? Neuroscientists themselves often argue that their theoretical framework can provide a way of integrating other sources of understanding. One possible interpretation of this claim is that it represents a reductionist view: the view that all behaviour and experiences can be explained by (or reduced to) neurobiological factors. This view sometimes goes hand in hand with the view that our behaviour is determined, and thus that we have no free will. If behaviour that is underwritten by processes in the brain is seen to be 'caused by the brain' (rather than intended by the person), and if the mental is thought to be reducible to neurobiological processes, then this threatens to eliminate free will and personal responsibility. However, while reducing complex problems to component parts that can be studied in isolation is often a constructive approach in science, this more fundamental type of reductionism is deeply problematic. If the social world is seen to be complex then it is likely to

have emergent properties, that is, properties that emerge from interactions of component parts of the system but are not predictable from them or reducible to them. On this view, the interaction between different agents – each influenced in turn by their individual personalities, beliefs, circumstances and backgrounds – is not something that can be reduced to individual brain activations. Moreover, as we have seen above, it is far from obvious how we are to conceptualise the relation between psychological and sociological (i.e., personal level) concepts and explanations and neurobiological (i.e. sub-personal level) concepts and explanations. Put differently, the explanatory framework of the social sciences differs from that of the neurosciences in terms of the scale and the concepts involved. However, both frameworks are legitimate insofar as they are explanatorily useful for the task at hand.

In studies of maltreatment, the most dominant theoretical framework at present is the ecological one (Cicchetti & Valentino, 2006; Department of Health, 1999; National Research Council, 1993; Sidebotham, 2001). The National Research Council's tome on child maltreatment puts the case for this as:

> "This perspective reflects the understanding that development is a process involving transactions between the growing child and the social environment or ecology in which development takes place. Positive and negative factors merit attention in shaping a research agenda on child maltreatment. We have adopted a perspective that recognizes that dysfunctional families are often part of a dysfunctional environment" (National Research Council, 1993, p.50).

The importance of the family and the wider social and physical environment (and the complex interplay between them) in both the causation and response to maltreatment points to the need for a framework that can capture the *interactions* between very different aspects of a child's experience. The neuroscientific contribution to our understanding can and has been fitted into this framework (see Cicchetti, 2013) with a key focus being on the multiple levels of analysis needed to understand the complex and dynamic process of child development. In this approach, no level is seen as foundational. Neuroscience can illuminate some of the factors that are involved, for instance by helping to uncover the neurobiological mechanisms that are involved in certain psychological processes (as exemplified by the hypothesis testing approach that we discussed). Professionals, however, need to recognize the complex and multifaceted nature of the aetiology of child abuse (Belsky, 1980). An ecological framework provides a scheme for systematically ordering the large body of data and integrating the various theoretical viewpoints (Sidebotham, 2001, p. 103).

In contrast, it is not clear how the neuroscientific framework can be used as a foundation that could readily accommodate the wealth of understanding we have of the social and psychological factors that contribute to the occurrence of maltreatment and to the consequences for the victim. Indeed, in the relevant neuroscientific publications there is little mention of the extensive social science literature or discussion of how it can be integrated into the neuroscientific one.

5. Conclusion

We began by questioning what contribution evidence from neuroscience (and in particular from neuroimaging studies) is making to public policy debates regarding the need for early intervention in cases of children suffering from maltreatment. In this article we focused on neuroimaging, as this technique – due to the images it is able to produce – seems to evoke particularly strong reactions (especially in those who are not familiar with the complexities and limits of the technique). The status of neuroscience as a 'hard science' leads to people giving great weight to its findings and using it as a powerful weapon as they argue for their preferred policies in relation to reducing maltreatment. For example, in Graham Allen's (2011) report for the UK government on early intervention, the second chapter 'Using our brains' anchors the argument on the development of the architecture of the brain, although the bulk of the chapter then deals with social science research. Neuroscience offers a powerful rhetorical device for persuading others of the rightness of your policy but it carries dangers when it overstates its case. Few feel competent to challenge neuroscientific assertions and so intelligent debate can be stifled. This article has aimed to give people a better understanding of the technical aspects of the subject with the hope of de-mystifying it somewhat. Simplistic and selective use of neuroscientific findings holds other dangers. At an anecdotal level, we are aware that for some, neuroscientific evidence leads to fatalism – the view that early harm is permanently embedded in brain development. Although this is contradicted by the findings on the continued ability of the brain to change into adulthood, such an assumption may distort policy and practice by taking attention away from helping children at later stages. Moreover, even if we accept that maltreatment causes long-lasting and developmentally course-altering changes to the brain, it is not at all obvious what follows from this in terms of policy implications. Some might conclude that children

suffering from maltreatment should be removed from their families as soon as possible. Others might conclude that families should be supported earlier and more strongly instead, especially in light of the fact that removal from the family is likely to cause severe stress as well, which could contribute further to the harm being done.

We showed, first, that neuroimaging studies need to be treated with caution insofar as their reliability and validity is concerned. This relates to general methodological issues in neuroimaging as a science that is still in the early stages of development as well as to more specific concerns regarding existing neuroimaging studies of victims of maltreatment. The latter are often based on small sample sizes and do not differentiate between the effects of maltreatment per se and those of co-morbid psychiatric conditions, or socio-economic variables. Moreover, due to large inter-individual differences, it is difficult to generalize from such studies. We then discussed the relationship between neuroscience and the social sciences and argued that many participants in the debate seem to subscribe either to an implicit version of dualism, or to reductionism. Both of these are problematic. Not only do they threaten the way in which we think about the agency and moral responsibility of victims of abuse (and, ultimately, of ourselves), but they are also misguided for conceptual reasons. Dualistic assumptions create a false dichotomy between brain processes on the one hand, and psychological processes on the other. Reductionism ignores that the social world is complex and has emergent properties that cannot be reduced to neuroscience.

We can conclude that evidence from neuroscience does have the potential to contribute to a better understanding of the effects of maltreatment, and hence to better decision making in public policy. But it cannot do so when looked at in isolation

from the social sciences. Rather, neuroscientific studies need to be based on, and developed in concert with, social-scientific theories if they are to truly advance our understanding. We suggest that this can best be achieved by integrating findings from neuroscience into the existing broader ecological framework for the study of maltreatment.

References

Allen, G. (2011). *Early intervention: the next steps*. London: HM Government.

Amodio, D. M. (2010). Can neuroscience advance social psychological theory? Social neuroscience for the behavioral social psychologist. *Social Cognition, 28*(6), 695-716.

Aspinwall, L. G., Brown, T. R., & Tabery, J. (2012). The Double-Edged Sword: Does Biomechanism Increase or Decrease Judges' Sentencing of Psychopaths? *Science, 337*(6096), 846-849.

Barlow, J., & Schrader-MacMillan, A. (2009). Emotional Maltreatment-what works? . London: Department for Children, Schools and Families.
Belsky, J. (1980). Child maltreatment: An ecological integration. *American psychologist, 35*(4), 320.
Bennett, C. M., & Miller, M. B. (2010). How reliable are the results from functional magnetic resonance imaging? *Annals of the New York Academy of Sciences, 1191*(1), 133-155.
Bennett, C. M., Wolford, G. L., & Miller, M. B. (2009). The principled control of false positives in neuroimaging. *Social cognitive and affective neuroscience, 4*(4), 417-422.
Bennett, M. R., & Hacker, P. M. S. (2003). *Philosophical foundations of neuroscience*. Oxford: Blackwell Publishing
Buxton, R. B. (2002). *Introduction to functional magnetic resonance imaging: principles and techniques*. Cambridge: Cambridge University Press.
Cacioppo, J. T., & Decety, J. (2011). Social neuroscience: challenges and opportunities in the study of complex behavior. *Annals of the New York Academy of Sciences, 1224*(1), 162-173.
Cawson, P., Wattam, C., Brooker, S., & Kelly, G. (2000). Child maltreatment in the United Kingdom: a study of the prevalence of abuse and neglect. *London: NSPCC*.
Center on the Developing Child at Harvard University. (2007). A Science-Based Framework for Early Childhood Policy: Using Evidence to Improve Outcomes in Learning, Behavior, and Health for Vulnerable Children Retrieved from http://www.developingchild.harvard.edu
Cicchetti, D. 2013. "Annual Research Review: Resilient functioning in maltreated children - past, present and future perspectives." *Journal of Child Psychology and Psychiatry* 54(4):402-422.
Cicchetti, D. & Valentino, K. (2005) An ecological-transactional perspective on child maltreatment: Failure of the average expectable environment and its influence on child development. *Developmental Psychopathology, 3, 129-201*.
Colombo, M. (2013). Constitutive relevance and the personal/subpersonal distinction. *Philos. Psychol.* 26, 547–570. doi:10.1080/09515089.2012.667623
Department of Health. (1999). *Framework for the Assessment of Children in Need and their Families*. London: Department of Health.
Fonagy, P. (2012). The neuroscience of prevention. *JRSM, 105*(3), 97-100.
Gillihan, S. J., & Farah, M. J. (2005). Is self special? A critical review of evidence from experimental psychology and cognitive neuroscience. *Psychological bulletin, 131*(1), 76.
Hart, H., & Rubia, K. (2012). Neuroimaging of child abuse: a critical review. *Frontiers in Human Neuroscience, 6*.
HM Treasury. 2003. *Every Child Matters*, London: HMSO.
Kriegeskorte, N., Simmons, W. K., Bellgowan, P. S., & Baker, C. I. (2009). Circular analysis in systems neuroscience: the dangers of double dipping. *Nature neuroscience, 12*(5), 535-540.
Lindsey, D. (2004). *The welfare of children*. Oxford: Oxford University Press.
Logothetis, N. K. (2008). What we can do and what we cannot do with fMRI. *Nature, 453*(7197), 869-878.

MacMillan, H. L., Wathen, C. N., Barlow, J., Fergusson, D. M., Leventhal, J. M., & Taussig, H. N. (2009). Interventions to prevent child maltreatment and associated impairment. *The Lancet, 373*(9659), 250-266.

McCrory, E., De Brito, S. A., & Viding, E. (2010). Research review: the neurobiology and genetics of maltreatment and adversity. *Journal of Child Psychology and Psychiatry, 51*(10), 1079-1095.

Monterosso, J., Royzman, E. B., & Schwartz, B. (2005). Explaining away responsibility: Effects of scientific explanation on perceived culpability. *Ethics & Behavior, 15*(2), 139-158.

National Research Council. (1993). Understanding Child Abuse and Neglect. Washington DC: The National Academies Press.

O'Connor, C., Rees, G., & Joffe, H. (2012). Neuroscience in the public sphere. *Neuron, 74*(2), 220-226.

Rose, N. (2010). 'Screen and intervene': governing risky brains. *History of the human sciences, 23*(1), 79-105.

Sabates, R., & Dex, S. (2012). Multiple risk factors in young children's development. London: Centre for Longitudinal Studies.

Shea, N. (2013). "Neural mechanisms of decision-making and the personal level," in *The Oxford Handbook of Philosophy and Psychiatry*, eds K. W. M. Fulford, M. Davies, R. Gipps, G. Graham, J. Sadler, G. Stanghellini, and T. Thornton (Oxford: Oxford University Press), 1063–1082

Sidebotham, P. (2001). An ecological approach to child abuse: a creative use of scientific models in research and practice. *Child Abuse Review, 10*(2), 97-112.

Teicher, M. H. (2000). Wounds that time won't heal: The neurobiology of child abuse. *Cerebrum, 2*(4), 50-67.

Vul, E., Harris, C., Winkielman, P., & Pashler, H. (2009). Puzzlingly high correlations in fMRI studies of emotion, personality, and social cognition. *Perspectives on Psychological Science, 4*(3), 274-290.

Wastell, D., & White, S. (2012). Blinded by neuroscience: social policy, the family and the infant brain. *Families, Relationships and Societies, 1*(3), 397-414.

Weisberg, D. S., Keil, F. C., Goodstein, J., Rawson, E., & Gray, J. R. (2008). The seductive allure of neuroscience explanations. *Journal of cognitive neuroscience, 20*(3), 470-477.

Wood, M., & Wehrli, F. (1998). Principles of Magnetic Resonance Imaging. In D. Stark & W. Bradley (Eds.), *Magnetic Resonance Imaging* (3rd ed.). Missouri: Mosby.