Neuroethics
DOI 10.1007/s12152-013-9181-0

ORIGINAL PAPER

Perceived Access to Self-relevant Information Mediates Judgments of Privacy Violations in Neuromonitoring and Other Monitoring Technologies

D. A. Baker · N. J. Schweitzer · Evan F. Risko

Received: 23 January 2013 / Accepted: 4 February 2013

Abstract Advances in technology are bringing greater insight into the mind, raising a host of privacy concerns. However, the basic psychological mechanisms underlying the perception of privacy violations are poorly understood. Here, we explore the relation between the perception of privacy violations and access to information related to one's "self." In two studies using demographically diverse samples, we find that privacy violations resulting from various monitoring technologies are mediated by the extent to which the monitoring is thought to provide access to self-relevant information, and generally neuromonitoring did not rate among the more invasive monitoring types. However, brain monitoring was judged to be more of a privacy violation when described as providing access to self-relevant information than when no such access was possible, and control participants did not judge the invasiveness of neuromonitoring any differently than those told it provided no access to self-relevant information.

D. A. Baker (✉) · N. J. Schweitzer
School of Social and Behavioral Sciences,
Arizona State University,
MC3051, Tempe, AZ 85287-3051, USA
e-mail: dabaker3@asu.edu

E. F. Risko
Department of Psychology, University of Memphis,
Memphis, TN 38152, USA

Keywords Privacy · Self · Neuroethics · Self-identity · Neuromonitoring

Privacy is an important part of our day-to-day lives and one that lies at the heart of recent debates in neuroethics. Nevertheless, we have a poor understanding of the basic psychological mechanisms underlying the perception of privacy violations. In other words, we lack a systematic science of the human experience of privacy. What leads one to feel like their privacy has been violated? When this state of affairs is juxtaposed with the seemingly constant barrage of new threats to our privacy (e.g., location tracking cell phones, unmanned aerial drones, data hungry internet service providers, RFID school badges, neuromonitoring), an acute need to improve our understanding in this area becomes apparent. Such an effort would be beneficial from both a basic science (e.g., providing insight into how we form and maintain psychological boundaries between ourselves and others) and applied perspective (e.g., providing aid in developing policy on how to deploy new technologies that have the potential to violate one's privacy). A critical task within this effort will be to identify the factors that modulate individuals' subjective response to potential privacy violations. Here, we investigate one such factor, specifically, we explore the relation between the perception of privacy violations and the extent to which a given form of monitoring (e.g., neuromonitoring) is perceived to provide access to information related to one's "self."

Published online: 15 February 2013

Privacy violations are often grounded in some form of monitoring by a third party and are, in many cases, facilitated through technology. The development and everyday use of sophisticated forms of monitoring technology is becoming increasingly prevalent (e.g., geographical positioning systems, unmanned aerial drones, brain imaging). The use of such technology is, of course, central to psychological science as it provides a record of an individual's overt behavior as well as the psychophysiological mechanisms supporting that behavior as in the case of sophisticated brain imaging technology. But, of course, the use of monitoring technology extends well beyond the psychological laboratory (e.g., RFID tags in school badges; [1]) underlining the need for a better understanding of the construct. Here we leverage the use of such technology as stimuli in an investigation of the perception of privacy violations. Specifically, we measure the extent to which individuals perceive different forms of monitoring as potential privacy violations.

The hypothesis tested in the present investigation is that the perception of a privacy violation is determined (at least in part) by the extent to which the information gathered by a given type of monitoring is perceived to be related to one's self. This hypothesis is grounded in philosophical and legal definitions of privacy that often emphasize control over information about oneself [2–4]. For example, according to Parker (1973) "Privacy is control over when and by whom the various parts of us can be sensed by others" (p. 281). The self represents a core psychological construct composed of a multitude of self-relevant elements (e.g., thoughts, beliefs, preferences) that can be thought of as coalescing into an integrated representation of what individuals would typically associate with the first person pronouns "I" and "me" [5, 6]. Critically, not all types of information about a person are perceived to provide information about that person's self [7, 8]. For example, Andersen and Ross (1984) demonstrated that information about thoughts and feelings are perceived to provide more information about a person's "real self" than information about that person's overt actions. Thus, the extent to which a given form of monitoring provides access to self-related information represents a viable dimension on which the perception of privacy violations could be based. To test this hypothesis we measured, in parallel with privacy perceptions, the extent to which individuals perceive different forms of monitoring as providing more or less information about the self. We find that individuals did indeed perceive different forms of monitoring to provide more or less access to self-related information, and we find that the extent to which this access is perceived to exist predicts the extent to which such monitoring reflects a privacy violation. Furthermore, when we control the type of monitoring and experimentally manipulate an individual's beliefs about the extent to which it provides access to self-related information we can alter individuals' privacy perception. In addition to examining the relation between privacy and the self in the context of monitoring technology, the present investigation also provides a detailed picture of how privacy varies across a number of current monitoring technologies including monitoring individuals' brain activity. This provides a unique opportunity to situate, with respect to privacy perceptions, brain based monitoring technologies amongst other forms of monitoring (e.g., is monitoring brain activity perceived as more or less of a potential privacy violation than other forms of monitoring such as monitoring the eyes; [9]) as opposed to considering the former in isolation. To our knowledge the present investigation represent the first (and most comprehensive) such effort.

Present Investigation

In Study 1 we presented participants with scenarios describing being monitored using a variety of forms of monitoring. Each of the described monitoring types was "real" in the sense that they each have, in fact, been used at one time or another to investigate some aspect of human behavior. We measured both the extent to which individuals perceived each form of monitoring as a potential privacy violation and the extent to which individuals perceived each form of monitoring as providing information related to the self. The hypothesis outlined above would predict that as the latter increases the former would increase as well. We were also interested in comparing how individuals perceive the monitoring of brain activity relative to other forms of monitoring. The potential for monitoring brain activity to act as a gateway to our inner, typically private, worlds has drawn attention to a host of ethical concerns [10–13], however, as noted above, little data actually exists regarding individuals' perceptions of such technology. Given research by Andersen and Ross [7, 8] we would expect that the monitoring of brain activity, as a putative index of ones' thoughts and feelings, to be perceived as a greater potential privacy violation than forms of

monitoring that provide information about overt behaviors [see also 14]. In Study 2 we manipulated the extent to which a given form of monitoring provided access to self-related information (rather than using natural variation across different monitoring technology). In this latter experiment we focused specifically on the monitoring of brain activity.

Experiment 1

Amazon Mechanical Turk (AMT) [15–17] was used to gather a sample of 362 US residents from 48 states (M_{age}=31; Female=43.9 %; 50.6 % college graduates). In a between subjects experiment, online participants were asked to consider the act of being monitored in a particular manner while ignoring any physical or logistical aspects of being monitored in that manner. Subsequently participants were randomly directed to a very brief description of one of 15 different monitoring scenarios ostensibly used by researchers (e.g. "…researchers have developed a system for monitoring a person's brain activity in a way that would provide them with a constant record of the which area of the brain are active") and asked the extent to which the monitoring would constitute a "violation of [the participant's] privacy" (i.e., using a Likert scale from 1 to 7, with 1 being "not at all" and 7 being "completely"). The 15 monitoring types are listed in Appendix A. On the following screen, participants were asked to rate the extent to which they believed the monitoring type that was described to them could provide information about a number of self-relevant constructs such as thoughts, desires, beliefs, intentions, values, etc. (using the same Likert scale as in the privacy rating).

Results & Discussion

Principal axis factor analysis was used to combine ratings of the self-relevant constructs into a standardized 'access-to-self' factor score,[1] which was then used to predict ratings of privacy violation across the different types of monitoring. Both privacy ratings and perceived access to self differed as a function of monitoring type, $F_{Privacy}(14,347)=10.703$, $p<.001$, $F_{Self}(14,340)=16.385$, $p<.001$. However, two more critical results emerged from this analysis. As illustrated in Fig. 1, the access-to-self factor score was highly predictive of individuals' privacy ratings (the correlation between mean privacy ratings and mean access-to-self factor scores was significant, $r(14)=.824$, $p<.001$). Thus, Study 1 clearly demonstrates that the perception of privacy is modulated by the extent to which a given form of monitoring provides access to self-related information.

In addition, the privacy ratings associated with the monitoring of brain activity did not rank particularly high relative to other forms of monitoring and this lower than expected ranking seemed to be at least partly due to individuals not rating the monitoring of brain activity particularly high with respect to providing access to self-relevant constructs. Thus, individuals do not view neuromonitoring as providing a particularly 'special' tool in terms of its ability to provide such information. Interestingly, this belief seemingly underestimates what some would claim is the current (and certainly future) capability of neuromonitoring to reveal such information [14, 18].

Experiment 2

In order to provide a further test of the relation between the perception of privacy violations and access-to-self, we manipulated the extent to which a given form of monitoring, specifically neuromonitoring, provided access to self-related information. We used AMT to gather a new sample of 288 US residents from 45 states (M_{age}=31; Female=53.8 %; 45.1 % college graduates).[2] In Study 2 participants were led to believe a "scientific advance" in the form of an unobtrusive real-time brain monitoring technique was now being used by researchers. Participants were randomly assigned to one of three conditions manipulating the description of the monitoring technique: In the first condition the brain monitoring was described as

[1] A principal axis factor analysis extracted a single factor, accounting for 63.71 % of the overall variability in the nine items. The resulting composite factor score represents a single measure of the extent to which participants believe that a monitoring type provides access to self-related information.

[2] We were able to track timing data for Experiment 2 therefore participants who completed the experiment in fewer than 240 s (a pre-determined minimum time cutoff based on pilot testing) were excluded from our analysis.

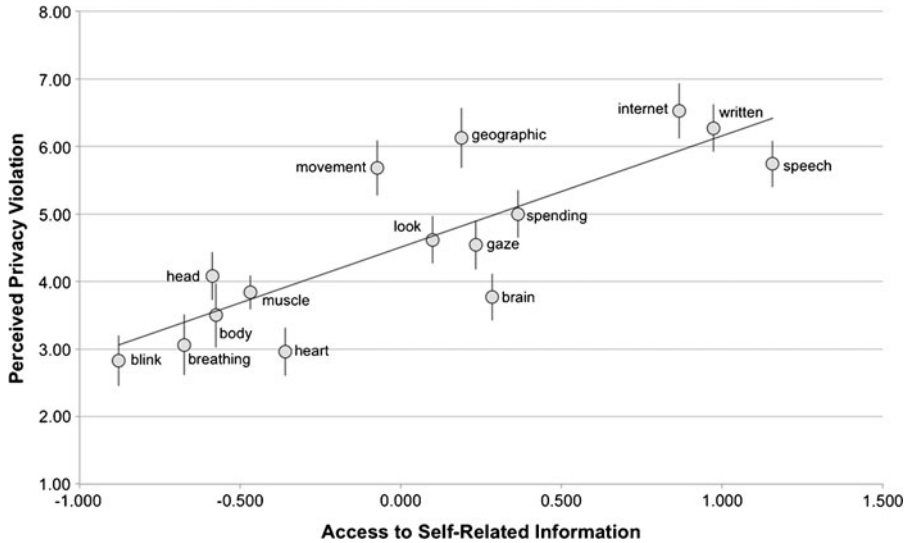

Fig. 1 Access to self and the extent of perceived privacy violation (+/− 1 SE) of Different Monitoring Types. Using a standardized 'access-to-self' factor score (see *Methods*) mean scores for ratings of privacy violations are plotted across the 15 different monitoring modalities. The strong positive correlation demonstrates that the perception of privacy is modulated by the extent to which a form of monitoring is thought to provide access to self-relevant information

allowing access to self-related information (Access), in the second it was described as being unable to allow such access (No Access), and in the third condition no information about the capabilities of the monitoring type, with respect to self-related information, was given (Control). As in Experiment 1, the participants were asked to rate the extent to which they judged the described form of neuromonitoring to be a potential privacy violation. As a manipulation check, we further asked participants to rate the extent to which they believed this monitoring could provide information about self-related constructs (thoughts, desires, beliefs, intentions, values, etc.)

Results & Discussion

We used univariate ANOVA to compare judgments of privacy violations between our three conditions, $F(2,285)=7.69$, $p<.001$, $\eta_p^2=.051$. Post-hoc tests with Tukey's HSD correction indicated that the Control condition ($M=3.84$) did not significantly differ from the No Access condition ($M=3.58$, $p=.611$). However, the Access condition ($M=4.66$) was judged to be a significantly greater privacy violation than either the Control ($p=.01$) or the No Access condition ($p<.001$). Participants judged the monitoring of brain activity to be more of a privacy violation when they were told that it could provide access to self-related information than when they were told that it could not or when they were told nothing about this particular capability (see Fig. 2). Thus, manipulating the extent to which an individual believes a particular type of monitoring provides access to self-related information has a marked impact on the extent to which they believe such monitoring reflects a privacy violation.

The lack of a difference between the no access and control conditions suggest that the general perception of neuromonitoring is that it does not provide (much) access to self-related information.[3] That is, without intervention, individuals do not seem to believe that monitoring brain activity provides a particularly efficacious means of providing access to self-related information. Critically, as in Study 1, Study 2 demonstrates that the perception of privacy is modulated by the extent to which a form of monitoring provides access to such information.

Discussion

The results are clear: individuals' privacy perceptions vary as a function of the specific technology proposed to do the monitoring. Critically, the differences between

[3] Prior to our manipulation participants were administered a brief pretest measure of preconceived notions about the capabilities of current neuromonitoring technology. A box-plot of this pretest suggested that participants were less likely to believe that neuromoniting technology has the capability to monitor self-related information–such as likes or dislikes–than they were to believe physiological information–such as signs of disease–could be monitored. (A repeated-measures ANOVA on these pretest items is not particularly useful as the large sample size resulted in nearly every item being significantly different from every other item.)

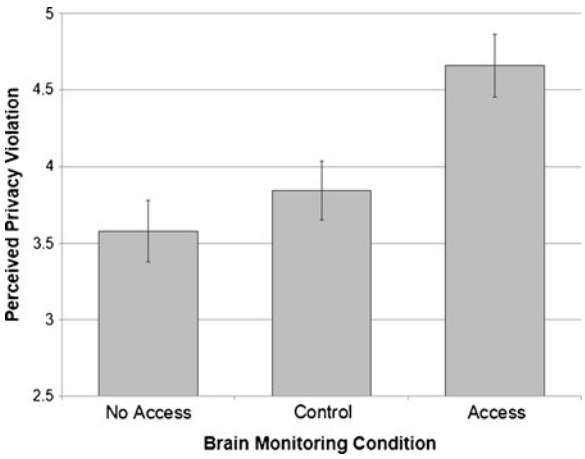

Fig. 2 Perceived privacy violation by brain access condition, +/− 1SE. Plotted are the mean scores of perceived violations of privacy across three neuromonitoring conditions; neuromonitoring described as not being able to derive information about the self (No Access), neuromonitoring described without providing reference to information about the self (Control), and neuromonitoring described as being capable of providing access to information about the self (Access). Perceived privacy violations did not significantly differ between the Control condition and the No Access condition (suggesting a possible baseline assumption of No Access), however, judgments of privacy violations increased significantly when participant were led to believe this type of monitoring could provide access to self-relevant information. See *Methods* for details and statistics

the various technologies appear to be explained to a large extent by the degree to which individuals believe that a given technology provides access to self-related information. Technologies that are perceived to provide greater access to self-related information are perceived as greater potential violators of ones' privacy. In addition, when the type of monitoring (i.e., monitoring brain activity) is held constant and the extent to which individuals are led to believe that a given technology provides access to self-related information is manipulated, we again see a clear demonstration of the importance of access to self to privacy perceptions. Specifically, when participants are led to believe that monitoring brain activity can provide access to self-related information they perceive the potential privacy violation as greater than if they are led to believe that it cannot. The present results also provide novel information regarding the perception of monitoring brain activity in relation to the other forms of monitoring with respect to both privacy and the extent to which it provides access to self-related information. On this front the results were particularly surprising. Namely, monitoring brain activity does not appear to rank particularly high (relative to other forms of monitoring) with respect to either its perceived potential to violate one's privacy or its ability to provide access to the self.

Privacy, the Self, and Beyond

The present results are consistent with the hypothesis that was derived from philosophical and legal conceptions of privacy. That is, the perception of a privacy violation is determined (at least in part) by the extent to which the information gathered by a given type of monitoring is perceived to provide access to self-related information. Access to self was able to explain differences between various monitoring technologies in terms of their privacy perceptions and an experimental manipulation of the construct was also able to alter privacy perceptions. With respect to conceptions of the self in the context of privacy these results also point to the importance of considering access to self on a continuum (i.e., different types of monitoring provide different levels of access to self-related information) rather than dichotomously (i.e., a given type of monitoring either does or does not provide access to self-related information).

While we have clearly provided support for the hypothesized relation between privacy and the self in the context of monitoring technology, examination of Fig. 1 indicates that access to self, as we have measured it here, is clearly not the only variable modulating the perception of privacy violations. For example, monitoring one's speech provides much greater access to self than monitoring one's physical location (i.e., geographical position system), but the latter is perceived as a greater potential violation of ones' privacy than the former. Future work investigating other potential contributors to the perception of privacy violations (e.g., ambiguity of the monitored information) will provide a more complete picture of the mechanisms underlying the perception of privacy.

Neuroprivacy

The present results also provide new insights into the ongoing debate regarding privacy and neuroimaging [14, 18, 19]. First, when monitoring brain activity is compared to other forms of monitoring it appears to raise a moderate level of concern over potential privacy violation. This can be viewed as somewhat

surprising given portrayals of the mind-reading potential of neuromonitoring [20–25]. In addition, at least in the sample tested, the data suggest that people are as concerned (and in many cases more concerned) about other forms of monitoring as they are about the monitoring of their brain activity. It is important to note that this observation concerns individuals' subjective perception and is independent of whether this view is "correct" in a legal or philosophical sense. As such, this observation should not be taken to minimize the important ethical dilemmas surrounding privacy that arise in the act of monitoring an individual's brain activity [10, 12, 18, 19]. Rather, it emphasizes the utility of viewing brain monitoring in the larger context of other forms of monitoring.

While brain monitoring's position amongst other forms of monitoring in terms of privacy perceptions is surprising, Experiment 2 provides a plausible explanation of this pattern. Specifically, when individuals are told that monitoring brain activity cannot provide access to internal mental states their privacy ratings were equivalent to when they were told nothing about the ability or lack thereof of the technology to provide such access. This result suggests that the sampled individuals, in general, do not believe that monitoring brain activity can provide particularly impressive access to self-related information, which at least partially explains the unimpressive showing of monitoring brain monitoring as a potential privacy violation. The observation that individuals do not, in general, view monitoring brain activity as providing much in terms of self-related information is interesting to consider in the context of people's folk beliefs about science [e.g., see 26]. In particular, this result might reflect individuals' dualist tendencies (i.e., the belief that mind and brain are not the same). For example, young children often do not perceive their personal identity to be tied to their brain such that (hypothetically) transplanting their brain would not lead to the transplantation of their identity [27]. Such beliefs are not limited to young children: Miresco and Kirmayer [28] found that mental health professionals viewed patients as more responsible and blameworthy for their behavior if the disorder was seen as having a "psychological" origin and less blameworthy if the behavior was seen as having a neurobiological origin. A dualistic belief would certainly be consistent with the observation here that individuals do not, in general, view monitoring brain activity as providing much in terms of self-related information. If the self were not located in the brain, then monitoring the brain would not provide much information about the self. It is important to note that such a dualistic belief is not a pre-requisite for not believing that monitoring brain activity can provide information about the self. For example, whereas someone might fully belief that the self is located in the brain, they might not believe that our current level of technological advancement is sufficient to provide such access. Indeed, this latter position is likely to be true, for the most part, at present. This interpretation of the present results suggests that as the potential of brain monitoring to provide self-related information becomes more widely known (or believed) the potential for perceived privacy issues associated with such monitoring will also increase.

Conclusion

Despite the centrality of privacy to our day-to-day lives we lack a deep understanding of the mechanisms underlying our perception of privacy violations. The present work represents an important step in remedying this state of affairs. Privacy perception is intimately tied to self-perception and further work aimed at elucidating this relation should provide novel insights into both constructs.

Appendix A

List of 15 possible monitoring types from Experiment 1

As part of a study, researchers have developed a system for monitoring…

1. a person's heart activity in a way that would provide them with a constant record of the person's heart rate.
2. a person's blink activity in a way that would provide them with a constant record of the person's blink rate.
3. a person's internet activity in a way that would provide them with a constant record of the person's internet usage.
4. a person's head orientation in a way that would provide them with a constant record of the direction a person's head is facing.

5. a person's spending activity in a way that would provide them with a constant record of the person's purchase history.
6. a person's movement in a way that would provide them with a constant record of how a person is moving their body.
7. a person's gaze orientation in a way that would provide them with a constant record of the direction a person's eyes are facing.
8. a person's gaze orientation in a way that would provide them with a constant visual record of what the person is looking at.
9. a person's written communication in a way that would provide them with a constant record of anything the person's writes or types.
10. a person's respiratory activity in a way that would provide them with a constant record of the person's breathing rate.
11. a person's body orientation in a way that would provide them with a constant record of the direction a person's body is facing.
12. a person's brain activity in a way that would provide them with a constant record of which areas of the brain are active.
13. a person's speech in a way that would provide them with a constant record of anything the person's says aloud.
14. a person's geographical position in a way that would provide them with a constant record of where a person is physically located.
15. a person's muscle activity in a way that would provide them with a constant record of which muscles a person is using.

References

1. A.H., et al. v. Northside Independent School District, et al. 5:12-cv-01113-LOG US (Western District of Texas San Antonio Division, 2013).
2. Moore, A.D. 2003. Privacy: its meaning and value. *American Philosophical Quarterly* 40: 215–227.
3. Parent, W.A. 1983. Privacy, Morality, and the Law. *Philosophy and Public Affairs* 12: 269–288.
4. Parker, R.B. 1973. A definition of privacy. *Rutgers Law Review* 27: 275–296.
5. Nowak, A., R.R. Vallacher, A. Tesser, and W. Borkowski. 2000. Society of self: the emergence of collective properties in self-structure. *Psychological Review* 107: 39–61.
6. Leary MR, Tangney JP. 2005. Handbook of self and identity. In eds. Leary MR, Tangney JP, New York: Guilford Press.
7. Andersen, S.M., and L. Ross. 1984. Self-knowledge and social inference: I. The impact of cognitive/affective and behavioral data. *Journal of Personality and Social Psychology* 46: 280–293.
8. Andersen, S.M. 2000. Self-knowledge and social inference: II. The diagnosticity of cognitive/affective and behavioral data. *Journal of Personality and Social Psychology* 78: 772–790.
9. Risko, E.F., and A. Kingstone. 2011. Eyes wide shut: implied social presence, eye tracking and attention. *Attention, Perception, & Psychophysics* 73: 291–296.
10. Parens E, Johnston J. 2007. Does it make sense to speak of neuroethics? Three problems with keying ethics to hot new science and technology. EMBO reports 8 Spec No: S61–4. Available: doi:10.1038/sj.embor.7400992. Accessed 6 July 2012.
11. Garland, B., and P.W. Glimcher. 2006. Cognitive neuroscience and the law. *Current Opinion in Neurobiology* 16: 130–134.
12. Wilfond BS, Ravitsky V. 2005. On the proliferation of bioethics sub-disciplines: do we really need "genethics" and "neuroethics"? The American Journal of Bioethics 5: 20–1; discussion W: 3–4. Available: doi:10.1080/15265160590960924. Accessed 16 June 2012.
13. Schweitzer, N.J., M.J. Saks, E.R. Murphy, A.L. Roskies, W. Sinnott-Armstrong, et al. 2011. Neuroimages as evidence in a mens rea defense: no impact. *Psychology, Public Policy, and Law* 17: 357–393.
14. Farah, M.J., M.E. Smith, C. Gawuga, D. Lindsell, and D. Foster. 2009. Brain imaging and brain privacy: a realistic concern? *Journal of Cognitive Neuroscience* 325: 119–127.
15. Rand, D.G. 2012. The promise of mechanical turk: how online labor markets can help theorists run behavioral experiments. *Journal of Theoretical Biology* 299: 172–179.
16. Buhrmester, M., T. Kwang, and S.D. Gosling. 2011. Amazon's mechanical turk: a new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science* 6: 3–5.
17. Paolacci, G., J. Chandler, and P.G. Ipeirotis. 2010. Running experiments on amazon mechanical turk. *Judgment and Decision Making* 5: 411–419.
18. Olson, S. 2005. Brain scans raise privacy concerns. *Science* 307: 1548–1550.
19. Räikkä, J. 2010. Brain imaging and privacy. *Neuroethics* 3: 5–12.
20. Leggett H. 2009. The next hacking frontier: Your brain? Wired.com. Available: http://www.wired.com/wiredscience/2009/07/neurosecurity/. Accessed 10 July 2012.
21. Wardlaw JM, O'Connell G, Shuler K, DeWilde J, Haley J, et al. 2011. "Can it read my mind?" - What do the public and experts think of the current (mis)uses of neuroimaging? PloS one 6: e25829. Available: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3186771&tool=pmcentrez&rendertype=abstract. Accessed 27 March 2012.
22. Covin R. 2012. Mind-reading advance lets brain scientists 'eavesdrop' on thoughts. Huffington Post.com. Available: http://www.huffingtonpost.com/2012/02/01/mind-reading-breakthrough_n_1246752.html. Accessed 8 August 2012.
23. Duncan DE. 2012, April 3. iBrain, a device that can read thoughts—NYTimes.com. The New York Times D2.
24. Paul, A. M. March 18, 2012. Your brain on fiction. The New York times SR6. Available http://www.nytimes.com/2012/03/18/opinion/sunday/the-neuroscience-of-your-brain-on-fiction.html?pagewanted=al. Accessed 6 July 2012

25. Venkataramanan, M. 2012. Brain-reading bike helmet shows how stress you are. Wired.com Available: http://www.wired.com/playbook/2012/11/mit-media-lab-bike-helmet/. Accessed 11 November 2012.
26. Bloom, P., and D.S. Weisberg. 2007. Childhood origins of adult resistance to science. *Science* 316: 996–997.
27. Johnson, C.N. 2008. If you had my brain, where would I be? Children's understanding of the brain and identity. *Child Development* 61: 962–972.
28. Miresco, M.J., and L. Kirmayer. 2006. The persistence of mind-brain dualism in psychiatric reasoning about clinical scenarios. *The American Journal of Psychiatry* 163: 913–918.