Episteme, 10, 2 (2013) 193–205 © Cambridge University Press
doi:10.1017/epi.2013.16

A DUTY OF IGNORANCE

DAVID MATHESON
david_matheson@carleton.ca

ABSTRACT

Conjoined with the claim that there is a moral right to privacy, each of the major contemporary accounts of privacy implies a duty of ignorance for those against whom the right is held. In this paper I consider and respond to a compelling argument that challenges these accounts (or the claim about a right to privacy) in the light of this implication. A crucial premise of the argument is that we cannot ever be morally obligated to *become* ignorant of information we currently know. The plausibility of this premise, I suggest, derives from the thought that there are no epistemically 'non-drastic' ways in which we can cause ourselves to become ignorant of what we already know. Drawing on some recent work in the epistemology and psychology of self-deception and forgetting, I seek to undermine this thought, and thus provide a defense against the challenging argument, by arguing that there are indeed such ways.

1

Just from the fact that I have no privacy[1] with respect to my spouse and information about my life savings, it doesn't follow that I have no privacy with respect to you and that personal information about me. Nor does it follow from the fact that I lack this privacy relative to my spouse that I lack other privacy relative to her – say, about the amount of cream I habitually consume with my coffee at the office, or about my plans to surprise her with a gift next month. Such considerations, as I have previously stressed (Matheson 2007), tell us something about the *explanandum* of accounts of privacy's nature: it is fundamentally a (concept of a) relation between two or more individuals[2] and one or more pieces of personal information[3] about at least one of those individuals. When theorizing about the nature of privacy, in others words, our central aim is not to provide illuminating truth

1 Throughout, I treat privacy and informational privacy as one and the same. But I am not an eliminativist about putatively non-informational forms of privacy (e.g. spatial privacy, decisional privacy); I see them as important species of informational privacy. See Matheson (2007: §VII).

2 Or groups. My focus here is on the privacy of individuals relative to other individuals, but, provided we are willing to countenance personal information about groups, I have no objection to the suggestion that groups can have privacy relative to other groups. Similarly, I'm happy to admit that groups can have privacy relative to individuals, and that individuals can have privacy relative to groups. The accounts of privacy's nature I canvass below are easily transformed, *mutatis mutandis*, into accounts of the nature of group privacy.

3 Personal information must of course be about specific persons (or groups). But it seems that more is required for information to count as personal. For distinct accounts of what more is required, see Parent (1983), Davis (2006), and Matheson (2008a, 2007: n. 1).

https://doi.org/10.1017/epi.2013.16 Published online by Cambridge University Press

conditions simply for the likes of *S has privacy*, but rather for something along the lines of *S_1 has privacy relative to p and to S_2*, where 'S_1' and 'S_2' range over distinct individuals and '*p*' over personal information about S_1.

Conjoined with the plausible claim that there is a moral right to privacy,[4] each of the major contemporary accounts of the nature of privacy, understood as accounts of an *explanandum* of this sort, implies a duty of ignorance.[5] This is perhaps most obvious in the case of the ignorance[6]-theoretic account I have defended (Matheson 2007; cf. Parent 1983; Davis 2006; Blaauw 2012; Le Morvan 2012b[7]), according to which an individual has privacy relative to personal information about her and to another individual just in case the other individual is ignorant of that information. That is:

Privacy(ignorance) S_1 has privacy relative to p and to S_2 iff S_2 does not know p.

If Privacy(ignorance) is true and S_1 has a moral right to privacy relative to p and to S_2, then S_1 is morally entitled to S_2's not knowing p. But rights generally entail corresponding obligations for those against whom right-bearers hold the rights. Hence if S_1 is morally entitled to S_2's ignorance of p, S_2 has a moral obligation of ignorance with respect to p.[8] Thus, if Privacy(ignorance) is true and there is such a thing as a moral right to privacy, there is such a thing as a (moral) duty of ignorance.

Or consider a control-theoretic account of the sort suggested by Westin (1967) and Fried (1968), which we may capture as follows:

Privacy(control) S_1 has privacy relative to p and to S_2 iff whether S_2 knows p depends on whether S_1 has decided to reveal p to S_2 (i.e. iff S_1 controls whether S_2 knows that p).

Here too, conjunction with the claim that there is a moral right to privacy entails a (conditional) duty of ignorance. If Privacy(control) is true and S_1 has a moral right to privacy

4 There have been eliminativist accounts of the right to privacy that call this claim into question, however; see e.g. Thomson (1975). In (2008b) I attempt to boost the plausibility of the claim by locating the moral right to privacy against a background of more familiar rights to person, liberty and property.

5 It's worth noting, as Ian Kerr has pointed out to me, that an analogous sort of duty implied by a putative 'right to be forgotten' in online contexts is a hotly debated topic in current European policy and legislation. See Rosen (2012). For more background on the digital issues see Mayer-Shönberger (2009) and Kerr (2010).

6 I take ignorance simply to be the state of not knowing, in accordance with what Le Morvan (2012a, 2011) aptly calls 'the standard view'.

7 Parent, Davis, Blaauw and Le Morvan offer accounts that differ slightly from my own, but are sufficiently similar to count as members of an ignorance-theoretic family. On Parent's account, unlike on my own, another's ignorance of one's personal information amounts to privacy with respect to that information only if the information is 'undocumented'. On Davis's account, unlike on my own, it is not quite another's ignorance of one's personal information that makes for privacy relative to that information, but rather a weaker negative epistemic state (viz., her not having warranted belief in that information). Whereas my own account allows for privacy in every case in which another is ignorant of one's personal information, Blaauw's allows for it only in cases where the ignorance has a certain modal strength. And although Le Morvan's account, unlike my own, allows that one can lose privacy simply by virtue of another entertaining false propositions about her, Le Morvan treats such privacy losses as involving a kind of attentional (or referential) knowledge on the part of the other.

8 Even those who deny the existence of epistemic obligations or duties (e.g. Wrenn 2007) can admit moral obligations with respect to our epistemic states.

https://doi.org/10.1017/epi.2013.16 Published online by Cambridge University Press

relative to p and to S_2, S_2 has a moral obligation to be ignorant of p unless S_1 decides to reveal p to her. So if Privacy(control) is true and S_1 has a moral right to privacy relative to p and to S_2, S_2 has a moral obligation to be ignorant of p if S_1 has not decided to reveal p to her. Hence, if Privacy(control) is true and there is such a thing as a moral right to privacy, there is such a thing as a duty of ignorance.

Limited-access accounts of privacy (Gavison 1980; Allen 1988) also imply a duty of ignorance when conjoined with the claim that there is a moral right to privacy. On such accounts, an individual's privacy derives not from her own control over others' knowledge of personal information about her, but from the limits on others' knowledge of this information imposed either by physical or technological barriers or by social (including legal) norms of access. Thus:

Privacy(limited-access) S_1 has privacy relative to p and to S_2 iff physical or technological barriers, or social access norms limit the ways in which S_2 can know that p.

Physical or technological barriers can obviously limit the ways in which others can know personal information about an individual: the walls of my home, like the drapes across my windows, dramatically reduce the ways in which others can know such personal information about me as my tastes in pyjamas or housekeeping habits; my clothes similarly reduce the ways in which people can acquire information about my bodily appearance; my email-account password limits the ways in which others can know information about my personal communications; and so on. But social access norms can also play the role of epistemic limitation. Consider a statute against warrantless wiretapping, for example: it is an access norm that limits the ways in which others can know personal information about an individual's telephone conversations because it eliminates warrantless wiretapping from among the legal (or unprosecutable) ways of knowing this information.

If Privacy(limited-access) is true and S_1 has a moral right to privacy relative to p and to S_2, then S_2 has a moral obligation to know p only within the limits imposed by the relevant barriers or norms; and if S_2 has such an obligation, she is morally obligated to be ignorant of p outside of those limits. So if Privacy(limited-access) is true and S_1 has a moral right to privacy relative to p and to S_2, there is yet again a duty of ignorance.

Little attention has been drawn in the privacy literature to the fact that all of the major accounts of privacy imply a duty of ignorance when they are conjoined with the claim about a right to privacy. In what follows I want to consider and respond to a prima facie plausible argument that might be mounted against these accounts (or against the claim that there is a moral right to privacy) in the light of this implication of a duty of ignorance.

2

So what is the argument? Note to begin with that every obligation (whether positive, negative, conditional or unconditional), as long as it stands, carries in its train at least two sorts of conditional obligation for the obligation-bearer – one pertaining to how the bearer ought to act on the condition that she currently satisfies the demands of the obligation, and another to how she ought to act if she does not. If, that is, one has an obligation

https://doi.org/10.1017/epi.2013.16 Published online by Cambridge University Press

to act in a certain manner or be in a certain state, then (a) if one is currently acting in this manner or currently in that state, one is obligated to act so as to continue, and (b) if one is not currently acting in this manner or currently in that state, one is obligated to act so as to begin – where the entailed, conditional obligations are presumably of the same sort (e.g. moral, epistemic, legal, professional) as the entailing obligation. To illustrate, consider my obligation as an academic not to plagiarize from the work of colleagues. The obligation applies to me regardless of whether I happen to be satisfying it; and because it applies, I am bound by two further conditional academic obligations: if I am not currently plagiarizing, then I ought to do what I can to ensure that I don't start; and if I am currently plagiarizing, I ought to stop.

The legal obligation to provide for the necessities of one's dependents likewise carries two conditional obligations. Given this obligation, if you have dependents and are currently providing for them, then you ought legally to take measures to maintain that laudable practice; and if you are not currently providing for them, you surely ought legally to start. Even an obligation so universal as the moral obligation not unduly to interfere with the liberty of others fits the pattern. Anyone not currently interfering in an undue manner with the liberty of others ought morally to do their best to ensure continued non-interference, and those who are currently interfering have a clear moral obligation to give it up.

One might suppose that we can have obligations such that once we fail to meet them it is impossible to bring about (a return to) the actions or states they require, and that such obligations do not fit this pattern.[9] I have an obligation to prevent the book I've recently borrowed from the library from being destroyed, for example; does this obligation fit the pattern? Once the obligation is not met, it may seem, it is impossible to bring about renewed compliance: I can't put the book back together once it's been (say) burned. Surely, then, one might suppose, my obligation to prevent the book from being destroyed does not carry in its train the obligation to bring about a return to the book's physical integrity if it has been destroyed (as well as the obligation to maintain the integrity of the book if it hasn't been destroyed), for that would be an obligation toward the impossible.

A more careful consideration of the content of such an obligation, however, suggests that it does fit the pattern. The obligation to prevent the destruction of the book is an obligation to act in ways conducive to the preservation of its integrity; and one can fail to meet this obligation without it being consequently impossible to effect a return to what it requires. If I fail to act in ways conducive to the preservation of the book's integrity – for example, if I leave it in the stack of scrap paper used for starting the fireplace – it doesn't follow that it has been destroyed; it follows only that I have acted inappropriately, by (perhaps) significantly raising the probability that it will be destroyed. But as long as it isn't destroyed despite my negligence, I do have an obligation to bring about a return to the behavior required – to bring it about that I do once again act in ways conducive to the preservation of the book's integrity. Of course it can be the case that my negligence results in the book's destruction. (My partner grabs it from the stack to light the fireplace, reasonably assuming that I meant it for scrap paper.) In that case my initial obligation to act in ways conducive to the preservation of the book's integrity no longer stands (even though

9 Thanks to Sanford Goldberg for discussion on this point.

https://doi.org/10.1017/epi.2013.16 Published online by Cambridge University Press

my culpability for failing to act in such ways presumably does). Still, as long as it does stand, it carries with it an obligation to begin (again) acting in such ways if I currently am not.

It seems, then, that every obligation of a given kind implies not just a similar kind of obligation to maintain or preserve what is required if what is required currently obtains, but also to bring about what is required if it currently does not obtain. Hence we have the general principle that if one has an obligation to *F*, then one can (if not-*F*) have a similar-in-kind obligation to bring it about that *F*. Applied to the putative duty of ignorance under consideration – the moral obligation not to know certain personal information about others – this principle yields the result that if we have a moral obligation not to know certain personal information about others, then we can (if we currently *do* happen to know that information) be morally obligated to bring it about that we don't know it – i.e. we can be morally obligated to *become* ignorant of the information. And it is here that the trouble enters, according to the argument, because the notion of being morally obligated to become ignorant of information we already know is absurd: whereas we may in certain situations have a moral obligation to avoid acquiring knowledge of information we don't currently possess, and whereas we may in other situations be morally obligated to acquire knowledge of information we don't currently possess, we can never be morally obligated to *become* ignorant of information we do currently possess – to take steps so as (somehow) to destroy the knowledge of information we already have. Hence there can't really be a duty of ignorance, according to the argument, which means that either accounts like Privacy(ignorance), Privacy(control), and Privacy(limited-access), or the claim that there is a moral right to privacy, must go.

This (call it) 'No Obligation to Become Ignorant' ('NOBI') argument can be succinctly captured as follows:

The NOBI Argument

P1 If any of the major accounts of privacy is true and there is a moral right to privacy, then we have a moral obligation to be ignorant of certain personal information about others.

P2 If we have an obligation to *F*, then we can (if not-*F*) have a similar-in-kind obligation to bring it about that *F*.

C1 If any of the major accounts of privacy is true and there is a moral right to privacy, we can have a moral obligation to become ignorant of information.

P3 We cannot ever have a moral obligation to become ignorant of information.

C2 Either none of the major accounts of privacy is true or there is no moral right to privacy.

3

The crucial premise of the NOBI argument is P3. I think it is false, but I acknowledge its first-glance plausibility. My aim in this section and the next is to mount a defense against the NOBI argument by providing a diagnosis of why P3 seems so plausible, and by calling into question a central element of what, according to this diagnosis, yields that plausibility.

Let's stipulate that a 'non-drastic' way of changing one's epistemic state is a way of changing one's epistemic state – of moving from a state of ignorance to a state of

https://doi.org/10.1017/epi.2013.16 Published online by Cambridge University Press

knowledge, or vice versa – that does not significantly diminish one's general epistemic competences or abilities. Non-drastic ways of changing one's epistemic state are ubiquitous, and include such mundane things as opening up the backyard window and listening carefully (so as to find out whether the chickadees are chirping, say), asking a knowledgeable friend (to learn a tax rate, for example) and consulting a product manual (to discover the source of a product's problem). Drastic ways of changing one's epistemic state, by contrast, would include such things as inflicting a serious brain injury on oneself, taking some sort of permanent memory-impairment pill[10] and the like.

It's clear from such examples that there are non-drastic ways in which we can cause ourselves to know information we don't currently know. Less clear is whether there are any non-drastic ways in which we can cause ourselves to become ignorant of information we do currently know. A plausible explanation for the plausibility of NOBI's P3, I think, derives from this comparative lack of clarity and our commitment to a suitable ought-implies-can principle: because we cannot so clearly see that there are any non-drastic ways in which we can cause ourselves to become ignorant of information we currently know, we are prompted to think that there are no such ways; and because we think that a moral obligation to become ignorant of information we currently know implies that there are some such ways – that if we are morally obligated to become ignorant of information we currently know, we can in some non-drastic manner cause ourselves to become ignorant of that information – we find ourselves inclined to accept the likes of P3.

But despite the comparative lack of clarity about whether there are any non-drastic ways in which we can cause ourselves to become ignorant of information we currently know, there are such ways. Or so I will now argue, by considering recent epistemological and psychological work in two areas of relevance: self-deception and forgetting. Work in both areas, I suggest, brings into sharper focus non-drastic ways in which we can indeed cause ourselves to become ignorant of information we currently know.

4

The occasional skeptic's complaints notwithstanding, the pervasiveness of self-deception has strong empirical support. As Sahdra and Thagard (2003) note, it accounts for a wide variety of everyday epistemic phenomena, from the useful illusions of those facing physical illness or stronger opponents on the battlefield, to the self-employed's unrealistic optimism, to driving practices among professional drivers, and various cooperative and social-defection behaviors (p. 214). There are also plausible evolutionary accounts of its emergence. Von Hippel and Trivers (2011), for example, have recently argued that self-deception developed primarily in the service of advantageous interpersonal deception: being self-deceived on a given matter can make you less likely to be detected when deceiving others about the matter, because it can significantly diminish or remove altogether the typical behavioral signals of deception.

On one prominent account of the state of self-deception – that of Bach (1981) – it involves a desire-driven interference with the activation of an underlying belief.

10 Given the non-selective nature of current, neuro-pharmacological memory modification technologies (see Liao and Sandberg 2008: 89–90), their use would likewise seem to fall into the 'drastic' category.

https://doi.org/10.1017/epi.2013.16 Published online by Cambridge University Press

Following Bach (1981: 354–7), let 'thinks that *p*' ('is thinking that *p*', 'has the thought that *p*', etc.) stand for what has traditionally been called occurrent belief that *p* (i.e. for the mental event of actively assenting to *p*), 'thinks of *p*' ('is thinking of *p*', 'has the thought of *p*', etc.) indicate the event of simply considering or entertaining *p*, and 'believes that *p*' ('has the belief that *p*', etc.) refer to a disposition to (among other things) think that *p* when prompted to think of *p*. On Bach's account, one is self-deceived that not-*p* just in case, although one believes that *p*, one also desires that not-*p* in such a way that one is successfully motivated to prevent one's belief that *p* from being activated regularly to yield the thought that *p* – either by being motivated successfully to prevent oneself from regularly thinking of *p* at all, or by being motivated successfully to prevent oneself from regularly thinking that *p* despite one's inclination to do so when the thought of *p* arises (Bach 1981: 362–5). In self-deception, in other words, one has a desire-based disposition to avoid thinking something that competes with one's underlying disposition to think it (1981: 361); the self-deceived individual 'does not [regularly] think what he believes' (1981: 369).

Bach discusses three common ways of deceiving oneself – of either putting oneself into, or keeping oneself in, the state of being self-deceived that not-*p*. *Rationalization* involves explaining away evidence one has for *p* (by e.g. constructing hypotheses to render the evidence consistent with the content of one's desire that not-*p*) so as to avoid regularly thinking that *p*. As such, rationalization can precede, and eventually lead to, the state of being self-deceived that not-*p*, because it can occur without having the belief that *p* at all (and, indeed, can occur even with the belief that not-*p*).

> A classic example is the case of a man who believes that a certain woman loves him even though he possesses (and realizes he possesses) strong evidence to the contrary. He is cognitively unmoved by the fact that she has never wanted to go out with him, always hangs up on him whenever he calls, returns his unsolicited gifts, plans to marry someone else, etc. He 'knows' there must be some explanation for all this: her mother poisoned her mind against him, her father wants her to marry someone respectable, she thinks he is too good for her, she does not realize how much he loves her, or whatever. (Bach 1981: 358–9)

Once the individual has spent a sufficient amount of time confronting the evidence for *p* (in this example, the proposition that the woman does *not* love him), and thus being regularly confronted with the thought of *p*, his belief that not-*p* (that she does love him) may be lost and replaced with the belief that *p*. But his desire-driven rationalization prevents him from regularly having the thought that *p* regardless of whether he has ('deep down', as it were (Bach 1997: 105)) come to believe that *p*. Over time, that is, the desire-driven practice of rationalization may turn one's desire that not-*p* into something that successfully competes with an underlying belief that *p*.

Evasion involves avoiding the regular thought of *p* so as to avoid the regular thought that *p*; and the evasive individual typically avoids the regular thought of *p* by quickly entertaining such things as thoughts about an entirely different matter, the thought that *p* isn't worth thinking about or a consideration against *p*, whenever the thought of *p* springs to mind, so as to forestall future occurrences of the thought of *p*. Like evasion, the third means of self-deception Bach identifies – *jamming* – 'occurs primarily in the service of an already established state of self-deception' (1981: 362). It involves forcing oneself to think that not-*p* whenever the thought of *p* arises, for example by focusing one's

https://doi.org/10.1017/epi.2013.16 Published online by Cambridge University Press

'attention on what it would be like if not-*p* were true and vividly imagin[ing] desirable consequences of not-*p*' (1981: 361); the thought that not-*p* thus 'jams' or (we might say) 'crowds out' the normal result of the activation of the belief that *p* – the thought that *p* – as one cannot think both that *p* and that not-*p* at the same time.

More recent research on self-deception has highlighted various ostensive means of self-deception, and most if not all can be comprehended under one or more of the general means that Bach identifies. Thus, for example, in addition to rationalization, von Hippel and Trivers (2011) discuss such means as biased information searching, biased interpretation, misremembering and efforts to convince oneself that falsehoods are true. Biased information searching involves limiting the amount of time one spends seeking, or domains in which one seeks, evidence on a given matter, or limiting the attention paid to certain sorts of evidence, so as to avoid the acquisition of evidence that is incompatible with what one desires to be the case about that matter (von Hippel and Trivers 2011: 8–9) Such practices seem easily classified under the evasion (and perhaps in some cases rationalization) heading. Biased interpretation is fairly clearly a form of rationalization, involving as it does the (mis)interpretation of evidence on a matter that goes against what one desires to be the case about that matter as irrelevant or misleading (von Hippel and Trivers 2011: 9). Misremembering occurs when one (conveniently) fails to retrieve, or retrieves incorrectly, stored information that would, if retrieved correctly, count against what one wants to be true about a given matter (von Hippel and Trivers 2011: 9–10); and, as such, is seems best understood as a variety of evasion (or perhaps as a kind of jamming, depending on the details). And efforts to convince oneself that a falsehood is true, for the purpose of protecting what one wants to be true about a matter (von Hippel and Trivers 2011: 10–12), would in their most rudimentary form fall under the category of jamming. If the efforts involved enough by way of explanatory manipulations, of course, they might reach the level of rationalization.

Practiced over a sufficiently long period of time, all three of the more general means of self-deception that Bach identifies seem likely to cause someone who is self-deceived that not-*p* eventually to lose her underlying belief that *p*: the competing, desire-driven disposition to avoid regularly thinking that *p* stands a good chance of rendering the individual who has it no longer disposed to think that *p* when the thought of *p* occurs to her. (So the means of deceiving oneself Bach identifies would only lead to or maintain a state of self-deception that not-*p* for a limited period of time; on his account, once the belief that *p* is lost, the individual no longer strictly counts as being in the state of self-deception that not-*p*.) This claim – that regular interference with, or inhibition of the activation of, one's belief that *p* will over time tend to eradicate that belief – finds support in plausible adaptationist accounts of the nature of forgetting information stored in long-term memory (e.g. Bjork 1989; Anderson and Schooler 2000). On such accounts, as Michaelian notes in his recent work on the epistemology of forgetting, the less frequently the stored information is retrieved from long-term memory, the greater the chances that it will be forgotten: 'the memory system renders records inaccessible (in part) according to their retrieval history' (2011: 420). On the broadly dispositionalist understanding of belief adopted here, to have a retrievable piece of information (or record) that *p* stored in long-term memory is to have a belief that *p;* to retrieve this information is to activate the belief by thinking that *p*; and to forget, or no longer be able to retrieve, the stored information is to lose the belief that *p* (while perhaps, to use Audi's terminology (1994), to retain 'a disposition to believe' that *p*: having such a disposition to believe would amount to having the record stored in,

https://doi.org/10.1017/epi.2013.16 Published online by Cambridge University Press

but not currently retrievable from, long-term memory). The adaptationist accounts thus imply that the less frequently the belief that *p* is activated or retrieved over time, the greater the likelihood that it will eventually be lost. Since regular, desire-driven retrieval inhibition of the belief that *p* has the effect of decreasing the frequency of the belief's retrieval, it should increase the likelihood of eventually losing the belief.

Indeed, the psychological research on directed forgetting suggests that such retrieval inhibition can induce forgetting, and hence the loss of belief, even if the inhibition is voluntarily effected, has been operative for a relatively short period of time, and pertains to episodic information of emotional significance to the subject.[11] In one well-known experiment (Joslyn and Oakes 2005), for example, subjects kept a personal diary of memorable events over a two-week period. At the end of the first week some of the subjects were instructed to forget the recorded events of the previous week and remember only those of the coming week; the other subjects were instructed to remember recorded events for the previous as well as for the coming week. At the end of the two-week period, subjects in the forget group performed notably worse, recalling approximately 24 percent fewer recorded events from the first week than subjects in the remember group (Joslyn and Oakes 2005: 581).[12] Barnier et al. (2007) report on a set of experiments in which subjects were asked to retrieve memories they had recently recalled and associated with words on cue lists. Some of the memories generated were of positive emotional significance to the subjects, others of negative emotional significance, and others yet of no particular emotional significance at all. In each experiment, after generating memories in connection with cues from a first list ('List 1') but before doing so in connection with cues from a second list ('List 2'), half of the subjects were asked to forget the first list memories and remember only the second list memories; the other half were asked to remember both first list and second list memories. Again, in all the experiments, subjects performed notably worse at recalling the memories they had been directed to forget. Although '[i]ndependent[ly] of directed forgetting instructions, participants recalled fewer unemotional (neutral) memories than emotional (either negative or positive) memories, ... the directed forgetting instruction was successful, irrespective of the emotional valence of the targeted memories' (Barnier et al. 2007: 315), with 'the forget instruction consistently and reliably impair[ing] recall of List 1 memories to approximately 11% below the baseline performance of the group instructed to remember' (p. 315). Given that in each of the experiments the forgetting effects occurred over a very short period of time and after just a single effort of voluntary retrieval-inhibition, it's difficult not to conclude that a more prolonged use of such retrieval-inhibition techniques would be likely to result in even more significant effects:

> [The] success in reliably inducing inhibition of recently recalled autobiographical memories in the laboratory with a simple, unrepeated procedure suggests that in everyday cognition, much more powerful effects might be present. Spontaneous and repeated use of a directed forgetting procedure

11 This, I take it, is the sort of information to which a right-to-privacy generated obligation of ignorance would most commonly apply.

12 A replication of the experiment resulted in forget group subjects performing better, but still significantly worse than remember group subjects, recalling approximately 14% fewer recorded events of the first week. The difference in results was largely explicable in light of the fact that in the repeated experiment a special holiday (Valentine's Day) occurred during the first week of recording, rendering the recorded events of that week of more emotional significance, and hence more memorable, even for subjects in the forget group (Joslyn and Oakes 2005: 584).

https://doi.org/10.1017/epi.2013.16 Published online by Cambridge University Press

> on the same knowledge may induce much stronger and enduring inhibitory effects. In everyday life, there may be many opportunities for repeated directed forgetting, and inhibitory control over autobiographical memories may be both common and effective. (Barnier et al. 2007: 319)

On the assumption that the memory or belief that *p* in many cases amounts to (or corresponds integrally with[13]) knowledge that *p*, the observation that retrieval-based inhibition tends to eradicate belief suggests that techniques like rationalization, evasion and jamming (along with whatever more specific techniques they comprehend, such as biased information searching, biased interpretation, etc.) count as ways of eventually causing oneself to become ignorant with respect to some piece of information. But would these techniques ever count as non-drastic in the relevant sense?

It's difficult to see why they wouldn't. They of course hold the potential for changing one's epistemic state from positive to negative – for moving one from a state of knowledge to a state of ignorance with respect to a piece of information – but there's no reason to think that they would have to diminish one's general epistemic competences.

Indeed, if Michaelian is right and forgetting can be epistemically virtuous in an important sense, they might in some cases even turn out to enhance one's general epistemic competences, for they can be plausibly characterized as techniques of forgetting – if, as I have argued, they are techniques likely to result in the eradication of underlying belief. Provided that forgetting something frees up the computational resources of the mind so as to allow for greater epistemic speed, power and reliability with respect to the mind's recall of other, useful information stored in long-term memory, and provided that it does not entail a diminishment of accessibility to too much useful information stored there – features that normal human forgetting often has (Michaelian 2011: 420–2) and that voluntarily effected, directed forgetting seems at least sometimes to have[14] – then it is epistemically virtuous. Techniques like rationalization, evasion and jamming, where they do result in the eventual eradication of underlying belief, and thus amount to successful techniques of forgetting, *can* have such features: they can free up the mind's computational resources in the relevant way, and because they are so limited in scope they needn't result in widespread diminishment of retrievability of useful information stored in long-term memory. Hence techniques like rationalization, evasion and jamming can at least in principle be epistemically virtuous in the relevant sense; and if they can be so virtuous, they needn't lead to the diminishment of one's basic epistemic competences and needn't be drastic ways of causing oneself to become ignorant of some piece of information one currently knows.

To the question of whether there are ways in which an individual might cause herself to become ignorant of information she currently knows, then, contemporary work on

13 In the sense that if one were to lose the positive attitude of belief, one would also lose whatever positive attitude knowledge amounts to. Some non-reductionist accounts of knowledge (e.g. Williamson 2000; White 1983) provide reason to doubt that knowledge (though a propositional attitude) is just a particular species of the attitude of belief.

14 Barnier et al. (2007: 301–2) discuss previous research showing this with respect to 'relatively simple', non-autobiographical information of no particular emotional significance to subjects. And they note that, although across all of their own experiments there was generally 'little evidence of enhanced recall of List 2 memories after the forget instruction', in one of the experiments 'the forget group recalled fewer positive List 1 memories but more positive List 2 memories than the remember group recalled' (p. 316).

https://doi.org/10.1017/epi.2013.16 Published online by Cambridge University Press

self-deception provides a plausible answer: yes. The general techniques of belief (or memory) recall-inhibition that Bach identifies as central means of self-deception – rationalization, evasion and jamming – along with whatever more particular strategies they comprehend, would seem to count: they can be effected by us (indeed, frequently are for self-deceivers) and they are likely if effected over a sufficient period of time eventually to result in the loss of currently possessed knowledge (in any cases in which the relevant beliefs amount to or correspond integrally with knowledge). To the question of whether such ways of causing oneself to become ignorant of information one currently knows could be non-drastic, Michaelian's insights on the epistemic virtues of forgetting provide a similarly positive answer: indeed they could. Not only might they not carry any degradation of the individual's general epistemic competences, they might in some cases carry some improvement.

5

In my defense against the NOBI argument, I have provided a diagnosis of why we might find its crucial premise plausible: from our inability to see clearly that there are non-drastic ways in which we can cause ourselves to become ignorant of information we currently know, we infer that there are no such ways, from which – together with the thought that there must be some such ways if we are ever to be morally obligated to become ignorant of the information – we further infer the likes of P3. I have argued that the main step in this line of thought in support of P3 is false: there are non-drastic ways in which we can cause ourselves to become ignorant of information we currently know, and so there is no principled reason to deny that we can live up to a right-to-privacy-entailed duty of ignorance in situations where we already know the relevant personal information about others.

I recognize that this defense against the NOBI argument might leave us feeling uncomfortable, however, because of its appeal to self-deception. I have pointed to typical means of self-deception – rationalization, evasion and jamming – as non-drastic ways in which we can cause ourselves to become ignorant of the relevant personal information. Doesn't this have the unpalatable implication that in situations where we currently have knowledge of this information but morally ought not to, we are morally obligated to become, repugnantly, self-deceivers with respect to that information?

I think the answer is 'no'. Robust support for this answer, and thus a more compelling attempt to deal with the discomfort that may be caused by my defense's appeal to self-deception, will have to wait for another time, but let me close by sketching two potential lines of support for a negative answer here.

First, we might note that, in typical cases of self-deception, what makes the adoption of techniques like rationalization, evasion and jamming repugnant is the fact that the adoption is motivated by something other than a desire to discharge one's moral duty. But, the thought would go, the adoption of these techniques can also be motivated by a desire to discharge one's moral duty – a motivation that may well be present in the sorts of cases wherein one recognizes that one has a right-to-privacy generated duty of ignorance with respect to the relevant personal information; and when so motivated, the adoption of the techniques amounts to an effort at self-deception but not a particularly repugnant one – no more repugnant than other attempts to respect moral obligations that preclude

https://doi.org/10.1017/epi.2013.16 Published online by Cambridge University Press

the maximization of our epistemic position (e.g. obligations not to acquire medical knowledge by unethical human or animal experimentation).[15]

Alternatively, if we are inclined to think that self-deception – morally obligated or not – is always repugnant, we might plausibly argue that what makes the adoption of techniques like rationalization, evasion, and jamming count as self-deception is motivation by an overweening desire that the facts be a certain way: by rationalizing, evading or jamming, one counts as deceiving oneself that not-*p* only if one's rationalization, evasion or jamming is motivated by a dominating desire that not-*p*. One can, however, be motivated to adopt such techniques by the desire *not to know that (or whether)* $\underline{p}$ rather than by the desire that not-*p*. And when so motivated, we might further note, the adoption of the techniques no longer counts as self-deception, but rather as another form of epistemic behavior – an epistemically moderate act of self-ignorance, we might say.[16] The motivating desire sets the aim of adopting the techniques as simply not knowing the relevant facts of the matter; for the adoption to count as self-deception, it would have to have the aim of falsely believing something about the facts.[17]

REFERENCES

Allen, A. 1988. *Uneasy Access: Privacy for Women in a Free Society*. Totowa, NJ: Rowman & Littlefield.

Anderson, J., and Schooler, L. 2000. 'The Adaptive Nature of Memory.' In E. Tulving and F. Craik (eds), *The Oxford Handbook of Memory*, pp. 557–70. Oxford: OUP.

Audi, R. 1994. 'Dispositional Beliefs and Dispositions to Believe.' *Noûs*, 28: 419–34.

Bach, K. 1981. 'An Analysis of Self-Deception.' *Philosophy and Phenomenological Research*, 41: 351–70.

——. 1997. 'Thinking and Believing in Self-Deception.' *Behavioral and Brain Sciences*, 20: 105.

Barnier, A., Conway, M., Mayoh, L., Speyer, J., Avizmil, O., and Harris, C. 2007. 'Directed Forgetting of Recently Recalled Autobiographical Memories.' *Journal of Experimental Psychology*, 136: 301–22.

Bjork, R. 1989. 'Retrieval Inhibition as an Adaptive Mechanism in Human Memory.' In H. Roediger III and F. Craik (eds), *Varieties of Memory and Consciousness: Essays in Honour of Endel Tulving*, pp. 309–30. Hillsdale, NJ: Lawrence Erlbaum Associates.

Blaauw, M. 2012. 'The Epistemology of Privacy.' Presentation for the *Episteme* conference in Delft.

Davis, S. 2006. 'Privacy, Rights, and Moral Value.' *University of Ottawa Law and Technology Journal*, 3: 109–31.

Fried, C. 1968. 'Privacy.' *Yale Law Journal*, 77: 475–93.

Gavison, R. 1980. 'Privacy and the Limits of Law.' *Yale Law Journal*, 89: 421–71.

Joslyn, S., and Oakes, M. 2005. 'Directed Forgetting of Autobiographical Events.' *Memory and Cognition*, 33: 577–87.

15 I'm indebted to an anonymous referee for pressing me to mention this potential line of support.

16 Thanks to Pierre Le Morvan for helpfully suggesting the language of moderation here, and for encouraging me to consider whether the adoption of the techniques, so motivated, should be described as a form of self-deception at all.

17 Thanks to Kirk Michaelian, Timothy Chambers, Vida Panitch, Ian Kerr, Pierre Le Morvan, Vanessa Davies, Sanford Goldberg, Nathan Cockram, Gabriele Contessa and two anonymous referees for their helpful comments and criticisms. An earlier version of this paper was presented at the 2012 *Episteme* conference in Delft, the Netherlands, and at a Carleton University Department of Philosophy colloquium; the present version has benefited from the insightful feedback of audience members at both events.

https://doi.org/10.1017/epi.2013.16 Published online by Cambridge University Press

Kerr, I. 2010. 'Review of Mayer-Shönberger's *Delete: The Virtue of Forgetting in the Digital Age*.' *Surveillance and Society*, 8: 261–4.
Le Morvan, P. 2011. 'On Ignorance: A Reply to Peels.' *Philosophia*, 39: 335–44.
——. 2012a. 'On Ignorance: A Vindication of the Standard View.' Forthcoming in *Philosophia*.
——. 2012b. 'Privacy, Secrecy, and Ignorance.' Presentation for the *Episteme* conference in Delft.
Liao, S., and Sandberg, A. 2008. 'The Normativity of Memory Modification.' *Neuroethics*, 1: 85–9.
Matheson, D. 2007. 'Unknowableness and Informational Privacy.' *Journal of Philosophical Research*, 32: 251–67.
——. 2008a. 'Deeply Personal Information and the Reasonable Expectation of Privacy in *Tessling*.' *Canadian Journal of Criminology and Criminal Justice*, 50: 349–66.
——. 2008b. 'A Distributive Reductionism about the Right to Privacy.' *The Monist*, 91: 108–29.
Mayer-Shönberger, V. 2009. *Delete: The Virtue of Forgetting in the Digital Age*. Princeton, NJ: Princeton University Press.
Michaelian, K. 2011. 'The Epistemology of Forgetting.' *Erkenntnis*, 74: 399–424.
Parent, W. 1983. 'Privacy, Morality, and the Law.' *Philosophy and Public Affairs*, 12: 269–88.
Rosen, J. 2012. 'The Right to be Forgotten.' *Stanford Law Review* (online), 88 (13 Feb.): 88–92.
Sahdra, B., and Thagard, P. 2003. 'Self-Deception and Emotional Coherence.' *Minds and Machines*, 13: 213–31.
Thomson, J. 1975. 'The Right to Privacy.' *Philosophy and Public Affairs*, 4: 295–314.
von Hippel, W., and Trivers, R. 2011. 'The Evolution and Psychology of Self-Deception.' *Behavioral and Brain Sciences*, 34: 1–16.
Westin, A. 1967. *Privacy and Freedom*. New York: Atheneum.
White, A. 1983. *The Nature of Knowledge*. Totowa, NJ: Rowman & Littlefield.
Williamson, T. 2000. *Knowledge and its Limits*. Oxford: OUP.
Wrenn, C. 2007. 'Why there are No Epistemic Duties.' *Dialogue: Canadian Philosophical Review*, 46: 115–36.

DAVID MATHESON is an Associate Professor in the Department of Philosophy at Carleton University, Ottawa, Ontario, Canada. His primary research interests and publications are in the areas of epistemology, axiology and theories of privacy.

https://doi.org/10.1017/epi.2013.16 Published online by Cambridge University Press