Rev.Phil.Psych. (2015) 6:455–468
DOI 10.1007/s13164-015-0261-4

Nudging as a Threat to Privacy

Andreas Kapsner[1] · Barbara Sandfuchs[2]

Published online: 20 May 2015

Abstract Nudges can pose serious threats to citizens' privacy. The essay discusses several examples of nudges that must appear problematic to anyone valuing privacy. The paper also re-draws a well established connection between privacy and autonomy and argues that insofar as nudges incur too great a loss of privacy, they are incompatible with the libertarianism that libertarian paternalism is committed to by virtue of its very name.

1 Introduction

A *nudge*, according to the term's inventors Cass Sunstein and Richard Thaler, is an "aspect of the choice architecture that alters people's behavior in a predictable way without forbidding any options or significantly changing their economic incentives".[1] Nudges have been used to positively influence behavior on issues as diverse as the consumption of healthy foods, retirement savings, and urine spillage in public toilets. Typically, nudges are small, even tiny[2] interventions, such as setting a default for a choice in a beneficial way. For example, vastly more people in Austria are organ donors than in Germany, simply because the default is set differently in the two countries. Austrians have to opt out of organ donation, Germans have to opt in. For a variety of reasons, people tend to stick with the defaults that are chosen for them. Sunstein and Thaler have given the name *libertarian paternalism* to the political philosophy that holds that such empirical findings should be exploited to drive citizens to make better choices.

[1] Thaler and Sunstein (2008), p. 6.

[2] The solution to the spillage problem is a tiny fly painted on the inside of the urinal. Apparently, men cannot resist the urge to aim at the fly.

✉ Andreas Kapsner
Andreas.Kapsner@lrz.uni-muenchen.de

[1] Munich Center for Mathematical Philosophy, LMU Munich, Munich, Germany

[2] GRF Research Training Group "Privacy", University of Passau, Passau, Germany

In this essay, we wish to draw attention to the fact that libertarian paternalistic measures, for all the good that they do, can pose serious threats to citizens' privacy. On the one hand, this might not seem surprising. The ability to process personal data brings great advantages for policy makers. The more one knows about people's preferences, past choices, social connections, daily routines, and so on, the easier it is to nudge them successfully, since many types of nudges can be more precisely targeted on the basis of such information.

On the other hand, one of the aims of libertarian paternalists is to nudge people to take better care of their personal data. Given this apparent concern for privacy, one might expect them to balk at the idea of harvesting personal data for their ends. However, we will see that this is not always the case. And even if libertarian paternalists did not express an interest in promoting privacy, we argue that to disregard privacy concerns in the design of nudges will lead to regulation that is incompatible with the sort of libertarianism the program propounds.

Our argument will proceed as follows: We start with some preliminaries on paternalism in Second section. In Third section, we will delineate the aspect of privacy that we focus on, namely *informational privacy*. In Fourth section, we ask what the value of privacy might be, focusing on autonomy and creativity. We argue that, since informational privacy is a precondition of autonomy, invasive measures cannot claim to be in the spirit of any truly libertarian program. In Fifth section, we consider nudges that are designed to steer people towards better privacy self-management, and display the incongruity that lurks in the idea of using personal information to achieve this aim. Sixth section takes us to a thought experiment involving an all-knowing government that seeks to avoid a particularly strong form of paternalism. In Seventh section, we broaden the scope of the discussion by considering privacy aspects of all kinds of nudges, not just those aimed at better privacy management. Finally, in Eighth section, we ask how privacy costs might be weighed against other benefits in assessing particular nudges.

We will for the most part concentrate on the works of Cass Sunstein, though we take our points to be important to the whole of the emergent libertarian-paternalistic industry. Sunstein suggests himself as a focal point for several reasons. Not only is he one of the inventors of libertarian paternalism and the most prolific writer on the subject, he has also had the chance to put his ideas into practice, having served as the head of the U.S. Office of Information and Regulatory Affairs (OIRA) from 2009 to 2012.[3] His views have already led to ample political change and will continue to do so, hence the study of these views is especially worthwhile.

[3] OIRA is one of the most powerful government agencies in the US, able to review, reject and delay many proposed regulations. Originally, however, its scope was more limited, in that it reviewed only regulations involving information collection from the public. This might suggest that it was set up as a data protection agency. However, while Sunstein in an executive memorandum mentions the protection of privacy among a cluster of topics to be kept in mind, it is not among the three main functions of the agency he singles out. One of these is to ensure that "the collection of information (…) maximizes the practical utility of and public benefit from information collected by or for the Federal Government." (Sunstein 2011, p.5) Privacy enhancing measures such as timely destruction of records are surely far from maximizing such practical utility.

2 Libertarian Paternalism

First, some clarification of the terminology that will play a role in the following, starting with the term *libertarian paternalism*. Sunstein and Thaler describe their measures as libertarian, since they allow the nudged citizens to disregard the nudges and follow their original intention, if that intention was strong enough. Much has been written about whether their program really is libertarian, and some of our comments will aim in a broadly similar direction. In this section, however, we want to concentrate on the term *paternalism*.

According to most definitions of paternalism, a rule is paternalistic if and only if it overrules the preferences of subjects solely for their own good.[4] The "solely" in the definition is important. A government[5] that explicitly rejects paternalistic measures might still override citizens' choices that it feels are bad for them, if some *other* good is achieved in addition.

For example, German law states that you must wear a seat-belt. According to the legal construction, this is to protect *other* passengers from getting hurt by your body, which might be flung about uncontrollably after a high-velocity impact.[6] The rule is therefore not seen to be paternalistic, since it is not only protects your own personal safety.

Note that according to such an understanding of paternalism, many nudges would not count as paternalistic. In early papers, Sunstein and Thaler made a distinction between *libertarian paternalism* and *libertarian benevolence*,[7] the latter applying, for example, to interventions increasing organ donorship, which cannot be said to primarily benefit the donor. As far as we can tell, this distinction has disappeared from the literature; it now seems that *libertarian paternalism* encompasses all instances of regulations based on nudges.

This might suggest that a broader understanding of the term *paternalism* than the one we are using here has become dominant, one that perhaps equates paternalism with any kind of deviation from Mill's harm principle. This principle says that the only legitimate reason for government to interfere with individuals' choices is to avoid harm to others. In their more limited senses, both libertarian paternalism and libertarian benevolence extend that principle. In the first case, the main alteration is from "others" to "others or themselves"; in the second case from "avoid harm to others" to "avoid harm to others or bestow benefits on others." The new understanding of libertarian paternalism seems to encompass both alterations. In the following, we draw attention to this distinction where it matters.

Another distinction that will be important below is the one made by Sunstein between *means paternalism* and *ends paternalism*.[8] The difference is whether the government accepts citizens' ends, concerns and goals, and only intervenes when it sees them to be employing the wrong means to achieve those ends (means paternalism), or whether it wishes to supplant new ends for those of the citizens (ends paternalism). It

[4] See Dworkin (2014).

[5] While other actors and institutions can act paternalistically, we are only concerned with paternalistic governments.

[6] Bundesverfassungsgericht (Federal Constitutional Court), Neue Juristische Wochenschrift 1987, 180.

[7] See, e.g., Sunstein and Thaler (2003), p.1193.

[8] See Sunstein (2013b), p. 1875–1878.

appears that Sunstein wants to be a means paternalist as far as possible, since ends paternalism is the much more troubling and controversial idea.

However, there is also a problem with means paternalism, and it is an epistemic one: How can a government come to know the real ends of its citizens? People are heterogeneous in their aims and values; in promoting some salient ends, government will make sweeping choices that override or underestimate the ends of minorities. We will see that this problem is the source of one of the greatest temptations for libertarian paternalists when it comes to collecting information about citizens and invading their private spheres.

3 Privacy

Before we can express our worries about endangered privacy, we should also make clear what we mean when we speak of *privacy*. There is general consensus that attempts to analyze the concept of privacy have so far failed to come up with a conclusive result. The various characterizations that have been suggested, such as privacy as *secrecy*, as *control over personal data*, as *limited access to persons* or simply as *intimacy*, all leave out important aspects, or are too broad, and in most cases both.[9]

That said, we find the tripartite distinction offered by Beate Rössler[10] to be useful. She distinguishes between local, informational and decisional privacy. While these areas might not encompass all aspects of privacy, and are not without significant overlap, they help to frame our present concern.

Local privacy[11] is the realm of privacy associated with the home and other spaces that are experienced as safe havens from unwanted attention. Though nudges that focus on behavior in the home or, for example, public toilets, should be designed with an acute awareness of the private functions of these spaces, we will not expand on this aspect of privacy here.

Instead, our main focus will be on issues of *informational privacy*. This concerns personal information, information the subject would like to be in control over but often is not. While dangers to informational privacy are not new, technological progress has hugely exacerbated them. Informational privacy is under threat by new (and old) types of information collection, information processing and information dissemination.[12]

Lastly, we are talking about what Rössler refers to as *decisional privacy* whenever we say: “This is my private decision.” We do not necessarily mean that no one should know about our decision, but that no one except for ourselves should have a say in it. Decisional privacy is a central issue in the question of abortion, for example. Indeed, the verdict in the seminal case *Roe v. Wade* was explicitly grounded in a woman’s right to privacy.[13]

[9] See Solove (2008) for an extensive critique of these conceptions and an account of privacy as a nexus of many concepts that share a Wittgensteinian *family resemblance*.

[10] See Rössler (2005) and, more concisely, Rössler (2008).

[11] Other taxonomies speak of “locational” privacy, which concerns access to information about your whereabouts. In Rössler’s framework, this would fall under informational privacy.

[12] These are main categories in the taxonomy of privacy harms in Solove (2006).

[13] See Rössler (2008) p. 703.

Raising the question of decisional privacy in connection with nudges would readily lead us to the famous question of whether "libertarian paternalism" is an oxymoron.[14] If, as on might well argue, any form of paternalism compromises our decisional privacy, then the clash between nudging and privacy is inevitable. This, however, is not the discussion we wish to take up in this essay. While we are concerned with the loss of autonomy in general and decisional privacy in particular, we are only interested in such a loss insofar as far as it is the result of a breach of informational privacy. We will expand on this idea in the next section.

4 Privacy, Creativity and Autonomy

Much has been written on the intrinsic and instrumental values of informational privacy. We focus on two issues[15] that we think should especially concern libertarians: Privacy[16] as an element that is essential to the creative development of new ideas and, as mentioned above, privacy as a precondition for autonomy.

To create genuinely new ideas, we need to be able to consider and reject a variety of possibly unorthodox alternatives, and to change our opinion freely until our thoughts are ready to be shared with the world. Fear of disapproval might mean that new ideas are smothered by conformism and thus trapped by the familiar and the mainstream. Additionally, we might refrain from sharing controversial ideas with trusted persons in order to avoid negative consequences for them. In so doing, we forfeit the chance for these ideas to be improved by constructive criticism. As Neil Richards puts it, the "engine of expression—the imagination of the human mind"[17] can be endangered by a lack of informational privacy.[18]

Apart form the danger of individual people becoming "boring", as Richards puts it, [19] this smothering of new ideas may seriously harm society at large, in that innovative solutions to social and other problems may not be found. A loss of informational privacy will therefore not only affect the individuals, but lead to whole "communities governed by apathy, impulse, or precautionary conformism."[20]

A related consideration, one that should be even more important for libertarians, is that autonomy itself presupposes privacy. The mere impression of being watched can alter behavior. If we cannot be sure who is watching us and what activities our watchers might perceive as negative, whether today or any time in the future, it is likely that we

[14] See Sunstein and Thaler (2003), Mitchell (2005).

[15] Rössler (2008) and the references therein are a good start for accounts of other valuable aspects of privacy.

[16] Unless indicated otherwise, we will from now mean *informational privacy* whenever we simply write *privacy.*

[17] Richards (2008), p. 404.

[18] An example that shows that great ideas sometimes require an immense amount of undisturbed privacy is Andrew Wiles's solution to Fermat's famous last theorem, which had exercised the best mathematical minds of the last 350 years. Wiles had been working on this problem for many years without telling anyone about it. After his triumphant disclosure of the proof, he stressed the impossibility of working on this famous problem "out in the open". He would have been unable to stomach the disbelieving interest, the possible comments about his intellectual hubris, the ridicule and the dissuasions of his colleagues (see Singh 1997).

[19] Richards (2008), p. 404.

[20] Cohen (2000), p. 1426–1427.

will limit ourselves to the least-risky behavior. This mechanism has found its definitive description in Foucault's famous discussion of the panopticon.[21]

One might expect a libertarian paternalist to want to give this idea a different spin: If privacy infringements lead to predictable changes in behavior, then this might be far from being a problem. On the contrary, it might constitute a resource for a whole new series of nudges. But this seems an all too uncharitable suspicion, one that is unwarranted by libertarian paternalist writings. At least, nothing they have said so far advocates instilling and exploiting fear and discomfort.

The arguments for privacy as a source of creativity and autonomy go right back to Mill's *On Liberty,*[22] and have been much developed since.[23] If they are right, as we believe they are, any movement that has the word "libertarian" in its title should seek to reduce privacy infringements in its policies as far as possible.

5 Nudging towards Privacy

With the sense of the importance of privacy in mind, we shall now turn to a discussion of specific nudges. We begin with those that aim to encourage people to take better care of their personal data. As Lauren Willis writes: "Consumer privacy on the internet appears poised to be the next arena [for libertarian paternalist interventions]." (Willis 2013, p. 1159)

There is ample evidence that most people are concerned about their privacy, but actually do little to protect it,[24] not unlike the many people who are positively disposed towards organ donorship but are not actually carrying donor cards. This similarity suggests that nudges that increase organ donorship might serve to protect private data, as well.

Given this initial diagnostic, what might a libertarian paternalistic strategy for privacy protection look like? An obvious proposal,[25] directly parallel to the organ donation case, would be to implement mandatory defaults in software services so that no data about users can be passed on to anyone, unless a clear and conscious choice is made to allow this.[26]

But what if a government wants to avoid *ends paternalism* and acknowledges that not all people care about their privacy all that much? In order not to reshuffle people's ends concerning privacy, would it not be better to set the defaults differently for those who would like to receive advertisements that really interest them, and would happily trade in their private data for the privilege? Would it not be better were the government to set the defaults in a way that respects the preferences of individual people, i.e., give the privacy

[21] See Foucault (1977).

[22] See Mill (1869).

[23] See also Allen (2011), de Bruin (2010), Rössler (2005), Rössler (2013) Slobogin (2002) and Solove (2008).

[24] See Nissenbaum (2011), p. 36.

[25] See Sunstein (2013a), p. 102. Willis (2013), p. 1202 gives a list of examples of actual privacy regulation proposals in this vein.

[26] In practice, it is a matter of some debate whether such defaults would work as well for privacy protection as they do for organ donorship (see Willis 2013). However, we are more concerned here with the question of whether a policy should be rejected on normative grounds rather than on the grounds of practical efficiency.

lovers one set of defaults and another set to those who do not care about privacy?

As Sunstein writes, such an idea presupposes knowledge about privacy preferences:

> If enough information is available about someone's past choices or personal situation, we could design, for that person, default rules with respect to health insurance, privacy, rental car agreements, computer settings, and everything else. (Sunstein 2013b, p. 1871, emphasis added.)

Technically, acquiring such information might well be feasible; if it is not, it will be soon. Governments already have collected huge amounts of data about their citizens, private businesses probably even more (who knows?). It would presumably not be too difficult to analyze big data in a way that enabled the categorization of people along the required lines. For example, we could create one category for those who value their privacy and the integrity of their private data above most other concerns, another for those who are willing to give away certain bits of their data if they get a good return for it, and yet another for those who enjoy any sort of attention they get and want their profiles to be disseminated as widely as possible.

Government could now for each group set appropriate defaults for information use. For example, for the group that values privacy extremely highly: "*Treat my personal data with the utmost confidentiality, even if that means that I will be unable to use certain services and other amenities*"; and for the group that gives the least importance to privacy: "*Make free use of my personal data*". In an article suggesting to personalize default rules for, amongst other things, privacy settings,[27] Ariel Porat and Lior Strahlevitz spell out such possibilities in much more sophisticated ways; however, our sketch here suffices for our purpose.

Now, that there is potential for conceptual pitfalls is clear. We find out that someone does not want his or her privacy invaded by invading his or her privacy. If this is not yet problematic in itself, then it surely becomes so if we insist that we are doing this while honoring his or her ends.

This is not lost on those who make such proposals. Porat and Strahlevitz call their own suggestion to invade privacy in order to further privacy protection paradoxical,[28] though they do not think that the paradox runs too deep. Likewise, Sunstein himself shows in an unpublished manuscript[29] that he is well aware of the problem:

With respect to privacy, there is a great deal of heterogeneity and a risk of self-interested judgments on the part of choice architects, both of which argue for active choosing.[30]

However, he also writes:

> [P]erhaps choice architects know that Jones is fiercely protective of her privacy and that in the face of any kind of doubt, she prefers to prevent other people from

[27] Porat and Strahilevitz (2013).

[28] Porat and Strahilevitz 2013, p. 1468.

[29] Sunstein (2013c). It is especially interesting to compare this manuscript with an earlier draft from 2012, which at the point of this writing is also available online. The sections on privacy are among the main changes, suggesting that the kinds of problems we discuss here have been on Sunstein's mind lately, as well.

[30] Sunstein (2013c), p. 40. Active choosing here refers to an alternative to default setting, namely to prompt individual users to actively make their privacy choices themselves.

> knowing about her behavior and her choices. If so, that very knowledge can be used to produce privacy protective default rules.[31]

There is a background assumption that Sunstein, Porat and Strahlevitz all seem to be making, though they do not quite articulate it. This assumption is that no one could be so concerned about their privacy that they wouldn't want government to know about this concern. That strikes us as false. Just consider the many reports of people being placed under surveillance merely because they use services such as the anonymity network *TOR*.[32] The unspoken assumption presupposes a level of trust that governments today simply do not enjoy.

One of the many issues on which such trust hinges concerns the sources of information a government accesses to personalize its nudges. How exactly do the choice architects come to know of Jones's status as "fiercely protective of her privacy"? Of course, it is conceivable that they might limit themselves to data that individual citizens have willingly and explicitly disclosed, a measure that would go some way to instill trust in the population. Sunstein's own sympathy, however, seems to lie with much more sweeping big data analyses, as we will see in the next section.

6 Big Data and Ends Paternalism

In an interesting passage, Sunstein claims that a government that seeks comprehensive knowledge about its citizens can still be libertarian. What is more, he thinks that a highly informed government might not even be paternalistic any more:

> Consider a thought experiment or perhaps a little science fiction. We should be able to agree that the government would focus only on means, and would not be paternalistic, if it could have direct access to all of people's internal concerns and provide them with accurate information about everything that concerns them. And perhaps in the fullness of time, government, or the private sector, will be able to do something like that. (Sunstein 2013b, p. 1857)

This quote occurs, again, in the context of answering to the complaint that ends paternalism is inevitable, i.e., that government will unwittingly override the ends of minorities. This problem, so the idea goes, would disappear if nudges could be tailored to the concerns of the individual, given sufficient information about those concerns.

Take Smith, who loves fast food. She also has a terminal illness (unrelated to her preference for unhealthy food) that makes it very unlikely that she will live long enough to suffer from the consequences of her diet. If government had access to Smith's medical file and could use that information in its design of interventions, it would not have to bother her by nudging her towards soybean products she neither needs nor wants.

This example, of course, is meant to make you feel uncomfortable, but we cannot see how it is precluded by anything Sunstein says about his fantasy. Comfortable or not,

[31] Sunstein (2013c), p. 40.

[32] See, e.g., http://daserste.ndr.de/panorama/aktuell/nsa230_page-1.html.

we hope that you will come to agree that it is far from uncontroversial that this fantasy would not constitute a form of paternalism, as he claims.

In fact, we take that claim to be false, and rather obviously so for those who have privacy concerns on their minds. A government that keeps tabs on its citizens to such a degree is surely paternalistic, because it overrides important concerns of those whose preference it is *not* to be spied on. We would guess that everyone has some concerns they would not want the government to be informed about.[33] However, for our argument it is not necessary that all, or even most people have such concerns. As long as in Sunstein's fictional future there is at least one such person left (they certainly exist today), their concern for government not knowing certain things will be overridden, maybe for their own good, maybe not. In any case, we argue that this is surely a form of paternalism.[34] More, going down Sunstein's science fiction route would constitute a perfect example for *ends* paternalism, not a way to avoid it.

7 Privacy Conscious Nudging

Let us then assume that a libertarian paternalistic administration has abandoned the dream, or nightmare, of a perfectly informed government (and maybe the dream of a pure means paternalism along with it). Let us also leave for a while the topic of pro-privacy nudges, as the administration has to make important choices in how far it will respect citizens' personal information in order to better their behavior on other issues, as well.

For example, governments might start to monitor the calorie intake of their citizens, by tapping into the data streams generated by various self-improvement apps. This would set the scene for some possibly very successful nudges towards better health, but it would entail collecting data that many would consider highly sensitive.

There seem to be two options for those libertarian paternalists who take our concerns seriously. They might abandon privacy endangering practices as far as possible, even if that means less effective regulation. Or they might choose to openly give up on privacy protection and use all the information they can get their hands on. However, as we argued above, this second option might well imply a loss of entitlement to the label "libertarian", and thus seems to be out of the question.

[33] On the factual falsity of most pronouncements of the claim "I've got nothing to hide", see Solove 2007.

[34] Our claim might be questioned if at issue is the more exclusive understanding of paternalism, according to which measures can only be considered paternalistic if they overrule the preferences of the individuals solely for their own good. Might there not be a secondary benefit to universally informed governments that might save Sunstein's fantasy from being paternalistic? Indeed, we know that governments collect vast amounts of information about their citizens in the name of national safety; what harm could there be in letting the "nudging unit" look over the information the NSA has collected? This sort of information sharing between government agencies, however, has already been shown to be far from unproblematic in the famous census decision of the German Federal Constitutional Court in 1983 (translation available at 5 Human Rights Law Journal 1984, p. 94–116 (100)). Even back then it was recognized that free flow of information between such agencies dramatically increases the threats to citizens' privacy as it increases uncertainty about which persons have access to sensitive information. It also makes data theft and leakage more likely and leads to other related problems. The preference of privacy conscious citizens would be that information about them, if it must be collected, should not be freely distributed to other agencies such as the nudging unit. Even if the collection of the data might itself not be paternalistic (in the narrow sense), as it might serve some purpose unrelated to the citizens' own good, using it to find out more about their ends in order to serve those ends surely is.

While the first option is thus the one we would hope legislators chose, it might turn out to be more cumbersome than expected. Let us pick two prominent nudges discussed in Thaler and Sunstein (2008) to see whether they might be problematic from a privacy perspective, their success in achieving laudable goals notwithstanding.

The first example involves adding an element of competition in order to get consumers to save energy.[35] Here, the first empirical finding is that electricity consumption goes down if those who consume above average are informed that they do so. This might be worrisome already from a privacy conscious perspective. What if someone glances over our shoulders as we open the letter from the company that informs us that we are reckless wasters of energy?

But there is another part to the story: Those who are informed that they consume *less* than average actually *increase* their consumption. It is only when that fact is presented as positive and directly compared to the households in the neighborhood that these electricity users further strive to lower their consumption. When comparisons are thus localized, it might not be hard to figure out who particularly wasteful members of the community are, and, given the right social climate, they might end up being ostracized for their behavior.

How about unpersonalized nudges, such as the simple setting of defaults across all users? At first sight, it seems nothing could go wrong here. No data about you is retrieved to set the appropriate default, rather it is uniformly chosen by the regulators in accordance with their aims. However, on second thought, it might not be so simple.

Let us return to organ donation, one of the most frequently cited examples of the power of “sticky” defaults. As was mentioned in the introduction, donor rates are massively higher in countries in which donorship is the default. This seems to be an impeccable success for libertarian paternalists who implement such a default, for their aim must be to get as many people who have no strong opinion about the subject to donate their organs, while enabling those who want to keep their organs to do just that.

But suppose you do not want to donate your organs after death, for reasons you are unwilling to discuss with anyone. In Germany, all you have to do is not opt in on organ donation. In other words, all you need to do is nothing at all. This is indistinguishable from not signing up out of laziness or being unaware of the possibility of giving away your organs after death.

In Austria, on the other hand, one has to actively opt out. This means that, should someone find out about your decision and press you on the matter, you will have less maneuvering space. It might mean that you picture the emergency surgeon muttering “What a selfish person!” as he fishes the “No-Donor”-card out of your wallet. Maybe you are so concerned about such a possibility that you decide not to opt out after all, even though this is against your aims. Even if opting out does not involve carrying a card, every such mechanism will need to keep some kind of record of individuals’ decisions, a record which they might be concerned and worried about, maybe to the point of not pursuing their true aims.

These examples are not meant to show that concern for privacy poses insurmountable problems for all kinds of nudges. Absolute protection of privacy is impossible in any form of regulation; privacy costs will always have to be balanced against other benefits. Maybe the impact of our privacy concerns will not tip the scales against the

[35] See Thaler and Sunstein (2008), p. 68 et seq.

measures. Maybe the worst squanderers of energy are apparent simply by inspecting their Christmas decoration. Maybe people who can be swayed by the thought of possible social discomfort do not, after all, want to keep their organs enough for that wish to deserve protection.

What the examples are meant to establish, however, is that explicitly adding privacy concerns to the check list for future nudges will involve some thorny decisions and in some cases the abandonment of otherwise promising programs. It will involve answering difficult questions about retention periods for data, levels of data security, impacts on government trust, and so on.[36]

8 The Weight of Privacy

We want to end by singling out the most fundamental of these problematic questions: How should privacy in practice be weighed against other concerns?

Much of the deliberation about a proposed nudge will involve some form of cost-benefit analysis. However, it is not easy to give a numerical value to such an intangible good as privacy. Sunstein mentions this difficulty in his recent book *Valuing Life*,[37] in which he gives an inside look into cost-benefit analyses he was engaged in as head of OIRA. In that book, he seems only to be thinking of privacy as it figures on the "benefit" side of the ledger. When assessing the weight of the kinds of risks we point to, of course, the problem of unquantifiability will crop up on the "cost" side, as well.

Indeed, we might have privacy featuring both on the "cost" and on the "benefit" side of the calculation, as is the case with the nudges towards better privacy control we discussed earlier. Since such calculations involve similar, though still hard to quantify, goods, they might be amongst the easiest to solve. Let us try our hand at some cases:

- There are apps and plugins that pose a number of challenging puzzles in critical hours to see whether you are sober enough to be allowed to post on social media sites.[38] The first step in setting some of these services up is to indicate the hours in the week when you are likely to be too drunk or high to trust yourself in your posting behavior. Taking inspiration from this, governments might force social media providers to hardwire such an app into their services. They might not be activated by user-set default timeslots, but, more efficiently, by tacitly tracking users' GPS-signals and activating the inhibiting device after a visit to the bar. Even if sensitive private data is spared from self-disclosure in this way, it seems to us that the privacy costs would clearly outweigh the benefits here.
- Similarly, an algorithm was actually used in a (non-representative) study[39] to monitor status updates for overly emotional language while they were written. Different types of warnings were shown, resulting in a decrease in posts that were later regretted. Suppose government wants to legally implement such algorithms in social media sites. Whether privacy gains outweigh privacy concerns will, among

[36] Solove (2006) gives a taxonomy of privacy harms that might serve as guide for such checks.
[37] Sunstein (2014).
[38] See e.g., https://www.facebook.com/pages/Drunk-Detector/251632671590218?sk=info&tab=page_info.
[39] See Wang et al. (2013).

other things, depend on how plausibly government can claim that the algorithms only displaying their results to the user and do not feed secondary applications.

- Consider a different and rather elaborate nudge: An interesting video produced on behalf of the Belgian Financial Sector Federation (Febelfin) shows a fake magician who is able to "mindread" his subjects and find out all sorts of sensitive information such as alcohol consumption, bank details, and so on. At the end of the video, it is revealed that it was not the magician but a team of hackers that was able to retrieve that information in real time, given apparently nothing more than the name of the subject.[40] The effect is like that of a benevolent hacker attack: to show that the possibilities of real harm are not just theoretical, and thus to overcome what behavioral economists call *availability-bias*.[41] Though the privacy invasion is substantial, the benefits could well be greater still. This is a question that only empirical studies of the post-assault privacy behavior of the victims could answer.[42]
- Lastly, what about our original problem, the personalized setting of privacy defaults on the basis of comprehensive data analyses suggested by Sunstein, Porat and Strahlevitz? Even if the government were able to instill warranted trust that it would not collect further information once it learned about Jones's strong preference for privacy and delete all earlier records, we believe that this is a case in which privacy costs are higher than benefits, if the initial information was obtained against her wish. If, on the other hand, the model were one in which citizens had to actively opt in on government tracking for the purpose of having their defaults set for them, the issue might depend on how successful the ensuing privacy nudges were, as there will still be worries about data security and the like on the cost side.

While these cases are not all easy to decide, it gets even harder when privacy has to be balanced with goods of completely other types, such as health, wealth or happiness. We can suggest no metric that would solve this problem; however, that does not mean that nothing can be said on the matter.

Some calculations will be obvious, even if privacy is weighed against a quite different sort of good. Consider improving further on urinal performance by strategically placing surveillance cameras that broadcast to screens out in the hall, or maybe the internet. Quite effective, we imagine. Also, we hope, out of the question.

And even before we engage in any specific comparison, it might be a good idea to draw some red lines at the outset: Which, if any, privacy infringements should be out of bounds for a libertarian paternalist, no matter the benefits they might bring? Imagine, for example, that we were to find strong correlations between certain genes and responsiveness to certain types of nudges. Imagine that if we know that someone has these particular genes, we can design highly effective nudges for them, and maybe even

[40] See https://www.youtube.com/watch?v=F7pYHN9iC9I.

[41] See Thaler and Sunstein (2008), p. 24 on availability biases.

[42] Also, the effect on those who merely watch the video would have to beb e factored in. We would guess that watching the video isn't near as effective as being a victim of the magician, in part because of the persistent overoptimism that behavioral economics has observed time and again: People are aware of risks, but they grossly underestimate the chance that they themselves could fall victim to those risks. Smokers think that smoking is causing cancer, but tend to believe that their personal chance of getting cancer isn't any higher than that of non-smokers, etc. On the other hand, many more people might watch the video than could feasibly be directly targeted by such a demonstration, so that the net benefit might be comparable in the end.

for their close relatives. Should a government seek and employ knowledge about the genetic makeup of its citizens in order to optimize its design of nudges, or should it refrain from such possibilities categorically?

9 Conclusion

We have highlighted the potential dangers to privacy inherent in many nudges, whether already employed, merely proposed or (plausibly, we maintain) extrapolated by us. Further, we have argued that these concerns should not be dismissed lightly, since they strike the core of the libertarian enterprise. This does not mean that nudges that infringe upon privacy may never be licit. However, the dangers to privacy will have to be considered and balanced against the goals reached by the implementation of each nudge. This calculation will often be difficult to make. Nonetheless, libertarian paternalists should monitor their ideas closely for potential privacy infringements and abandon those that fail muster, even if that means abstaining from highly effective measures of regulation.

Acknowledgments The research of the first author was funded by the German Research Foundation (Project PI 1082/1-1 0228 885 2134).

We would like to thank Simon Garnett, Prof. Gerrit Hornung, three anonymous referees and the editors for helpful comments on this paper.

Compliance with ethical standards The work complies with all ethical standards as listed on the ROPP webpage.

References

Allen, A.L. 2011. *Unpopular privacy.* Oxford: Oxford University Press.

Cohen, J.E. 2000. Examined lives: informational privacy and the subject as object. *Stanford Law Review Online* 52: 1373–1438.

de Bruin, B. 2010. The liberal value of privacy. *Law and Philosophy* 29: 505–534.

Dworkin, G. 2014. Paternalism, *The Stanford Encyclopedia of Philosophy* (Summer 2014 Edition), <http://plato.stanford.edu/archives/sum2014/entries/paternalism/>. 22nd January 2015

Foucault, M. 1977. *Discipline and punish: the birth of the prison*. New York: Random House LLC.

Mill, J. S. 1869. *On liberty.* London: Longmans, Green, Reader, and Dyer.

Mitchell, G. 2005. Libertarian paternalism is an oxymoron. *Northwestern University Law Review* 99: 1245–1277.

Nissenbaum, H. 2011. A contextual approach to privacy online. *Daedalus, the Journal of the American Academy of Arts and Sciences* 140: 32–48.

Porat, A., and L. Strahilevitz. 2013. Personalizing default rules and disclosure with big data. *Michigan Law Review* 112: 1417–1478.

Richards, N. 2008. Intellectual privacy. *Texas Law Review* 87: 387–445.

Rössler, B. 2005. *The value of privacy.* Cambridge: Polity.

Rössler, B. 2008. New ways of thinking about privacy. In *The Oxford Handbook of political theory*, ed. Dryzek, Honig, and Phillips, 694–712. Oxford: Oxford University Press.

Rössler, B. 2013. Autonomy, paternalism, and privacy: some remarks on Anita Allen. *APA Newsletter* 13(1): 14–18.

Singh, S. 1997. *Fermat's last theorem*. London: Fourth Estate.

Slobogin, C. (2002). Public Privacy: Camera Surveillance of Public Places and the Right toAnonymity. *Mississippi Law Journal*, 72, p. 213-299.

Solove, D. 2006. A taxonomy of privacy. *University of Pennsylvania Law Review* 154: 477–560.

Solove, D. 2007. 'I've got nothing to hide' and other misunderstandings of privacy. *San Diego Law Review* 44: 745–772.

Solove, D. 2008. *Understanding privacy.* Cambridge: Harvard University Press.

Sunstein, C. 2011. Memorandum for the heads of executive departments and agencies, and of independent regulatory agencies: Executive Order 13563, "Improving Regulation and Regulatory Review".

Sunstein, C. 2013a. *Simpler: the future of government*. New York: Simon and Schuster.

Sunstein, C. 2013b. The storrs lectures: behavioral economics and paternalism. *Yale Law Journal* 122: 1826–1899.

Sunstein, C. 2013c. Impersonal default rules vs. active choices vs. personalized default rules: a triptych. Manuscript available at SSRN: http://ssrn.com/abstract=2171343 or doi:10.2139/ssrn.2171343, earlier version mentioned in the text available at http://dash.harvard.edu/handle/1/9876090. 22nd January 2015.

Sunstein, C. 2014. *Valuing life*. Chicago: Chicago University Press.

Sunstein, C., and R. Thaler. 2003. Libertarian paternalism is not an oxymoron. *The University of Chicago Law Review* 70: 1159–1202.

Thaler, R., and C. Sunstein. 2008. *Nudge: improving decisions about health, wealth, and happiness*. New Haven: Yale University Press.

Wang, Y., et al. 2013. From facebook regrets to facebook privacy nudges. *Ohio State Law Journal* 74: 1307–1344.

Willis, L. 2013. When nudges fail: slippery defaults. *The University of Chicago Law Review* 80: 1155–1229.