EXPLORATIONS
DIGITAL
CULTURES
BURKHARDT
SHNAYIEN
GRASHÖFER

# Personalized Concepts of Privacy in Ever-Changing Computing Environments: Challenges and Critique

Martin Degeling and Jasmin Degeling

**Bibliographic information:**

Degeling, Martin and Jasmin Degeling. 2021. "Personalized Concepts of Privacy in Ever-Changing Computing Environments: Challenges and Critique." In *Explorations in Digital Cultures*, edited by Marcus Burkhardt, Mary Shnayien, and Katja Grashöfer. Lüneburg: meson press. DOI: 10.14619/1716 <Online first version>.



PRIVACY

PERSONAL DATA

UBIQUITOUS COMPUTING

COMPUTER SCIENCE RESEARCH

DIGITAL MEDIA HISTORIES

# Personalized Concepts of Privacy in Ever-Changing Computing Environments: Challenges and Critique

**Martin Degeling and Jasmin Degeling**

**Although privacy is multifaceted, being researched in various scientific fields such as legal, social and technological sciences, the recent discourse privileges a technological perspective, in which privacy concepts are shaped by user-study-driven technology developments. Scientific research in this area is concerned with privacy-enhancing technologies like privacy assistants. These systems are supposed to help users cope with the problem of how to handle their personal data within current information systems and ubiquitous computing environments. This same research also tries to shed light on the shifting conceptions of privacy. In our chapter, we want to map out and discuss the epistemological implications of current privacy research that attempts to measure and mediate social norms by implementing them in technology. The challenges of this are especially**

**pressing, because privacy-enhancing technologies are bound to the same technological hypes as computer science in general: machine learning and big data are used to develop privacy assistants, entering the realm of prediction and automatic decision making and enforcing old norms on new problems.**

Privacy and data security continue to be among the main issues in IT-related debates of the last decade. Challenged by the rapid development of new technologies and surveillance by private and state actors, privacy has become a privileged research subject in computer science research. While there are some differences between the European and the Northern American perspective on the legal definition of "personal information"[1], the underlying concept of "informational self-determination" – the idea of the individual being in control of what information is available about *her – is similar on both sides of the Atlantic.

Researchers have recognized the limitations of this individualized perspective on privacy that puts the rational being at its center. However, informational self-determination remains prevalent and efforts are made to integrate this concept of privacy into new technical developments. Currently, computer science adopts data mining techniques to predict privacy norms that are derived from scenario-based quantitative studies to support or replace individual privacy decisions.

In this chapter, we want to examine how notions of privacy have been shaped by privacy research in computer science in recent years. Being from different disciplines ourselves, namely computer science (and involved in some research described in this paper) and media and gender studies, we are interested in describing and discussing the emergence of this research field. By trying to outline at least some important examples of recent research approaches, we aim to shed light on the particular disciplinarity of privacy concepts. We are interested in understanding how they deal

1 In this paper, we do consider account the differences between the European and US-American discourse on data protection and privacy. Still, it should at least be noted that data protection in Europe is more proactive and puts more responsibilities on those who collect personal information, whereas it is less strictly governed in the US, where regulations are negotiated within markets and (technical) developments.

with our evolving media environments and their complex and sometimes opaque technical operations.

## Challenges for Technological Privacy Research

As more and more processes get digitalized, computer science research is focusing on (technical) concepts of regulating the disclosure of personal information in a rapidly changing technological environment. Technical advances take place at all levels of software and hardware development, ranging from new devices and sensors capturing more data and exchanging it in the "Internet of Things", to new algorithms capable of extracting more information from data streams like facial recognition and voice assistants. The narrative of "data as the new oil" fuels debates about ownership, access rights and the question of what data is considered anonymous and what is personal.

The majority of research on privacy and security in computer science focuses on technical means that prevent (personal) data from being stolen, leaked or processed for purposes other than the one it was initially collected for. This "privacy as confidentiality" (Gürses, Preneel, and Berendt 2009) approach often overlooks the socio-technical contexts in which the sharing of information is considered desirable. For example, while cryptographers have created numerous algorithms to exchange encrypted messages between two parties, the question of how and if these algorithms can be adopted to group messaging, which has become an increasingly important mode of communication, has barely been solved. Similarly, to decide whether a data point should be considered personal data and what implication this has on how or by whom the data can be stored, processed or accessed, is a hard problem. In some cases, legal frameworks like the European General Data Protection Regulation (GDPR) have set guidelines, e.g. data collection is, in most cases, allowed if the data subject has given consent and when there is a specific purpose for the data being collected. For example, smartphone applications are by default not allowed to access the location of the device, but access can be granted, if the developers indicate that the data is necessary for the application to fulfill its purpose. The responsibility to decide what types of data disclosure are ok and which are not is shifted to the user by requesting permission through the user interface, assuming that s*he can execute *her right to informational self-determination.

As a result, users are confronted with a growing number of privacy decisions that, in turn, have inspired a line of privacy and security research

that tries to lift the burden of an increasing number of complex privacy decisions from the user. Researchers are looking for common norms in disclosure practices that can guide and render manageable what a user might find acceptable and what not.

But these privacy norms have become problematic[2] – not only for users but also for researchers and technology developers. Referring to the Foucauldian concept of "problematization" as a means of determining when and under what (societal, technical) conditions something has become an object of knowledge and social regulation in a specific historical situation, we would suggest that "privacy" has become such a problem, as a symptom of shifting discourse: research approaches rely on empirical – quantitative as well as qualitative – user studies to gain knowledge about supposedly existing privacy understandings and practices and then implement those into software solutions that one day may automatically manage the individual's digital privacy using machine learning. The mediation between technical developments, legal, social and technical privacy norms, and what is called "user preferences", i.e. individual privacy preferences that might deviate from the norm, has become a field for intervention by developers and researchers. The complexities and virtualities of media settings are simulated to elicit and model norms for barely used or even not-yet-existing technologies.

2 Foucault explains his notion of problematization as follows: "What I tried to do [...] was to analyze the process of 'problematization' - which means: how and why certain things (behavior, phenomena, processes) became a problem. Why, for example, certain forms of behavior were characterized and classified as 'madness' while other similar forms were completely neglected at a given historical moment; the same thing for crime and delinquency, the same question of problematization for sexuality. [...] I have tried to show that it was precisely some real existent in the world which was the target of social regulation at a given moment. The question I raise is this one: how and why were very different things in the world gathered together, characterized, analyzed, and treated as, for example, 'mental illness?' What are the relevant elements for a given 'problematization?' [...] For I think there is a relation between the thing which is problematized and the process of problematization. The problematization is an 'answer' to a concrete situation that is real. [...] In fact, however, I have tried to show, for instance, that the new problematization of illness or physical disease at the end of the 18th Century was very directly linked to a modification in various practices or to the development of a new social reaction to diseases, or the challenge posed by certain processes, and so on." (cf. Foucault 2006,cf. Foucault 1996, 178f.)

## Privacy as Contextual Integrity: Technology vs. Social Norms

As technological environments are becoming increasingly complex, with individuals, businesses, and governments having interests in personal information, privacy research has started taking into account what is described as the "data subject".[3] Instead of fully embracing the concept of a rational individual, research has shown that the actual decision of whether to disclose information or not is often a secondary concern to the main question of whether one wants to use a service (and therefore disclose personal information) or not (to protect privacy, but also prevent oneself from receiving whatever benefit the service might offer). In light of this discussion, the discourse has adopted a more complex view of privacy as "contextual integrity" (Nissenbaum 2004). Nissenbaum's approach has become very influential in privacy research in more recent years, building on the assumption that privacy decisions are a result of different norms of privacy itself. Contextual integrity, as Nissenbaum framed it, is dependent on a set of various "factors". Some of those factors are regulated by data protection laws (e.g. data receivers, purposes and retention times have to be defined in privacy policies in Europe), and some other contextual factors are shifting, such as social context and transmission principles (who else is using the application and who can access the data).

Nissenbaum challenges the most common model of regulating digital data flows, which is called "transparency and choice," or more often "notice and consent," meaning that users have to decide whether to opt-in or opt-out of given information-flow practices: "Transparency-and-choice appears to model control because it allows individuals to evaluate options deliberately and then decide freely whether to give or withhold consent" (Nissenbaum 2011). This model treats users as free actors in a free market and, according to Nissenbaum and many others, has failed. With respect to even more recent technological developments, it is clear that her diagnosis of 2011 stills holds up: not only is "online activity … deeply integrated into social life and is radically heterogeneous", it is simply impossible for users to be provided with the necessary information on data flows. Even more so, data practices

3 The European General Data Protection Regulation defines 'data subject' as "an identifiable natural person … one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person" (GDPR 2016).

have begun operating on a whole new level, gathering and seeking new and different types of information with data mining techniques that often make it difficult to even hold up the notion of "personal information".

The contextual integrity approach does not elaborate much on the technological dimension of privacy practices and concepts. Although Nissenbaum acknowledges digital and online practices as being an "online life" which is "heterogeneous and thickly integrated with social life" (Nissenbaum 2011), the notion of "contextuality" tries to model so called online life based on the already existing norms of so-called social life: consulting a search engine is in this regard "akin to … searching a library catalogue" (Nissenbaum 2011). Therefore, existing social (mainly meaning: legal) norms are to be referred back to norms for "online life". The internet, in this regard, remains a public good and not a political economy. To Nissenbaum, privacy remains a concept shaped by shifting social practices, yet, she somehow does not integrate technology into her notion of sociality. With regard to the urgent problem of "online cases without straightforward social precedents", Nissenbaum argues, their "ends, purposes, and values" need to be considered by "working from there back to norms".

Therefore, the approach of "contextual privacy" does not consider technology itself to be the driving force constructing notions and practices of privacy. Though Nissenbaum's approach successfully shows how social norms are flexible and historically changing[4], it refers to a rather conservative notion of social norms. Privileging legal norms in influencing social contexts results in this approach in a concept of norm that is non-technological, or at least non-digital.

Other privacy theories often referenced in computer science research try to determine the shifts in privacy concepts in contemporary socio-technical settings. For example, Palen and Dourish (2003) argue that technology constantly shifts disclosure, identity and temporal boundaries. Those boundaries have to be considered when designing socio-technical systems and evaluating their impact.

Still, Nissensbaum's theory, which stresses the importance of multiple factors of privacy – social context, senders, receivers, transmission principles –, has seen wide adoption in emerging technologies. It has also led to an explosion of privacy settings allowing individual users to define what they think is appropriate to disclose and to decide whom they want

4 As Tobias Matzner emphasized, Nissenbaum's "contextual integrity" was one of the early approaches that successfully challenged the idea of a presupposed "autonomous subject" (Matzner 2018, 83).

to disclose it to and for what reasons and purposes. Despite Nissenbaum's critique, these more fine-grained access control systems leave the notice-and-consent-model intact, as well as the liberal bias that assumes a free and rational individual.

This contradiction can, to some extent, be explained by the fact that this type of access control is easy to operationalize: systems are ultimately still designed according to the notice-and-consent paradigm. While early smartphone operating systems requested a one-time consent to the privacy policy during installation time of a mobile app, companies have changed the consent processes to a more contextual consent scheme.[5] Today users are asked each time (or at least the first time) an application requests access to an information (e.g. an app that wants to take a photo) – respecting the context of the decision. And while previous versions of Android and iOS only asked the user to decide based on technically described access permissions (e.g. "Allow app to write to external storage"), current operating systems allow developers to describe senders, receivers and purposes of a data flow before the user is asked to permit it (e.g. "This app needs access to your phones SD card to store photos").

Still, privacy is often not the main driving factor behind the decision and no one wants to or can make a complex decision every time a status update is posted or a favorite app is used. Therefore, the following research questions emerged: which social norms are ultimately influencing individual decisions, and how can those decisions be made easier? In this context, current IT trends like personalization, machine learning, and automation come into play.

## Nudging Users: Technically Implementing Privacy Norms

One example of personalization in privacy is nudging, an approach from behavioral economics which aims to help users make the "right" decision. Behavioral economics is actively engaging in privacy research, drawing from Nissenbaum's concept of "privacy factors", and from psychology and different economic approaches. Yet, contrary to Nissenbaum, "Privacy Nudging" is particularly interested in "normative design research" (Acquisti 2009), which integrates research "results about cognitive and behavioral

5 In 2015 Android changed the permission model to runtime permissions, which can be controlled more easily.

biases in privacy and security decision making" by users into the design of privacy-enhancing technologies.

Behavioral research has tried to show that participants in studies, especially when confronted with technical systems, are easily convinced or tricked into disclosing sensitive information and to act contrary to their beliefs. In the early days of online shopping, Berendt et al. (2005) discussed that participants in their shopping experiment would easily give away personal details in a dialog with a virtual avatar to receive tailored recommendations, regardless of how privacy-sensitive they claimed to be. Besides the differences between claimed privacy preferences and actual behavior, other research has highlighted participant's irrationality with respect to privacy protection assurance. John et al. (2009) have shown that the mere presence of a consent form and notifications that the data will be handled carefully actually decreases the willingness to disclose information (and to be honest).

Not only might the task of having to decide individually what (not) to disclose in various different media settings and contexts be overwhelming, but often individual privacy seems disconnected and distant from the immediate benefit of disclosing data. Therefore, privacy researchers have started to look into "privacy nudging". Nudges refer to short information fragments that are given when a privacy relevant decision has to be made; they are thought to convince people to make more privacy-friendly decisions (Acquisti 2009). Nudges have been tested on social networks where users have been made aware of the actual audience or the perceived sentiment before posting (Wang et al. 2013), and on Android, where participants were "nudged" into restricting the access rights of apps by showing them how often they used certain information like the user's location (Almuhimedi et al. 2015).

Nudging therefore deals with a problem that has already become apparent in discussing Nissenbaum's approach: privacy research is challenged by having to acknowledge that the modern idea of a "free individual" fails to translate to rational and responsible decision making. It therefore suggests an approach that Acquisti calls the "soft paternalism solution" (Acquisti 2009, 84), which insists on a mild solution, rejecting a "strong paternalism" which would install technical measures to ban users from disclosing information (one might very well think of children's mobile operating systems which ban specific applications, for example). Calling for a "soft paternalism" instead might "provide [additional] context to aid the user's

decision" or "it might alter the system's default settings" so that some data would not be disclosed unless "the individual sets them that way" (ibid.).

Rejecting the behaviorist idea of a "homo oeconomicus" and the methodologies of decision calculus that come along with it, Acquisti tries to take into account the less stable, more "emotional components" (ibid., 82) of decision making, indicating that calculating human behavior is dependent on shifting social desires that, because of immanent irrationality, need to be measured in crowd-sourced studies. Taking into account methodologies from psychology, Acquisti hopes to "be able to reconcile the human need for publicity with our ostensible desire for privacy" (ibid., 82). Yet, such anthropological notions – the naturalization of "human needs" and "desires" – do indeed inform research methodologies. That becomes clear with regard to Acquisti's illustration of his arguments: with depreciation, or maybe even articulation of slight technophobia, he refers to a Facebook group that became popular in 2007 and whose name clearly indicates the practices of self-documentation it triggers: "30 Reasons Girls Should Call it a Night". It is a group with 150,000 members in which (mostly) young woman uploaded pictures of (presumably) themselves depicting intoxicated persons.

Yet the example of such Facebook groups might help us to re-frame the problem: rather than conceptualizing media technologies as the platform for the articulation of natural human desires – and therefore suspending media technologies from shaping such desires – Facebook groups like the one mentioned above illustrate vividly how media techniques and media practices actually produce new meanings of privacy. Apparently, in privacy research, there is some obliviousness involved to the intrinsic digital media effects which remediate privacy.

When privacy is framed using the idea that users need to be nudged, it shows how privacy research struggles with a sort of vacuum that remains of the failed concept of a "free and rational individual". Behavioral privacy constructs a notion of privacy as decision making which demands trade-offs in a "privacy calculus" (cf. Dinev 2006). In other words, it renders apparent the liberal bias of this research field and its continuing struggle to deal with it. The liberal bias – the assumption of a free decision being taken – is intertwined with the behavioral, a mechanistic concept of the user needing to be nudged, or regulated, into a specific behavior. Both views come together in the idea of privacy nudges, which try to leverage the behavioristic view for the sake of a positive effect on privacy.

## Eliciting Privacy Norms – Modeling Users

Drawing from both research approaches – Nissenbaum's suggestion of identifying specific "factors" of privacy like actors, information types and transmission principles, together with "nudging" to implement norms into technology design – one contemporary take on privacy-enhancing technologies in computer science envisions modeling and predicting privacy preferences.

Privacy and security research often demonstrates how a technology has violated or is undermining privacy. But some researchers have started to anticipate and envision future ways of data collection and processing to get ahead of the problem. Years before the invention of the smartphone, Lederer et al. (2003) asked, in a crowd-sourced study, what information about their location and their current activity participants would share with specific parties in different scenarios, using a cell phone. They found that *who* is collecting information is more relevant than *what* information is to be disclosed. Nowadays, the focus of studies relies on what is called the "Internet of Things" (IoT), where a myriad of different devices is, in theory, able to capture a variety of different data points (see e.g. Naeini et al. 2017, Lee and Kobsa 2017). Actual IoT-applications are, however, not yet in widespread use. The scenarios in those studies were drafted by describing specific factors that are also considered by Nissenbaum, for example, who is collecting which types of information for what purposes, and who else may get access. Scenarios are then modeled to ask participants about their privacy preferences, with respect to, for example, video cameras in different locations (at home, at work, in public restrooms, in public libraries). The scenarios also describe different types of data being collected (cameras can simply record the scenery or apply facial recognition algorithms to the image stream). By asking participants how comfortable they would feel in each situation and whether they thought it was appropriate, such studies also aimed to gain knowledge about the evolving social norms that are shaping privacy notions in emerging technological settings.

The number of possible combinations of scenarios is not only overwhelming for users but for technology designers and researchers as well. Being overwhelmed sometimes leads to the adoption of stereotypical ideas about how to handle the lack of knowing what "preserving" privacy could mean: Li et al. (2017) tried to globalize privacy preference prediction by including what they called a "cultural" factor, measured by six "cultural values" which included what was called the "acceptance of inequality" as well as "masculinity".

Such examples, as well as the call for a "soft paternalism" mentioned above, show that empirical research methodologies and technical design processes can reproduce historically situated cultural, sexual and racial differences. Moreover, in the discourse between computer and communication science as well as legal and public policy it is critically discussed that "profiling"[6] as a method and as a cultural technique tends to fall for the same biases and stereotypes that privacy researchers have also criticized in data mining and big data (see e.g. Gutwirth and Hildebrandt 2010).

Still, data mining and big data are the most influential approaches in contemporary privacy research in computer science, which is paradoxical considering that the ever-evolving digital environments raised the problem of what privacy could mean in the first place. Thus, research suggests that even more information needs to be collected about users to "preserve" their privacy.

More importantly, it comes with a risk: identified privacy profiles can be mistreated as target groups thereby allowing for the economic exploitation of privacy assistants. In this sense, data mining and big data techniques create the whole problem of privacy, producing a sheer never-ending feedback of data that needs to be mined for privacy purposes (see also Kapsner and Sandfuchs 2015). "Privacy" becomes the problematic object of knowledge of data mining technologies itself. Rather than trying to "preserve" privacy, the discourse actively produces it, undermining its original intent.

## Automating Privacy Assistants

Having attempted to model privacy preferences, the resulting profiles can be used to adjust privacy settings semi-automatically. As outlined in the "Privacy Research Roadmap for the Computing Community" (Cranor et al. 2016), "user-oriented machine learning techniques" should be used "to help users refine their privacy settings, leveraging user feedback to suggest modifications to these settings." Researchers are trying to build privacy

6 Categorizing consumers by their privacy preferences has a long history in the US. Alan Westin consistently categorized the (US) population for over 30 years into three groups: the fundamentalists, unconcerned and so-called privacy pragmatists, the latter presenting the majority (Kumaraguru and Cranor 2005). According to Draper (2017) the notion that a majority of the population does not have a fixed opinion on questions regarding the disclosure of their data, but makes decisions based on the facts presented to them, has been prevalent in policy making in the US.

assistants as agents for users, which would run on their phone and mediate the complex task of managing privacy settings.

The easily controllable context of a mobile operating system has also been used by researchers to create prototypes of those assistants. Liu et al. (2016) created an application that puts users in one out of four categories based on their answers to three to five questions about whether they want to allow apps from different categories to access data on their phone. Similarly, Wijesekera et al. (2017) "built a classifier to make privacy decisions on the user's behalf by detecting when context has changed and, when necessary, inferring privacy preferences based on the user's past decisions and behavior." Both claim to accurately predict user preferences, proving their point by showing that the recommendations their systems make are not rejected, or are in line with previous decisions made by the user. Accepting recommendations is therefore read as proof of constructing realistic profiles. Yet critique on big data has highlighted that the construction of such profiles tends to take correlations for causalities: studies barely explain factors and categories which construct specific profiles (Degeling 2014).

But looking at the results of such sophisticated data analysis of factors and profiles, it becomes clear that many of the results are not always surprising. For example, the studies mentioned in the previous section, on future data collection scenarios, found that the most influential factor was whether the data was being collected in a public or a private space: Few participants felt comfortable with video cameras, controlled by a third party, being set up in their home. Here the studies (or the participants) are in line with Nissenbaum's approach of applying known norms to new contexts and therefore drawing implicitly on conventional modern dichotomies of public and private, thereby disregarding their reconfiguration by digital techniques.

Other scenarios lead to inconsistent results: for example, whether it is appropriate to measure health data at the workplace, or to continuously track users in buildings to provide information in case of emergencies. Users do not quite know how to deal with such media environments in terms of privacy issues. Consequently, Naeini et al. (2017) emphasized those scenarios which anticipate privacy issues emerging in unknown media environments as those where individual preferences seem to differ, and therefore profiling with privacy assistants would be especially helpful. That means developing privacy assistants aims particularly at implementing norms for such purposes. But it might very well be that this approach fails when norms have not yet been sufficiently discussed or when norms simply not seem realistic yet because they require some

investment of imagination for anticipating the media environments that elicit them. This carries the risk that uncertainties in the results of a crowdsourced survey are modeled away in diverse profile categories, instead of acknowledging the need for a debate on what data practices are acceptable as Nissenbaum approached it. Those categories tend to reproduce and remediate conventional, modern privacy practices. Paradoxically, this can lead to the enforcement of non-technical notions of social norms via digitally automated norm prediction.

## Summary

Privacy research is struggling to keep up with the rapid pace of technological development. It is trying to adapt to the dynamic environments of apps and the IoT that overwhelm users as well as researchers and developers. As a consequence, research addresses the problematic opacity of privacy by translating it into empirical methodologies which favor crowdsourcing and data mining techniques to create privacy profiles, and yet do not meet the complexities of contemporary relations of digital environments and social practices. Instead of experimenting with methodologies that conceptualize privacy as an entanglement of actors, media practices and media environments, such empirical approaches tend to ignore that technologies and media practices actually produce privacy practices, as well as processes of subjectivation and individuation.

While trying to adapt to new technological environments, current research aims at eliciting new norms by modeling privacy profiles with the help of scenario based crowdsourcing studies. But by relying on data mining techniques, instead of debating what privacy could potentially mean, empirically found preferences result in technically implemented norms.

Since the idea of informational self-determination relies on individual decision making, privacy assistants are created that are thought to nudge users to heed to their stated privacy preferences which would not conform to their actual media practices. Yet these assistants make use of the same (privacy invasive) data mining techniques that create the problem. Additionally, it needs to be taken into account, that digital media techniques and practices are reconfiguring modern notions of the social distribution of public and private spheres.

## References

Acquisti, A. 2009. "Nudging Privacy: The Behavioral Economics of Personal Information." *IEEE Security and Privacy.*

Almuhimedi, H., F. Schaub, N. Sadeh, I. Adjerid, A. Acquisti, J. Gluck, L. Cranor, and Y. Agarwal. 2015. "Your Location Has Been Shared 5,398 Times!: A Field Study on Mobile App Privacy Nudging." In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, 787–796. CHI '15. New York, NY, USA: ACM. doi:10.1145/2702123.2702210.

Berendt, B., O. Günther, and S. Spiekermann. 2005. "Privacy in E-Commerce: Stated Preferences vs. Actual Behavior." *Commun. ACM* 48 (4): 101–6. doi:10.1145/1053291.1053295.

J. David Bolter, Richard A. Grusin, Remediation: understanding new media, Cambridge, Mass (MIT Press) 1999.

Cranor, L., T. Rabin, V. Shmatikov, S. Vadhan, and D. Weitzner. 2016. "Towards a Privacy Research Roadmap for the Computing Community." *ArXiv:1604.03160 [Cs].*

Degeling, M. 2014. "Profiling, Prediction und Privatheit: Über das Verhältnis eines liberalen Privatheitbegriffs zu neueren Techniken der Verhaltensvorhersage." In *Medien und Privatheit*, edited by S. Garnett, S. Halft, M. Herz, and J. M. Mönig, 69–92. Medien, Texte, Semiotik 7. Passau: Verlag Karl Stutz.

Dinev, T., and P. Hart. 2006. "An Extended Privacy Calculus Model for E-Commerce Transactions." Information Systems Research 17 (1): 61–80. https://doi.org/10.1287/isre.1060.0080.

Draper, N. A. 2017. "From Privacy Pragmatist to Privacy Resigned: Challenging Narratives of Rational Choice in Digital Privacy Debates." *Policy & Internet* 9 (2): 232–51. doi:10.1002/poi3.142.

EU General Data Protection Regulation (GDPR), Available at http://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32016R0679.

Foucault, M. 1996. *Diskurs und Wahrheit: die Problematisierung der Parrhesia. Sechs Vorlesungen, gehalten im Herbst 1983 an der Universität von Berkeley/Kalifornien.* Berlin: Merve.

———. 2006. *Discourse and Truth: the Problematization of Parrhesia: 6 lectures given by Michel Foucault at the University of California at Berkeley, Oct-Nov. 1983.* Available at: foucault.info/downloads/discourseandtruth.doc : 66/67.

Gürses, S., B. Preneel, and B. Berendt. 2009. "PETs under Surveillance: A Critical Review of the Potentials and Limitations of the Privacy as Confidentiality Paradigm." In *Proceedings of HOTPETS 2009*. Seattle, Washington, United States.

Gutwirth, S., and M. Hildebrandt. 2010. "Some Caveats on Profiling." In *Data Protection in a Profiled World*, edited by Serge Gutwirth, Yves Poullet, and Paul De Hert, 31–41. Springer Netherlands.

John, L.K., A. Acquisti, and G. Loewenstein. 2009. "The Best of Strangers: Context Dependent Willingness to Divulge Personal Information." https://dx.doi.org/10.2139/ssrn.1430482

Kapsner, A., and B. Sandfuchs. 2015. "Nudging as a Threat to Privacy." *Review of Philosophy and Psychology* 6 (3): 455–68. doi:10.1007/s13164-015-0261-4.

Kumaraguru, P., and L. Faith Cranor. 2005. "Privacy Indexes: A Survey of Westin's Studies." *ISRI Technical Report*.

Lederer, S., J. Mankoff, and A. K. Dey. 2003. "Who Wants to Know What When? Privacy Preference Determinants in Ubiquitous Computing." In *CHI'03 Extended Abstracts on Human Factors in Computing Systems*, 724–725. ACM.

Lee, H., and A. Kobsa. 2017. "Privacy Preference Modeling and Prediction in a Simulated Campuswide IoT Environment." In *Proceedings 15th IEEE Conference on Pervasive Computing and Communications*. Kona, HI.

Li, Y., A. Kobsa, Bart P. Knijnenburg, and M-H. Carolyn Nguyen. 2017. "Cross-Cultural Privacy Prediction." *Proceedings on Privacy Enhancing Technologies* 2017 (2). doi:10.1515/popets-2017-0019.

Liu, B., M. Schaarup Andersen, F. Schaub, H. Almuhimedi, S. Zhang, N. Sadeh, Y. Agarwal, and A. Acquisti. 2016. "Follow My Recommendations: A Personalized Privacy Assistant for Mobile App Permissions." In *Proceedings of the 12th Symposium on Usable Privacy and Security (SOUPS 2016)*, 27–41.

Matzner, Tobias, Der Wert informationeller Privatheit jenseits von Autonomie, in: Steffen Burk et al. (Ed.), *Privatheit in der digitalen Gesellschaft*, Issue 10, Berlin (Duncker & Humblot) 2018 (Internetrecht und Digitale Gesellschaft), 75–94.

Naeini, P. E., S. Bhagavatula, H. Habib, M. Degeling, L. Bauer, L. Cranor, and N. Sadeh. 2017. "Privacy Expectations and Preferences in an IoT World." In *Proceedings of the 13th Symposium on Usable Privacy and Security (SOUPS 2017).* Santa Clara, CA: Usenix Association.

Nissenbaum, H. 2004. "Privacy as Contextual Integrity." *Washington Law Review* 79: 119.

———. 2011. "A Contextual Approach to Privacy Online." *Daedalus* 140 (4): 32–48. doi:10.1162/DAED_a_00113.

Palen, Leysia, and Paul Dourish. 2003. "Unpacking 'Privacy' for a Networked World." In *Proceedings of the SIGCHI Conference on Human Factors in Computing*, 129. ACM Press. doi:10.1145/642611.642635.

Wang, Y., P. G. Leon, K. Scott, X. Chen, A. Acquisti, and L. Cranor. 2013. "Privacy Nudges for Social Media: An Exploratory Facebook Study." In *Proceedings of the 22nd International Conference on World Wide Web*, 763–770. WWW '13 Companion. Geneva, Switzerland.

Wijesekera, P., A. Baokar, L. Tsai, J. Reardon, S. Egelman, D. Wagner, and K. Beznosov. 2017. "The Feasibility of Dynamically Granted Permissions: Aligning Mobile Privacy with User Preferences." *ArXiv:1703.02090 [Cs]*, March.