

Towards an Engineering Model of Privacy-Related Decisions

Joachim Meyer^(✉)

Department of Industrial Engineering,
Tel Aviv University, Tel Aviv-Yafo, Israel
jmeyer@tau.ac.il

Abstract. People make numerous decisions that affect their own or others' privacy, including the decisions to engage in certain activities, to reveal and share information or to allow access to information. These decisions depend on properties of the information to be revealed, the situation in which the decision is made, the possible recipients of the information, and characteristics of the individual person. System design should ideally protect users from unwanted consequences by allowing them to make informed decisions, at times blocking users' ability to perform certain actions (e.g., when the user is a minor). The development of alerting and blocking mechanisms should be based on predictive models of user behavior, similar to engineering models in other domains. These models can be used to evaluate different design alternatives and to assess the required system specifications. Predictive models of privacy decisions will have to combine elements from normative decision making and from behavioral, descriptive research on decision making. Some major issues in the development and validation of such models are presented.

Keywords: Privacy · Decision making · Models · Cognitive engineering

1 Introduction

Privacy has become a major concern in people's interaction with technologies. The storing of vast amounts of information and the possible access to this information by other people, by governmental agencies, or by companies and other organizations expose people to the threat of others gaining information about them on almost all aspects of their lives. The people who access the information are usually unknown to the individual, may use the information against the individual's interest, and the individual generally has no way to redress the issue.

At the same time, people also gain benefits from revealing information. They receive personalized services, such as adapted product offerings on websites, they may have access to location-related recommendations, they can get emergency support when they are in an accident (if they are connected to a system that monitors their status and location), etc. The rapidly blooming field of social networks is based entirely on people's willingness, and even desire, to share personal information. Thus sharing information and having others access one's information are not necessarily bad, nor are they necessarily good. Rather, as is usually the case, they have both positive and negative sides.

1.1 Privacy Decision Making

The notion that providing access to one's personal information can have advantages and disadvantages for a person has been known for a long time. It implies that people may want to weigh the advantages and disadvantages and choose whether to reveal information. This idea is central in the definition of privacy, proposed by Westin (1967), as "the claim of individuals, groups or institutions to determine for themselves when, how, and to what extent information about them is communicated to others." He recognizes the dynamic nature of these choices by also stating that "... each individual is continually engaged in a personal adjustment process in which he balances the desire for privacy with the desire for disclosure and communication ..."

Thus one can analyze a person's privacy related actions as the result of decision processes. The active sharing of information, the engagement in activities that generate information, or the failure to prevent private information from becoming public, can all be seen as results of decision processes. According to economic normative models of decision making (such as the Expected Utility Model), the decisions should be made, based on the expected outcomes when information is revealed and when it is not. However, for privacy decisions, as for decisions in most other domains, people's actual decision making deviates from the prescriptions of classic economic models (e.g., Acquisti and Grossklags 2005). Furthermore, privacy-related decisions are inherently difficult to analyze, even with simple economic models, since the consequences (costs and benefits) occur at different points in the future, they occur with some (largely unknown) probabilities, and they are in most cases not directly translatable into monetary values.

Privacy-related decisions have a variety of outcomes that have very different importance and meaning for different people. Basically, there are three major categories of outcomes (see Table 1):

Social. Privacy-related decisions can affect the relations a person has with other people. Communicating with others, by, for instance, posting on social networks, can provide various benefits. These include communicating about a person's status, creating and managing the impressions others might have about the person, maintaining relationships with others, etc. These actions may also have negative consequences, such as offending certain people, or information reaching people who were not supposed to see it (e.g., the boss seeing an employee intoxicated).

Economic. Sharing of information may be motivated by economic benefits a person receives when agreeing to share the information. Examples are people joining customer loyalty programs, where they receive minor benefits for agreeing to reveal their identity (e.g., swipe their card) whenever they perform a purchase. Revealing information may also have negative economic implications. For instance, if an insurance company obtains information showing that a person is at an increased risk for some chronic disease, the company may raise the person's insurance rates.

Functional. Sharing of information may provide functional benefits. For instance, one must share location information to receive location-dependent services or recommendations. Sharing one's identity with a website allows the site to customize the information to the individual's characteristics, etc. However, the shared

information may also be misused, as happens in the most extreme case when it is used by a criminal, for instance to perform identity theft.

Ideally people should make privacy-related decisions after considering all possible consequences. This is obviously problematic, and it is unrealistic to expect that people explicitly evaluate and weigh each of the consequences (and there may be very many), their probability, and their utility in some common measure. However, it may be possible to predict to some extent which possible consequences people consider, depending on the prior information they have and the display of relevant information by the system.

Table 1. Some types of costs and benefits related to privacy

	Benefits	Dangers and costs
Social	Communicate with others, impression management, maintain relationship	Unintended consequences of information reaching people
Economic	Incentives from sharing information	Possible negative effects (increased insurance rates, etc.)
Functional	Improved services when functions are shared (location based recommendations)	Possible misuse of information (identity theft, etc.)

2 Privacy Engineering

The design of systems that take privacy into account has to deal with numerous aspects of privacy, including the encryption of information, the protection of information from unwanted access, the limitation of information collection, etc. Eventually these boil down to technical decisions made by the people who develop, deploy and maintain systems. These are part of the engineering of systems, and hence the engineering of privacy may be a relevant term. Spiekerman and Cranor (2009) published an analysis of the development of privacy-sensitive systems, with the title “engineering privacy”. They describe two approaches in the engineering of privacy. One, which they name “privacy by architecture”, is the prevention of privacy violations by designing the system so that the data collection will be minimal or privacy violations will ideally be impossible. The other approach, “privacy by policy”, deals with cases in which the possibility of privacy violations still exists. Then system designers need to inform users about possible privacy risks and must leave users the choice whether to expose themselves to such risks or not (the “notice and choice” approach).

Gurses (2014) points out that building systems that cope appropriately with the plethora of legal and societal aspects of privacy is a “bewilderingly complex” task. She describes three major approaches in privacy research in computer science, which can form the basis of the engineering of privacy: (1) Privacy as confidentiality, which means limiting the amount of information collected and the possibility that information can be revealed to others; (2) privacy as control, which means creating mechanisms that allow people to control the collection and use of data about them; and (3) privacy as

practice, which considers privacy as part of social interactions in which people exchange information and signals about the use of the information. Gurses doubts that it will be possible to engineer privacy. Rather, this may be a, perhaps unattainable, ideal towards which engineers should strive.

3 Cognitive Engineering

The design of systems that allow people to take adequate control over their privacy requires the understanding of people's decision making process. This includes observing how people obtain information on which they base their decisions, how they use this information to evaluate different alternative actions, and how they choose a particular course of action. The information and the available actions are often displayed by computers, and action implementation is mediated by a computer. Thus, in the context of privacy, a computer may (or may not) tell a person what information is collected if he or she grants a program a specific permission. The computer may also inform the person (correctly or incorrectly) what will be done with this information and how it will be protected. The person's decision should eventually be based on the evaluation of this information, together with some evaluation of the expected benefits from providing the information.

"Cognitive engineering" studies systems in which people and computers interact to perform some task, or as Vicente (1999) defined it, "Cognitive engineering is a multidisciplinary endeavor concerned with the analysis, design, and evaluation of complex systems of people and technology". The field emerged from the attempt to understand and predict human performance in complex systems, such as advanced aircraft cockpits or the control rooms in nuclear power plants. It encompasses a variety of different approaches, ranging from qualitative, descriptive analyses to highly quantitative predictive and analytical models.

3.1 Quantitative Models

Among the different approaches in cognitive engineering the attempt to create an engineering process of the specification and design of human computer systems might be particularly valuable in the context of privacy. In this engineering process (as in engineering in general), the decisions and actions should be based on quantitative models of systems and operator actions in the systems. The American Institute of Aeronautics and Astronautics (1998) defined a model as a "Conceptual / mathematical / numerical description of a specific physical scenario, including geometrical, material, initial, and boundary data."

Such work should aim to generate models of people's decisions, given specific system properties and usage conditions. These models can be used for a number of purposes. For one, they can support design decisions, and they can help develop specifications for the system. For instance, they can be used to decide which functions to automate, so that the computer will perform them and which to leave to the human operator (a process named "function allocation").

In addition, the models can be used for interface design, including the decisions what to display to the users and what actions users should be able to perform (which will affect the choice of displays and input devices for a system). At times exist regulations that specify which information must be provided to the user, such as Article 10 in the EU Data Regulation Directive. It states that people about whom data is collected must be informed about the collection of the data, who collects it, for what purpose is it collected, and other relevant information. The model can be used to predict the conditions that will provide optimal presentation of this information, so that people will become aware of it without too strongly disrupting their interaction with the system.

The analysis can also help in the development of training and simulation facilities by supporting various decisions, such as to understand what skills and knowledge are required for a particular task? What situations should be trained? How frequently should refresher training take place? Etc.




Fig. 1. The relation between models, observations and controlled experiments.

Models are part of the continuous attempt to observe, describe, analyze and predict phenomena. As shown in Fig. 1, models are closely related to observations of the world and to controlled experiments. They are based on intuitions and observations, and they inform interventions, which should be based on conclusions drawn from models. They also generate hypotheses that can be tested in controlled experiments, and they are adjusted, based on the results of the experiments. Finally, the design of the experiment should be informed by observations of the world and ideally resemble the conditions that exist in the world (to ensure the generalizability of the results from the experiment to situations outside the lab, the so-called “external validity”). The process of model development, adjustment and validation is a continuous process, which can never end. There is never a “correct model” that has been reached. Rather, as the statistician George E. Box stated “Essentially, all models are wrong, but some are useful” (Box and Draper, 1987; p. 424).

The engineering model is first a model of the system, incorporating the functions, properties and the behavior of the system that is modeled (see Fig. 2). For instance, the engineering model of a privacy component of a system will describe the way information in the system is stored, who can access it, and how this access is done. The model should not only model the system, but it also needs to model the environment in




Fig. 2. Aspects to be addressed by a cognitive engineering model.

which the system is deployed. Thus, properties of the environment, such as the likelihood of privacy transgression, the severity of the expected outcomes, the information available to detect privacy problems, etc., need to be incorporated in the model. Finally, in a human-computer model the model also needs to describe and specify the user. It needs to specify relevant stable, general user properties that may affect the user's behavior, such as age, gender, education, cultural background, etc. In addition, it may include situation-specific user variables, such as the individual user's experience with a given situation, the information the user received (perhaps through word of mouth from others) about the situation, etc. The three domains – system, environment and user – are not independent, and properties of the system may alter the user and to some extent also affect the situation in which the system is used (because users may, for example, choose to avoid certain situations).

3.2 The Triad of Privacy-Related Behaviors

The modeling of privacy-related decisions is complicated by the fact that these are not single decisions. Rather, they are decisions that are part of an ongoing sequence of actions, where decisions made at earlier points in time will affect the set of alternatives among which people choose in the future, the available information for making decisions, and the expected outcomes of decisions. Essentially, as in cyber-security decisions (Ben-Asher et al. 2010; Möller et al. 2011), there are three different time perspectives of decision making, which all need to be considered, and which affect the decisions a person makes at some moment in time:

1. *Precautionary actions.* These actions are done in advance, and often only once, when a person begins to use a system. They include the choice of system settings, the installation (or disabling) of functions and services, and the installation and setting of protective mechanisms. These decisions are made, based on the information available to the user at the time at which she or he begins to use the system, and they may often

not be adjusted when the use of the system, the system itself or the environment change. One of the major problems, inherent in these actions, is the fact that people are not very good in deciding how to adjust a system and its settings. These decisions depend on the available information, and they will often deviate greatly from optimal settings, if such can be computed (Botzer et al., 2010).

2. *Exposure to a privacy risk.* These are decisions to engage in activities that make privacy risks possible (e.g., posting information on social networks, providing identifying information when signing up for a service, allowing a mobile app to collect information, etc.). These decisions are made continuously, whenever people encounter situations in which they might reveal information that may be sensitive at some point in time and under certain conditions.

3. *Actions when a negative event occurs.* These decisions are made at the moment when possibly sensitive information is revealed to somebody who is not supposed to have access to this information. Unfortunately, in most cases, this will happen without the person about which the information is revealed having control over the event or even knowing about it (thus he or she will usually not be able to make a decision and take action at this point). In systems in which people are somehow involved (e.g., are alerted when someone tries to access their information or tries to download personal material), the involvement is tied to alerts and warnings. These can be followed, or, quite often, especially if they occur frequently, they will be ignored. Overall, people's responses to alerts differ systematically from the optimal responses to such information, if these can be computed (e.g., Meyer, Wiczorek and Günzler, 2014).

The three behaviors, related to these three times, can be considered a “triad of privacy-related behaviors”, in parallel to the “triad of risk-related behaviors” we describe in the context of cyber-security (Ben-Asher and Meyer, 2015).

An additional challenge in the modeling of privacy-related decisions is that the decisions are not necessarily the ones a classic economic decision-making model would prescribe. The expected value maximization can perhaps be a starting point for developing a model, but it needs to be adjusted to properties of the decision process, such as risk aversion and non-linear utility functions, and bounded rationality due to limited time and cognitive abilities. It must also take into account characteristics of behavioral decision making, such as deviations from classical probability theory in the estimation of the likelihood of outcomes and the computation of preferred alternatives.

4 Discussion

The design of information systems and the mechanisms involved in having people manage their personal information should be based on systematic analytical tools, similar to the tools used in other engineering disciplines. Such models should combine an understanding of the technical, as well as the behavioral aspects of a system and a situation.

To develop such models, we should adapt the standards and views of engineering modelers. For one, in contrast to scientific models which often strive to be as accurate as

possible, we aim to develop as simple models as possible. Models need not be complete (and actually never can be). Rather, they should provide sufficient information to be used to make the decisions for which they were developed. If the outcomes from choosing either of a number of different design alternatives are very similar, it will not particularly matter which of the alternatives one chooses. Thus the accuracy of model predictions needs not to be very high.

Also, models should be easy to develop. The development should be doable with relatively simple tools, and even people with limited experience in modeling should be able to develop models. To do so, it may be necessary to develop software tools that can support and guide the modeling process. Models should also be easy to communicate. The modeler should be able to present the outcome of the modeling in a simple and convincing way to stakeholders for whom the model predictions are relevant. And finally, models need to be verifiable, so that people who want to inspect the model can relatively easily see if model predictions were computed correctly and can recreate the computations leading to these predictions.

Although we strive to develop simple models, modeling of privacy decisions is inherently difficult. There are large individual differences in people's preferences and the factors they consider when evaluating outcomes, in most cases consequences of choices are not known when choices are made, and the values of consequences may change over time (a person who prides himself of an active social life as a student, may be less happy to reveal this information after accepting an executive position). At the moment of the decision, the situational characteristics, the information that is salient, and the recent experiences a person has (or events the person heard about) may all affect the decisions. Thus the timing and context in which people make decisions will all have to be taken into account.

Gurses (2014) expresses some doubt about the possibility to engineer privacy. If by engineering privacy, one means that there will be full control of privacy including all its aspects, this statement is certainly correct. However, design decisions, and in particular decisions regarding the information provided to people regarding the implications of their choices, should be based on scientifically validated models, rather than on the intuitions, gut feelings, and impressions of software developers, designers or project managers. Even if such models currently often cannot provide predictions, their development helps to structure the thinking about the design decisions and can point to a subset of alternatives among which one may find an adequate solution to the problem. This may help us develop systems that are better adapted to protect users' privacy and that provide people with the ability to make the choices that are required to manage the collection and exposure of information about themselves.

References

- Acquisti, A., Grossklags, J.: Privacy and rationality in individual decision making. *IEEE Secur. Priv.* **3**(1), 26–33 (2005)
- AIAA (American Institute of Aeronautics and Astronautics): Guide for the verification and validation of computational fluid dynamics simulations, AIAA-G-077-1998. American Institute of Aeronautics and Astronautics Reston, VA (1998)

- Ben-Asher, N., Meyer, J.: The triad of risk related behaviors. Unpublished Manuscript (2015)
- Ben-Asher, N., Meyer, J., Parmet, Y., Möller, S., Englert, R.: An experimental microworld for evaluating the tradeoffs between usability and security. Usable Security Experiment Reports (USER) Workshop in the Symposium on Usable Privacy and Security (SOUPS), Redmond, WA, USA, 14–16 July 2010
- Botzer, A., Meyer, J., Bak, P., Parmet, Y.: Cue threshold settings for binary categorization decision. *J. Exp. Psychol. Appl.* **16**, 1–15 (2010)
- Box, G.E.P., Draper, N.R.: *Empirical Model Building and Response Surfaces*. Wiley, New York (1987)
- Gurses, S.: Privacy and security: Can you engineer privacy? *Commun. ACM* **57**(8), 20–23 (2014)
- Meyer, J., Wiczorek, R., Günzler, T.: Measures of reliance and compliance in aided visual scanning. *Hum. Factors* **56**(5), 840–849 (2014)
- Möller, S., Ben-Asher, N., Engelbrecht, K.-P., Englert, R., Meyer, J.: Modeling the behavior of users who are confronted with security mechanisms. *Comput. Secur.* **30**(4), 242–256 (2011)
- Spiekerman, S., Cranor, L.F.: Engineering privacy. *IEEE Trans. Softw. Eng.* **35**(1), 67–82 (2009)
- Vicente, K.: *Cognitive Work Analysis*. Lawrence Erlbaum Associates, Mahwah (1999)
- Westin, A.: *Privacy and Freedom*. Atheneum, New York (1967)

Privacy and Identity Management for the Future

Internet in the Age of Globalisation

9th IFIP WG 9.2, 9.5, 9.6/11.7, 11.4, 11.6/SIG 9.2.2

International Summer School, Patras, Greece,

September 7-12, 2014, Revised Selected Papers

Camenisch, J.; Fischer-Hübner, S.; Hansen, M. (Eds.)

2015, X, 267 p. 44 illus. in color., Hardcover

ISBN: 978-3-319-18620-7