Chapter 7
Risks of Profiling and the Limits of Data Protection Law

Bart Schermer

Abstract. Profiling and automated decision-making may pose risks to individuals. Possible risks that flow forth from profiling and automated decision-making include discrimination, de-individualisation and stereotyping. To mitigate these risks, the right to privacy is traditionally invoked. However, given the rapid technological developments in the area of profiling, it is questionable whether the right to informational privacy and data protection law provide an adequate level of protection and are effective in balancing different interests when it comes to profiling. To answer the question as to whether data protection law can adequately protect us against the risks of profiling, I will discuss the role of data protection law in the context of profiling and automated decision-making. First, the specific risks associated with profiling and automated decision-making are explored. From there I examine how data protection law addresses these risks. Next I discuss possible limitations and possible drawbacks of data protection law when it comes to the issue of profiling and automated decision-making. I conclude with several suggestions to for making current data protection law more effective in dealing with the risks of profiling. These include more focus on the actual goals of data processing and 'ethics by design'.

7.1 Introduction

Profiling, the application of profiles to individuate and represent a subject or to identify a subject as a member of a group or category (Hildebrandt 2008), is commonplace in our data-driven information society. While profiling may have many benefits for businesses, the government and citizens themselves, there are also potential risks for data subjects attached to profiling. To mitigate these risks, traditionally the right to (informational) privacy is invoked. However, given the

Bart Schermer
eLaw, Institute for Law in the Information Society, Leiden University, The Netherlands
e-mail: schermer@considerati.com

B. Custers et al. (Eds.): Discrimination & Privacy in the Information Society, SAPERE 3, pp. 137–152.
springerlink.com

rapid technological developments in the area of profiling and automated decision-making, it is questionable whether the right to informational privacy and more specifically data protection law (still) provide an adequate level of protection and whether they balance the interests of the actors involved effectively.

In this chapter I explore the possible risks associated with profiling and examine whether the current legal framework can mitigate these risks effectively.[1] I shall do so by seeking answers to the following questions:

- What risks does profiling pose for individuals (and groups)?
- How are these risks addressed by the current data protection framework?
- Does the current legal framework for data protection provide adequate protection whilst also taking into account the legitimate interest of profilers?

After answering these questions I examine what changes might be necessary in order to mitigate the risks posed by profiling.

7.2 Risks Associated with Profiling

While profiling can be a valuable aid for businesses and governments, profiling may also entail risks. Risks commonly associated with profiling are: discrimination, de-individualisation, stereotyping, information asymmetries, inaccuracy and the abuse of profiles.

7.2.1 Discrimination

Classification and division are at the heart of profiling. As such, discrimination is part and parcel of profiling. However, there are situations where discrimination is considered unethical and even illegal. This can occur for instance when a profiling exercise is focussed on characteristics such as ethnicity, gender, religion or sexual preference. But even without a prior desire to judge people on the basis of particular characteristics, there is the risk of inadvertently discriminating against particular groups or individuals.

7.2.2 De-individualisation

In many cases profiling is in large parts concerned with classification and thus there is the risk that persons are judged on the basis of group characteristics rather than on their own individual characteristics and merits (Vedder 1999). Group profiles usually contain statistics and therefore the characteristics of group profiles may be valid for the group and for individuals as members of that group, though not for individuals as such. For instance, people who live in a particular neighbourhood may have a 20% higher chance of defaulting their loan than the

[1] In discussing data protection legislation, I shall focus exclusively on the EU framework for data protection.

average person. This characteristic goes for the group (i.e., people living in that particular neighbourhood), for the individuals as members of that group (i.e., randomly chosen people living in the neighbourhood), but not necessarily for the individuals as such (i.e., for John, Mary and William who all live in the same neighbourhood). When individuals are judged by group characteristics they do not possess as individuals, this may negatively affect them (Custers 2010).

Group profiling may not only have direct negative effects on individuals, but may also lead to stigmatisation of group members. Moreover, divisions into groups can damage societal cohesion. When group profiles, whether correct or not, become public knowledge, people may start treating each other accordingly. For instance, when people start believing that individuals from a particular neighbourhood default their loans more often, they may conclude that those individuals live in a 'bad' neighbourhood.

7.2.3 Stereotyping

Closely related to the risk of de-individualisation and stigmatisation is that of stereotyping. A profile casts us on the basis of predetermined categories (e.g., 'valuable customer', 'young urban professional', but also 'security risk' or 'dodgy debtor'). For a profiling exercise to remain effective and efficient there are a finite number of general categories. These profiles are, almost by definition, incapable of accurately reflecting all the nuances of our personality. As such, the profile we fit will become a stereotype on the basis of which we are judged. Moreover, these profiles can also make it more difficult for a person to 'escape' the stereotype.

7.2.4 Information Asymmetries

A fourth risk associated with profiling is that it can lead to information asymmetries. In other words, through profiling, the position of the data controller improves with regard to the data at his disposal, whereas that of the data subject remains the same. This is a particular issue when the data subject is unaware of the profiling exercise, or does not have complete information about the profiling exercise. Information asymmetries may lead to an imbalance in the playing field between government and citizens, and between businesses and consumers, upsetting the current balance of power between different parties.

In the context of the relation between government and citizens, information asymmetries can also affect individual autonomy. If data mining indeed yields information the government can act upon, the government will have more power. Moreover, the fear of strong data mining capabilities on the part of the government may 'chill' the willingness of people to engage in political activities, given the fear of being watched. For this fear to materialise, profiling does not even have to be effective (Schermer 2007, p. 137).

In the context of the relation between businesses and consumers, information asymmetries may lead to unfair economic practices and discriminatory pricing. For instance, certain goods or services may be withheld from individuals, solely on the basis of them fitting or not fitting a particular profile. It is also possible to

adjust prices of goods and services on the basis of the profile of the individual. Charging different prices on the basis of particular characteristics (e.g., race, sex, or sexual preference) is likely a violation of anti-discrimination legislation.

7.2.5 Inaccuracy

A fifth risk associated with profiling is that profiles might be inaccurate. In particular there is the problem of 'false positives' and 'false negatives'. This means that people that in fact do not fit the profile are fitted within it (a false positive), or people that fit the profile are left outside of it (false negative). False positives and false negatives occur for various reasons, for instance because insufficient data is available, or the data is inaccurate. False positives and false negatives are a particular problem in automated decision making since there is no human intervention and it is not an adversarial process where both sides are heard. This is troublesome as it places the burden of proof on the side of the data subject: they must prove that they do or do not fit the profile.

7.2.6 Abuse

A final risk associated with profiling is that data controllers or third parties (for instance hackers) abuse profiles and/or the information contained therein. Possibilities for abuse arise in particular when the profile can be linked to an identified individual. A profile could for instance be made public leading to reputational damage for the data subject (e.g., the data subject is exposed as a dodgy debtor), or the (personal) data contained in the profile could be used for fraudulent purposes.

7.3 Privacy and Data Protection in Light of Profiling

To mitigate the risks mentioned in the previous paragraphs, traditionally the right to (informational) privacy is invoked. The right to informational privacy acts as a boundary against the free flow of information and thus ensures the protection of personal information. An important aspect of informational privacy is personal data protection. In particular in the context of the private sector, data protection legislation has become the most important aspect of informational privacy protection. Van den Hoven (2008, p. 311) lists four different moral reasons for protecting personal data. They are: 1) protection against information based-harm, 2) protection against informational inequality, 3) protection against informational injustice and, 4) the protection of moral autonomy.

Information based-harm
Because information can be used to cause harm (e.g., identity theft, fraud) or other serious disadvantages to data subjects, personal data needs to be protected from access by parties who wish to cause harm using personal data. Data protection sets

forth rules on access and security to personal data, thwarting the efforts of those who wish to cause harm.

Informational inequality
A second moral reason for the protection of personal data is that it reduces the negative effects of informational inequality. Since consumers are not always (fully) aware of the economic opportunities their personal data may present, and/or not in a position to trade their identity-relevant information in a fair and transparent market, they may be disadvantaged in the marketplace for identity-relevant information. Constraints on the flow of personal data need to be put in place in order to guarantee economic equality of arms, transparency and fairness (Van den Hoven 2008, p. 313).

While van den Hoven only describes the issue of informational inequality from a private sector perspective, it is also relevant in the context of the relationship between governments and citizens. In this realm, informational inequality is closely associated with personal autonomy. If the government knows a great deal about its citizens, but is not equally transparent, the balance of power is upset.

Informational injustice
A third moral reason for data protection is to avoid informational injustice. Informational injustice occurs when the boundaries of the 'spheres of access' are disrespected. People do not mind when there data are being processed for a legitimate goal (e.g., their medical data being used for their treatment). But if a sphere of access is disrespected (e.g., the medical data is being used in a job application procedure) informational injustice takes place.

Moral autonomy and moral identification
A fourth reason to invoke data protection rules is that they allow us to set a 'distance' between the outside world and ourselves. This distance is crucial for what van den Hoven calls 'shaping our own moral biographies' (Van den Hoven 2008, p. 316). Without the observing gaze of others we can freely develop our thoughts and our identity. Furthermore, it allows us to present ourselves to the outside world as we see fit. When the outside world can readily access personal data across a number of different contexts, the individual's freedom to shape our own moral biography is reduced.

These moral foundations for protecting personal data are also relevant when we observe the possible risks of profiling. For instance, stereotyping and de-individualisation encroach upon our sense of moral autonomy, informational inequality may occur when profiling is surreptitious or when profiles become too rich, and informational injustice may occur when profiles cross the boundaries of spheres of access. Therefore, the right to informational privacy and data protection law are also relevant in the context of profiling.

7.4 Data Protection Law

In Europe there are two main bodies of law that address profiling for purposes other than national security and law enforcement.[2] They are the Data protection directive (1995/46/EC) and the ePrivacy directive (2002/58/EC), which was amended in 2009 by Directive 2009/136/EC. The Data protection directive deals with the use of 'personal data' in general, whereas the ePrivacy directive deals with the use of unique identifiers and tracking technologies that can be used to facilitate profiling (e.g., cookies).

European data protection law has its roots in the OECD principles on privacy protection and the transborder flow of personal data and the Council of Europe treaty on personal data protection.[3] It aims to strike a balance between the (informational) privacy of the data subject and the free flow of information. The Data protection directive does this by providing a harmonised framework for the secure and legitimate exchange of personal data throughout Europe.[4]

The Data protection directive states that personal data must be processed fairly and lawfully and only for specified, explicit and legitimate purposes. To ensure fair and lawful processing the data protection sets a number of rules for the processing of personal data. These include –amongst others- obligations to keep the data secure, ensure its quality, inform the data subject, register the process in a public register, and grant the data subject access to the data.

In order for the provisions of the Data protection directive to be applicable, data must first be qualified as 'personal data'. Personal data is described in article 2(a) as:

"any information relating to an identified or identifiable natural person ('data subject'); an identifiable person is one who can be identified, directly or indirectly, in particular by reference to an identification number or to one or more factors specific to his physical, physiological, mental, economic, cultural or social identity"

An individual is considered 'identified' when that individual can be distinguished from all other members of a group.[5] Identification is commonly achieved through

[2] The use of profiling techniques for law enforcement purposes is governed –for the most part- by the law of criminal procedure, which differs from member state to member state. Though they differ from country to country, all laws that govern profiling must be in accordance with the rules set forth in article 8 of the European Charter of Human Rights (ECHR).

[3] Council of Europe Convention for the Protection of Individuals with regard to Automatic Processing of Personal Data (Convention ets. no. 108, Strasbourg 28-1-1981).

[4] Early December 2011, a draft version of a new general Regulation on data Protection prepared by the European Commission leaked (version 56, 29 November 2011). Relevant provisions include more strict rules on profiling (article 18) and the inclusion of online identifiers such as cookies in the definition of personal data. Given the fact that this Regulation is still in the drafting phase it is not discussed further in this chapter.

[5] Opinion Nº 4/2007 on the concept of personal data, Article 29 Working Party, p. 12.

the combination of certain 'identifiers' that hold a particularly privileged and close relationship to the individual.[6] Common identifiers are: name, physical appearance, and certain unique numbers such as a social security number. The extent to which certain identifiers are sufficient to achieve identification depends on the context of the particular situation.[7]

Van den Hoven explains that the data used in an identification process are *referential*, meaning that the data refers to a specific person, not just any person (Van den Hoven 2008, p. 309) This means that personal data always need an identity-relevant context. Without such context data have no meaning and are just *attributive*: they would describe a situation or fact without reference to any specific individual. One could argue that attributive data are *conditionally referential*; they can become personal data if another identity-relevant condition occurs, for instance because the raw data are placed in an identity-relevant context or combined with a piece of identity relevant information (Terstegge 2009). So while an individual might not be directly identified on the basis of a unique identifier such as a name, he or she may nonetheless be 'identifiable'. In the context of profiling three situations may occur that would render attributive data referential, personal data. They are: 1) adding identifying data to a profile, 2) spontaneous identification based on the uniqueness of the profile, 3) linking the profile to an individual by means of unique identifiers.

The first situation occurs when referential data (personal data) is added to attributive data. By adding information that is considered uniquely identifying (e.g., full name, date of birth, address) to a profile, all the data in that profile will become personal data.

The second situation occurs when the data contained in a profile leads to the spontaneous identification of the data subject. This is the case when the constellation of data is considered so unique, that the profile can only fit a single person and that person can be identified on the basis of the profile.[8] A well-known example of this is the case of 'AOL searcher 4417749'. In 2006 AOL published an anonymised dataset consisting of search queries for research purposes. But it did not take researchers long to trace back the search queries of an anonymous user (4417749) to her real name: Thelma Arnold (Barbaro and Zeller 2006). The combinations of search queries were so unique for each individual that they were able to trace back the queries to Ms. Arnold.

The third situation occurs when the profile of a data subject is associated with unique identifiers that are closely associated with him. Apart from identification on the basis of unique attributes such as for instance name, address and/or social security number, a profile may also be linked to a natural personal via other means. Most often this will be the case with profiling, since profiling is only effective if the data subject can be somehow be linked to the relevant profile. Apart from using identifiers that are unique to the data subject (e.g., name, address, place and date of birth), other types of identifiers may also be used. A common method is to link a data subject to a profile through the terminal equipment (e.g., mobile

[6] Opinion Nº 4/2007 on the concept of personal data, Article 29 Working Party, p. 12.
[7] Opinion Nº 4/2007 on the concept of personal data, Article 29 Working Party, p. 12.
[8] Opinion Nº 4/2007 on the concept of personal data, Article 29 Working Party, p. 13.

phone, computer) used by the data subject. For instance, a profile may be linked to a unique number associated with the terminal equipment. Identifiers that can function in this manner are IP-addresses, IMEI-numbers and MAC-addresses. Another option is to read from and write information to the terminal equipment, for instance by means of a cookie.

According to the article 29 Working Party when these indirect links are sufficient to single out a person, the associated profile should be considered personal data:

"Without even enquiring about the name and address of the individual it is possible to categorise this person on the basis of socio-economic, psychological, philosophical or other criteria and attribute certain decisions to him or her since the individual's contact point (a computer) no longer necessarily requires the disclosure of his or her identity in the narrow sense."[9]

The applicability of the Data protection directive to profiling in this manner is further confirmed in the Article 29 Working Party opinion on behavioural advertising. In this opinion, the Article 29 Working Party explains why it feels that personal data is processed in the context of behavioural advertising:

"This is due to various reasons: i) behavioural advertising normally involves the collection of IP addresses and the processing of unique identifiers (through the cookie). The use of such devices with a unique identifier allows the tracking of users of a specific computer even when dynamic IP addresses are used. In other words, such devices enable data subjects to be 'singled out', even if their real names are not known. ii) Furthermore, the information collected in the context of behavioural advertising relates to, (i.e. is about) a person's characteristics or behaviour and it is used to influence that particular person."[10]

The key element is that the identifier enables the person to be singled out for a specific treatment on the basis of the associated profile. On the basis of this reasoning by the Article 29 Working Party, most if not all, profiling exercises will fall under the scope of the Data protection directive. In those cases where an identifier is used that reads and/or writes information to terminal equipment, article 5(3) of Directive 2002/58/EC also applies. We may thus conclude that profiling in most if not all current forms, falls within the scope of the Data protection acquis in Europe.

7.5 Drawbacks to the Current Approach to Data Protection in the Context of Profiling

We have established that there are moral reasons for the protection of (personal) data in the context of profiling. If we follow the (broad) interpretation of the

[9] Opinion Nº 4/2007 on the concept of personal data, Article 29 Working Party, p. 13.
[10] Opinion Nº 2/2010 on online behavioural advertising, Article 29 Working Party, p. 9.

concept of personal data as set forth by the Article 29 Working Party we may also conclude that data protection law apply to many profiling practices.

The goal of profiling is to individualise and give a representation of a subject, or to identify that subject as a member of a group or category (Hildebrandt 2008, p. 17). For profiling to be effective it is unnecessary to know the data subjects actual identity. Furthermore, it is most often unnecessary to distinguish an individual from other members of a group. Rather, profiling calls for categorising an individual, for instance by fitting that individual into a certain predetermined target group. While it might be possible to identify an individual on the basis of a profile, more often than not, the data controller is not interested in actually identifying the data subject.[11] Nevertheless, as signalled by the Article 29 Working Party a user can be singled out on the basis of a profile in combination with a unique identifier such as a cookie or an IP-address. The idea is that because a person can be singled out and the information contained in the profile is used to make decisions about the person, data protection law should apply. While this is understandable from a privacy perspective, it is questionable whether data protection law is always the most effective mechanism for dealing with the risks posed by profiling. This question is important, as there are possible drawbacks to applying the current data protection law to profiling. Below I shall discuss several drawbacks that may prompt us to rethink the current approach to data protection in the light of profiling.[12]

7.5.1 The 'Binary' Nature of Data Protection Law

The first drawback of current data protection law is its binary nature. The Data protection directive only applies to the processing of personal data. While this sounds logical, it leads to difficulties in practice. The difficulty with the binary nature of data protection law is that it is oftentimes difficult to establish when data should be considered personal data. A combination of different pieces of data may all of a sudden become personal data when a referential piece of information is added, or when the different pieces of data by themselves spontaneously identify an individual. Moreover, while the profiling exercise itself may not be aimed at identifying a data subject, the possibility of identification is always present. For instance, pieces of information (oftentimes outside of the control of the data controller) may be linked to the profile, enabling identification. The question then becomes: at what point are all the protection mechanisms and legal obligations of the Data protection directive exactly to come into play. This leads to a great deal of uncertainty for (potential) data controllers. But the binary nature of data protection law is also problematic for data subjects, as it is unclear to whom they should turn for redress when a profile is misused or abused.

---

[11] While the data controller might not be interested in identifying or re-identifying the individual other parties may wish to do so. However, given the limited space available, we shall not address this issue in this chapter.

[12] It is relevant to note that in this discussion the focus is mainly on the use of profiling for commercial purposes.

The solution to this issue as set forth by the Article 29 Working Party seems to be to err on the side of caution, and consider any form of profiling the processing of personal data. The difficulty with this is that once a dataset or a profile is considered personal data, all the rules of the Data protection directive apply. In practice, this leads to a substantial administrative burden for data controllers (see paragraph 7.2). Moreover, it dilutes the effectiveness of enforcement (see paragraph 7.3).

7.5.2 The Procedural Nature of Data Protection Law

The second drawback, which ties in with the binary nature of the current data protection law, is the procedural nature of the law. Data protection legislation in its current form is primarily aimed at the *ex-ante* protection of privacy and personal data. This means that data controllers need to ensure that their processing of data is compliant with all the demands set forth by the Data protection directive. Though this should ensure the privacy of the data subject, in practice the effect of this ex-ante approach is often limited. In practice, privacy protection is for data controllers mainly an issue of compliance and following the procedural rules of the Directive (e.g. registering the processing in a public register, informing the data subject), rather than a discussion on what is considered a sustainable, ethical and responsible (business) process.

For the data subject there are also possible drawbacks. A significant drawback is that the data subject has limited options for redress in case there is a misuse or an abuse of personal data.[13] The reason is that the Data protection authorities are in charge of the enforcement of the law, leaving less room for individual redress.

7.5.3 Inflation of the Personal Sphere

Another issue with the application of the Data protection directive in the context of profiling is that it further expands the scope of the Data protection directive. The risk this brings with it is that as more and more activities fall under the header of personal data protection, the protection the law can provide actually decreases. Zwenne (2010, p. 335) for instance argues that a law that applies to essentially everything applies to effectively nothing. Blok (2002) also warns for this problem, calling the expansion of the concept of personal data 'an inflation of the personal sphere'. The main problem that might arise as a result of this inflation is that key privacy interests get heaped up with less important infringements of privacy, leading to an overstretched enforcement apparatus, confusion on how the law should apply, and possibly a degradation of the importance of privacy as a human right and the underlying values which it aims to protect.

A further problem with the inflation of the personal sphere is that the data protection authority becomes the *de facto* judge of what is considered the ethical use

[13] Whether there are options for individual redress is dependent on the actual implementation of the Data protection directive and associated privacy laws in national law. For the most part though we can say that the Data protection authority is in the lead when it comes to the enforcement of privacy rules, rather than the data subject.

of ICT. As soon as something is considered personal data, the data protection authority can decide whether or not a particular use of technology is acceptable or not. While the data protection authority can provide valuable input for discussion, oftentimes other institutions are more suited to this task. In the context of discrimination for instance, an equal treatment committee is probably more suited to determine when data processing is discriminatory. Moreover, when it comes to balancing different interests, it is important to involve all relevant stakeholders. This is particularly relevant in the context of profiling as most data protection authorities seem more 'privacy oriented' than 'data controller oriented', which could lead to an unfair balancing of different interests.

7.5.4 Data Minimisation

The Data protection directive states that the personal data processed must be adequate with regards to the goal of the processing. In essence, this means that no more data may be processed than is necessary (data minimisation). But the converse is also true: since the data processed must be adequate for the specified goal, processing too little personal data is also undesirable. In theory, the more attributes that are added to a profile, the more accurate a profile will be. So from the perspective of accuracy it could be argued that data *maximisation* rather than data *minimisation* should be a goal.[14] This leads to an interesting paradox when it comes to privacy and profiling. The goal of privacy and data protection legislation is to minimise the amount of data being processed. However, this may lead to profiles being less accurate, which in turn may engender the risks mentioned in paragraph (e.g. false positives/negatives, stereotyping, de-individualisation and discrimination). Furthermore, data minimisation is not necessarily a guarantee against discriminatory effects. Research in this area suggests that even by eschewing sensitive data (race for instance) altogether, discrimination may still occur (Verwer and Calders 2010). For a possible alternative to the current approach of data minimisation see Chapter 15 on data mini*mum*misation.

7.6 Is Data Protection Law an Adequate Solution?

The binary nature of data protection legislation has led to a situation whereby more and more data are considered personal data, in turn leading to an inflation of the personal sphere. As discussed in the previous paragraphs, this inflation is troublesome for several reasons. First of all, it leads to legal uncertainty and unnecessary burdens for data controllers. Second, the inflation of the personal sphere dilutes the effectiveness of enforcement and places too much emphasis on the role of the Data protection authority. Thirdly, the role of privacy and data protection legislation in addressing societal issues associated with profiling will become too big. This last point requires some further explanation. Informational privacy,

[14] Of course in saying this we must take the constraints of computer science into account, since an excessive amount of data may be detrimental for the efficiency and effectiveness of the algorithms used.

while an important human right in itself, is oftentimes more a means than an end. By limiting access and use of data via the right to informational privacy and data protection (the *means*), we limit the possibilities for misuse and abuse of these data, thus protecting interests such as personal autonomy, reputation and equal treatment (the *ends*). For instance, the right to privacy in the context of government surveillance is aimed at protecting personal autonomy: because knowledge is power, less information about citizens means less power for governments. In the context of processing data about an individual's race or religion the primary interest is not privacy protection, but rather equal treatment and/or avoiding discrimination: by not allowing racial information to be processed, it will be impossible to discriminate on the basis of these data. So by regulating the use of personal data, we mitigate possible risks and protect underlying goals (i.e., the moral reasons for data protection).

While this approach has proven useful, it also has its limitations. The binary nature of data protection law (it either applies, or it does not) also means that there are few possibilities to differentiate in the application of data protection legislation. On the one hand this may mean that too strict a regime is applied to 'mundane' privacy issues, while serious issues such as discrimination do not get the attention they deserve and are only treated as data protection issues. Moreover, too strong a focus on data protection may draw away our attention from alternative (legislative) solutions that provide more protection for individuals and groups as well as take into account the interests of the profiler.

7.7 Shifting the Focus in Data Protection Law

We have seen that the EU Data protection directive is quite expansive in its scope because of the broad interpretation of the concept of personal data, which may be troublesome. As such, we may conclude that there are limits to the effectiveness of data protection law in the context of profiling. Nonetheless, privacy and data protection law provide an important barrier against privacy intrusions and there are compelling moral reasons for protecting personal data. Therefore it is worthwhile to explore how we can make privacy and data protection law more effective, particularly in the context of profiling.

7.7.1 Differentiation in Data Protection: Data Centric Approach

A first option to make data protection law more effective is to differentiate in the application of data protection law based upon the data being processed. Depending on factors like the type of data processed, the likelihood of identification, and the scope of the data processing exercise we could set different standards of protection. When it comes to the appropriate (legal) safeguards, we could for instance employ a light regime that focuses on transparency, data quality and data security, possibly linked with stronger *ex-post* protection mechanisms, for data that is not easily identifiable; and a stronger regime that employs all the (ex-ante) safeguards of the data protection directive for data with a clear link to an identified person.

Ohm (2010) proposes a differentiation based on the roles different entities can play in the identification or re-identification process. He argues that because identification or re-identification is made possible (or easier) by combining different data sets from different entities, entities that process large amounts of (personal) data (what Ohm calls 'large entropy reducers'), should have a higher duty of care (e.g., companies like Google, Microsoft and Choicepoint).

Schwarz & Solove (2011) propose a differentiated system based on the difference between 'identified data' and 'identifiable data'. They divide the use of data into three risk categories: identified, identifiable, and non-identifiable. Rather than defining these categories in law, they opt for a more flexible, standards based approach to determine under which circumstances what regime should apply.

7.7.2 Focus on the 'Why' Instead of the 'What': Goal Oriented Approach

While a more fine-grained data centric approach will, to some extent, remedy the issues associated with the binary and procedural nature of data protection legislation, it does not necessarily deal effectively with the possible risks of profiling. Therefore, we should also look towards other mechanisms to function alongside data protection legislation.

An alternative (or an addition) to the data centric approach is a more goal oriented approach. Depending on the actual goal of the data processing and the possible risks involved, the most effective protective measures may be chosen.

Purpose specification and purpose binding already form key elements of the structure of the current Data protection directive. Data controllers need to have a specified, explicit and legitimate purpose for collecting personal data and any further processing may not be incompatible with the specified purpose (see article 6 of the Directive). However, the goal of the data processing does not determine which rules should apply. Rather, the general rules of the Data protection directive apply, regardless whether they are the most effective protective measures.

7.7.3 Revisiting the Moral Reasons for Data Protection

In a goal-oriented approach the type and level of protection would be based primarily on the goal of the profiling exercise and the risks associated with this goal, rather than on the basis of the fact that certain data is considered personal data. By looking more closely at the risks involved with a particular type of processing we can ascertain whether data protection law should apply, and to what extent. A more goal-oriented approach makes data protection rules more context-sensitive, opening up the possibility for other legal protection mechanisms (such as consumer protection, equal treatment, and unfair commercial practices legislation) that might be more effective or suitable.

A goal-oriented approach to data protection and profiling would therefore place more emphasis on the moral grounds for data protection than is currently the case. This may also entail that other types of legislation (anti-discrimination legislation for instance) may come into play in addition to data protection law. In some cases

these rules may even supersede data protection legislation. For instance, data minimisation rules and prohibitions on the processing of sensitive data may be overruled if they undermine the accuracy of a profiling exercise, or if they deny us the possibility to detect discrimination in a profiling exercise.

A goal-oriented approach will likely mean less focus on ex-ante protection and more focus on ex-post protection mechanisms. A positive effect of this shift is that it will force data controllers to actually make an assessment of the risks involved in their data processing and profiling activities, rather than reducing privacy and data protection to a mere compliance issue, as is currently often the case.

7.7.4 From 'Privacy by Design' to 'Ethics by Design'

Apart from the application of data protection and the rules associated with it, we should also examine other means of regulation. In particular the 'code as code' solution of *privacy by design* should be taken into consideration (Lessig 2006). 'Privacy by design' refers to the notion that we must incorporate privacy-protecting measures into the architecture of information systems (Borking 2010). In this way we can 'hardwire' the rules into the system. While privacy by design is an important measure when it comes to the protection of the individual in the context of profiling we should also be cognisant of the limitations and drawbacks of such an approach. In particular, we should take into account the limitations of privacy and data protection law in drafting functional and legal requirements for IT systems. Rather than a narrow focus on privacy and data protection we should look towards the actual goal of the profiling exercise and determine whether we can apply appropriate safeguards. This requires a broader focus than privacy, so instead of privacy by design we should focus on *ethics by design.* While closely related to the idea of privacy by design, ethics by design allows for a more focused, context sensitive approach to dealing with the possible risks of profiling.

7.8 Conclusion

Profiling is becoming an increasingly important tool for the public and private sector. While an effective tool, profiling might also entail risks for individuals and groups. These risks include stereotyping, inaccuracies in the application of profiles, stigmatisation, de-individualisation and abuse of profiles.

Currently these risks are addressed mainly through the application of data protection law. However, it is questionable whether this legal framework in its current form and application effectively mitigates the risks of profiling.

By stretching the definition of personal data to include profiles and the identifiers that link these profiles to individuals (e.g., IP-addresses and cookies) the protection mechanisms of the Data protection directive apply. While such an approach is understandable, there are some drawbacks. Because of the binary and procedural nature of data protection law there is no way to differentiate in the application of the Data protection directive. Labelling all data as personal data either because there are moral reasons to have some form of protection, or because there

is a risk of identification or re-identification, will lead to an inflation of the personal sphere. In some cases, applying data protection law may be ineffective and even counterproductive.

To counter the drawbacks that currently flow forth from the application of data protection legislation, a first option could be to differentiate between different types of data (identified, identifiable, non-identifiable). While this would make data protection law more flexible and practical it will not necessarily address all the issues associated with profiling, nor remedy all the drawbacks of applying data protection law in the context of profiling. Therefore, in addition to the data centric approach, we should focus more on the actual goal of the profiling exercise and determine on the basis of the actual risks associated with this goal which safeguards should apply. Not only would such an approach likely provide more protection to the individual, it would also allow for a better balancing of the interests of the data subject and those of the data controller.

References

Barbaro, M., Zeller, T.: A face is exposed for AOL searcher no. 44177179. New York Times (2006), online version via: http://www.nytimes.com/2006/08/09/technology/09aol.html (last visited: December 28, 2011)

Borking, J.J.F.M.: Privacyrecht is code. Kluwer, Deventer (2010) (in Dutch)

Blok, P.: Het Recht op Privacy. Boom Juridische uitgevers, Den Haag (2002) (in Dutch)

Bygrave, L.: Minding the machine: article 15 of the EC data protection directive and automated profiling. Computer Law & Security Report 17, 17–24 (2001)

Custers, B.H.M.: Data Mining with Discrimination Sensitive and Privacy Sensitive Attributes. In: Proceedings of ISP 2010, International Conference on Information Security and Privacy, Orlando, Florida, July 12/14 (2010)

Hildebrandt, M.: Defining Profiling: A New Type of Knowledge? In: Hildebrandt, M., Gutwirth, S. (eds.) Profiling the European Citizen, Cross-Disciplinary Perspectives. Springer Science (2008)

van den Hoven, J.: Privacy and the Varieties of Informational Wrongdoing. Australian Journal of Professional and Applied Ethics, Special Issue 1 (June 1999)

van den Hoven, J.: Information Technology, Privacy and the Protection of Personal Data. In: van den Hoven, J., Weckert, J., et al. (eds.) Information Technology and Moral Philosophy, pp. 301–332. Cambridge University Press, Cambridge (2008)

Lessig, L.: Code 2.0. Perseus Book Group, New York (2006)

Ohm: Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization. UCLA Law Review 57, 1701–1777 (2010)

Schermer, B.W.: Software Agents, Surveillance, and the Right to Privacy: a Legislative Framework for Agent-enabled Surveillance, PhD. Thesis. Leiden University (2007)

Schwarz, Solove: The PII Problem: Privacy and a new concept of personal identifiable information. New York University Law Review 86, 1814 (2011)

Terstegge, J.: Back to basics: privacy ethics (2009), http://jeroenterstegge.blogspot.com (last visited: December 28, 2011)

Vedder, A.: KDD: The challenge to individualism. Ethics and Information Technology 1(4) (December 1999)

Verwer, Calders: Three Naive Bayes Approaches for Discrimination-Free Classification. In: Data Mining: Special Issue with Selected Papers from ECML-PKDD 2010. Springer (2010)

Zwenne, G.J.: Over persoonsgegevens en IP-adressen, en de toe komst van privacywetgeving. In: Mommers, L., Franken, H., Klaauw, F., van der Herik, H., van der Zwenne., G.-J. (eds.) Het Binnenstebuiten, Liber Amicoru m Aernout Schmidt, pp. 321–341. Universiteit Leiden (2010) (in Dutch)