Tilburg University

Medical Data, New Information Technologies, and the Need for Normative Principles other than Privacy Rules

Vedder, A.H.

Published in:

Law and Medicine : Current Legal Issues 2000 : Volume 3

Publication date:

2000

[Link to publication in Tilburg University Research Portal](#)

Citation for published version (APA):

Vedder, A. H. (2000). Medical Data, New Information Technologies, and the Need for Normative Principles other than Privacy Rules. In M. Freeman, & A. Lewis (Eds.), *Law and Medicine : Current Legal Issues 2000 : Volume 3* (pp. 441-459). (Current Legal Issues; No. 3). Oxford University Press.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Medical Data, New Information Technologies and the Need for Normative Principles Other than Privacy Rules

Anton Vedder¹

1. Introduction

Recent technological developments invite us to reconsider our traditional conceptions of privacy and to extend the normative framework for the moral and legal assessment of new information technologies. I will propose a new account of the conceptualisation of privacy and I will introduce some additional normative principles, which – taken together – are better fit to capture the social problems relating to certain technologies. The technologies on which I will focus are those of profiling through data mining, especially profiling through data mining in epidemiology.

In the years to come, profiling through data mining will become an important and powerful set of techniques in the hands of epidemiologists and policymakers in the field of health care. The techniques of data mining will enable them to predict the developments in the health condition and health risks of populations more accurately and, more importantly, much easier than ever before. Applying the techniques results in generalisations about groups of persons, rather than about individuals. For this reason, existing privacy norms do not apply. Nevertheless, some of the ways in which these generalisations are used or can be used may lead to serious social and moral problems. In this paper, I will first present some preliminaries about profiling through data mining. Next, I will explain in some detail why current conceptions of privacy, as they occur in law, legal theory and ethics, are too weak as to their normative and descriptive evaluation and distinction potential to capture the problems relating to profiling through data mining. Subsequently, I will propose a new conceptual scheme for the privacy notion.

This approach equips us with a better understanding of the descriptive core and the basic points of the value notion. In addition, it enables us to understand and articulate the problems arising in relation to profiling through data mining that cannot be adequately and completely formulated in terms of other evaluative notions and principles. Nevertheless, invoking these additional principles and values is necessary. I will expound the ones needed to extend the normative framework for assessing information technologies. Finally, I will put forward and discuss some practical solutions to the problems at issue.

2. Profiling through data mining

'Data mining' is the term I use to refer to the whole set of techniques certain experts call 'knowledge discovery in databases (KDD)'. This set of techniques is applied in a process that is traditionally divided into three phases, i.e. the data warehousing phase, the data mining phase proper and the interpretation phase. In the data-warehousing phase, data is collected, enriched, checked, and coded. The data is analysed in what the aforementioned experts call the data-mining phase in the strict sense. Finally, the results are interpreted. During these phases, a search hypothesis is used to guide the process.² In this paper, I will use the perhaps somewhat sloppy language that has become customary, and *not* restrict the use of 'data mining' to refer to the middle phase of the KDD process only.

The basic aims and greatest opportunities of data mining are description and prediction through the discovery of significant patterns in the relationships of whatever kind between data.³ By 'profiling through data mining', I refer to those data-mining processes that result in profiles of groups of persons, i.e. characterisations of groups that can be assigned to those groups and to the members of those groups. The process of profiling through data mining aims at tracking down significant relationships between characteristics that define and identify a group of persons on the one hand and whatever properties or characteristics on the other.

In recent years, profiling through data mining has been acknowledged as an important set of techniques in analysing data for purposes such as direct marketing and credit scoring. Other applications include checking for patterns in criminal behaviour for

forensic and judicial purposes and defining high-risk groups for tax evasion in order to improve rateability. A very inviting and – at least from the perspectives of public interest and public health – very useful utilisation is analysing medical data or data about the use of medical services in combination with demographic data for epidemiological purposes. So, for instance, data about the use of medical services in the files and databases of health-insurance companies, or data from medical files or electronic care records maintained by health-care providers, could be merged and analysed. These analyses can result in the description and prediction of the incidence and prevalence of diseases. They also enable epidemiologists to ascertain and find high-risk groups, and to determine relations between chances of recovery from diseases and other – until now as yet unknown – influencing factors, etc.

It is primarily the utilisation for epidemiological or public health purposes I will be concerned with here. I will not go into the details of these specific applications. Much of what I will say equally applies to profiling through data mining for other purposes – be they fiscal, judicial, or commercial. One of the reasons to select data mining and profiling for epidemiological purposes is that it is a salient example of applying the techniques for undoubtedly good and legitimate primary purposes, while many of the secondary uses of the resulting profiles may turn out to be highly questionable.

Throughout this paper, I will treat of profiling and data mining only insofar as it involves the use of personal medical data in the so-called 'warehousing phase'. Personal medical data is personal data directly or indirectly – e.g. *via* data concerning the use of medical services and medication – referring to health condition and health prospect. Of course, I will include profiling through data mining insofar as it involves the use of medical personal data *in combination with* other data, be it personal or of another kind. I will, however, not address profiling through data mining involving no personal data at all.

3. Current conceptions of privacy

For an understanding of present-day conceptions of privacy, some acquaintance with the legal notion of privacy is necessary, especially with what might be called the legal conception of informational privacy. Over the last decades, informational privacy has

become more and more important in law and regulation. Informational privacy in law and regulation is mostly interpreted as data protection. There is one kind of data that is considered exclusively eligible for protection by privacy law and regulation: personal data. Personal data is commonly defined as data and information relating to an identified or identifiable person. A clear illustration of this rather narrow starting point can be found in the highly influential European Directive 95/46/EC of the European Parliament and of the European Council of 24 October 1995 'on the protection of individuals with regard to the processing of personal data and on the free movement of such data'. The Directive poses some basic principles with regard to the processing of personal data. As might be expected from the definition of 'personal data', most of these principles lean heavily on the idea that there is some kind of direct connection between a particular person and his or her data.

First, there are some principles regarding data quality. Personal data should only be collected for specified, explicit, legitimate purposes and should not be further processed in a way incompatible with these purposes. No excessive amounts of data should be collected, relative to the purpose for which the data is collected. Moreover, the data should be accurate and, if applicable, kept up to date. It must be guaranteed that inaccurate or incomplete data is either rectified or erased. Also, personal data should be kept in a form which permits identification of data subjects for no longer than is necessary for the purpose for which the data was collected.

Second, some principles apply to the processing of personal data. Data processing is legitimate if an individual has unambiguously given his or her consent. Furthermore, processing may be legitimate if the processing is needed for compliance with a legal obligation, for the protection of vital interests of the data subject, for some public interests, or for some legitimate interests pursued by the data controller or by third parties to whom the data is disclosed.

Third, the data subject has some specific rights with regard to 'his or her' personal data. Among these rights are the right of knowing what data is being stored and whether the data relating to the data subject is being processed, the right of rectification, the right to know to whom the data has been disclosed, and the right to object to the processing of data relating to the data subject.

The definitions and principles formulated in the European Directive are implemented in the national privacy laws and regulations of the European Union member states. For this reason, the impact of the Directive's definition and principles should not be underestimated.

The conceptual and normative individualism of the Directive reflects ideas about informational privacy currently held amongst legal and ethical theorists.⁴ Sometimes, these ideas on informational privacy are not much more than implicit assumptions. However, things are different and more articulate where theorists define privacy as being in control over (the accessibility of) personal information, or where they indicate some kind of personal freedom, such as the preference freedom in the vein of John Stuart Mill's individuality, as the ultimate point and key value behind privacy.⁵ These theorists consider privacy to be mainly concerned with information relating to individual persons. They also tend to advocate protective measures in terms of safeguarding an individual's control and consent *vis-à-vis* certain dispersions of personal information.

Applying the narrow definition of personal data and protective measures such as the Directive's and some philosophers' to the profiles discussed here is problematic. As long as the process involves personal data in the strict sense of data relating to an identified or identifiable individual, the definition and the principles apply without reservation. This will mainly concern some of the source material in the so-called warehousing phase of data mining. However, as soon as the data has ceased to be personal data in the strict sense, it is not at all clear how the principles should be applied. For instance, the right of rectification applies to the personal data in the strict sense; it does not apply to information derived from this data. The same goes for the requirement of consent. Once the data has become anonymous, or has been processed and generalised, an individual cannot exert any influence on the processing of the data at all. In sum, the definition and the rights and requirements – at least as they are traditionally formulated and interpreted – make no sense regarding anonymous data and group profiles.

Before I can elaborate further on this problem, it seems advisable to have a closer look at current conceptions of privacy in the legal and ethical debate over the last decades. Doing so will enable us to be more accurate when dealing more specifically

with the normative and descriptive evaluation and distinction potential of the notion of privacy.

In the last decades, there has been ample discussion on the subject of privacy, particularly on the descriptive and normative aspects of the notion. One line of discussion has focused on the normative political problem of the (right to) privacy of citizens versus the duty of government or those same citizens to serve the public interest or the common good.⁶ The conception of privacy that played a part in this environment could perhaps be best characterised as a conception of *decisional* privacy or an *autonomy* conception of privacy. The debate concentrated on the political question of how much liberty should be granted to an individual regarding his or her intimate personal domain, particularly concerning his or her sexual orientation, the use of contraceptives, abortion, etc. In this approach, privacy boils down to a kind of personal freedom within a restricted area.

Other discussions ran along the borders of the partially overlapping fields of semantics, epistemology and value theory. So, for instance, there has been discussion on the role of conventions, traditions and the total network of social practices in a given culture regarding the identification of spheres of life or aspects of persons which are called private or personal.⁷ Other debates have been concerned with the question whether the right to privacy could be derived from the autonomy of individual persons or be better interpreted as protective of certain kinds of interests of individuals. There have been debates about the question whether privacy was concerned with the inaccessibility of personal spheres as such or rather with an individual's control over the access to his or her personal sphere. Again others have been about the question whether privacy is just about physical accessibility of personal spheres or about access to information about an individual's personal sphere as well. In this debate, the information approach, with a conception of privacy as informational privacy in its wake, remained dominant.⁸

Furthermore, important work has been done explicitly on the normative point of privacy. Fried and Rachels situated the point of privacy norms in respect for the need of individuals to be able to control the spread of information about their personal spheres in order to define the character of relationships which they (would like to) establish and maintain with other individuals and institutions.⁹ In addition to this, Deborah Johnson emphasised privacy's empowering role. Privacy as control over personal information is

an important instrument to strengthen the position of citizens and consumers *vis-à-vis* governmental institutions and private organisations.¹⁰ Stanley Benn and Jeffrey Johnson highlighted the possibilities, offered by privacy norms, to be immune from the judgement of others and to be respected as a (now and then erring and failing) plan-making and plan-implementing being, instead of an object without will, ready to be fixed and pinned down in the judgement of others.¹¹

4. Privacy: A contextual-functional approach

Although privacy is an often and easily used notion, its precise meaning is far from clear, and liable to stimulate controversies.¹² In order to avoid digressions, I will stipulate an account of the meaning of privacy. This account serves two purposes. First, it indicates a defensible unifying core concept that is common to most of the conceptions of privacy currently used. Together with some broader considerations concerning the contextuality and functionality of meaning in the account, the core concept enables us to clarify and systematise the seemingly diffuse and chaotic family of privacy conceptions. Second, the account of privacy put forward will help to understand and articulate the problems in relation to profiling and data mining that cannot be formulated adequately and completely in terms of other evaluative notions and principles.¹³

My account of privacy leans heavily on Jeffrey Johnson's conception of privacy as 'immunity from the judgements of others'. I think that Johnson is right in assuming that the idea of immunity from the judgement of others is in some way or other the *trait d'union* between all the instantiations of privacy. The way in which it performs the function of this unifying feature is important. Unfortunately, Johnson is not very clear about this. I would say that we are in fact talking about two levels: the level of the phenomena in which privacy is instantiated and the level of language in which these instantiations are articulated. I would claim that apart from all other characteristics that may be present, common to most of the phenomena in which privacy is instantiated is the feature of an individual's immunity from the judgement of others (as being present or as being absent). At the level of linguistic articulations or propositional renderings of instantiations of privacy, I would say that the absence or presence of immunity from the

judgement of others is the basic description frame common to the correct articulations of all possible instantiations of privacy. By calling the individual's immunity from the judgement of others a basic description frame, I mean that in the articulations of instantiations of privacy, somehow a positive or negative account of a description of this immunity to the judgement of others is present or at least presupposed.

The feature of immunity to the judgement of others can be found in most instantiations of privacy – with the important exception of decisional privacy, but to this issue, I will return later. Privacy is often rightfully considered as the condition of an individual in which others do not have access to his or her personal sphere as such or without the individual's permission. 'Having access', however, should be understood as being near to or present in an individual's personal sphere. This nearness or presence may be physical (with one's body) or instrumental (by bugging, using binoculars, etc.), or by apprehending information that has already been abstracted by oneself or by others from an individual's personal sphere. A common feature of all these variations of access is that another person or other persons become *informed* about certain aspects of the individual's personal sphere. Through this information, others are enabled to judge and evaluate the person in some respect. A person's privacy is being directly infringed when others have physical or instrumental access to his or her personal sphere, thereby gathering information without permission. A journalist infringes on the privacy of a person, for instance, if he or she uninvitedly intrudes into the person's home and takes a picture of this person while he or she is asleep. Furthermore, the journalist infringes on this person's privacy indirectly by publishing the photograph in a newspaper to inform the readers about the person's personal sphere. The readers who are thus informed, infringe on this person's privacy directly, albeit possibly non-voluntarily.

This articulates sufficiently the basic description frame of privacy. The basic description frame is the common denominator of privacy conceptions. The normative point for (the importance of) articulating problems and problematic situations in terms of privacy and for valuing privacy is a different subject. It is to be found in a complex amalgam and interplay of contextual elements and different values, and not only, as Johnson would have it, in conventions, on the one hand, and freedom or autonomy, on the other.

Why is it important to protect or to safeguard the immunity of persons from the judgement of others? Why is infringing on privacy, why is compromising this immunity, a serious thing to do? Many writers have suggested that there is one value or one set of some familiar values at issue for which the value of privacy is instrumental or constitutive. This single value or small set of familiar values is thought to underlie all instantiations of privacy. Johnson has suggested that personal freedom is behind privacy. Benn suggested that we are dealing with specific kinds of freedom: the freedom of self-presentation and moral autonomy. Again others have put forward other connected values. Interestingly enough, they all seek the unifying factor of the apparent diversity of privacy conceptions in a single or in some familiar values that are thought to be behind privacy. A possible reason for this is that they have all been trying to find the conceptual link between decisional privacy and privacy in terms of access or information *via* some common underlying value that was supposed to bridge the gap between these notions.

The unity of privacy, however, is not to be sought in the value(s) behind privacy, but precisely in what I have called the basic description-frame. This should not lead us to underestimate the importance of all the different values, associated with privacy. Let me start to clarify this point by commenting on Jeffrey Johnson's view that conventions determine what is to be considered the personal sphere, the domain supposed to be protected by privacy:

'Privacy is a conventional concept. What is considered private is socially or culturally defined. It varies from context to context, it is dynamic, and it is quite possible that no single example can be found of something which is considered private in every culture. Nevertheless, all examples of privacy have a single common feature. They are aspects of a person's life which are culturally recognised as being immune from the judgement of others.... The function of privacy, therefore, is to isolate certain limited and culturally defined aspects of the individual's life as being morally and legally protected from the evaluative judgement of others.'¹⁴

By an individual's personal sphere we usually mean the domain of a person, consisting of his mind and body, possessions, letters, home, certain activities, etc. It should be kept in mind that the notion of personal sphere is originally a metaphor. It should not be interpreted exclusively in terms of spatial dimensions or territory. The

personal sphere also includes *aspects* and *dimensions* of persons, lives and actions, as seen from certain perspectives. Jeffrey Johnson is right in claiming that what is considered to belong to the personal sphere is in several respects defined by the conventions and traditions of one's culture and community. One should not, however, feel tempted to equate conventions and traditions too hastily with either arbitrary tendencies or with purely contingent social arrangements. In many cases, conventions are only part of the story. The rationale for including certain elements in the personal sphere (e.g. nudity of the body, the performance of certain biological functions) in a Western community nowadays seems sometimes hard to find. Nevertheless, for some of these there certainly has been a reason and for most of the others there certainly have been good reasons to count them as belonging to the personal sphere.

Conventions also play a part in combination with the prevention of harm and offence, or some conception of individual well-being. The point of defining certain items or aspects of persons as 'personal' often lies in the fact that persons with respect to these items and aspects just tend to be vulnerable, and are therefore in need of protection. Conventions and traditions are important in the definition of the personal sphere to the extent that they themselves or the social arrangements accompanying them create or keep up the conditions and occasions for special kinds of harm or offence. So, for instance, one's medical condition and prospects are rightfully considered to belong to the personal sphere. There is a certain tradition, convention or at least a seemingly ineradicable tendency in Western societies towards harmful discrimination and stigmatisation of and scorn for the ill and the suffering. Although this tendency is mostly suppressed, it sometimes comes to head in the rejection of persons with certain diseases. Recent examples are AIDS and other communicable diseases.

Conventional aspects of privacy like these are unfortunately rarely acknowledged by contemporary scholars, impressed as they are by the apparent irrationality of some conventions designating certain features as belonging to the personal domain. The scholars who do have an open eye for conventions are a little fuzzy about the precise role of conventions. So, for instance, it is not in the least clear whether Jeffrey Johnson, when talking about the conventional aspects of privacy, envisages the aspects which I have distinguished above. He may as well be making a statement about the development of the

value of privacy in the traditions of a community when he wrote the passage quoted above.

Both these interpretations are important to privacy, but they should not be confused. They explain part of the constitution of the personal sphere and part of the constitution of the value of privacy. Conventions determine in part under which aspect or from which perspective the disclosure of a feature of a person or a person's life might hurt or harm this person. These features – and these not necessarily *per se* but under a certain aspect or in a certain perspective – are part of the personal sphere. They explain indirectly why such disclosure would be bad: the person would be hurt or harmed because of the stigma, taboo, prejudice or (social) exclusion that is in the normal course of events in a society or culture like the one involved attached to such disclosure. The value which in such situations is to be protected ultimately is the well-being of individuals. It is protected by avoiding the possibility that information about persons is disclosed that might harm them by exposing them to prejudice, stigmatisation and exclusion, etc.

Conventions are not the whole story about the personal sphere. Neither are they – in combination with some ideal conception of personal well-being – the whole normative point of privacy. Economic and social arrangements and technical and technological developments play their role as well, as do important values other than the one of individual well-being.

Certain social, economic and technical arrangements – themselves partially based on conventions and traditions – provide the opportunity, and sometimes even the institutional necessity, to harm persons on the grounds of features of their personal sphere, such as their medical condition and prospects. They do so, for instance, by hindering access to insurance, credit facilities, jobs and offices, some of the few social institutions in which discrimination on the grounds of health aspects of persons is still accepted in Western societies and cultures. Here, disclosing information about the health condition and prospects of persons is reckoned to be private or belonging to the personal sphere because people may be harmed by it in that they are excluded from certain social provisions and amenities. Again, an ideal of individual well-being is the normative point behind protecting the privacy of individuals in this respect.

In addition to and, often, in combination with conventions, social, economic and technical arrangements and the general value individual well-being, other values play their part as well. Fried and Rachels have shown that privacy enables individuals to control the spread of information about their personal spheres. Fried certainly has a case in claiming that privacy is constitutive of friendship and instrumental to veracity. Rachels understands the control over the dissemination of personal information as an important constituent of autonomy, enabling persons to define themselves the character of relationships which they (choose to) establish or maintain with other individuals and institutions. Jeffrey Johnson considers privacy to be a guarantee against the undesired evaluative judgement of others and therefore as a constituent of freedom. Benn shows the further point of such immunity: abstaining from judgements about an individual's personal sphere allows him, at least in a certain area, to be respected as a person. 'The value of privacy', says Benn, 'hinges on a person's interest in forms of self-presentation, as part of his self-awareness as a maker of projects, without which the individual is depersonified, reduced to an object.... The ends of privacy rights are derived from the needs of human beings to be safe from persecution, to develop intimate personal relations, etc.'¹⁵

The unity in the family of privacy conceptions lies hidden in a common description frame, and not in some value or values underlying privacy. There are many values behind privacy. For which value privacy exactly functions as an instrument, depends on the context of social, economical and technical circumstances and conventions. This is why the meaning of privacy is determined by context and function.

A potential objection against this account of privacy could be that immunity to the judgement of others can be discovered in all instantiations of privacy, with the exception of instantiations of decisional privacy. The main thesis of my account is, however, that all privacy *conceptions* have the common feature of using one basic description frame. There is a difference between the instantiations of privacy in the phenomena and the linguistic articulations of such instantiations of privacy. Immunity to the judgement of others is a common basic description frame for all articulations of instantiations of privacy; it is *not* a common feature to all instantiations of privacy in the phenomena. I would say that instantiations of 'decisional privacy' as they occur in the phenomena do

not share this characteristic. That we label these occurrences privacy phenomena has other reasons. The domain of life and actions which decisional privacy or an autonomy conception of privacy covers – in the sense that it requires respect for or forbears interference with decisions and actions in this domain – is essentially part of the sphere or the domain in which persons normally are thought to have immunity from the judgement of others. In addition to this, there seem to have been strategic and practical reasons to bring the kind of decisions that decisional privacy actually protects under the label of privacy. The idea of decisional privacy is strongly associated to a certain interpretation of some of the amendments to the American Constitution. The notion of privacy – as relating to physical access and information about persons – already played some role in the interpretations of these. It lay in the natural course of events that the defendants of what I have labelled 'decisional privacy' would turn to the privacy vocabulary. A more appropriate juridical linguistic framework just seems to have been missing. Apparently, other constitutional starting points were not at hand. Finally, although there might be better ways of formulating the instantiations and the general values of decisional privacy, convention and custom have kept us from abandoning the conception in favour of a better one. Needless to say, that, although immunity to the judgement of others is not a feature of instantiations of decisional privacy, it is a basic description frame of articulations of instantiations of decisional privacy because the idea of immunity from the judgement of others is conceptually, *indirectly*, presupposed by the articulation. To put it differently: the conception of decisional privacy feeds on the conception of privacy that generally articulates instantiations of privacy in which immunity to the judgement of others is at issue.

In sum: Privacy is a servant of many master values. In addition to the values that have been brought to the fore there are still others, such as individuality, fairness in judgement and treatment, etc. In the next section, I will show how these values are active behind instantiations of privacy. In this section, I have put forward an account of privacy in which the basic description frame of 'immunity to the judgement of others' is the unifying factor in the family of privacy conceptions. The variety, which I labelled decisional privacy, stands somewhat apart from the other varieties because it is only indirectly connected to the basic description frame. I have also explained that privacy

should not be considered to be derived somehow from just one value or from some familiar values. Behind different instantiations of privacy, in different contexts, different values may be active, mostly in combination with certain conventions, traditions and technical, economic and social-institutional arrangements.

5. Categorical privacy, individuality, fairness and justice

Most conceptions of individual informational privacy currently put forward in law, regulation and ethical debate have one feature in common that is important to my point. Not only do they assume that the personal data with which privacy is concerned, *originally* consist of statements about states of affairs or aspects accompanied by indicators of individual natural persons, they also assume that the data during processing *continue* to contain those identifiers of individual natural persons. This feature of many current privacy conceptions has two significant consequences: it makes it difficult to label the problematic aspects of using data abstracted from personal data and producing and applying group profiles and generalisations; it also makes it difficult to fathom the seriousness of these problems in practice.

It should be observed that group profiles and generalisations may occasionally be incompatible with respect for individual privacy and laws and regulations regarding the protection of personal data, as it is traditionally conceived of. In order to understand this, we must distinguish between distributive profiles and non-distributive profiles.

Distributive profiles assign certain properties to a data or information subject, consisting of a group of persons however defined, in such a way that these properties are actually and unconditionally manifested by all members of that group. Distributive generalisations and profiles are put in the form of down-to-earth, matter-of-fact statements. As opposed to this, *non*-distributive profiles are framed in terms of probabilities and averages and medians, significant deviances from other groups, etc. They are based on comparisons of members of the group with each other and/or on comparisons of one particular group with other groups. Non-distributive profiles are, therefore, significantly different from distributive profiles. The properties in non-distributive generalisations apply to individuals as members of the reference group, whereas these individuals taken as such

need not in reality exhibit these properties. For instance, an applicant may be refused a life insurance on the basis of a non-distributive generalisation of certain health risks of the group (e.g. defined by a postal code) to which he happens to belong, whereas he or she is a clear exception to the average risks of his or her group. In all such cases, the individual is primarily judged and treated on the basis of belonging to a group or category of persons and not on his or her own merits and characteristics.

Distributive generalisations and profiles may sometimes be rightfully thought of as infringements of (individual) privacy when the individuals involved can easily be identified through a combination with other information available to the recipient or through spontaneous recognition. In the case of non-distributive profiles, the information remains attached to an information subject constituted by a group. It cannot be traced back to individual persons in any straightforward sense. The groups which are the information subjects of non-distributive profiles can often only be identified by those who defined them for a special purpose. From the perspectives of people other than the producers and certain users of the profiles and generalisations, the definition of the information subject will remain hidden because they do not know the specific purposes of the definition. When accidentally found out by the people in the reference group, they will probably think of the definition as being arbitrarily chosen. Most importantly, however, the information contained in the profile envisages individuals as members of groups; it does not envisage the individuals as such. Supposing for the sake of argument that the profile has been produced in a methodically sound and reliable way, it only tells us some truth about individual members of those groups in a very qualified, conditional manner. Therefore, privacy rules, as they are traditionally conceived of, do not apply. The information in non-distributive profiles cannot be traced back to individual persons.

One might think that perhaps we could be saved by some notion of collective privacy. However, collective privacy will not do the job properly. The notion of collective privacy is too easily associated with the concept of collective rights. The subjects of collective rights are groups or communities. In order to make sense of the idea of collective rights, these subjects are often treated as beings analogous to persons or moral agents, or at least as conglomerates having certain characteristics which cannot ultimately and exhaustively be explained by the input of the individual members.

Furthermore, they are often thought to be structured or organised in some way so as to be able to exercise their rights or let their rights be advocated by vicarious agents.¹⁶ All these properties do not apply to the reference groups of the profiles we are considering. From the perspective of their members, these groups are mostly randomly defined. Their members do not have any special ties of loyalty among one another. Nor do they have organisational structures. Therefore, they are not able to take decisions or to act as collectivities.

In order to remove the deficiencies of current conceptions of privacy as regards analytical and distinctive evaluative potential, we would be better off utilising a concept, which I have elsewhere labelled, 'categorical privacy'.¹⁷ I suggest that we conceive of categorical privacy not as an independent concept, but as a dimension of privacy. This dimension relates to data or information to which the following conditions apply. (1) The information was originally taken from the personal sphere of individuals, and – after aggregation and processing according to statistical methods – is no longer accompanied by identifiers of individual natural persons, but, instead, by identifiers of groups of persons. (2) When attached to identifiers of groups and when disclosed, the information is apt to carry with it the same kind of negative consequences for the members of those groups as it would for an individual natural person if the information were accompanied by identifiers of that individual.

Categorical privacy is strongly connected with individual privacy. It uses the same basic description frame as individual privacy, i.e. immunity from the judgement of others. The values which – in combination with contextual factors such as conventions and social arrangements – oppose infringements on individual privacy, such as individual well-being, personal autonomy, individuality and certain social interests, equally oppose infringements of categorical privacy. But there are more values at issue, such as fair judgement and treatment, and respect for the individuality, i.e. the individual merits and characteristics, of persons. Unlike collective privacy, categorical privacy has its points in respecting and protecting the individual rather than in respecting and protecting some group to which the individual belongs. Furthermore, the conception of categorical privacy presented here – just like many current conceptions of individual privacy – builds on a conception of the personal sphere that is partially predefined by conventions and

social, economic and technical arrangements. Categorical privacy, however, is different from its individual counterpart in that it draws attention to the attribution of generalised properties to members of groups, which may result in the same effects as the attribution of particularised properties to individuals as such. In this respect, infringements of categorical privacy resemble stereotyping and wrongful discrimination on the basis of stereotypes.¹⁸

We began reconsidering privacy norms by questioning the possibilities of profiling through data mining for epidemiological purposes. In section 2, I emphasised that many of these applications are undoubtedly very useful. Earlier on, I also stated that it might be difficult to come to understand precisely the darker side of these applications, the current conceptions of privacy in law, regulation and ethical theory standing in the way. By now, things may be a little less complex. Problems accompanying distributive profiles may at least in part be articulated in terms of a traditional privacy conception. Problems related to non-distributive profiles can be partially articulated in terms of the notion of categorical privacy. It is important to see, however, that the significance or seriousness of these drawbacks depends heavily on the context in which the profiles are used. Suppose that the profiles are produced and used only and strictly for epidemiological purposes, in such a way that there are guarantees that access to them is only permitted to some researchers who do not pass the information on to others. Against this background the privacy objections have a seriousness that is of a rather academic kind. Things are different, however, as soon as the guarantees mentioned are absent. Then the information in the profiles may become available to others and thus become part of the body of public knowledge in society or the information may be used for completely other purposes, such as selection procedures for jobs, insurance, loans, etc. Of course, if the latter happens, then not only privacy is at stake, but also values of social justice and fairness. Social justice is at stake where the distribution of provisions and amenities in society is based on health criteria. Fairness is at stake where non-distributive profiles as such are used. When such profiles are applied, an individual is judged and treated on the basis of characteristics he or she did not acquire voluntarily, such as a bad health condition or a bad health prospect. More importantly, however, an individual as

such will often not exhibit the characteristic at all, since such a characteristic is one of the group and not necessarily also of the individual.

Medical profiles may carry yet another problem with them. Through data mining and profiling, information may be produced about an individual which is unknown to this individual. For instance, one may think of a profile indicating a health risk of a group of persons, without these persons knowing the risk. In such a case, disclosure of the profile to the members of the group may confront these individuals with medical information about themselves which they have not sought freely. It might even confront them with information about a risk which they as a matter of fact do not have. Of course, in situations like these much depends on the ways in which people are confronted with the information. Are they, for instance, informed about the methodical and methodological aspects of the profiles and risks indicated, so that they may conclude that they do not necessarily run the risk? Is it clear to those involved what kind of disease the risk is about? Nevertheless, in certain cases divulging profiles about health condition and health risk may confront persons with information that they had no desire to have.

6. From the protection of privacy to the protection of persons: practical solutions

The problems surrounding profiling through data mining cannot be dealt with in ways similar to those in which individuals are protected against possible infringements of individual informational privacy. The application of principles and rights of, for instance, rectification and consent to potential infringements on categorical privacy is to a large extent impossible. Even if it were possible, it would nevertheless be unacceptable for obvious reasons. First, as has been explained above, the reference group of the profile will only rarely be able to reach and enact collective decisions because of its lack of organisational structure and personal or social ties. Second, if one were to turn from the group as such to the individual members of the group, then an individual's possibility of refusal or of opting out could be harmful. It would be harmful to other members of the reference group as well as to the very person refusing to allow personal information to be used in producing the profile. For, actual refusal will reduce the reliability of the profile or generalisation. Nevertheless, all members of the reference group, including the

individual who opted out, are at risk of being judged and treated on the basis of just this profile with reduced reliability. Of course, the possibility of opting out may also, in some respects, benefit the members of the reference group. If, in the case of application in selection procedures, only people with bad risks actually refuse the use of their information this may turn out to be rather advantageous for the healthy. This, however, does not diminish the wrongfulness of, for instance, judging and treating persons on the basis of properties which they do not, if only with a decreased probability, instantiate.

Perhaps then the only way to protect individuals against the possible negative consequences of the use of generalisations and profiles based on personal information in the broad sense lies in a careful case-by-case assessment of the ways in which the group profiles are in fact used and can be used. By meticulously investigating and evaluating these applications one may hope to find starting points for restrictions of the purposes for which the profiles are produced and applied. An elaborate proposal concerning such acceptable and unacceptable purposes cannot be provided here. It is important, however, to keep in mind that solutions will not be found only in forbidding the production and application of profiles for certain purposes. In many cases, it may be more appropriate to reconsider those purposes themselves. Sometimes it may be easier or even morally more desirable to do something about social and economic arrangements that induce wrongful applications of information technologies than abolishing those applications. This is the case especially where, for instance, profiles can be used for desirable purposes and for undesirable purposes at the same time. Also, in such situations where there is a possibility of good use and bad use of the same newly produced information, doubtless some help is to be expected from encryption and authentication techniques. It may turn out to be possible, for instance, to protect databases against certain types of queries, or to make certain information accessible only to certain persons for certain purposes. Together with agreements in the legal or contractual sphere, through which the behaviour of these persons is bound and can be controlled, technical solutions may turn out to be the best practicable solutions to the problems mentioned.

Touching on practical solutions, it should be noted that data mining can in some way complicate the problems that have been discussed so far even further. Data mining gives us the opportunity of producing profiles easier, quicker and in greater numbers than

was possible before the techniques, called data mining, came to hand. It also enables us to discover many more correlations between phenomena, such as characteristics of individuals and groups, than were known until today. For instance, data mining may show correlations between characteristics that are trivial in a certain context and characteristics that are significant in that same context. It may be possible to establish a statistical relationship between the ownership of a certain type of car, on the one hand, and a certain health risk, on the other (without there necessarily being any natural causal relationship between the two). If this were the case, then the use of the health-risk profile could be hidden, as it were, 'behind' the use of the profile of car owners. This is especially complicating in situations where we would not have other possibilities of controlling what people in a certain institution or enterprise are doing with profiles than to wait until an individual consumer or patient discovers what is happening and starts to complain. The possibility of linking significant, potentially harmful, profiles to trivial ones makes it difficult to uncover dubious applications.¹⁹

7. Concluding remarks

Our current moral and legal vocabularies and conceptual frameworks for dealing with information technologies should be revised and extended in order to adjust them to the problems that arise. On the occasion of problems relating to data mining and profiling for epidemiological purposes, I have suggested a reconceptualisation of privacy. In my account, the basic description frame of 'immunity to the judgement of others' is the unifying factor in the family of privacy conceptions, while the importance of privacy in different situations depends on contextual factors such as conventions, social and economic arrangements and technical developments on the one hand and a large set of diverging values, ranging from individual well-being and personal autonomy to fairness and individuality, on the other. Practical problems relating to new technologies such as those lying at the heart of data mining urge the moral and legal assessment of information technologies to be no longer exclusively focused on privacy norms and related norms such as those of confidentiality. Instead, the scope of privacy norms should be interpreted more broadly, so as to cover a wider range of personal data than data relating to

identifiable individual persons only. In addition, the criteria for assessing the technologies in the field of medical information should also include norms of social justice, fairness and respect for the individuality of persons. Only by making these adjustments will we be able to understand the significance of the problems occurring and will we be in a position to solve conflicts between public-health interests and the protection of individuals and groups against stigmatisation, discrimination and violations of their dignity.

¹ The research for this paper was partially funded by the Netherlands' Organisation for Scientific Research. Thanks are due to Peter Blok and Eric Schreuders of Tilburg University for their constructive suggestions and criticisms.

² Cf. U. Fayyad, G. Piatetsky-Shapiro and P. Smyth, 'Knowledge Discovery and Data Mining: Towards a Unifying Framework', in: E. Simoudis, J. Hian and U. Fayyad, eds., *Proceedings of the Second International Conference on Knowledge Discovery & Data Mining* (Menlo Park, Cal., 1996); W.J. Frawley, G. Piatetsky-Shapiro and C.J. Matheus, 'Knowledge Discovery in Databases: An Overview', in: G. Piatetsky-Shapiro and W.J. Frawley, eds., *Knowledge Discovery in Databases* (Menlo Park, Cal. / Cambridge, Mass. / London, 1991).

³ Of course, the significance of patterns is determined by contextual factors such as the aims of the analyst, the analyst's superiors' aims, their institutional environment, social and political situation, natural predicament, etc.

⁴ Anita Allen and Helen Nissenbaum have also drawn attention to the problem of restricting data protection to the protection of personal data, be it from other perspectives: A. Allen, *Uneasy Access* (Totowa, N.J., 1988); H. Nissenbaum, 'Protecting Privacy in an Information Age: The Problem of Privacy in Public', (1998) 17 *Law and Philosophy*, 559-596.

⁵ See, for instance, W.A. Parent, 'Recent Work on the Concept of Privacy', (1983) 20 *American Philosophical Quarterly*, 341-356; S.I. Benn, *A Theory of Freedom* (Cambridge, 1988), 264-305; J.L. Johnson, 'Privacy, Liberty and Integrity', (1989) 3 *Public Affairs Quarterly* 15-34.

⁶ See, for example: P. Devlin, *The Enforcement of Morals* (Oxford, 1959); H.L.A. Hart, *Law, Liberty and Morality* (London, 1963); A. MacIntyre, 'The Privatization of the Good, An Inaugural Lecture', (1990) 52 *The Review of Politics*, 344-361; and, more recently, A. Etzioni, *The Limits of Privacy* (New York, 1999).

⁷ J.J. Thomson, 'The Right to Privacy' (1975) *Philosophy and Public Affairs* 295-315; T.M. Scanlon, 'Thomson on Privacy' (1975) *Philosophy and Public Affairs*, 315-322.

⁸ W.A. Parent, 'Recent Work on the Concept of Privacy', (1983) 20 *American Philosophical Quarterly* 341-356.

⁹ C. Fried, *An Anatomy of Values* (Cambridge, Mass., 1971); J. Rachels, 'Why Privacy is Important' (1975) *Philosophy and Public Affairs* 323-333.

¹⁰ D. Johnson, *Computers and Privacy* (Upper Saddle River, 1994) 81-102.

¹¹ J.L. Johnson, 'Privacy and the Judgement of Others' (1989) 23 *Journal of Value Inquiry* 157-168; 'Privacy, Liberty and Integrity', (1989) 3 *Public Affairs Quarterly* 15-34; 'A Theory of the Nature and Value of Privacy', (1992) 6 *Public Affairs Quarterly* 271-288; S.I. Benn, *A Theory of Freedom* (Cambridge, 1988), 264-305.

¹² One should not say that privacy is an essentially contested concept. Gallie's idea of essential contestedness ultimately tends to suggest that essentially contested concepts are nonsense; W.B. Gallie, 'Essentially Contested Concepts', (1955-56) 56 *Proceedings of the Aristotelian Society* 167-198.

¹³ For some methodology regarding the ways I expound and defend the conceptual scheme of privacy, see: A. Vedder, 'Considered Judgements: Meaning, Community and Tradition' in W. van der Burg and Th. van Willigenburg (eds.), *Reflective Equilibrium*, (Dordrecht, 1998) 55-72.

¹⁴ J.L. Johnson, 'Privacy and the Judgement of Others' (1989) 23 *Journal of Value Inquiry*, 157

¹⁵ S.I. Benn, *A Theory*, n. 11 above, 264, 305.

¹⁶ M. Hartney, 'Some Confusions Concerning Collective Rights' (1991) 4 *Canadian Journal of Law and Jurisprudence*, 293-314.

¹⁷ A. Vedder, 'Privatization, Information Technology and Privacy: Reconsidering the Social Responsibilities of Organizations' in Geoff Moore (ed.), *Business Ethics: Principles and Practice*. (Sunderland, 1997), 215-226.

¹⁸ J. Harvey, 'Stereotypes and Group-Claims: Epistemological and Moral Issues and Their Implications for Multiculturalism in Education', (1991) 24 *Journal of Philosophy and Education*, 39-50.

¹⁹ Vedder, n. 17 above.

Contact address

Dr Anton Vedder

Schoordijk Institute, Faculty of Law, U 31

Tilburg University

PO Box 90 153

5000 LE Tilburg, The Netherlands

Phone: ** 31 13 466 83 85

Fax: ** 31 13 466 80 45

email: anton.vedder@kub.nl

Biographical Note on the Author

Anton Vedder teaches ethics and is a research fellow at the Faculty of Law of Tilburg University. His research interests are: moral epistemology and semantics, value theory and applied ethics (relating to information technology and medical ethics). Recent publications include articles and papers on Wittgenstein, Kovesi and coherence theories, on ethical aspects of data mining and on medical confidentiality.