Privacy Online†

Herman T. Tavani
Philosophy Department, Rivier College
htavani@rivier.edu

In the recent literature on privacy and technology, considerable attention has been paid to online privacy issues and concerns. For example, a number of books and scholarly journal articles, as well as reports in the popular press, have examined the impact of certain online activities involving the Internet and the World Wide Web for personal privacy. In this essay, I use the expression "online activities" to refer to those activities involving both the Web in particular and the Internet in general. And I use the tern "Internet" in its most generic sense to include protocols such as File Transfer Protocol (FTP) and Gopher as well as the Web (HTTP). Our concern in this study is with privacy issues arising from online activities involving any of these Internet protocols. Note, however, that we will not be concerned with privacy issues peculiar to privately owned computer systems and networks (such as privacy issues involving employee email). Nor will we be concerned with privacy issues involving the use of digital telephony, including devices such as cell phones. Instead, we will limit our analysis of privacy issues to those concerns arising from certain user activities involving Internet resources.

One interesting question to consider—and one that is often overlooked in the contemporary privacy literature—is: To what extent has the Internet itself generated new privacy issues, and to what extent has it merely exacerbated existing ones? It will be argued that in most cases, the privacy issues frequently associated with the Internet are issues which have their genesis in technologies that predate the Internet. It will also be shown, however, that at least two of the privacy issues currently associated with the Internet are such that they would not have arisen without the availability of certain Internet tools and techniques. Issues of the latter type are labeled *Internet-specific* and are contrasted with *Internet-enhanced* privacy issues. Another interesting, and perhaps more important, question is: How can we best resolve Internet-related privacy issues, regardless of whether such issues are unique to, or merely exacerbated by, Internet technology? In the final section of this study, we consider some policy proposals recently put forth to address current online privacy issues involving certain Internet activities. We begin our study with a brief look at the concept of privacy.

The Concept of Privacy

Privacy is a concept that is neither clearly understood nor easily defined. We often hear remarks to the effect that one's privacy has been "lost," "diminished," "intruded upon," "invaded," "violated," "breached," and so forth. Each of these descriptions, in turn, reflects the insights and biases of one or more models or theories of privacy. For example, some theories see privacy as an "all-or-nothing" concept—i.e., privacy is something that one either has (totally) or does not have. Other theories view privacy as something that can be diminished, e.g., as a repository of information possessed by an individual, which can be eroded gradually. Still others see privacy in terms of a spatial metaphor such as a zone that can be intruded upon or invaded by individuals and organizations. And other theories view privacy in terms of confidentiality that can be violated, or trust that can be breached. There are several interesting theories or models of privacy that either directly correspond to, or approximately fit, one or more of the metaphors described above. Elsewhere (see Tavani 1999c), I have critically reviewed and analyzed various philosophical and legal theories of privacy. However, I will not repeat that discussion here, since our primary objective in this work is to examine specific online privacy issues and policy proposals.

It is perhaps worth noting that at least some authors believe that it is more useful to view privacy as either a presumed or a stipulated interest that individuals have with respect to protecting personal information, personal property, or personal space than to think about privacy as a moral or legal right. For example, Posner (1978) has suggested that privacy can be viewed in terms of an economic interest and that information about individuals might be thought of in terms of personal property that could be bought and sold in the commercial sphere. Clarke (1999) has recently suggested that privacy can be thought of as an "interest individuals have in sustaining personal space free from interference by

† Major portions of this essay are extracted from "Privacy and Security," Chap. 4 in *Internet Ethics* (ed. D. Langford). London, UK: Macmillan Press, forthcoming (December 1999). I am grateful to Macmillan Press for permission to reprint that material in the present essay.

other people and organizations." Many Western European nations have preferred to approach issues related to individual privacy as issues of "data protection" for individuals—i.e., as an interest in protecting personal information—rather than in terms of a normative concept that needs philosophical analysis. In the U.S., on the other hand, discussions involving the concept of privacy as a legal right are rooted in extensive legal and philosophical argumentation, including debate in both Constitutional and tort law. However, as stated in the preceding paragraph, our concern in this study is not in examining privacy theories; nor is it in determining whether privacy is a right—moral, legal, or otherwise. Rather, our aim is to consider some specific online-privacy issues as well as some specific policy proposals for addressing those issues.

How is Privacy Threatened by Internet Activities?

A number of recent U.S. studies and surveys suggest that online users are very concerned about their privacy while engaged in activities on the Internet. According to a study conducted by the Boston Consulting Group in 1997 (cited in Wright and Kakalik, 1997), over 70% of 9,300 users responded in an online survey that they were "more concerned about privacy on the Internet" than they were about privacy threats from any other medium. A recent survey conducted by Business Week/Harris survey (cited in Benassi, 1999) indicated that privacy is the "number one consumer issue"—ahead of ease of use, spam, security, and cost—facing the Internet. Wang, Lee, and Wang (1998), who cite a recent survey undertaken by Equifax and Harris Associates which determined that over two-thirds of online consumers surveyed are very concerned about their privacy on the Internet, believe that the most crucial issue currently facing the emerging e-commerce market is the "fear and distrust regarding loss of personal privacy."

Some authors now use the expression "Internet privacy" to refer to privacy concerns associated with activities involving the Internet (see, for example, Cranor, 1999). The use of such an expression might cause us to ask whether we need a separate category of privacy called "Internet privacy." To answer this question, we should determine whether any—and if so, which—privacy threats currently associated with the Internet are uniquely attributable to that medium. However, even if it is determined that some privacy issues are indeed unique or even special to the Internet, we can still consider whether such privacy issues might be adequately handled under certain existing categories of privacy. A brief look at some of those categories would perhaps be useful.

Two Categories of Privacy

Some authors draw a distinction between *information privacy* and *communications privacy* (see, for example, Johnson and Nissenbaum, 1995; and Regan, 1995). Johnson and Nissenbaum differentiate information- or database-privacy issues from issues in communications privacy by placing in the latter category the set of privacy concerns related to technologies such as electronic surveillance, encryption, email, and digital telephony. Information privacy issues, on the other hand, include those associated with personal information contained in large databases. On this two-fold distinction, it might seem that privacy issues related to the Internet would fall under communications privacy. After all, the Internet is among other things a communications medium.

Although certain activities on the Internet have raised privacy concerns that would seem to fall under the heading of communication privacy, other Internet activities have raised privacy concerns that seem more closely related to information privacy. For example, concerns regarding Internet access to personal information that resides in databases have also been raised. So it would seem that Johnson and Nissenbaum's distinction, which generally works very well in parceling out pre-Internet privacy issues into two useful categories, will not help us to sort out Internet-related privacy issues from those privacy issues associated with earlier technologies. It might also seem that because there is no easy way to differentiate these privacy issues, a separate "Internet privacy" category is needed. Of course, we have not yet determined whether any privacy issues currently associated with the Internet are qualitatively or even significantly different from pre-Internet privacy concerns. In deciding whether a separate privacy category is warranted, then, it would be useful to contrast certain Internet and pre-Internet privacy concerns.

Privacy Threats Involving Pre-Internet Technologies

First, we should note the obvious but relevant point that concerns about personal privacy existed long before the era of the Internet. For that matter, they existed long before the advent of either modern information or communications technologies. Prior to the computer era, for example, certain uses of technologies such as the camera and the telephone raised concerns about personal privacy. Earlier forms of record keeping also raised privacy concerns. Since privacy controversies regarding the use of technology to collect and communicate personal data predate the Internet era, it might seem that there is nothing at all new with respect to the privacy issues currently associated with the Internet. For example, it might appear that Internet technologies have merely intensified the debate over concerns already introduced by the use of earlier information and communications technologies. However, even if certain privacy concerns currently identified with the Internet were introduced in certain pre-Internet technologies, we should not underestimate the magnitude of the impact that the Internet itself has had for personal privacy. And more importantly, we should not infer, based on what we have seen thus far, that *no* new or special privacy threats have been introduced by the use of Internet tools and resources.

Activities on the Internet have contributed to the ongoing privacy-and-computers debate in at least two ways. First, the Internet has made it possible for certain existing privacy threats to occur on a scale that would not have been possible, i.e., in a practical sense, with pre-Internet technologies. This set of privacy concerns, while not original to the Internet, can be said to be *enhanced* by the Internet. Secondly, it has made possible, by virtue of certain tools and techniques unique to the Internet itself, specific privacy threats that were not possible with earlier information and communications technologies. These latter privacy concerns could be said to be *specific* to the Internet. We next consider both kinds of concerns, which we shall henceforth refer to as *Internet-specific* and *Internet-enhanced* privacy concerns, beginning with an analysis of the latter.

Internet-Enhanced Privacy Threats

Privacy concerns which may have arisen because of certain uses of earlier information and communications technologies, but which are now also inextricably associated with and exacerbated by the Internet, can be analyzed under two general headings: (i) dataveillance and data gathering, and (ii) data exchange and data mining.

Dataveillance and Data-Gathering Activities on the Internet

Some authors suggest that the Internet can be viewed as a new "surveillance medium." Clarke (1988) uses the term *dataveillance* to refer to the systematic use of systems, which would include the Internet, in the "monitoring of people's actions or communications." As noted earlier, privacy issues related to surveillance and data-monitoring (as well as data-gathering) techniques are hardly new to the Internet. Nonetheless, concerns over surveillance have, in more recent times, been aggravated by certain Internet activities. And because of its surveillance and data-recording capabilities, many privacy advocates believe that the Internet poses a threat to privacy on a scale that could not have been realized in the pre-Internet era.

In the early days of computing, when computers were owned and operated mostly by large public agencies, it was feared that strong centralized governments would be able to monitor the day-to-day activities of their citizens. Today, however, the threat of surveillance comes not so much from governments and their agencies, at least not in Western democratic societies, as it does from surveillance by online businesses and corporations in the private sector who can now monitor the activities of persons who visits their Web sites, determine how frequently those persons visit particular sites, and draw conclusions about which preferences those visitors have when accessing a certain site. Even the number of "clickstreams"—key strokes and mouse clicks—entered by a Web site visitor can be monitored and recorded.

Data about Internet users can be gathered either directly or indirectly from their online activities. One direct method of information gathering involves the use of *Web forms*, a data-gathering mechanism into which Web users enter information online. Forms technology, often used to collect information about Web visitors (such as a user's name and address) can, as Kotz (1998) notes, also be used to track the sequence of pages one visits within a given Web site. For the most part, the use of Web forms would seem to be uncontroversial since online users voluntary submit the requested information. However, information gained from forms can be combined with other directly gathered information, such as information about the items a user purchases online, and can then be combined with online information about individuals that is gathered indirectly. One indirect method for gathering personal information involves the use of Internet *server log files*. Kotz points out that because Web browsers transmit certain kinds of data to a Web server, such as the Internet address of the user's computer system as well as the brand name and version number of the user's Web browser and operating system software, Internet server logs can be used to gather personal data in an indirect manner. Information gathered indirectly from server logs can also be combined with information gained directly from Web forms, which can then eventually be used by advertising agencies to target specific individuals.

Although Web forms and Internet server logs can be used in ways that pose significant threats to the privacy of Internet users, it is worth noting that neither of these two technologies is unique to the Internet since the development and use of both technologies predates the Internet era. However, the scale on which dataveillance and data gathering can now be carried out on the Internet has increased dramatically because of the use of forms- and server-log technologies.

Data-Exchanging and Data-Mining Activities on the Internet

Whereas the dataveillance and data-gathering tools described in the preceding section are used mainly to monitor and record activities of online users, other tools are used to exchange that data across databases accessible the Internet. This exchange of online personal information often involves the sale of personal data to third parties, which has resulted in commercial profits for certain online entrepreneurs, often without the knowledge and consent of individuals about whom the data is exchanged.

Techniques for exchanging personal data online are hardly new to the Internet. Data-exchange techniques such as the merging (see Tavani, 1996a) and matching (see, Tavani, 1996b) of electronic records stored in databases occurred before the arrival of the Internet. However, Internet technology has facilitated the exchange of online personal information at a rate that was not possible in the pre-Internet era. In response to earlier data-exchange practices involving privately owned computer networks, certain privacy laws and data-protection guidelines have been enacted and implemented. These laws and principles, which specifically address the

exchange of personal information between databases in privately owned computer networks, might also seem by extension to apply to the exchange of personal data on the Internet as well.

Although earlier privacy legislation involving the exchange of personal information in databases have centered mainly on the transfer of such information *between* databases in computer networks, some recent privacy concerns have arisen because of the kind of personal information that can now be extracted or "mined" from *within* a single database. These concerns arise from the use of a certain technique commonly referred to as data mining. Data-mining technology, which combines research in artificial intelligence (AI) and pattern recognition, is defined by Cavoukian (1998) as a "set of automated techniques used to extract buried or previously unknown pieces of information from large databases." Using data-mining techniques, it is possible to unearth patterns and relationships, which were previously unknown, and to use this "new" information—i.e., new "facts" and relationships in the data—to make decisions and forecasts. Through the use of data-mining algorithms, individual pieces of data about an online user's activities, which in themselves might seem innocuous, can be recorded, combined, and recombined in ways to construct profiles of individuals. Moreover, the profiles or groups in which individuals end up may be ones that are not at all obvious to those individuals. As a result of data-mining applications, for example, an individual might eventually discover that he or she belongs to some consumer category or some "risk group," the existence of which he or she had been previously unaware (see Tavani, 1998a, 1998b).

The mining of personal data in the pre-Internet era, which depended on the use of large commercial databases called *data warehouses* to store the data, focused mainly on transactional information. Personal data mined from the Internet, however, need not be (and frequently is not) transactional. For example, information typically included in and mined from personal Web pages, as well as noncommercial Web sites, is nontransactional. Because of Internet commerce, however, much transactional information can now also be gained from the Web as well. When an individual orders a book from Amazon.com (an online book store), for instance, transactional information is recorded about the purchase, and information about that particular transaction can be (and frequently is) used for future business decisions. What distinguishes the Internet as a mining resource from the large databases or data warehouses used in data mining, however, is the vast amount of nontransactional, personal information currently available for mining on the Internet.

Cavoukian (1998), who points out that one of the purposes of data mining is to "map the unexplored terrain of the Internet," notes that the Internet is becoming an "emerging frontier for data mining." She also notes that with access to an Internet server, it is possible to FTP (file transfer protocol) the data from the client's server and then conduct various data mining activities. Fulda (1998) points out that because data-mining software employs certain AI techniques, it can "learn" about the Web by coming to understand the content associated with common HTML tags. Eisenberg (1996) notes that *intelligent agents* can "sift through" the potential wealth of data on the Internet, and O. Etzioni (1996) describes the use of "learning techniques" or systems such as *softbots* (intelligent software robots or agents that use tools on a person's behalf) and metasearch engines (such as MetaCrawler and Ahoy) to uncover general patterns at individual Web sites and across multiple Web sites. So data-mining techniques that previously raised privacy concerns only at the database or data-warehouse level now raise concerns on the Internet as well.

Although data-mining technology is currently used on the Internet, we have also seen that the use of that technology predates the Internet era. So in this sense, privacy issues associated with data mining on the Internet are similar in kind to privacy concerns associated with Web forms and Internet server log files. Like the privacy concerns raised by the use of the latter tools, concerns raised by data-mining technology are not unique to the Internet. Instead, they are instances of what we have earlier identified as Internet-enhanced privacy concerns.

Internet-Specific Privacy Threats

We next examine privacy issues that are specific to, rather than merely enhanced by, activities on the Internet. Privacy concerns that are attributable to tools and techniques provided by the Internet itself, which we described earlier as "Internet-specific" concerns, arise mainly from the use of two new types of online tools. One technique involves the gathering of personal data from users who visit certain Web sites, whereas the other tool can be used to locate information about persons via the Internet.

Internet Cookies

Through the use of a data-gathering technique called Internet cookies, online businesses and Web-site owners can store and retrieve information about a user who visits their Web sites, typically without that user's knowledge or consent. Cookies technology has generated considerable controversy, in large part, because of the novel way in which certain information about Internet users can be collected and stored. Information about an individual's online browsing preferences can be "captured" while that user is visiting a Web site, and then stored on a file placed on the hard drive of the user's computer system. The information can then be retrieved from the user's system and resubmitted to a Web site the next time the user accesses that site. Cookies technology is the only data-gathering technique that actually stores the data it gathers about a user on the user's computer system.

It is perhaps important to note that owners and operators of one Web site cannot access cookies-related information pertaining to a user's activities on another Web site. However, information about a user's activities on different Web sites can, under certain circumstances, be gathered and compiled by online advertising agencies. For example, online advertising agencies such as DoubleClick.net, who pay to place advertisements on Web sites, include a link from a host site's Web page to the advertising agency's URL. So when a user accesses a Web page that contains an advertisement from DoubleClick.net, a cookie is sent to the user's system not only from the requested Web site but also from that online advertising agency. The advertising agency can then retrieve the cookie from the user's system and use the information it acquires about that user in its marketing advertisements. The agency can also acquire information about that user from cookies retrieved from other Web sites the user has visited, assuming that the agency advertises on those sites as well. The information can then be combined and cross-referenced in ways which enable a marketing profile of that user's online activities to be constructed and used in more direct advertisements.

Several privacy advocates have argued that because cookies technology involves the monitoring and recording of a user's activities while visiting Web sites (without informing the user) as well as the subsequent downloading of that information onto a user's computer system, that technology violates the user's privacy. Defenders of cookies, who are usually owners of online businesses and Web sites, maintain that they are performing a service for repeat users of a Web site by customizing a user's means of information retrieval and by providing the user with a list of preferences for future visits to that Web site. Despite any alleged advantages provided by cookies to users who frequently visit one or more Web sites, most users surveyed indicated that they are more concerned about not losing their privacy while visiting Web sites than they are with gaining customized retrieval preferences for their favorite sites. According to a 1996 Equifax/ Harris Internet Consumer Privacy Survey (cited in Wright and Kakalik 1997), 64% of respondents believed that providers of on-line services and Web sites should not be able to track users' activities on the Internet, including the Web sites they visit, regardless of whether that information is eventually used by online advertisers.

It should be noted that in the newer versions of most Web browsers, users have an option to "disable" cookies. With that capability, users can either opt-in or opt-out of cookies, assuming that they are aware of cookies technology and assuming that they know how to enable/disable that technology through their Web browsers. So it would seem that many users now have more control over whether or not to accept cookies. Nonetheless, cookies technology continues to remain a point of contention for many privacy advocates. Because the privacy concerns raised by cookies are unique to Internet technology, and thus did not exist prior to the Internet, these particular concerns can be categorized as Internet-specific privacy issues.

Internet Search Engines

Internet technology has also provided tools that support new techniques for locating and retrieving information about persons. Wright and Kakalik (1997) note that a certain kind of information about individuals, which was once difficult to find and even more difficult to cross-reference, is now readily accessible and collectible through the use of automated search facilities on the Internet. These facilities are called Internet search engines. Included in the list of potential topics on which search-engine users can inquire is information about individual persons. By entering the name of an individual in the search-engine program's entry box, search engine users can potentially retrieve information about that individual. However, because an individual may be unaware that his or her name is among those included in a search-engine database, or because he or she might be altogether unfamiliar with search-engine programs and their ability to retrieve information about persons, questions concerning the implications of search engines for personal privacy have been raised (see Tavani, 1997). As in the case of privacy concerns associated with Internet cookies, privacy issues involving search engines did not exist prior to the Internet era.

A search for a person's name will often return the addresses of Web pages written by that person or the addresses of Web sites that include information about that person. Kotz (1998) points out that since many email-discussion lists are stored and archived on Web pages, it is possible for a search engine to locate information that users contribute to electronic mailing lists or listservers. Search engines can also search through archives of news groups, such as Usenet, on which online users also post and retrieve information. One such group, DejaNews, is set up to save permanent copies of new postings and thus provides search engines with a comprehensive searchable database. DejaNews also provides "author profiles" which include links to all of the online articles posted by a particular person. Because the various news groups contain links to information posted by a person, they can provide search-engine users with considerable insight into that person's interests and activities.

It could be argued that information currently available on the Internet, including information about individual persons, is, by virtue of its residing on the Internet, public information. We can, of course, question whether all of the information currently available to Internet users should be viewed as public information. The following scenario (see Tavani, 1997) may cause us to question whether at least some information about individual persons, such as personal information stored on a Web server or in a database that is accessible to Internet users, should be viewed as public in-

formation. Consider a case in which an individual contributes to a cause sponsored by a homosexual organization. That individual's contribution is later acknowledged in the organization's newsletter (a hardcopy publication that has a limited distribution). The organization's publications, including its newsletter, are then converted to electronic format and included on the organization's Internet Web site. The Web-site is "discovered" by a search-engine program and an entry about that site's address is recorded in the search engine's database. Suppose that you enter this individual's name in the entry box of a search-engine program and a "hit" results, identifying that person with a certain homosexual organization. Since the individual in question might have no idea that such publicly available information about his or her activities exists in a database or on a server accessible to Internet search-engine programs, or have no control over how that information is subsequently used, the use of search-engine technology in this case could indeed raise privacy concerns for that individual.

Some Proposals for an Online-Privacy Policy

Regardless of whether the above online-privacy issues are labeled Internet-specific or Internet-enhanced, the level of controversy that has recently been generated by these issues is such that an online-privacy policy, or perhaps even a set of online-privacy policies, is clearly needed. Numerous proposals that address privacy concerns involving activities on the Internet have recently been put forth. Some proposals call for stricter privacy laws on the part of governments and for the formation of privacy oversight commissions to enforce those laws. Others call for more serious self-regulatory measures by those in the commercial sector. Still other proposals call for technological solutions through the use of privacy-enhancing tools by Internet users. We begin with a look at proposals involving privacy-enhancing technologies.

Privacy-Enhancing Technologies and Industry Self-Regulation Initiatives

A number of tools and techniques for enhancing individual privacy are currently available to Internet users. One type of application allows online users to be *anonymous* while they interact with the Internet. Certain "anonymizing" programs provide online users with a considerable degree of anonymity in so far as they make it extremely difficult to identify individual users by linking the Internet activities of those users to a particular Internet (or IP) address. Although many privacy advocates have championed the use of anonymity programs because they allow for greater privacy for online users, it must be also be noted that these tools have raised concerns for security on the Internet. Since anonymity tools also make it difficult to identify and trace the activities of those individuals who engaged in unauthorized activities on the Internet, some government agencies and law enforcement organizations have been critical of these tools.

One of the best known anonymity tools is a program called the *Anonymizer* (www.anonomizer.com). It is important to note, however, that although Anonymizer users enjoy anonymity while visiting Web sites, they are not anonymous to the Anonymizer itself or to their Internet Service Providers (ISPs). As we saw earlier in the section on Internet data-gathering tools, most users' activities on Web sites are recorded in server log files and thus, in certain cases, can be traced back to a specific ISP or IP. To enjoy complete anonymity on the Internet, online users need tools that do not require them to place their trust in a single "third party." Two currently available anonymity programs that do not require such trust are *Crowds* and *Onion Routing*. Crowds is an anonymity tool based on the idea that people can be anonymous when they "blend into a crowd." When Crowds users submit a request such as the URL of a certain Web site, the request is sent to a "randomly selected member of their crowd." So neither the server at the end of the destination (e.g., a particular Web site) nor any member of the "crowd" can determine where the request originated (see Reiter and Rubin, 1999). A slightly different solution to the problem of needing to trust in a third party is offered by the Onion Router program, through which users submit encrypted requests. An "onion"—i.e., a "layered data structure" that specifies cryptographic algorithms and keys—is sent to the intended recipient. On this scheme, one "layer of encryption" is removed according to the "recipe" contained in the onion, as the data passes between each onion-router along the way. Only the IP address of the last onion-router on the path is revealed to the requested Web site (see Goldschlag, Reed, and Syverson, 1999).

While anonymity programs are useful for Internet activities in which users have neither a need nor a desire to be identified, they are not useful for certain online activities which require identification of users. For example, when Internet users wish to make online purchases, they often need to provide some identifying information. To attract online consumers who might be inclined to avoid e-commerce activities because of privacy worries, many online entrepreneurs have supported the need for a set of Internet-wide privacy standards. The World Wide Web Consortium (W3C), an international industry consortium, was commissioned to provide these standards. In 1997, W3C announced its Platform for Privacy Preferences (P3P), a standard through which Internet users would be able to specify their own individualized Web privacy preferences. For example, users would be able to enter their privacy preferences into their Web browsers, choosing from a range of options. Users would also have the ability to change their privacy preferences each time they accessed a Web site. On the P3P standard, data exchange between a user and Web site would occur only when a user's indicated privacy preferences match that Web site's stated privacy policy. Although W3C was charged with establishing privacy standards on the Web, it was not set up

to enforce the actual protection of personal data. Thus, additional measures are needed to assure users that any personal information that they release to a Web site will be used only in the ways they had specified. Various technologies, sometimes referred to as "negotiation agents" and "trust engines" have been developed to assist users in this process.

One industry-backed privacy initiative, called TRUSTe (www.trustee.org), which is self-regulatory in nature, was designed to ensure that Web sites would indeed adhere to the privacy policies they advertise. TRUSTe uses a branded system of "trustmarks" (graphic symbols) which represent a Web site's privacy policy regarding its use of personal information. Trustmarks provide consumers with the assurance that a Web site's privacy practices accurately reflect its stated policies, and online users are assured that there will be an audit of the Web site and a means of recourse for them if the Web site does not abide by its stated policies (see Benassi, 1999). Any Web site which bears the TRUSTe mark must satisfy several conditions regarding the disclosure of personal information about users who visit that site. Wright and Kakalik (1997) point out that such a Web site must also clearly explain in advance its general information-collecting practices, including which personally identifiable data will be collected, what the information will be used for, and with whom the information will be shared. In addition, the Web site must disclose whether the user will be able to correct and update personally identifiable information, and whether the user information will be removed from that Web site's database upon request.

Critics point out that the practical application of TRUSTe and similar tools may prove difficult. For example, certain critics worry that the amount of information users are required to provide may discourage some users from carefully reading and thus adequately understanding what is expected from them. Other critics point out that the various "warnings" these tools display may appear "unfriendly" and thus work against the ideal of easy Web site access and use. They also note that, unfortunately, "friendlier" trustmarks or graphic icons might result in online users being supplied with less direct information that is important to protecting their privacy. Despite the concerns raised by critics, advocates of tools such as TRUSTe argue that since online users will be better able to make informed choices regarding online transactions, including electronic purchases, those users would clearly benefit from programs like TRUSTe. Certain critics, however, worry that such programs do not go far enough, and some argue that what is also needed is stronger privacy legislation and enforcement of privacy laws.

Privacy Legislation and Data-Protection Principles

In response to recent concerns about privacy concerns on the Internet, some nations have recently passed or are currently in the process of seriously considering strong privacy legislation. The U.S., however, has not been at the forefront of this movement. In 1974 the U.S. Congress passed the Privacy Act, which has been frequently criticized for containing far too many loopholes and for lacking adequate provisions for enforcement. This Act is also restricted in its application to records in federal agencies and thus is not applicable in the private sector. Subsequent privacy legislation in the U.S. has resulted in a "patchwork" of individual state and federal laws that are neither systematic nor coherent. Generally, the U.S. government has resisted requests from the public for stronger privacy laws, siding instead with various business interests in the private sector whose concerns are based on a belief that such legislation would undermine economic efficiency and thus adversely impact the overall U.S. economy. Critics point out, however, that many of those U.S. businesses who also have subsidiary companies or separate business operations in countries with strong privacy laws and regulations, such as nations in Western Europe, have found little difficulty in complying with the privacy laws of the host countries, and that profits for those American-owned companies have not suffered because of their compliance. In any event, there is now increased pressure on the U.S. government to enact stricter privacy laws and on American businesses to adopt stricter privacy polices and practices because of global e-commerce pressures, especially from Canada and the European Union.

Many European nations have, through the implementation of strict data-protection principles, been far more aggressive than the U.S. in addressing privacy concerns of individuals. In 1980, most Western European nations signed on to the Organization for Economic Cooperation and Development (OECD) Principles, and in the early 1990s the European community began to consider proposals for synthesizing the data-protection laws of the individual European nations. The European Community has recently instituted a series of directives, including EU Directive 95/46/EC of the European Parliament and of the Council of Europe of 24 October 1995, which protects the personal data of its citizens by prohibiting the transborder flow of such data to countries that lack "adequate" protection of personal data. As in the case of Canada, which has set up privacy oversight agencies with a Privacy Commissioner in each of its provinces, many European countries have also established data-protection agencies to enforce data-protection laws.

Two Comprehensive Proposals

We have seen some of the strengths and limitations of privacy legislation and data-protection principles, industry self-regulation initiatives, and privacy-enhancing technologies as separate or distinct online-privacy proposals. Although each type of proposal seems to offer an important corrective measure, none by itself seems fully adequate. Perhaps, then, some combination of these individual solutions can be integrated into a more comprehensive and robust privacy policy for the Internet. Clarke (1999) and Wang, Lee, and Wang (1998)

have recently suggested some possible ways of combining the individual pieces. Arguing for a "co-regulatory" model, Clarke believes that a successful online-privacy policy must include strong legislation, a privacy oversight commission, and industry self-regulation, and that these provisions must also be accompanied by privacy-enhancing technologies that individuals can use to ensure their privacy. He further believes that a "privacy watchdog agency" and sanctions are also both needed for the overall privacy scheme to work.

Wang, Lee, and Wang also suggest that governments, businesses, and individuals each have a key role to play in any successful privacy policy. For example, the role of government would be to promote strong privacy laws in both the public and private sectors, to establish independent privacy commissions to oversee the implementation and enforcement of those laws, and to educate the public about privacy issues. Whereas businesses would be responsible for promoting self-regulation for fair information practices and for educating consumers about their online-privacy policies, individuals themselves would be responsible for using privacy-enhancing technologies and security tools. The authors admit, however, that developing an adequate privacy policy for the Internet is "one of the most challenging public policy issues of the information age."

It must be noted that not everyone agrees that stronger privacy-protection policies are needed. In fact, some argue that we have already moved too far in the direction of privacy protection, and that as a result we are in danger of jeopardizing the interests of the larger public good (see, for example, A. Etzioni, 1999). From what we have seen thus far in our analysis of certain activities involving the Internet, however, it would seem that a comprehensive policy similar to the models proposed by Clarke and by Wang, Lee, and Wang are clearly needed to ensure adequate privacy protection for individuals who engage in Internet activities.

Concluding Remarks

We conclude this study with some thoughts on how the competing parties and interests in the current privacy debate might be able to reach a solution that is satisfactory to those on both sides. A standard way of framing the debate over interests involving individual privacy and the implementation of a new technology has been in terms of a need to *balance* competing interests (see, for example, Johnson 1994). Because of recent activities on the Internet, some believe that there is an urgent need to balance the privacy interests of individuals against the economic interests of online businesses as well as against the interests of the greater public good—i.e., to balance the interests of those groups (e.g., government agencies and corporations) who claim to have a legitimate need for information about individuals against the needs or rights of those individuals about whom the information is collected. Others, however, believe that simply using a balancing scheme based on the tradeoff of interests involving the individual good vs. the larger social good misses an important point, because such a decision procedure fails to take into account the significance of privacy as a social (as well as an individual) value.

Regan (1995) and Blanchette and Johnson (1998) note that when we frame the debate simply in terms of how to balance privacy interests as an individual good against interests involving the larger social good, support for those interests believed to benefit the latter good will generally override concerns regarding individual privacy. For example, if evidence were put forth to show that the use of a certain Internet technology would increase the number of jobs in a community or would raise that community's standard of living, then a decision to use that technology would likely be perceived as yielding a greater overall good than would a decision to forgo that technology for the sake of protecting the privacy of individuals. If, however, privacy is understood not merely as a value involving the good of individuals but as one that also contributes to the broader social good—i.e., a value that is essential for democracy and freedom—then concerns for individual privacy might have a greater chance of receiving the kinds of consideration they deserve in those debates involving the balancing of competing interests. Keeping in mind this understanding of the importance of privacy for our social institutions and human values, it would seem that in future debates involving privacy and the Internet, especially in those nations where strong democratic institutions are valued, there are good reasons to proceed from a position that DeCew (1997) and others describe as a "presumption in favor of privacy."

Throughout this essay, there has been a presumption on my part that privacy is a positive value which is worth protecting. Essentially, I agree with DeCew (1997) that we should presume in favor of privacy and then develop ways that would "allow individuals to determine for themselves how and when that presumption should be overridden." Moor (1997) has recently put forth a proposal that describes how such a process might be carried out. He begins by arguing that privacy is best understood when viewed within the parameters of certain contexts or what he calls "situations." Moor goes on to note that we can then decide, via a rational dialog and debate, which situations to declare *normatively private*—i.e., which ("naturally private") situations to grant moral, legal, or perhaps some other kind of normative protection. One virtue of Moor's privacy theory is that it employs a "Publicity Principle" which requires that debates over whether a given "situation" should be protected as a normatively private one must be open to all interested parties affected by that situation. Through a rational discussion and debate of relevant issues, the interested parties can decide, by way of a referendum, whether a given situation will be declared normatively private. And since all of the parameters regarding the disputed situation must be specified in a

manner that is open and public, individuals affected by that situation can make "informed choices" in determining whether that situation should be protected as a normatively private one. To see how Moor's theory can be applied to privacy issues involving the use of Internet search engines and data mining, as well as other Internet activities, see Tavani (1997, 1998b, and 1999a).

Admittedly, many of the points raised in the concluding section of this study are in need of further development, and they require supporting argumentation. To provide that development and support here, unfortunately, would take us beyond the intended scope of the present work. My principal aim in this essay has been to show how personal privacy is currently threatened by certain online activities that are either made possible through or exacerbated by the Internet, and to consider some policy proposals for addressing those privacy concerns. ◆

Acknowledgments

I am especially grateful to Duncan Langford for his helpful comments on a longer version of this essay (included as Chap. 4 in his book *Internet Ethics*). This essay has also benefited from some valuable suggestions I received from Deborah Johnson and Jim Moor on a related project on privacy and the Internet, while I was a Fellow at Dartmouth College's Humanities Institute during the summer of 1998. The sections of this essay pertaining to privacy concerns involving data mining have benefited from discussions with Joseph Fulda and Steve Cooper.

References

Benassi, P. (1999) "TRUSTe: An Online Privacy Seal Program," *Communications of the ACM*, Vol. 42, No. 2, February, pp. 56-59.

Blanchette, J-F. and D. G. Johnson. (1998) "Data Retention and the Panopticon Society: The Social Benefits of Forgetfulness." In *Proceedings of the Conference on Computer Ethics: Philosophical Enquiry: CEPE '98.* (ed. L.D. Introna). London, UK: University of London Press, pp. 94-105.

Cavoukian, A. (1998) *Data Mining: Staking a Claim on Your Privacy.* Ontario, Canada: Information and Privacy Commissioner's Report.

Clarke, R. (1988) "Information Technology and Dataveillance," *Communications of the ACM*, Vol. 35, No. 5, May, pp. 498-512.

Clarke, R. (1999) "Internet Privacy Concerns Confirm the Case for Intervention," *Communications of the ACM*, Vol. 42, No. 2, February, pp. 60-67.

Cranor, L. F. (1999) "Internet Privacy," *Communications of the ACM*, Vol. 42, No. 2, February, pp. 28-31

DeCew, J. W. (1997) *In Pursuit of Privacy: Law, Ethics, and the Rise of Technology.* Ithaca, New York: Cornell University Press.

Eisenberg, A. (1996) "Privacy and Data Collection on the Net," *Scientific American*, March, p. 120.

Etzioni, A. (1999) *The Limits of Privacy.* New York: Basic Books.

Etzioni, O. (1996) "The World Wide Web: Quagmire or Gold Mine?" *Communications of the ACM*, Vol. 39, No. 11, November, pp. 65-68.

Fulda, J. (1998) "Data Mining and the Web," *Computers and Society*, Vol. 28, No. 1, March, pp. 42-43.

Goldschlag, D., Reed, M., and P. Syverson. (1999) "Onion Routing for Anonymous and Private Internet Connections," *Communications of the ACM*, Vol. 42, No. 2, February, pp. 39-41.

Johnson, D. G. (1994) *Computer Ethics.* 2nd ed. Englewood Cliffs, NJ: Prentice Hall.

Johnson, D. G. and H. Nissenbaum, eds. (1995) *Computers, Ethics & Social Values.* Englewood Cliffs, NJ: Prentice Hall.

Kotz, D. (1998) "Technological Implications for Privacy." In *Proceedings of the Conference on The Tangled Web: Ethical Dilemmas of the Internet* (ed. J.H. Moor). Hanover, NH: Dartmouth College, forthcoming.

Moor, J. H. (1997) "Towards a Theory of Privacy in the Information Age," *Computers and Society*, Vol. 27, No. 3, September, pp. 27-32.

Posner, R. A. (1978) "An Economic Theory of Privacy," *Regulations*, May-June, pp. 19-26.

Regan, P. M. (1995) *Legislating Privacy: Technology, Social Values, and Public Policy.* Chapel Hill, North Carolina: The University of North Carolina Press.

Reiter, M. K. and A. D. Rubin. (1999) "Anonymous Web Transactions With Crowds," *Communications of the ACM*, Vol. 42, No. 2, February, pp. 32-38.

Tavani, H. T. (1996a) "Personal Privacy in the Information Age: An Ethical Dilemma," *InSight*, Vol. 3, No. 1, pp. 155-169.

Tavani, H. T. (1996b) "Computer Matching and Personal Privacy: Can They Be Compatible?" In *Proceedings of the Symposium on Computers and the Quality of Life: CQL '96.* (ed. C. Huff). New York: ACM Press, pp. 97-101.

Tavani, H. T. (1997) "Internet Search Engines and Personal Privacy." In *Proceedings of the Conference on Computer Ethics: Philosophical Enquiry: CEPE '97.* (ed. M.J. van den Hoven). Rotterdam, The Netherlands: Erasmus University Press, pp. 214-223.

Tavani, H. T. (1998a) "Informational Privacy, Data Mining, and the Internet." In *Proceedings of the Conference on The Tangled Web: Ethical Dilemmas of the Internet.* (ed. J.H. Moor). Hanover, NH: Dartmouth College. Reprinted in *Ethics and Information Technology*, Vol. 1, No. 2, pp. 137-145.

Tavani, H. T. (1998b) "Data Mining, Personal Privacy, and Public Policy." In *Proceedings of the Conference on Computer Ethics: Philosophical Enquiry: CEPE '98.* (ed. L.D. Introna). London, UK: University of London Press, pp. 113-120. Reprinted in *Ethics and Information Technology*, Vol. 1, No. 4, in press.

Tavani, H. T. (1999a) "Privacy and the Internet." In *Proceedings of the Fourth Annual Conference on Ethics and Technology.* (ed. R.A. Spinello). Chestnut Hill, MA: Boston College Press, pp. 114-125.

Tavani, H. T. (1999b) "Internet Privacy: Some Distinctions Between Internet-Specific and Internet-Enhanced Privacy Concerns." In *Proceedings of the ETHICOMP99 Conference* (eds. A. D'Atri, et al.). Rome, Italy: LUISS Guido Carli University, in press.

Tavani, H. T. (1999c) "Philosophical Foundations of Privacy: Some Implications for the World Wide Web." In *Proceedings of the 1999 Northern New England Philosophical Association (NNEPA) Conference.* Manchester, NH: St. Anselm College, forthcoming.

Tavani, H. T. (in press) "Privacy and Security," Chap. 4. in D. Langford, ed. *Internet Ethics.* London, UK: Macmillan Press.

Wang, H., Lee, M., and C. Wang, (1998) "Consumer Privacy Concerns About Internet Marketing," *Communications of the ACM*, Vol. 41, No. 3, March, pp. 63-70.

Wright, M. and J. Kakalik. (1997) "The Erosion of Privacy," *Computers and Society*, Vol. 27, No. 4, December, pp. 22-25.