

Specifying Personal Privacy Policies to Avoid Unexpected Outcomes¹

George Yee and Larry Korba
Institute for Information Technology
National Research Council Canada
{George.Yee, Larry.Korba}@nrc-cnrc.gc.ca

Abstract

The growth of the Internet is increasing the deployment of e-services in such areas as e-commerce, e-learning, and e-health. In parallel, the providers and consumers of such services are realizing the need for privacy. The use of P3P privacy policies on web sites is an example of this growing concern for privacy. Managing privacy using privacy policies is a promising approach. In this approach, an e-service provider and an e-service consumer each have separate privacy policies. Before an e-service is engaged, the provider's policy must be "compatible" with the consumer's policy. However, beyond compatibility, the policies may lead to unexpected outcomes. This can result in the lost of privacy and even lead to serious injury in certain cases. This paper gives examples of how such outcomes can arise and suggests how the consumer's personal privacy policy can be modified to avoid such outcomes.

Keywords: privacy, personal privacy policy, specification, e-service, unexpected outcomes

1. Introduction

The rapid development of the Internet has been accompanied by a growth in the number of e-services available to consumers. E-services, and in particular, web services, are available for banking, shopping, learning, healthcare, and Government Online. However, each of these services requires a consumer's personal information in one form or another. This leads to concerns over privacy. Indeed, the public's awareness of potential violations of privacy by online service providers has been growing. Evidence affirming this situation include a) the use of P3P privacy policies [1] by web server sites to disclose their treatment of users' private information, b) the

enactment of privacy legislation and directives by major jurisdictions as a sort of owners' "bill of rights" concerning their private information, and c) the appointment of privacy commissioners or officials who can assist the consumer in addressing violations of privacy (Canada has a federal privacy commissioner as well as provincial level privacy commissioners). In order for e-services to be successful, privacy must be protected. An effective and flexible way of protecting privacy is to manage it using privacy policies. The objectives of this paper are a) to show that such use of privacy policies can lead to unexpected outcomes and b) to propose ways to eliminate or mitigate these bad outcomes.

1.1. Privacy legislation and directives

In Canada, privacy legislation is enacted in *PIPEDA (Personal Information Protection and Electronic Documents Act)* [2] and is based on the Canadian Standards Association's *Model Code for the Protection of Personal Information* [3] recognized as a national standard in 1996. This Code consists of ten Privacy Principles that for convenience, we label as CSAPP. We will examine the CSAPP below.

Data privacy in the European Union is governed by a very comprehensive set of regulations called the Data Protection Directive [4].

In the United States, privacy protection is achieved through a patchwork of legislation at the federal and state levels. Privacy legislation is largely sector-based [5]. At the Federal level there are presently more than a dozen privacy laws. Some of these laws are: Privacy Act of 1974 as amended (5 USC 552a), Electronic Communications Privacy Act of 1986, and Right to Financial Privacy Act of 1986. Laws applicable to the private sector include: Family Educational Rights and Privacy Act of 1978, Privacy Protection Act of 1980, and Video Privacy Protection Act of 1988. As can be seen, the laws typically apply to specific technologies or privacy threats to, for example, bank records,

government databases, or video rental history. The laws serve as operational boundaries rather than requirements and there is no national all encompassing code for privacy protection. As such, the US laws are less effective at protecting personal privacy than either the legislations of the European Union or Canada. The United States is not the leader in privacy protection [5, 6, 7].

1.2. E-Service model

It is useful to describe what we mean by e-services. An e-service is a service that is offered by a provider to a consumer across a computer network. This includes web services that are characterized by the use of XML and SOAP in a Service Oriented Architecture. A stock quotation service is often used as an example of an e-service. Here a consumer would logon to the service from a computer, and after appropriate user authentication, would make use of the service to obtain stock quotes. Accessing one's bank account through online banking is another example of an e-service. Here the provider is the bank and the service consists of allowing the consumer to check the balance, transfer funds, or make bill payments. The network is usually the Internet, but could also in principle be a private enterprise network. At any point in time, one provider may be serving many consumers and many providers may be serving one consumer. For the purposes of this paper, the business relationship between provider and consumer is always one-to-one, i.e. the service is designed for one consumer and is provided by one provider, payment for the service is expected from one consumer. In addition, service providers may also be service consumers, and service consumers may also be service providers.

1.3. Privacy management model

An effective and flexible way to protect privacy is to manage it using privacy policies. A provider has a privacy policy stating what private information it requires from a consumer and how the information will be used. A consumer has a privacy policy stating what private information the consumer is willing to share, with whom it may be shared, and under what circumstances it may be shared. An entity that is both a provider and a consumer has separate privacy policies for these two roles. A privacy policy is attached to a software agent that acts for a consumer or a provider as the case may be. Prior to the activation of a particular service, the agent for the consumer and the agent for the provider undergo a privacy policy exchange, in which the policies are examined for compatibility. Each agent examines the

other's policy to determine if there is a match between the two policies. If each agent finds a match, the agents signal each other that a match has been found, and service is initiated. If either agent fails to find a match, that agent would signal a mismatch to the other agent and service would then not be initiated. In this case, the consumer (provider) is free to exchange policies with another provider (consumer).

In our model, the provider requires private information from the consumer for use in its e-service and so reduces the consumer's privacy by requesting such information. This reduction in consumer privacy is represented by the requirements for consumer private information in the provider's privacy policy. The consumer, on the other hand, would rather keep her private information to herself, and so tries to resist the provider's attempt to reduce her privacy. This means that the consumer would only be willing to have her privacy reduced by a certain amount, as represented by the privacy provisions in her privacy policy. There is a *match* between a provider's privacy policy and the corresponding consumer's policy where the amount of privacy reduction allowed by the consumer's policy is at least as great as the amount of privacy reduction required by the provider's policy. Otherwise, there is a *mismatch*. Where time is involved, a private item held for less time is considered less private. A privacy policy is considered *upgraded* if the new version represents more privacy than the prior version. Similarly, a privacy policy is considered *downgraded* if the new version represents less privacy than the prior version. Figure 1 illustrates our privacy management model. For the purposes of this paper, it is not necessary to consider the details of service operation. However, the provider is expected to comply with the consumer's privacy policy if service is initiated.




Figure 1. Exchange of privacy policies (PP) between consumer agent (CA) and provider agent (PA)

Section 2 examines the specification of privacy policies by identifying some attributes of private information collection, using the CSAPP as a guide. Section 3 looks at how privacy policies can produce unexpected negative outcomes and presents some examples of such outcomes. Section 4 proposes an

approach to ensure that privacy policies are “well-formed” thereby avoiding negative outcomes. Section 5 describes related work. Section 6 gives conclusions and future research.

2. The Specification of Privacy Policies

2.1. Requirements from privacy principles

In this section, we identify some attributes of private information collection using the CSAPP as a guide. We use the CSAPP because it is representative of principles behind privacy legislation in many countries, including the European Union. We will then apply these attributes to the specification of privacy policy contents. Table 1 shows the CSAPP.

Table 1. CSAPP - The Ten Privacy Principles from the Canadian Standards Association [3]

<i>Principle</i>	<i>Description</i>
1. Accountability	An organization is responsible for personal information under its control and shall designate an individual or individuals accountable for the organization's compliance with the privacy principles.
2. Identifying Purposes	The purposes for which personal information is collected shall be identified by the organization at or before the time the information is collected.
3. Consent	The knowledge and consent of the individual are required for the collection, use or disclosure of personal information, except when inappropriate.
4. Limiting Collection	The collection of personal information shall be limited to that which is necessary for the purposes identified by the organization. Information shall be collected by fair and lawful means.
5. Limiting Use, Disclosure, and Retention	Personal information shall not be used or disclosed for purposes other than those for which it was collected, except with the consent of the individual or as required by the law. In addition, personal information shall be retained only as long as necessary for fulfillment of those purposes.
6. Accuracy	Personal information shall be as accurate, complete, and up-

	to-date as is necessary for the purposes for which it is to be used.
7. Safeguards	Security safeguards appropriate to the sensitivity of the information shall be used to protect personal information.
8. Openness	An organization shall make readily available to individuals specific information about its policies and practices relating to the management of personal information.
9. Individual Access	Upon request, an individual shall be informed of the existence, use and disclosure of his or her personal information and shall be given access to that information. An individual shall be able to challenge the accuracy and completeness of the information and have it amended as appropriate.
10. Challenging Compliance	An individual shall be able to address a challenge concerning compliance with the above principles to the designated individual or individuals accountable for the organization's compliance.

We interpret “organization” as “provider” and “individual” as “consumer”. In the following, we use CSAPP.n to denote Principle n of the CSAPP. Principle CSAPP.2 implies that there could be different providers requesting the information, thus implying a *collector* attribute. Principle CSAPP.4 implies that there is a *what* attribute, i.e. what private information is being collected. Principles CSAPP.2, CSAPP.4, and CSAPP.5 state that there are *purposes* for which the private information is being collected. Principles CSAPP.3, CSAPP.5 and CSAPP.9 imply that the private information can be disclosed to other parties, giving a *disclose-to* attribute. Principle CSAPP.5 implies a *retention time* attribute for the retention of private information. Thus, from the CSAPP we derive 5 attributes of private information collection, namely *collector*, *what*, *purposes*, *disclose-to*, and *retention time*.

The Privacy Principles also prescribe certain operational requirements that must be satisfied between provider and consumer, such as identifying purpose and consent. Our service model and the exchange of privacy policies automatically satisfy some of these requirements, namely Principles CSAPP.2, CSAPP.3, and CSAPP.8. The satisfaction of the remaining operational requirements depends on compliance mechanisms (Principles CSAPP.1,

CSAPP.4, CSAPP.5, CSAPP.6, CSAPP.9, CSAPP.10) and security mechanisms (Principle CSAPP.7). Security and compliance mechanisms are outside the scope of this paper.

2.2. Privacy policy specification

Based on the above examination of CSAPP, the contents of a privacy policy should, for each item of private data, identify a) *collector* - who wishes to collect the information, either an organization or a specific person, b) *what* - the nature of the information, c) *purposes* - the purposes for which the information is being collected, d) *retention time* - the amount of time for the provider to keep the information, and e) *disclose-to* - the parties to whom the provider will disclose the information. A privacy policy can be considered as a machine-readable document that lists each item of private information with corresponding description of *collector*, *what*, *purposes*, *retention time*, and *disclose-to*.

The attribute grouping <collector, what, purposes, retention time, disclose-to> is called a *privacy rule*. A privacy policy then consists of a header section followed by one or more privacy rules. This header consists of the fields: *Policy Use* (for what e-service?), *Owner* (name of the provider or consumer who owns the policy), *Proxy* (Yes or no - yes if a proxy will act for the consumer to give the information), and *Valid* (period of time during which the policy is valid). Figure 2(a) shows an example provider privacy policy for e-learning; Figure 2(b) shows the corresponding consumer privacy policy.

The above method of specification results in a minimal policy in the sense that the policy is the minimum required to satisfy privacy legislation. Additional provisions can be added but are not necessary for the purposes of this work.

As an aside, with reference to the Privacy Management Model of Section 1.3, the policies in Figure 2 match, since 2 years retention time offered by the consumer for course marks is less private than the 1 year required by the provider, i.e. the consumer allows a privacy reduction that is greater than the reduction required by the provider. Privacy policies need to be expressed in a machine-readable policy language such as APPEL [8] (XML-based). The investigation of suitable policy languages for privacy policies is ongoing research and outside the scope of this paper.

3. Unexpected outcomes

We are interested in unexpected outcomes that result from the matching of consumer and provider

policies. Unexpected outcomes result from a) how the matching policy was obtained, and b) the content of the matching policy itself. We examine each of these sources in turn.

3.1. Outcomes from how the matching policy was obtained

The matching policy can be obtained through policy upgrades or downgrades. These policy changes can occur while the policy was being formulated for the first time or after a mis-match had occurred in an attempt to obtain a match (e.g. during policy negotiation [11, 12]). Recall from Section 1.3, that an upgraded policy reflects a higher level of privacy. On the other hand, a downgraded policy reflects a lower level of privacy.

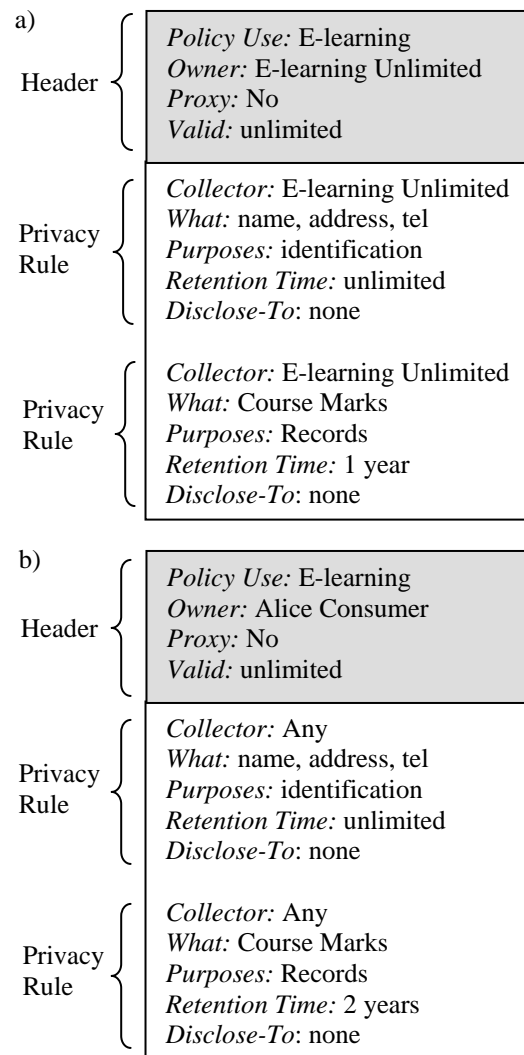


Figure 2. Example provider (a) and corresponding consumer (b) privacy policies

Policy upgrades

Given our privacy management model that the provider reduces the consumer's privacy, it is possible that the policy non-match was due to the provider's policy requiring too much privacy reduction. Suppose then that the match occurred after the provider upgraded its privacy policy to represent more privacy, i.e. require less privacy reduction. This could mean that the provider is requiring less information that is private. In this case, the provider or consumer may not realize the extra costs that may result from not having access to the private information item or items that were eliminated through upgrading. For example, leaving out the social insurance number may lead to more costly means of consumer identification for the provider. As another example, consider the provider and consumer policies of Figure 3. In this figure, suppose All Books Online upgraded its privacy policy by eliminating the credit card requirement. This would lead to a match with Alice's privacy policy, but may cost Alice longer waiting time to get her order, as she may be forced into an alternate slower means of making payment (e.g. mail a cheque), if payment is required prior to shipping.

<i>Policy Use:</i> Book Seller <i>Owner:</i> All Books Online <i>Proxy:</i> No <i>Valid:</i> unlimited	<i>Policy Use:</i> Book Seller <i>Owner:</i> Alice Consumer <i>Proxy:</i> No <i>Valid:</i> December 2004
<i>Collector:</i> All Books Onl. <i>What:</i> name, address, tel <i>Purposes:</i> identification <i>Retention Time:</i> unlimited <i>Disclose-To:</i> none <i>Collector:</i> All Books Onl. <i>What:</i> credit card <i>Purposes:</i> payment <i>Retention Time:</i> until payment complete <i>Disclose-To:</i> none	<i>Collector:</i> any <i>What:</i> name, address, tel <i>Purposes:</i> identification <i>Retention Time:</i> unlimited <i>Disclose-To:</i> none

Figure 3. Example online book seller provider (left) and consumer (right) privacy policies

Policy downgrades

Since the consumer resists the provider's privacy reduction, it is possible that the policy non-match was due to the consumer's policy allowing too little privacy reduction. Suppose then that the match occurred after the consumer downgraded her privacy policy to represent less privacy, i.e. allow for more privacy reduction. This could mean that the consumer is willing to provide more information that is private.

Then the provider or consumer may not realize the extra costs that may result from having to safeguard the additional private information item or items that were added through downgrading. For example, the additional information might be a critical health condition that the consumer does not want anyone else to know, especially her employer, which could result in loss of her employment. The provider had better add sufficient (costly) safeguards to make sure that the condition is kept confidential. The provider may not have fully realized the sensitivity of the extra information.

3.2. Unexpected outcomes from the content of the matching policy

We give here some example unexpected outcomes due to the content of the matching policy. We examine the content of the header and privacy rules in turn, as follows.

Proxy

If the consumer uses a proxy to provide her private information, the consumer needs to make sure that the proxy is trustworthy. Otherwise, the consumer's private information may be divulged against her wishes.

Valid

If the *valid* field of the consumer's policy is not carefully specified, the provider may become confused upon expiry if there is not another consumer policy that becomes the new policy. In this state of confusion the provider could inadvertently disclose the consumer's private information to a party that the consumer does not want to receive the information.

Collector

Specification of who is to collect the consumer's private information needs to consider what happens if the collector is unavailable to receive the information. For example, consider the privacy policies of Figure 4. The policies are not compatible, since Alice will reveal her medical condition only to Dr. Smith whereas the provider would like any doctor or nurse on staff to take the information. Suppose the provider upgrades its policy to satisfy Alice by allowing only Dr. Smith to receive information on Alice's condition. Then an unexpected outcome is that Alice cannot receive help from Nursing Online because Dr. Smith is not available (he might have been seriously injured in an accident), even though the policies would match and the service could theoretically proceed. There are various ways to solve this particular situation (one

way is simply to have an overriding condition that in an emergency, Alice *must* give her condition to any doctor or nurse on staff) but our point still holds. An improperly specified collector attribute can lead to unexpected serious consequences.

<i>Policy Use:</i> Medical Help <i>Owner:</i> Nursing Online <i>Proxy:</i> Yes <i>Valid:</i> unlimited	<i>Policy Use:</i> Medical Help <i>Owner:</i> Alice Consumer <i>Proxy:</i> No <i>Valid:</i> December 2004
<i>Collector:</i> Nursing Online <i>What:</i> name, address, tel <i>Purposes:</i> contact <i>Retention Time:</i> unlimited <i>Disclose-To:</i> pharmacy	<i>Collector:</i> any <i>What:</i> name, address, tel <i>Purposes:</i> contact <i>Retention Time:</i> unlimited <i>Disclose-To:</i> pharmacy
<i>Collector:</i> Nursing Online <i>What:</i> medical condition <i>Purposes:</i> treatment <i>Retention Time:</i> 1 year <i>Disclose-To:</i> pharmacy	<i>Collector:</i> Dr. A. Smith <i>What:</i> medical condition <i>Purposes:</i> treatment <i>Retention Time:</i> unlimited <i>Disclose-To:</i> pharmacy

Figure 4. Example medical help provider (left) and consumer (right) privacy policies

Retention time

Care must also be taken to specify the appropriate retention time for a particular information item. The responsibility for setting an appropriate retention time lies with both the provider and the consumer. For example, consider once again the policies of Figure 4. Suppose Alice changes her privacy rule for medical condition from “Dr. A. Smith” to “any” and from “unlimited” to “2 years”. Then the policies match and the service can proceed. At the end of 2 years, the provider complies with the consumer’s privacy policy and discards the information it has on Alice’s medical condition. But suppose after the 2 years, medical research discovers that Alice’s condition is terminal unless treated with a certain new drug. Then Nursing Online cannot contact Alice to warn her since it no longer knows that Alice has that condition. Poor Alice! Clearly, both the provider and the consumer are responsible for setting the appropriate retention time. One could conclude that in the case of a medical condition, the retention time should be “unlimited”. However, “unlimited” can also have its risks, such as retaining information beyond the point of when it no longer applies. For example, Alice could one day be cured of her condition. Then, retention of Alice’s condition could unjustly penalize Alice if it somehow leaked out when it is no longer true.

Disclose-To

If the *disclose-to* attribute of the consumer’s privacy policy is not specified or improperly specified, providers can share the consumer’s private information with other providers or consumers with resulting loss of privacy. Consider the following examples.

Suppose Alice has a critical health condition and she does not want her employer to know for fear of losing her job (the employer might dismiss her to save on sick leave or other benefits – this really happened! [9]). Suppose she is able to subscribe to Nursing Online as in the above examples. Then through the execution of the service, Nursing Online shares her condition with a pharmacy in order to fill her prescription. Suppose the company that Alice works for is a pharmaceutical supplier, and needs to know contact information of patients in the area where Alice lives in order to directly advertise to them about new drugs effective for Alice’s condition. Suppose the pharmacy with which Nursing Online shared Alice’s condition is a consumer of the pharmaceutical supplier and the pharmacy’s privacy policy does not restrict the sharing of patient information that it receives second hand. Then the pharmaceutical supplier, Alice’s employer, can learn of her health condition from the pharmacy and Alice could lose her job – an unexpected outcome with serious consequences. A possible solution to this situation is for Alice to specify “pharmacy, no further” for *disclose-to*. Then to comply with Alice’s policy, Nursing Online, as a consumer of the pharmacy, in its privacy policy with the pharmacy would specify “none” for the *disclose-to* corresponding to Alice’s condition, thus preventing Alice’s employer from learning of her condition and preserving her privacy.

As another example, suppose Alice, as a consumer, uses graphics services from company A and company B. Her privacy policy with these companies stipulate that the rates she pays them is private and not to be disclosed to any other party. Suppose she pays company A a higher rate than company B. Now suppose companies A and B are both consumers of company C, which provides data on rates paid for graphics services. In order to use company C’s services, companies A and B must provide company C with de-identified information regarding rates they are paid. This does not violate the privacy policies of consumers of companies A and B because the information is de-identified. However, company B now learns of the higher rate paid company A and seeks a higher rate from Alice. There does not appear to be any solution to this situation, since Alice has already specified *disclose-to* as “none”. This example shows that there can be unexpected outcomes that may not be preventable.

We have presented a number of unexpected outcomes arising from how the policy match was obtained and how the content of the policy was specified. Our outcomes are all negative ones because they are the ones we need to be concerned about. There are also, of course, positive unexpected outcomes, which are outside the scope of this work.

4. Preventing unexpected negative outcomes

The problem at hand is how to detect and prevent the unexpected outcomes that are negative or dangerous. Since all unexpected outcomes derive from the privacy policy (at least in this work), it is necessary to ensure “well-formed” policies that can avoid unexpected negative outcomes. Further, if a non-well-formed policy matches the first time and leads to negative outcomes, it is too late to do anything about it. Based on the discussion of Section 3.2, we define a “well-formed” privacy policy after defining “unexpected negative outcome”:

Definition 1

An *unexpected negative outcome* is an outcome of the use of privacy policies per the Privacy Management Model (Section 1.3) such that a) the outcome is unexpected by both the provider and the consumer, and b) the outcome leads to either the provider or the consumer or both experiencing some loss, which could be private information, money, time, convenience, job, and so on, even losses that are safety and health-related.

Definition 2

A *well-formed (WF)* privacy policy (for either consumer or provider) is one that does not lead to unexpected negative outcomes. A *near well-formed (NWF)* privacy policy is one in which the attributes *valid*, *collector*, *retention time*, and *disclose-to* have each been considered against all *known* mis-specifications that can lead to unexpected negative outcomes.

In Definition 2, the mis-specifications can be accumulated as a result of experience (e.g. trial and error) or by scenario exploration (as above). We have already presented a number of them in Section 3.2. A NWF privacy policy is the best that we can achieve at this time. Clearly, such a policy does not guarantee that unexpected negative outcomes will not occur – it just reduces the probability that an unexpected negative outcome will occur. Also, we do not include “proxy” in the above definition. We assume that if a

proxy is used, that proxy is reliable and will represent the interest of the consumer to the best of her ability.

4.1. Rules for specifying near well-formed privacy policies

Let us consider once more the content of a privacy policy by looking at the header and the privacy rules.

The header (Figure 2) consists of *Policy Use*, *Owner*, *Proxy*, and *Valid*. *Policy Use* and *Owner* serve only to identify the policy and assuming they are accurately specified, they are unlikely to lead to unexpected negative outcomes. That leaves *Proxy* and *Valid*, and we have already disposed of *Proxy* in the previous section. As discussed in Section 3.2, *Valid* must be specified so that it is never the case that the provider is in possession of the consumer’s private information without a corresponding valid consumer policy (i.e. with the policy expired). Another way to look at this is that it must be true that the provider is no longer in possession of the consumer’s information at the point of policy expiration. Hence a rule for specifying *Valid* is the following:

Rule for specifying *Valid*

The time period specified for Valid must be at least as long as the longest retention time in the privacy policy.

This rule ensures that if the provider is in possession of the consumer’s private information, there is always a corresponding consumer privacy policy that governs the information, which is what is needed to avoid the unexpected outcomes from an improperly specified *Valid*.

Let us now consider the content of a privacy rule. The privacy rule consists of the attributes *Collector*, *What*, *Purposes*, *Retention Time*, and *Disclose-To* (Figure 2). *What* and *Purposes* serve only to identify the information and the purposes for which the information will be put to use. Assuming they are accurately specified, they are unlikely to lead to unexpected negative outcomes. That leaves *Collector*, *Retention-Time*, and *Disclose-To*, which we discussed in Section 3.2. We now formulate specification rules for them, based on the discussion of Section 3.2.

Rule for specifying *Collector*

When specifying an individual for Collector, the consequences of the unavailability of the individual to receive the information must be considered. If the consequences do not lead to unexpected negative outcomes (as far as can be determined), proceed to specify the individual. Otherwise, or if there is doubt,

specify the name of the provider (meaning anyone in the provider's organization).

Rule for specifying Retention Time

When specifying Retention Time, the consequences of the expiration of the retention time (provider destroys corresponding information) must be considered. If the consequences do not lead to unexpected negative outcomes (as far as can be determined), proceed to specify the desired time. Otherwise, or if there is doubt, specify the length of time the service will be used.

Rule for specifying Disclose-To

When specifying Disclose-To, the consequences of successive propagation of your information starting with the first party mentioned in the Disclose-To must be considered. If the consequences do not lead to unexpected negative outcomes (as far as can be determined), proceed with the specification of the Disclose-To party or parties. Otherwise, or if there is doubt, specify "none" or "name of receiving party, no further".

These rules address the problems discussed in Section 3.2 that lead to unexpected negative outcomes. Except for *Valid*, in each case we require the consumer or provider to consider the consequences of their intended specification, and propose specification alternatives, where the consequences lead to unexpected negative outcomes or there is doubt. By definition, application of these rules to the specification of a privacy policy will result in a near well-formed policy. Undoubtedly, mathematical modeling of the processes at play together with state exploration tools can help to determine whether or not a particular specification will lead to unexpected negative outcomes. Such modeling and use of tools is part of future research.

4.2. Approach for obtaining near well-formed privacy policies

We propose an approach to obtain NWF privacy policies, applied during the initial privacy policy specification process and during the subsequent privacy policy negotiation process when there is a mismatch.

Initial specification

We propose that the above rules for obtaining near well-formed policies be incorporated during initial policy specification. This is best achieved using an automatic or semi-automatic method for specifying privacy policies, such as the approach for consumer

policies given in [10]. The Rule for *Valid* is easy to implement. Implementation of the remaining rules may employ a combination of artificial intelligence and human-computer interface techniques to assist the human specifier to reason out the consequences. Alternatively, the rules may be applied during manual policy specification in conjunction with a tool for determining possible consequences of a particular specification, as noted above.

Policy mismatch

Near well-formed policies may still not match. At a policies mismatch, the consumer or the provider upgrades or downgrades her/its individual policy to try to get a match. In so doing, each could inadvertently introduce new values into the policy or remove values from the policy that result in unexpected negative outcomes or loss of NWF-ness. We propose the use of privacy policy negotiation [11, 12, 17] between consumer and provider agents to guide the policy upgrading or downgrading to avoid undoing the values already put in place for NWF-ness in the initial specification. Alternatively, negotiation may expose a needed application of the above rules. This is also a consequences exploration, but here both provider and consumer do the exploration while negotiating in real-time. For example, in the All Books Online example of Section 3.1 where Alice does not need to provide her credit card, negotiation between Alice and All Books Online could have identified the consequence that Alice would need to wait longer for her order and direct her to another more viable alternative, such as agreeing to provide her credit card. Similarly, in the example of Section 3.1 where the provider has to introduce more costly safeguards to protect the consumer's added highly sensitive information, negotiation could have uncovered the high sensitivity of the new information and possibly result in a different less costly alternative chosen (e.g. the new information may not be needed after all).

Table 2 illustrates how negotiation can detect and prevent the unexpected negative outcome of Alice having no access to medical service when it is needed (read from left to right and down). The result of this negotiation is that Nursing Online will be able to provide Alice with nursing service whenever Alice requires it, once she makes the change in her privacy policy reflecting the results of negotiation. If this negotiation had failed (Alice did not agree), Alice will at least be alerted to the possibility of a bad outcome, and may take other measures to avoid it. This example shows how negotiation may persuade the consumer to resolve a mismatch by applying the above rule for specifying "collector".

Table 2. Preventing Unexpected Negative Outcomes – Nursing Online

Nursing Online (Provider)	Alice (Consumer)
<i>OK if a nurse on our staff be told your medical condition?</i>	<i>No, only Dr. Alexander Smith can be told my medical condition.</i>
<i>We cannot provide you with any nursing service unless we know your medical condition.</i>	<i>OK, I'll see Dr. Smith instead.</i>
<i>You are putting yourself at risk. What if you need emergency medical help for your condition and Dr. Smith is not available?</i>	<i>You are right. Do you have any doctors on staff?</i>
<i>Yes, we always have doctors on call. OK to allow them to know your medical condition?</i>	<i>That is acceptable. I will modify my privacy policy to share my medical condition with your doctors on call.</i>

Table 3 gives another example of negotiation at work using Alice's Book Seller policy in Figure 3. This policy mismatched because Alice did not want to provide her credit card information. At the end of the negotiation, Alice modifies her privacy policy and receives service from All Books Online.

Table 3. Preventing Unexpected Negative Outcomes – All Books Online

All Books Online (Provider)	Alice (Consumer)
<i>OK if you provide your credit card information?</i>	<i>No, I do not want to risk my credit card number getting stolen.</i>
<i>If you do not provide your credit card information, you will need to send us a certified cheque before we can ship your order. This will delay your order for up to 3 weeks.</i>	<i>I still don't want to risk my credit card number getting stolen.</i>
<i>Your credit card information will be encrypted during transmission and we keep your information in</i>	<i>OK, I will modify my privacy policy to share my credit card information.</i>

<i>secure storage once we receive it. You need not worry.</i>	
---	--

We have assumed that either the consumer or the provider will want to inform the other about changes to a policy that could lead to unexpected negative outcomes. We believe this is a reasonable assumption given that it is in their mutual interest to avoid the negative outcomes.

5. Related Work

Negative outcomes arising from privacy policies may be regarded as a feature interaction problem, where policies "interact" and produce unexpected outcomes [13]. Traditionally, feature interactions have been considered mainly in the telephony or communication services domains [14]. More recent papers, however, have focused on other domains such as the Internet, multimedia systems, mobile systems [15], and Internet personal appliances [16]. In this work, we have chosen not to frame negative outcomes from privacy policies as a feature interaction problem. In so doing, we have obtained new insights and results.

Apart from feature interactions, other possible related work has to do with resolving conflicts in access control and mobile computing (e.g. [18, 19]). However, it is believed that these methods and similar methods in other domains will not work for privacy due to the subjective nature of privacy, i.e. personal involvement to consider each privacy rule is necessary. Looking at these other methods for possible application to privacy is another topic for future research.

6. Conclusions and future research

The Privacy Principles impose legislative conditions to protect the rights of individuals (consumers) to privacy. They imply that the collection of private information may be done under the headings of *collector*, *what*, *purposes*, *retention time*, and *disclose-to*. Privacy policies may be constructed using these headings to specify each private information item to be shared. In an online community, consumers and providers of electronic services specify their privacy preferences using privacy policies. The provider specifies the private information items it requires from the consumer. The consumer specifies the items she is willing to share with the provider. Software agents for consumers and providers

exchange and compare these policies to see if they match. The service is initiated only if the policies match. The content of the matching privacy policies can lead to unexpected negative outcomes. Examples of unexpected negative outcomes arising from policy upgrades or downgrades, and from malformed policy content were given. An approach was proposed to avoid unexpected negative outcomes, involving (1) the application of rules to construct near well-formed policies at initial policy specification time, and (2) the negotiation of mismatched policies between consumers and providers to agree on private information items and to construct near well-formed policies through consequences exploration and the application of the rules of Section 4.1.

We have based our work on our particular formulation of a privacy policy. An obvious question is whether our approaches apply to other formulations of privacy policies. We believe the answer is yes, for the following reasons: a) privacy policy formulations (i.e. contents) cannot differ too much from one another since they must all conform to privacy legislation and our policy is a minimal policy that so conforms, and b) if necessary, we can fit our approaches to any formulation by applying the same logic we used in this work.

As future research, we plan to look at answering the following questions:

- What are other unexpected negative outcomes?
- What tools can be designed to help consumers and providers do consequences exploration and identify the seriousness of each consequence?
- What other methods may work to avoid or mitigate unexpected negative outcomes from privacy policies?

We also plan to implement the approach to ensure near well-formed-ness at initial policy specification time in a prototype for semi-automatically deriving privacy policies.

Acknowledgements

The authors acknowledge the support of the National Research Council Canada for this work. The authors also extend their thanks to the anonymous reviewers for their valuable and constructive comments.

References

- [1] W3C, "The Platform for Privacy Preferences", <http://www.w3.org/P3P/>
- [2] Government of Canada, "Personal Information Protection and Electronic Documents Act", available as of February 28, 2005 at: http://www.privcom.gc.ca/legislation/index_e.asp
- [3] Canadian Standards Association, "Model Code for the Protection of Personal Information", retrieved Sept. 5, 2003 from: <http://www.csa.ca/standards/privacy/code/Default.asp?articleID=5286&language=English>
- [4] European Union, "Directive 95/46/EC of the European Parliament and of the 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", unofficial text retrieved Sept. 5, 2003 from: <http://aspe.hhs.gov/datacncl/eudirect.htm>
- [5] D. Banisar, "Privacy and Data Protection Around the World", 21st International Conference on Privacy and Personal Data Protection, September 13, 1999.
- [6] D. Hurley, "A Whole World in One Glance: Privacy as a Key Enabler of Individual Participation in Democratic Governance", 21st International Conference on Privacy and Personal Data Protection, September 13, 1999.
- [7] S.J. Milberg, S.J. Burke, H.J. Smith, and E.A. Kallman, "Values, Personal Information, Privacy, and Regulatory Approaches", *Communications of the ACM*, Vol. 38, No. 12, December 1995.
- [8] W3C, "A P3P Preference Exchange Language 1.0 (APPEL1.0)", W3C Working Draft 15 April 2002, <http://www.w3.org/TR/P3P-preferences/>
- [9] KumeKawa, J.K., "Health Information Privacy Protection: Crisis or Common Sense?", retrieved Sept. 7, 2003 from: http://www.nursingworld.org/ojin/topic16/tpc16_2.htm
- [10] G. Yee and L. Korba, "Semi-Automated Derivation of Personal Privacy Policies", Proceedings, The IRMA International Conference 2004 (IRMA 2004), New Orleans, May 23-26, 2004.
- [11] G. Yee, L. Korba, "Bilateral E-services Negotiation Under Uncertainty", accepted for publication, The 2003 International Symposium on Applications and the Internet (SAINT2003), Orlando, Florida, Jan. 27-31, 2003.
- [12] G. Yee and L. Korba, "The Negotiation of Privacy Policies in Distance Education", accepted for publication, 14th IRMA International Conference, Philadelphia, Pennsylvania, May 18-21, 2003.
- [13] G. Yee and L. Korba, "Feature Interactions in Policy Driven Privacy Management", Proceedings, Seventh International Workshop on Feature Interactions in Telecommunications and Software Systems, Ottawa, Canada, June 11-13, 2003.
- [14] D. Keck and P. Kuehn, "The Feature and Service Interaction Problem in Telecommunications Systems: A Survey", *IEEE Transactions on Software Engineering*, Vol. 24, No. 10, October 1998.
- [15] L. Blair, J. Pang, "Feature Interactions – Life Beyond Traditional Telephony", Distributed Multimedia Research Group, Computing Dept., Lancaster University, UK.

- [16] M. Kolberg et al, "Feature Interactions in Services for Internet Personal Appliances", University of Stirling, UK, Telcordia Technologies, USA.
- [17] L. Korba, "Privacy in Distributed Electronic Commerce", Proc. of the 35th Hawaii International Conference on System Science (HICSS), Hawaii, January 7-11, 2002.
- [18] T. Jaeger, R. Sailer, X. Zhang, "Resolving Constraint Conflicts", Proceedings of the ninth ACM symposium on Access control models and technologies, June 2004.
- [19] L. Capra, W. Emmerich, C. Mascolo, "A micro-economic approach to conflict resolution in mobile computing", Proceedings of the 10th ACM SIGSOFT symposium on Foundations of software engineering, November 2002.

¹ NRC Paper Number: NRC 48250