NRC Publications Archive
Archives des publications du CNRC

Determining Internet Users' Values for Private Information
Buffett, Scott; Fleming, Michael; Richter, Michael; Scott, Nathan; Spencer, Bruce

This publication could be one of several versions: author’s original, accepted manuscript or the publisher’s version. / La version de cette publication peut être l’une des suivantes : la version prépublication de l’auteur, la version acceptée du manuscrit ou la version de l’éditeur.

Publisher’s version / Version de l'éditeur:

The Second Annual Conference on Privacy, Security and Trust (PST'04) [Proceedings], 2004

NRC Publications Archive Record / Notice des Archives des publications du CNRC :
https://nrc-publications.canada.ca/eng/view/object/?id=0dd7da53-f68f-4346-9d90-04e3f1a0a8c4
https://publications-cnrc.canada.ca/fra/voir/objet/?id=0dd7da53-f68f-4346-9d90-04e3f1a0a8c4

Access and use of this website and the material on it are subject to the Terms and Conditions set forth at
https://nrc-publications.canada.ca/eng/copyright
READ THESE TERMS AND CONDITIONS CAREFULLY BEFORE USING THIS WEBSITE.

L’accès à ce site Web et l’utilisation de son contenu sont assujettis aux conditions présentées dans le site
https://publications-cnrc.canada.ca/fra/droits
LISEZ CES CONDITIONS ATTENTIVEMENT AVANT D’UTILISER CE SITE WEB.

Questions? Contact the NRC Publications Archive team at
PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca. If you wish to email the authors directly, please see the first page of the publication for their contact information.

Vous avez des questions? Nous pouvons vous aider. Pour communiquer directement avec un auteur, consultez la première page de la revue dans laquelle son article a été publié afin de trouver ses coordonnées. Si vous n’arrivez pas à les repérer, communiquez avec nous à PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca.

National Research Council Canada
Institute for Information Technology

Conseil national de recherches Canada
Institut de technologie de l'information

NRC · CNRC

*Determining Internet Users' Values for Private Information**

Buffett, S., Fleming, M.W., Richter, M.M., Scott, N., and Spencer, B.
October 2004

* published in The Second Annual Conference on Privacy, Security and Trust (PST'04). Fredericton, New Brunswick, Canada. October 14-15, 2004. pp. 79-88. NRC 47408.

Copyright 2004 by
National Research Council of Canada

Permission is granted to quote short excerpts and to reproduce figures and tables from this report, provided that the source of such material is fully acknowledged.

Canada

Determining Internet Users' Values for Private Information

Scott Buffett, Nathan Scott, Bruce Spencer
National Research Council Canada
Institute for Information Technology - E-business
Fredericton, New Brunswick,
Canada, E3B 9W4
{scott.buffett, nathan.scott,
bruce.spencer}@nrc.gc.ca

Michael M. Richter
University of Kaiserslautern,
Kaiserslautern, Germany
richter@informatik.uni-kl.de

Michael W. Fleming
University of New Brunswick
Fredericton, New Brunswick,
Canada, E3B 5A3
mwf@unb.ca

***Abstract*— We examine the problem of determining a user's value for his/her private information. Web businesses often offer rewards, such as discounts, free downloads and website personalization in exchange for information about the user, such as name, phone number and e-mail address. We present a technique that helps the user determine whether such an offer is acceptable by computing its value in terms of the consequences that could occur as a result of such an information exchange. Bayesian networks are used to model dependencies in the user's utilities for such consequences, and utility elicitation is used to reduce the uncertainty of these utilities. We also derive a "bother cost", which is used by the elicitation engine to determine the optimal time to stop the question process. A simple example experiment demonstrates the effectiveness of the technique by significantly improving the user's expected utility in a simple privacy negotiation.**

I. Introduction

A growing concern in today's Internet is the privacy and protection of one's personal information. Obtaining personal data from visitors and customers is of huge importance among e-commerce websites. Such information helps businesses to better serve their customers by updating their websites to meet changing demands and demographics, to attract new business by learning how to effectively target advertisements, and also to run more efficiently simply by retaining pertinent information about their customers in their own database. In addition, some businesses may transmit this data to third parties, perhaps for financial gain or simply as necessary information sharing with integral partners that need to make use of the data. Regardless of use, personal information of users has a great value on the Web today.

Unfortunately for these websites, users do not generally freely give away their private information. One issue is users' trust of the security of information transferred. Security has always been an issue in communications, since there is a desire to keep certain information from an interceptor that may act to the disadvantage of the sender. For centuries, even millennia, this was the domain of cryptography. The reason was obvious: The information had value for both the sender and the interceptor. The exact value for such information was not discussed; it was usually clear, particularly in military situations. In the context of transmitting private information over the Internet, users typically value keeping their information private because of potentially annoying or damaging repercussions, ranging from misuse of e-mail addresses resulting in spam, to misuse of credit card information or other data that could result in identity theft.

While most businesses do all they can to protect customers from such acts, and state so in their privacy policies, there are still a few problems. Number one, customers often do not read privacy policies, believing them to be too time consuming and not completely understandable. Second, users often do not trust businesses to follow their policies, particularly the less familiar websites. Third, businesses often indicate that they may share certain information with their partners but do not post their partners' policies, leaving the users at their mercy. Finally, people often fear the unknown consequences that could result, even if policies are followed properly and information is not misused. For example, consider clothing stores A and B, who each purchase from distributor D. An employee of A secretly buys clothes online from B. B's privacy policy indicates that certain information is shared with D (who in turn may share it with A), but all identifiable information such as name and address is kept secret. The only information that D requires is age, occupation and city (which B perceives to be unidentifiable information), to identify target demographics for various items. The CEO of A could then receive this information, and may be able to conclude that this particular employee shops with the competitor, perhaps because the CEO knows that this employee is the only one that fits the description. This could thus hinder the employee's advancement in the company, or possibly even lead to termination.

Due to these inhibitions on the part of the users, potential e-commerce transactions are lost. To combat this, websites often make offers to customers in exchange for their private information in the form of discounts, website personalization, free memberships or software downloads. The idea is that if something has value, then you can buy or sell. From an economic point of view a central question is to determine the market value of the product to be sold. For this, we assume a simple scenario that contains a seller of information (the user) and a buyer (usually a company). The company makes the first step by offering some compensation and advantages to the user if some information is given to them. For such a transaction to take place, the two sides must agree on a price. Here we face the additional problem that both seller and buyer have very different views on the value of the information under consideration and the usual mechanisms of supply and demand do not work. In particular, the seller needs to get advice on what the buyer can do with the obtained information and on what an appropriate compensation might be.

In a sales situation where there are no fixed prices, negotiation can take place. In mathematical economy, such questions are associated with the concept of utility [7], [9]. A mathematical approach for negotiating exchanges of private information for compensation is given by Buffett et al. [2]. The results are formulated in such a way that the Platform for Privacy Preferences Project (P3P) [6] can be used.

The problem is that users need to determine their personal value for private information as a function of both the market demand for such information as well as the possible consequences to the individual of giving it away. A basic assumption made by Buffett et al. is that all the involved utilities are known. In fact, otherwise a negotiation makes little sense. It is, however, a non-trivial step to get access to these utilities. To make matters worse, the utilities are highly individual and context dependent. We discuss a framework for such situations with the main target to support the user for determining the price for the information. Because we cannot expect a universally valid formula for utility, this is not only a computational problem; it rather requires a careful analysis of the whole situation.

Our approach to the problem utilizes the concept of utility elicitation, as proposed by Chajewska et al. [3]. In this framework, a prior probability distribution over the true utility for each alternative is assumed, and questions are then asked in an effort to determine the user's true utilities with more certainty. The next question to ask is chosen as the one that will provide the highest increase in the expected utility of the chosen strategy when answered. Such a question will thus provide an answer that will significantly reduce uncertainty about an aspect that is important to decision-making in the particular strategy. To account for dependencies in the data, we model the set of information statements as a Bayesian network. For example, if the user has a surprisingly low utility for the consequence of telemarketing, then it is more likely that her utility for any other consequence that involves her phone number is low. Thus the answer to any single question may reduce the uncertainty of our beliefs about the user's utility for several alternatives.

In utility elicitation there are two central problems: deciding which question to ask next and deciding when to stop asking. Users will be unlikely to use a system that asks several questions each time a new website is visited. In many cases, users' attitudes toward privacy change very little between websites, and thus perhaps very few (or zero) questions need to be asked. We examine a notion referred to as the "bother cost" of asking a question, and use this to determine whether the information gain from the next question is worth bothering the user further.

```
<STATEMENT>
   <PURPOSE>
      <telemarketing/> <admin/>
   </PURPOSE>
   <RECIPIENT><ours/></RECIPIENT>
   <RETENTION><indefinitely/></RETENTION>
   <DATA-GROUP>
      <DATA ref="#user.name.given"/>
         <CATEGORIES><physical/></CATEGORIES>
      </DATA>
   </DATA-GROUP>
</STATEMENT>
```

Fig. 1. An example P3P statement

II. Platform for Privacy Preferences Project (P3P)

There are various ways to collect personal information from the Internet: asking the user directly, accessing legitimate user agents, sharing "cookie" files between websites, or illegitimately invading private storage. We do not consider illegitimate uses here, but many users do not understand the extent of access to their private data even via legitimate methods. An important first step is to make policies on collecting data explicit, so W3C has developed the Platform for Privacy Preferences (P3P) [6]. P3P is used by websites to express their privacy practice. A computerized agent, acting on behalf of the user, can fetch and read the P3P policy file, can inform the user about the site's privacy practices and can make an automatic or semi-automatic decision on behalf of the user.

The P3P policy file is an XML file that is defined for certain regions of a website or the entire website. Each P3P file contains at least one statement, and each statement describes what data will be collected, with whom it will be shared, for how long it will be retained and for what purpose. Figure 1 shows an example P3P statement requesting an element of the "physical" category (specifically the user's given name), and indicates that this information will be used for telemarketing and website administration purposes, that there are no intended recipients other than the requestor itself and that it will be retained for an indefinite period of time.

In this paper, we refer to private information in the context of a P3P statement. Specifically, we denote a P3P statement as a tuple $\langle d, r, p, \tau \rangle$ denoting a set d of data, a set r of recipients, a set p of purposes and a real-valued retention time τ. The techniques described in this paper are however quite general, and should be applicable in any language for information exchange.

III. The Value of Information

The value of some information (as of anything traded) is not an internal property of the information, but is rather defined by the actions that can be made possible by using it. These actions have two aspects, from the buyer's view and from the seller's view, since they have a very different impact on these two parties.

First it has to be described what kind of information is the subject of the trade. This can be of a different nature, but usually the information is concerned with personal or business aspects of the seller. Next the possible actions for which the information will or might be used must be determined. Such actions are to the benefit of the buyer but may create costs (in terms of money but also of other unwanted circumstances) for the seller. In addition, the information may also fall into the hands of an unforeseen third party.

A. The Buyer's View

The buyer is interested in the information because it can improve or simplify future business. Such information helps businesses to better serve their customers by updating their websites to meet changing demands and demographics, to attract new business by learning how to effectively target advertisements, and also to run more efficiently simply by retaining pertinent information about their customers in their own database. In addition, some businesses may transmit this data to third parties, perhaps for financial gain or simply as necessary information sharing with integral partners that need to make use of the data.

B. The Seller's View

The value of the benefits obtained from the buyer are usually quite clear because they can mostly be formulated directly in financial terms. We will not be concerned with this matter. What we have to consider is how much compensation the seller should get from the buyer in return. This is up for negotiation, but in order to negotiate, the seller has to have an opinion on what could be a fair equivalent to the disadvantages she might experience by giving away such information.

The difficulty in estimating the value of the information to be sold in terms of the expected disadvantages is that this value is usually of indirect nature for the seller. In the first place, it is given in the form of costs, where these costs are mostly not specified in terms of money. The indirect nature is due to the fact that other people (not only the buyer) may use the information in order to perform actions that have a negative consequence for the person. Therefore, the possible actions that can be done are the primary objects to be investigated. The consequences of such actions have a wide variety. They range from discomfort to serious financial risks. Because the different costs determine the utility, we deal with a multi-dimensional utility. The central problem is then to determine or at least to approximate this utility.

An important aspect is that such negative consequences of actions have a probabilistic character: Certain actions need not be performed and may not have the worst consequences. Therefore, we deal with an expected utility. The probabilities involved are often subjective because statistics are rarely available. Hence we can distinguish between consequences that are intended and agreed upon and those that arise from abusing the obtained information. Typical examples for the first kind are that the address is given to other companies who send mail or email or that certain amounts of money are charged to the credit card. Unintended consequences are, for example, that third parties have access to the credit card or that the name of the seller is used in unforeseen contexts. Here the reliability of the buyer plays a large role. The problem for the seller is that such unwanted consequences are difficult to foresee and, in an actual situation, important aspects may be overlooked. For this reason a systematic investigation of actions is necessary.

C. Actions

We are interested in actions that can be performed by the buyer or some other agent because of the knowledge of information bought. For a systematic description we extend the notation for actions by defining them in terms of their preconditions and postconditions. The extended description is as follows:

1) Type of action (e.g., giving address to another institution);
2) Intended use, the purpose (e.g., shipping ordered products);
3) Restrictions (e.g., only non-commercial institutions);
4) People or company performing the action;
5) Type of possible abuse (e.g., giving information to third party, violating restrictions);
6) Preconditions in the ordinary sense, i.e., those that make the performance of the action possible.

Only information units are used as preconditions for actions. For the postconditions of the action we have to distinguish between direct postconditions and those that arise as indirect consequences of the action. The nature of the indirect consequences depends often not only on the action itself but more on the purpose of the action. In principle, we consider two actions as different if they differ in at least one of the above description elements.

First we give a list of typical information units that may be sold:

- Name
- Address
- Phone number
- Credit card number
- Bank account
- Social security number
- Employer
- Position in the company
- Financial status
- Health insurance company, number
- Age
- Sex
- Diseases

Next we give a list of typical actions, which we group according to their degree of seriousness:

Without problems (or even useful) :

- Getting ordered products
- Paying a bill

Inconvenient:

- Spam
- Unintended phone calls
- Unintended post mail
- Listed as customer somewhere

More serious:

- Information sharing with Boss and company
- Information sharing with Legal authorities (government, tax, etc)
- Information sharing with Banks
- Information sharing with Insurances (car, etc.)
- Information sharing with Family
- Misuse of credit cards

The severity (i.e., postconditions) of such unwanted consequences depends not only on the actions in the ordinary sense but more often also on the purpose of the actions. For this reason we have introduced above the extended description of actions. We associate with each action:

1) The costs that can often only be estimated and are in different dimensions, mainly:
   - Money
   - Time
2) The probability that the action is performed. In order to estimate this we have to observe that each action has in addition an actor. With this actor some agreement can be reached that the information is used for certain actions only. Unfortunately this creates additional problems:
   - How to precisely formulate for what the information is used and what is excluded from the usage?
   - How reliable is the buyer to keep such an agreement?
   - How secure is the environment that no third party can intercept the message containing the information and what are the possible consequences?

The relation between information units iu and actions a is that the latter have information units as preconditions (using prolog-style notation):

$$Pre(a) :- \ iu_1, iu_2, \ldots, iu_n$$

This is understood as follows: in order to perform a, all conditions of the list have to be known. Of course, an information unit may occur in several precondition lists.

Notation:

- For each information unit iu, denote the list of all actions that have iu as a precondition with $A(iu)$.
- For a set IU of information units, the set $A(IU)$ of actions affected by IU is the set of all $a \in A$ where IU is subset of $Pre(a)$.
- The set $ExA(IU)$ is the set of all actions that can be executed on the basis of IU. For all $a \in ExA(IU)$, $Pre(a) \subseteq IU$.
- For each IU and iu, iu is critical for all a such that $IU \subset Pre(a)$ and $IU \cup iu = Pre(A)$.

As long as an information unit is not critical for some action, it can be given away without negative consequences resulting from that action.

Buyer known to be reliable?	yes	no	
If buyer unknown:	Expected lack of security	Expect secure	No info
Misuse Expected?	Maybe (probability?)	No	No info
⋮	⋮	⋮	⋮

Fig. 2. Offers and resulting actions in the example

D. Utility and Calculating the Costs of Information Units

Determining the costs (and the utility) of the given information units for the seller is still a difficult problem. They cannot be determined and calculated in a straightforward way, and they are often not given in financial terms. This is due to the fact that the seller often does not have enough knowledge in order to predict what kind of trouble might result. For this purpose the seller has to be supported. The support could consist of a menu where the seller initially enters some information about the situation. Then there could be a response where the seller is asked for more information. Finally, some advice is given to the seller. A partial example of such a menu is given in Figure 2.

The utility (cost) of giving away a set of information units is calculated based solely on the expected utility of the actions that can be performed as a result. Non-critical information units are given away freely; if $ExA(IU) = ExA(IU \cup iu)$ then the utility of IU equals the utility of $IU \cup iu$. To compute the expected utility of a set IU of information units, we (a) determine the set $ExA(IU)$ of actions that can be executed as a result of giving IU away, (b) determine the user's utility for each action in $ExA(IU)$ (to be discussed in section IV), and (c) compute the expected utility of IU as a function of the utilities for the actions and the probability of each action occurring, as estimated by the menu responses.

E. Negotiation

On the basis of the above, the negotiation can be started. There are various ways to organize negotiation between the seller and buyer (e.g., direct communication, communication via legitimate agents, etc.). Eventually these have to be summarized in a legally admitted contract. There are two types of restrictions that are up for negotiation: 1) Positive ones: Information can only be used for a specific purpose 2) Negative ones: Information must not be used for the following purposes P, Q, ...

P3P is a standard for web sites to express their privacy policy in a machine and human-readable format. This means in particular that the web is the communication medium; there may be, however, other ways to communicate. APPEL [5] is another language developed by W3C. It allows the user to set up privacy preference rules in a machine and human-readable format. The intention of P3P and APPEL is to help the user to protect her privacy when the user browses the Internet. This is done in such a way that a policy can be formulated. APPEL allows an information owner to express conditions under which specific information will be released to the information seeker. Users can use this language to express their preferences for what information can be collected by the web sites and what cannot be collected. An APPEL preference file is composed of a set of rules that express the user's privacy concerns. There are three behaviours of APPEL rules: "request", "limited" and "block". a) Request means that the provided evidence is acceptable. b) Limited means that the provided evidence is somewhat acceptable. If a URI is provided, the resource at that URI SHOULD be accessed. However, the request made to access the resource should be limited; that is, all but absolutely necessary request headers should be suppressed. c) Block means that the provided evidence is not acceptable. APPEL does not allow the information owner in a general way to express in terms of rules, for example, that some categories of information but not others are to be divulged. This means that it allows only positive restrictions that could be formulated but not negative ones. As a consequence, a loss of precision could result in a conflict problem between rules if two rules are satisfied by one website's policy but the rule with the unintended result is activated. The formalism of APPEL by Wang [10] at the element level was extended so that negative restrictions could also be handled by introducing a form of Boolean logic. Hence, an evaluation engine was obtained that accepts APPEL with negation. This was based on a translation of APPEL to RuleML (see Boley et al. [1]). The approach was, however, limited to the syntactic level. On the semantic level a new difficulty arises because the meanings of certain terms (e.g., "not for marketing purposes") are often not clear since they are not universally defined. This has to do with the fact that the terms used in the protocol are not the ones that the seller is directly interested in; the latter are rather a more or less indirect consequence. In order to be precise one would need some kind of a dictionary to which a reference can be made.

IV. Determining User Utilities

A. Utility Elicitation

To make effective, rational decisions, one needs a full understanding of her utilities for the various choice alternatives, as well as the possible consequences that can arise from those choices. In domains where there might be several such possibilities, accurately determining the decision-maker's utilities for each possible outcome can be extremely difficult. Research in utility elicitation aims to mitigate the burden of this task. In the approach taken by Chajewska et al. [3], utility is treated as a random variable that is drawn from a known distribution. Given the distributions D and a strategy π, which is the sequence of decisions to be made in the current problem and the criteria used for making those decisions, the expected utility $Eu[\pi|D]$ of the strategy given the distributions is computed. The goal is then to determine the question to ask the decision-maker that will provide the most valuable information. Let q be such a question with n possible answers. If the user gives the ith answer with probability $p(i)$ and the resulting distributions given this new information are D_i, then the posterior expected utility after asking q is

$$\sum_{i=1}^{n} p(i) Eu[\pi|D_i]$$

Typical questions follow the standard gamble pattern. Let x_1, x_2 and x be outcomes such that $u(x_1)$ and $u(x_2)$ are known, and $x_1 \succeq x \succeq x_2$ (x_1 is preferred over x which is preferred over x_2). The user could then be asked the question "Given a choice between receiving alternative x for sure and a lottery which gives alternative x_1 with probability s and alternative x_2 with probability $1 - s$, which would you choose?" If the user chooses x, then we know that $u(x) > u(x_1)s+u(x_2)(1-s)$, otherwise $u(x) < u(x_1)s+u(x_2)(1-s)$. The probability distribution function for $u(x)$ is then updated accordingly.

The utility elicitation process has three main challenges: 1) composing the prior distributions, 2) determining the next question to ask, and 3) deciding when to stop asking questions. We examine each of these challenges in turn, focusing on the special challenges inherent when performing utility elicitation for private information.

B. Our Model

We use the technique for utility elicitation described above to derive the user's utilities as follows. Let b be the business (i.e., the website) with which a negotiation is to take place, let A be the set of possible actions that can be executed as a result of the private information exchange, and let D_b be the set of probability distribution functions for the user's true utility for each action in the set A, given the business b. Thus our model allows for the user to have different utilities for the same action across different websites. Also, let π be the chosen strategy in the current negotiation, and let $Eu[\pi|D_b]$ be the expected utility of executing π, given the beliefs D_b about the user's utilities. Note that in some cases π might be quite simple, such as in take-it-or-leave-it negotiations where an offer o is proposed by the business and thus the strategy π is simply $\pi(o) =$ "accept" if and only if we believe that the user's utility for o is greater than or equal to some threshold. In these cases, $Eu[\pi|D_b]$ can be quite simple to compute. In other scenarios, such as multi-round bilateral negotiations, π can be quite complex and thus computation of $Eu[\pi|D_b]$ might require more sophisticated game-theoretic or decision-theoretic techniques. In either case, we assume that its computation, or at least its estimation, is feasible.

Each time a website is visited with which a privacy negotiation is to take place, the question q that maximizes the expected utility of π is determined. Let u_q denote the expected increase in utility associated with asking q. Also, a *bother* *cost* bc is determined, indicating how much the user will be bothered if another question is asked. This bother cost can be viewed as the reduction in utility that will be realized by the user. Thus, the question q is asked if and only if $u_q > bc$. Once there is no q that satisfies this inequality, the question period for this website visit is terminated and negotiation can begin.

C. Composing Prior Utility Distributions

As in many domains, building probability distributions for a user's likely utilities for the various outcomes in private information exchange is realistic and fairly straight-forward. For example, it is safe to assume that most people will have very low utility for consequences such as credit card fraud and identity theft, and likely higher utility for less meaningful events such a company storing the user's name in a database. Perhaps for some outcomes there is more uncertainty associated with the user's preferences and thus these distributions will have higher variances, but these distributions can still be estimated nonetheless. The difficulty in constructing these distributions instead lies in modeling the dependencies between utilities. Often a user's attitude towards privacy of some information is not independent of her attitude towards the privacy of other information. For example, someone who has more than average aversion to the possibility of being hassled by telemarketers is more likely to have a more than average reluctance towards spam. Even stronger assumptions of dependency hold when two consequences have a common aspect, such as 1) being notified of future offers by phone, and 2) telemarketing from third parties. Since both consequences have a common aspect (being bothered at home by phone), then perhaps a user's utilities for these are not independent.

To account for these dependencies, we model the system of outcomes as a Bayesian Network. Figure 3 gives a simple example of such a network, where five consequences are shown:

N: Company A gets name, keeps in database
P: Company A gets phone number, keeps in database
T: Receive unsolicited phone calls from company A for 1 year
E: Company B will be provided with e-mail address
S: Receive spam

The network in Figure 3 indicates that the user's utility for telemarketing is dependent on her utilities for giving up her name and her phone number, and her utility for spam is dependent on her utility for telemarketing and for having a third party receive her e-mail address. While some of these dependencies may be moderate, such as spam's dependency on telemarketing, there is likely a high interdependency between two very similar consequences, such as third party e-mail and spam. It is likely that if we can ascertain the user's utility for having a third party receive her e-mail address, we can be quite certain about her utility for receiving spam.

Judging these conditional probabilities can be difficult. Utilities for telemarketing almost certainly depend on those for name and phone number, but the degree of these dependencies can be difficult to estimate. Very low utilities for name and phone number would certainly imply low utility for telemarketing, but high utilities for name and phone number (indicating that the user would not mind very much giving either piece of information away separately) certainly does not necessarily imply that the utility for telemarketing is high. We utilize the approach proposed by Cooper and Herskovits [4], which builds a Bayesian Network completely from data. In this way, starting with some initial sample utilities, the network can essentially be learned from the true utilities that are derived from each user. Each time a new user's utilities are elicited, these can be added to the database (perhaps maintained by a web service), and the Bayes' net and corresponding conditional probability distributions can then be updated.

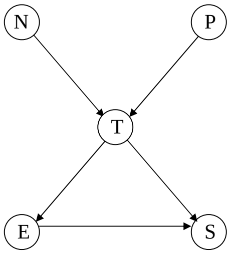

Fig. 3. An example Bayesian Network modeling dependencies in utilities for privacy consequences

D. Determining the Next Question

Utility elicitation questions take the form of standard gambles, as proposed by Chajewska et al. [3]. Let C be the set of consequences that can arise from the impending information exchange, let $u(c)$ represent the possibly unknown utility of a consequence $c \in C$, and let $Eu[c]$ be the expected value of the utility. A question $q = \langle c^-, c^+, c, s\rangle$ is a *candidate* if s is a probability in [0,1] (called a split point), c^-, c^+ and c are in C, $u(c^-)$ and $u(c^+)$ are known, and $u(c^+) \geq u(c) \geq u(c^-)$ is known to be true. Typically c^-, c^+ are chosen to be the worst and best option, respectively, and thus their utilities are known to be $u(c^-) = 0$ and $u(c^+) = 1$ (if utility is normalized to [0,1]). The question q is asked: "Given the option of having consequence c occur for certain or the lottery ℓ_q in which c^+ occurs with s probability or c^- occurs with $1-s$ probability, which would you choose? Since we know the expected utility of ℓ_q is $Eu(\ell_q) = u(c^+)s + u(c^-)(1-s)$, then the user's answer to q will indicate whether $u(c)$ is greater or less than $Eu(\ell_q)$. Moreover, since we have a probability distribution function for $u(c)$, then the probability $p_{<\ell_q}$ that $u(c) < Eu(\ell_q)$ (and thus $p_{>\ell_q} = 1 - p_{<\ell_q}$) is known. The net result of asking q is that we will be informed either that $u(c) < Eu(\ell_q)$ with $p_{<\ell_q}$ probability, or that $u(c) > Eu(\ell_q)$ with $p_{>\ell_q}$ probability.

The value of a question q is determined as the increase in expected utility $Eu[\pi]$ of executing the strategy π that will be realized as a result of obtaining this new information. Let $Eu[\pi| < \ell_q]$ be the expected utility of π given that $u(c) < Eu[\ell_q]$, and $Eu[\pi| > \ell_q]$ be the expected utility of π given that

$u(c) > Eu[\ell_q]$. Then the expected utility of π after asking q is

$$Eu[\pi|q] = Eu[\pi|{<}\ell_q]\cdot p_{<\ell_q} + Eu[\pi|{>}\ell_q]\cdot p_{>\ell_q} \qquad (1)$$

and thus the expected increase in utility is $Eu[\pi|q] - Eu[\pi]$. The question q that maximizes (1) is deemed the next question to ask.

For example, consider the following simple privacy negotiation where the website offers the user a discount on a product in exchange for the user's name and home address, which may be given to a third party. The negotiation is simply a take-it-or-leave-it type. We estimate that the consequence "receive junk mail" will occur with .8 probability as a result of this information sharing, and that the user's utility for this consequence lies somewhere in the range $[0.3, 0.6]$ with uniform probability. Assume that the discount being offered is such that the user will accept the offer if and only if the utility for the associated consequences is 0.54 or higher. This is equivalent to stating that the user achieves utility equal to the utility of the consequences if the deal is accepted, and otherwise achieves 0.54. Thus the strategy for the offer o is $\pi(o) =$ accept iff $u(o) > 0.54$. Since the user's expected utility given that the consequence occurs is 0.45, and thus the expected utility given that it has a 0.8 chance of occurring is $Eu[o] = 0.45\cdot .8 + 1\cdot 0.2 = 0.56$ (we assume that the utility of no consequence is 1), then the user should accept the offer and expect 0.56 utility.

However, there is a chance that the user's true utility is below 0.54, and thus there is a possibility that we have wrongly advised her. If we could further ascertain the user's true utility, we can reduce the chance of such ill-advised suggestions, and consequently increase the expected utility of the transaction. Let $q = \langle telemarketing, nothing, junkmail, 0.5\rangle$ be the question that asks whether the user would prefer receiving junk mail for sure or a lottery ℓ_q in which she would receive nothing with probability 0.5 and telemarketing (assumed to have utility 0) with probability 0.5. Since $Eu[\ell_q] = 1\cdot 0.5 + 0\cdot 0.5 = 0.5$ and the utility for junk mail lies uniformly in $[0.3, 0.6]$, then we know with probability $p_{<\ell_q} = 2/3$ that ℓ_q will be chosen, indicating that the utility for junk mail lies in the uniform range [0.3,0.5] (with expected utility 0.4), and with probability $p_{>\ell_q} = 1/3$ that junk mail will be chosen, indicating that the utility for junk mail lies in the uniform range [0.5,0.6] (with expected utility 0.55). Thus,

$$\begin{aligned} Eu[\pi|{<}\ell_q] &= \max\{Eu[junk|{<}\ell_q]\cdot 0.8 + 1\cdot 0.2, 0.54\} \\ &= \max\{0.4\cdot 0.8 + 1\cdot 0.2, 0.54\} \\ &= 0.54 \end{aligned} \qquad (2)$$

$$\begin{aligned} Eu[\pi|{>}\ell_q] &= \max\{Eu[junk|{>}\ell_q]\cdot 0.8 + 1\cdot 0.2, 0.54\} \\ &= \max\{0.55\cdot 0.8 + 1\cdot 0.2, 0.54\} \\ &= 0.64 \end{aligned} \qquad (3)$$

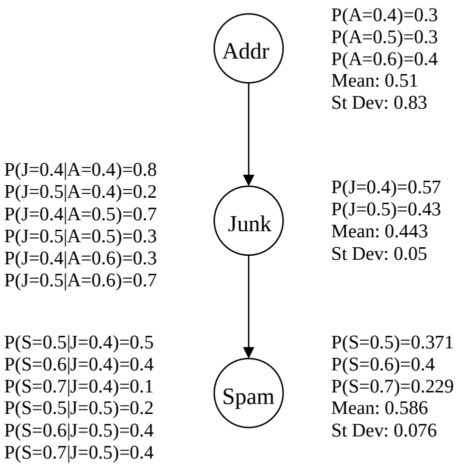

Fig. 4. A Bayesian network modeling dependencies in utilities for giving away address, receiving junk mail, and receiving spam

Thus the expected utility of the strategy after asking q is

$$\begin{aligned} Eu[\pi|q] &= Eu[\pi|{<}\ell_q]\cdot p_{<\ell_q} + Eu[\pi|{>}\ell_q]\cdot p_{>\ell_q} \\ &= 0.54\cdot 0.67 + 0.64\cdot 0.33 \\ &= 0.573 \end{aligned} \qquad (4)$$

Asking q thus increases the expected utility of π by 0.033.

In more complex situations such as bilateral negotiations, the utilities of several possible consequences may need to be derived. Since we model the dependencies in user utilities, some simple questions may effectively increase the utility certainties over several consequences. Consider the partial Bayes' net in Figure 4. In this network, the user's utility for spam is dependent on her utility for receiving junk mail, which is dependent on her utility for a third party receiving her address. Conditional probability distributions are given to the left of the dependent nodes, while the overall probability distributions are given to the right as well as the means and standard deviations. In Figure 5 we see the effect of eliminating an outcome for third party receiving address (i.e., that the user's utility is 0.6). Notice how the subsequent increased certainty in the user's utility for address (we now know it is either 0.4 or 0.5) has increased certainty in our beliefs about the dependent nodes. Before the elimination of this outcome, the standard deviations for the utility distributions for junk mail and spam were 0.05 and 0.076, respectively, while after the elimination they fell to 0.043 and 0.073. Therefore, when determining the effect of asking a question on a strategy that may consider several privacy outcomes and thus several consequences, considering the indirect effect on other utility distributions can be extremely advantageous.

E. Stopping Criterion

In information-gathering scenarios such as the one discussed in this paper, it is crucial to find the right balance between obtaining potentially useful information and avoiding bothering the user unnecessarily. To this end, we propose that the agent should stop asking questions when the most promising next

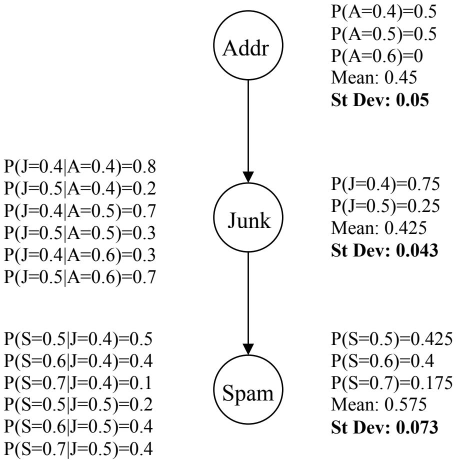

Fig. 5. The network from Figure 4 demonstrating the reduced standard deviation of each dependent action after a possibility for address utility (i.e., that $u(A) = 0.6$) is eliminated

question has an expected value that falls below some threshold. As described in Section IV.B, we achieve this by terminating a question period as soon as there is no question q that satisfies the inequality $u_q > bc$, where u_q is the increase in utility associated with asking question q and where bc is the bother cost.

Furthermore, because different users have different attitudes toward privacy and different tolerances for interruptions from semi-autonomous systems, we propose that this threshold should be user-dependent. In this section, we outline how the notion of "bother cost" presented by Fleming and Cohen [8] can be modified to suit our needs. Throughout the section, we define many variables that should in fact be defined as functions that depend on the user u. For example, the bother cost bc should be represented as a function $bc(u)$. However, to simplify the notation, we omit the parameter u and write this simply as bc.

When a particular user first employs our system, she will be informed that she will be asked questions occasionally regarding her attitudes toward giving away certain pieces of personal information. She will then be asked to indicate her willingness to field such questions, on the integer scale from 0 to 10. A user who is very willing to interact with such a system might choose a value of 10 (perhaps because she is very concerned about privacy issues), while a user who is unwilling to interact might choose a value of 0 (perhaps because privacy is less important to her or because she is particularly annoyed by unexpected interruptions). For a particular user, we use the notation w to represent the willingness of the user to interact with the system. For future work, we intend to implement a learning technique for refining this value based on observations of the user's subsequent behaviour. However, for the purpose of this paper, we assume that the user can provide a reasonably accurate estimate.

We will define the bother cost as a function that depends on the user's willingness to interact and on the past bother to which this user has been subjected in recent sessions.

First, we define $\alpha = 1.26 - 0.05w$. This yields a value very close to 1 for users with a moderate willingness level ($w = 5$) and a value below or above 1 for users who are more or less willing to interact, respectively.

We also keep track of the degree to which the user has been bothered recently. Let S be the set of the ten most recent sessions initiated by the user, where a session is defined simply as a visit to a website. For each $s \in S$, let $nq(s)$ be the number of questions that were asked of the user during session s, and let $t(s)$ be the time that has elapsed since session s, measured in minutes.

We define the recent bother experienced by the user as

$$rb = \sum_{s \in S} \beta^{t(s)} nq(s) \tag{5}$$

where β is a discount factor between 0 and 1. The intuition behind this formula is to measure the number of times the user has been asked questions in past sessions, but to give recent questions more weight in this calculation than questions that happened further in the past. The factor β is used to control how "forgetful" a user is of previous interruptions, and should simply be determined by the system designer.

Finally,

$$bc = \frac{\gamma \cdot (1 - \alpha^{1+rb})}{(1 - \alpha)} \tag{6}$$

where γ is a scaling factor that keeps the bother costs consistent with the scale of utilities. The result is that users with moderate willingness values have a nearly linear bother curve, while patient and impatient users will have more logarithmic or exponential curves, respectively.

As an example, suppose a user u_i has indicated a willingness level of $w = 9$ on a scale of 0 to 10. Then, $\alpha = 1.26 - 0.05w = 1.26 - 0.05(9) = 0.81$.

Suppose also that $\beta = 0.95$, and that, in this user's last ten visits to websites, she was asked one question 10 minutes ago, one question 30 minutes ago and one question 60 minutes ago. Then,

$$rb = (0.95)^{10}(1) + (0.95)^{30}(1) + (0.95)^{60}(1) = 0.86.$$

Finally, using a value of $\gamma = 0.02$,

$$\begin{aligned} bc &= \tfrac{\gamma \cdot (1-\alpha^{1+rb})}{(1-\alpha)} \\ &= \tfrac{0.02(1-0.81^{1.86})}{(1-0.81)} \\ &= 0.034 \end{aligned} \tag{7}$$

If the most promising next question has an expected utility gain u_q that is greater than 0.034, then the agent will ask the question. Otherwise, it will stop asking.

Offer	Resulting Actions
1. n, a, c, ai	J, B, Sa
2. n, a, p	T, J, Sa
3. e, p, ai, s	G, Sp
4. n, c, ai, s	G, B
5. n, p, e, c	T, Sp, B

TABLE I
OFFERS AND RESULTING ACTIONS IN THE EXAMPLE

V. EXPERIMENTATION

To illustrate the effect of our method on the expected utility of a negotiation strategy, we tested the technique on an example private information exchange. In this example, the website seeks the following information from the user: name (n), address (a), e-mail address (e), company (c), academic institution (ai), student number (s), and phone number (p). Thus the entire set of information units being sought is denoted by $IU = \{n, a, e, c, ai, s, p\}$. The set A of actions that could occur as a result of releasing the information in IU is

Spam	Sp	:- e
Junk mail	J	:- n, a
Telemarketing	T	:- n, p
Get grades	G	:- ai, s
Notice to boss	B	:- n, c
Visit from salespeople	Sa	:- a

Negotiation in this example comes in the form of a single round of offers from the website, where the user can select which one she prefers. The offers and the consequential actions that can arise are given in Table I.

The task is to determine this particular user's utilities for each of the possible actions, and subsequently for each offer. We begin with the prior conditional probability distributions given by the Bayesian network in Figure 6. After consultation with the user, we give the worst action (get grades) utility 0, and consider the lack of an action to have utility 1. Table II gives the expected values of the user's utilities, based on this network. The utility of an offer is computed based on the user's utilities for the actions that may result. For simplicity in this example, we assume that (a) all action probabilities are 1 and thus probabilities are ignored, and (b) utility is additive, thus the function

$$Eu(A') = \sum_{a \in A'} \frac{u(a)}{|A|} + \frac{|A| - |A'|}{|A|} \tag{8}$$

is used to determine the utility of a set $A' \subseteq A$ of actions. The utility is simply the average over the set A, where an action not appearing in A' gets utility 1. The utility of each offer is given in Table III. The user should choose offer 3, expecting a utility of 0.7767.

The question that produces the highest increase in expected utility is $\langle G, \{\}, J, 0.6\rangle$. Table IV shows the resulting expected utility of each offer for each answer to q. If the user prefers the lottery ℓ_q (receiving no consequence with 0.6 probability, get grades with 0.4), then the updated utilities suggest that offer 3 is still the best with expected utility 0.7767. If the user prefers J, then the updated utilities indicate that offer 2 is best with utility 0.8429. Since $p_{<\ell_q} = 0.7$ and $p_{>\ell_q} = 0.3$, the overall expected utility of asking the question is $Eu[\pi| < \ell_q]\cdot p_{<\ell_q} + Eu[\pi| > \ell_q]\cdot p_{>\ell_q} = 0.7767\cdot 0.7 + 0.8429\cdot 0.3 = 0.7965$, giving an increase of $0.7965 - 0.7767 = 0.0198$.

Given that the user chooses ℓ_q, the next question that produces the highest increase in expected utility is $q' = \langle G, \{\}, Sp, 0.5\rangle$. If the user prefers the lottery then the updated utilities suggest that offer 3 is still the best with expected

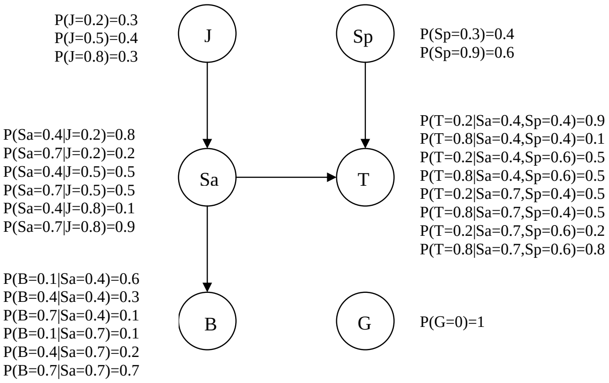

Fig. 6. The Bayesian network for the example, indicating the conditional probability distributions

Action a	$Eu[a]$
Sp	0.66
J	0.5
T	0.5121
B	0.4249
Sa	0.559
G	0

TABLE II
INITIAL EXPECTED UTILITY OF EACH ACTION

Offer o	$Eu[o]$
1	0.7473
2	0.7619
3	0.7767
4	0.7375
5	0.7662

TABLE III
INITIAL EXPECTED UTILITY OF EACH OFFER

Offer o	$Eu[o\|<\ell_q]$	$Eu[o\|>\ell_q]$
1	0.7092	0.8362
2	0.7271	0.8429
3	0.7767	0.7767
4	0.7288	0.7578
5	0.7521	0.7991

TABLE IV
NEW EXPECTED UTILITIES GIVEN THE TWO ANSWERS FOR QUESTION q

o	$Eu[o\vert<\ell_{q'}]$	$Eu[o\vert>\ell_{q'}]$	$Eu[o\vert<\ell_{q''}]$	$Eu[o\vert>\ell_{q''}]$
1	0.6702	0.7752	0.8362	0.8362
2	0.6959	0.7799	0.7783	0.8783
3	0.7767	0.7767	0.7767	0.7767
4	0.7083	0.7633	0.7578	0.7578
5	0.7190	0.8080	0.7345	0.8345

TABLE V

New expected utilities after either question q' or q'' is asked

Action	Prior St dev	Post St Dev
Sp	0.2939	0.2939
J	0.2324	0.1039
T	0.2997	0.1978
B	0.2591	0.2070
Sa	0.1497	0.0270

TABLE VI

Standard deviations of utilities before and after the two questions are asked

utility 0.7767 (see column 2 in Table V). Otherwise, offer 5 is best with utility 0.808 (see column 3). So the overall expected utility of asking the question is $0.7767 \cdot 0.629 + 0.808 \cdot 0.371 = 0.7883$, giving an increase of 0.0116.

Otherwise, if the user does not choose ℓ_q in the first question, the next question that produces the highest increase in expected utility is $q'' = \langle G, \{\}, T, 0.5\rangle$. If the user prefers the lottery, then the updated utilities suggest that offer 1 is best with expected utility 0.8362 (see column 4 in Table V). Otherwise, offer 2 is best with utility 0.8783 (see column 5). So the overall expected utility of asking q'' is $0.8362 \cdot 0.354 + 0.8783 \cdot 0.646 = 0.8634$, giving an increase of 0.0205.

Just by asking two questions in the elicitation process, the expected utility rises from 0.7767 to $0.7767 \cdot 0.4403 + 0.808 \cdot 0.2597 + 0.8362 \cdot 0.1062 + 0.8783 \cdot 0.1938 = 0.8108$. To further illustrate the effect on uncertainty, the expected values of the standard deviations for the utility of each action, before and after the questions are asked, are given in Table VI.

VI. Conclusions and Future Work

In this paper we examine a problem of practical relevance, namely determining and quantifying the possible negative aspects when exchanging personal information in an (internet) sales situation. The question is to determine the price for this information. We reduce the problem to determining the possible costs to the customer of actions that could be performed using this information.

We point out that this problem has a formal mathematical side as well as an informal and context dependent side. Therefore a combination of mathematical techniques and informal user interaction is presented. Mathematically, we rely on utility theory and Bayesian nets. For the user interaction we have developed a scenario in which the problem can be presented in terms that are directly related and understandable to the actual situation of the customer. We also relate this directly to P3P, which is a W3C recommendation for the automated exchange of privacy policies. Experiments show that, in domains such as this where a user's utilities over the set of consequences contain several interdependencies, uncertainty can be reduced for several outcomes by asking relatively few questions. This reduced uncertainty leads to better decision-making, and consequently higher utility.

For future work, we plan to further pursue research in utility elicitation with the goal of determining more effective techniques that exploit the interdependent nature of private information utilities. One problem with the standard technique used here is that the expected utility of a question is computed myopically. This means the value of the question is computed without considering the value of subsequent questions. Consider the example presented in section V. While we twice found the best question to ask, together these questions may not have been the most effective pair. More consideration should be given to determining which question will lead to more good questions.

Also requiring further examination is the determination of the bother cost. In particular, refining the user's willingness value over a number of sessions is necessary since (a) the user is not likely to accurately determine her own willingness value, and (b) the user's true value may change over time. Machine learning techniques could be used to constantly observe the user's behaviour (e.g., when questions are answered or ignored) and update the willingness value accordingly.

We also plan to examine how a user's utilities may change during a negotiation. If the website is very insistent on obtaining a particular type of data, one may conclude that this data is worth more than what was initially expected. Thus the user may lower her utility for giving it away, and perhaps demand more compensation in return. It may also be worth investigating the construction of negotiation strategies that can be used to elicit the website's utilities for the user's data.

References

[1] H. Boley, B. Grosof, M. Sintek, S. Tabet, and G. Wagner. RuleML design version 0.87. http://www.dfki.uni-kl.de/ruleml, 2004.

[2] S. Buffett, K. Jia, S. Liu, B. Spencer, and F. Wang. Negotiating exchanges of P3P-labeled information for compensation. *Computational Intelligence*, 20(4), 2004. To appear.

[3] U. Chajewska, D. Koller, and R. Parr. Making rational decisions using adaptive utility elicitation. In *AAAI-00*, pages 363–369, Austin, Texas, USA, 2000.

[4] G. F. Cooper and E. Herskovits. A bayesian method for the induction of probabilistic networks from data. *Machine Learning*, 9:309–347, 1992.

[5] L. Cranor, M. Langheinrich, and M. Marchiori. A P3P Preference Exchange Language 1.0 (APPEL1.0). http://www.w3.org/TR/P3P-preferences/, 15 April 2002. W3C Working Draft.

[6] L. Cranor, M. Langheinrich, M. Marchiori, M. Presler-Marshall, and J. Reagle. The Platform for Privacy Preferences 1.0 (P3P1.0) Specification. http://www.w3.org/TR/P3P/, 16 April 2002. W3C Recommendation.

[7] P.C. Fishburn. The axioms of subjective probability. *Statistical Science*, 1:335–358, 1986.

[8] M. W. Fleming and R. Cohen. A decision procedure for autonomous agents to reason about interaction with humans. In *the AAAI 2004 Spring Symposium on Interaction between Humans and Autonomous Systems over Extended Operation*, pages 81–86, 2004.

[9] J. von Neumann and O. Morgenstern. *Theory of games and economic behaviour, 2nd ed.* Princeton University Press, Princeton NJ, USA, 1947.

[10] F. Wang. Detecting and preventing rule conflicts in APPEL. Master's thesis, University of New Brunswick, Canada, 2003.