Use of a P3P User Agent by Early Adopters

Lorrie Faith Cranor
AT&T Labs-Research
180 Park Ave.
Florham Park, NJ 07932
+1 973-360-8607
lorrie@research.att.com

Manjula Arjula
AT&T Labs
200 Laurel Ave. S.
Middletown, NJ 07748
+1 732 420 3587
manjula@att.com

Praveen Guduru
AT&T Labs
200 Laurel Ave. S.
Middletown, NJ 07748
+1 732 420-3588
praveen@babel.lz.att.com

ABSTRACT

The Platform for Privacy Preferences (P3P), developed by the World Wide Web Consortium (W3C), provides a standard computer-readable format for privacy policies and a protocol that enables web browsers to read and process privacy policies automatically. P3P enables machine-readable privacy policies that can be retrieved automatically by web browsers and other user agent tools that can display symbols, prompt users, or take other appropriate actions. We developed the AT&T Privacy Bird as a P3P user agent that can compare P3P policies against the user's privacy preferences. Since P3P was adopted as a W3C recommendation in April 2002, little work has been done to study how it is being used and, especially, its impact on users. Many questions have been raised about whether and how Internet users will make use of P3P, and how to build P3P user agents that will prove most useful to end users. In this paper we first provide a brief introduction to P3P and the AT&T Privacy Bird. Then we discuss a survey of AT&T Privacy Bird users that we conducted in August 2002. We found that a large proportion of AT&T Privacy Bird users began reading privacy policies more often and being more proactive about protecting their privacy as a result of using this software. Unfortunately, the usefulness of P3P user agents is severely limited by the number of web sites that have implemented P3P. Our survey results also suggest that if it becomes easier to compare privacy policy across e-commerce web sites, a significant group of consumers would likely use this information in their purchase decisions.

Categories and Subject Descriptors

K.4.1 [**Computers and Society**]: Public policy issues – *privacy*

General Terms

Design, Human Factors, Standardization

Keywords

Privacy, P3P, user agent, user study, survey

1. INTRODUCTION

Internet users are becoming increasingly concerned about what personal information they may reveal when they go online and where that information might end up. These privacy concerns are making consumers nervous about going online, but current web site privacy policies are so long and difficult to understand that consumers rarely read them. Many consumers say they would prefer if privacy policies were presented in a standard, easy-to-read format [11]. The Platform for Privacy Preferences (P3P), developed by the World Wide Web Consortium (W3C), addresses this problem by providing both a standard computer-readable format for privacy policies, and a protocol that enables web browsers to read and process privacy policies automatically [4]. P3P enables machine-readable privacy policies that can be retrieved automatically by web browsers and other user agent tools that can display symbols, prompt users, or take other appropriate actions. We developed the AT&T Privacy Bird as a P3P user agent that can compare P3P policies against the user's privacy preferences and assist the user in deciding when to exchange data with web sites.

Since P3P was adopted as a W3C recommendation in April 2002 it has been adopted by numerous web sites and incorporated into two popular web browsers as well as into other software. However, little work has been done to study how P3P is being used and, especially, its impact on users. Many questions have been raised about whether and how Internet users will make use of P3P, and how to build P3P user agents that will prove most useful to end users.

In this paper we first provide a brief introduction to P3P and the AT&T Privacy Bird. Both P3P and the AT&T Privacy Bird are described in more detail in [3]. This paper focuses on a survey of AT&T Privacy Bird users that we conducted in August 2002. We present our survey methodology, describe our results, and discuss our findings.

2. THE PLATFORM FOR PRIVACY PREFERENCES

P3P provides a standard way for web sites to communicate about their privacy policies. Privacy policies are intended to describe a company's *data practices*—what information they collect from individuals (usually customers and potential customers, but sometimes also employees and others) and what they do with it. The P3P specification includes a standard *vocabulary* for describing these data practices and a *base data schema* for describing the kinds of information collected. A P3P *policy* is an XML-encoded collection of vocabulary and data elements that describe the data practices of a particular web site (or section of a

web site). A P3P policy is composed essentially of the answers to a number of multiple-choice questions, and thus does not always contain as much detailed information as a human-readable privacy policy (i.e., a policy written in English or another spoken language that is intended for people, rather than computers, to read).

The P3P 1.0 Specification [4] provides detailed definitions for all of the XML elements that comprise the P3P vocabulary. Figure 1 provides an overview of these elements.

ENTITY – contact information for the business, organization, or person who owns the site

ACCESS – whether individuals can find out what personal data a site keeps about them in its databases (6 types of access policies are specified)

DISPUTES – how to resolve privacy-related disputes with the site (customer-service desk, privacy seals, relevant privacy laws, etc.); also includes REMEDIES sub-element

DATA – the kinds of data collected (17 data CATEGORY elements and dozens of specific data elements are specified)

PURPOSE – how collected data is used (11 types of purposes and an "other-purpose" are specified), and whether individuals can opt-in or opt-out of any of these uses

RECIPIENT – whether and under what conditions data may be shared and whether there is an opt-in or opt-out (6 types of recipient policies are specified)

RETENTION – policies for periodic purging of collected data (5 types of retention policies are specified)

CONSEQUENCE – human-readable explanation of site's data practices

Figure 1. P3P policy elements overview

The P3P specification also includes a protocol, built on the HTTP protocol, for requesting and transmitting P3P policies. P3P user agents use standard HTTP requests to fetch a P3P policy reference file from a well-known location on the web site to which a user is making a request. The policy reference file indicates the location of the P3P policy file that applies to each part of the web site. There might be one policy for the entire site, or several policies that each cover a different part of the site. The user agent can then fetch the appropriate policy, parse it, and take action according to the user's preferences.

P3P also allows sites to place policy reference files in locations other than the well-known location. In these cases, the site must declare the location of the policy reference file using a special HTTP header or by embedding a LINK tag in the HTML files to which the P3P policies apply. Special HTTP headers are also used to transmit an optional P3P compact policy whenever cookies are set. Compact policies are very short summaries of full P3P policies that describe only the data practices related to cookies. They do not have the full expressive capabilities of P3P policies.

Figure 2 shows a plain English example of a relatively simple privacy policy typical of many non-commercial web sites. This privacy policy is represented as a P3P policy in Figure 3.

We do not currently collect any information from visitors to this site except the information contained in standard web server logs (your IP address, referer, information about your web browser, information about your HTTP requests, etc.). The information in these logs will be used only by us and the server administrators for website and system administration, and for improving this site. It will not be disclosed unless required by law. We may retain these log files indefinitely. Please direct questions about this privacy policy to privacy@p3pbook.com.

Figure 2. Example privacy policy

```
<POLICIES
xmlns="http://www.w3.org/2002/01/P3Pv1">
<POLICY
discuri="http://p3pbook.com/privacy.html"
        name="policy">
 <ENTITY>
 <DATA-GROUP>
  <DATA
    ref=
   "#business.contact-info.online.email">
       privacy@p3pbook.com
  </DATA>
  <DATA
    ref=
   "#business.contact-info.online.uri">
       http://p3pbook.com/
  </DATA>
  <DATA ref="#business.name">
   Web Privacy With P3P
  </DATA>
 </DATA-GROUP>
 </ENTITY>
 <ACCESS><nonident/></ACCESS>
 <STATEMENT>
  <CONSEQUENCE>
   We keep standard web server logs.
  </CONSEQUENCE>
  <PURPOSE>
   <admin/><current/><develop/>
  </PURPOSE>
  <RECIPIENT>
   <ours/>
  </RECIPIENT>
  <RETENTION>
   <indefinitely/>
  </RETENTION>
  <DATA-GROUP>
    <DATA ref="#dynamic.clickstream"/>
    <DATA ref="#dynamic.http"/>
  </DATA-GROUP>
  </STATEMENT>
</POLICY>
</POLICIES>
```

Figure 3. Example P3P policy

E-commerce sites on which customers can make purchases typically have lengthier privacy policies that describe their more complicated data practices. These policies often run several pages and include legalese that is difficult for most consumers to understand.

The P3P 1.0 Specification says very little about how P3P user agents should communicate with users. The P3P working groups left user interface issues mostly unspecified because a standardized user interface is not necessary for P3P to work [4]. The complexity and flexibility of the P3P vocabulary creates both opportunities and challenges for user interface developers. P3P user agent developers are charged with the task of making P3P both usable and useful, despite its complexity.

The P3P vocabulary was designed to distinguish a wide range of data practices; however, as the vocabulary was developed there was not a lot of attention paid to expressing the vocabulary in language readily understandable to the average user. Many of the P3P vocabulary terms borrow terminology from privacy laws and fair information practice principles. While these terms are well known to privacy experts, they are foreign to almost everyone else. In addition, some of the distinctions made in the vocabulary are unlikely to be important to most users—although it is quite likely that the distinctions users find most important will change over time and perhaps even vary across different regions of the world. Thus it is a challenge for user agent implementers to simplify the language of the P3P vocabulary and find ways of packaging it that will resonate with their users.

P3P user agents generally allow users to specify their privacy preferences so that they can compare a web site's policies to these preferences automatically. P3P user agents can also provide tools that make it easier for users to quickly assess a site's privacy practices for themselves. Some user agents display symbols that summarize a site's privacy policy or indicate that it has a privacy seal or is bound by certain privacy laws. Some user agents also include buttons that lead to a site's human-readable privacy policy without having to search for it on the site. P3P user agents are already built into the Microsoft Internet Explorer 6 [8] and Netscape Navigator 7 web browsers. Other P3P user agents are available as browser add-ons or proxies [3].

3. THE AT&T PRIVACY BIRD

We developed a P3P user agent, called the AT&T Privacy Bird, which works as a "browser helper object" [7] with the Microsoft Internet Explorer 5.01, 5.5, and 6.0 web browsers on Microsoft Windows 98/2000/ME/NT/XP operating systems. The beta 1.1 version is available as a 1.4MB self-extracting file that includes an installation wizard. The user interface design for this software benefited from our experience with four prototype P3P user agent implementations (our design process is discussed in more detail in chapter 14 of [3]).

The AT&T Privacy Bird displays a bird icon in the browser title bar that changes color and shape to indicate whether or not a web site's P3P policy matches a user's privacy preferences. As shown in Figure 4, the bird icon also functions as a button for accessing the AT&T Privacy Bird menus. The bird icon brings to mind phrases such as "a little bird told me" and images of a little bird whispering in one's ear. As shown in Figure 5, a happy green bird indicates a site that matches a user's preferences, the same green bird with an extra red exclamation point indicates a site that matches a user's preferences but contains embedded content[1] that does not match or does not have a P3P policy, an uncertain yellow bird indicates a site that does not have a P3P policy, an angry red bird indicates a site that does not match a user's preferences, and a sleeping gray bird indicates that the tool is turned off. The birds also have distinct "bubbles" that are distinguishable by colorblind users and users who do not have color displays. Sounds associated with the red, green, and yellow birds serve to reinforce the visual icons (users can choose whether or not they want to hear the sounds).

Figure 4. AT&T Privacy Bird menus

Figure 5. AT&T Privacy Bird icons

3.1 Policy Summary

Users can click on the bird icon and select Policy Summary from the About This Site menu to view a summary of the site's privacy policy that is generated automatically from the site's P3P policy. At sites that do not match a user's preferences, the policy summary also explains where the policy differs from the user's preferences. Figure 6 shows a policy summary for a site that has a policy that does not match the user's preferences. The policy summary begins with a Privacy Policy Check, which indicates any points where a site's policy does not match a user's preferences. Below the check is a summary derived from the site's P3P policy. It includes a bulleted summary of each statement in the policy, as well as information from the P3P ACCESS, DISPUTES, and ENTITY elements, including images of any privacy seals referenced. Rather than using the full definitions of each PURPOSE, CATEGORY, RECIPIENT and ACCESS element from the P3P specification, we developed abbreviated descriptions using language likely to be more readily accessible to users. We append the words "unless you opt-out" to purposes for which an opt-out is available, and provide a hyperlink to the site's instructions for opting out. We append the words "only if you request this" to purposes that occur only if a user opts-in.

[1] Embedded content includes images, sounds, frames, and other objects embedded in a web page. Any object that can be addressed by a URL can have a P3P policy.

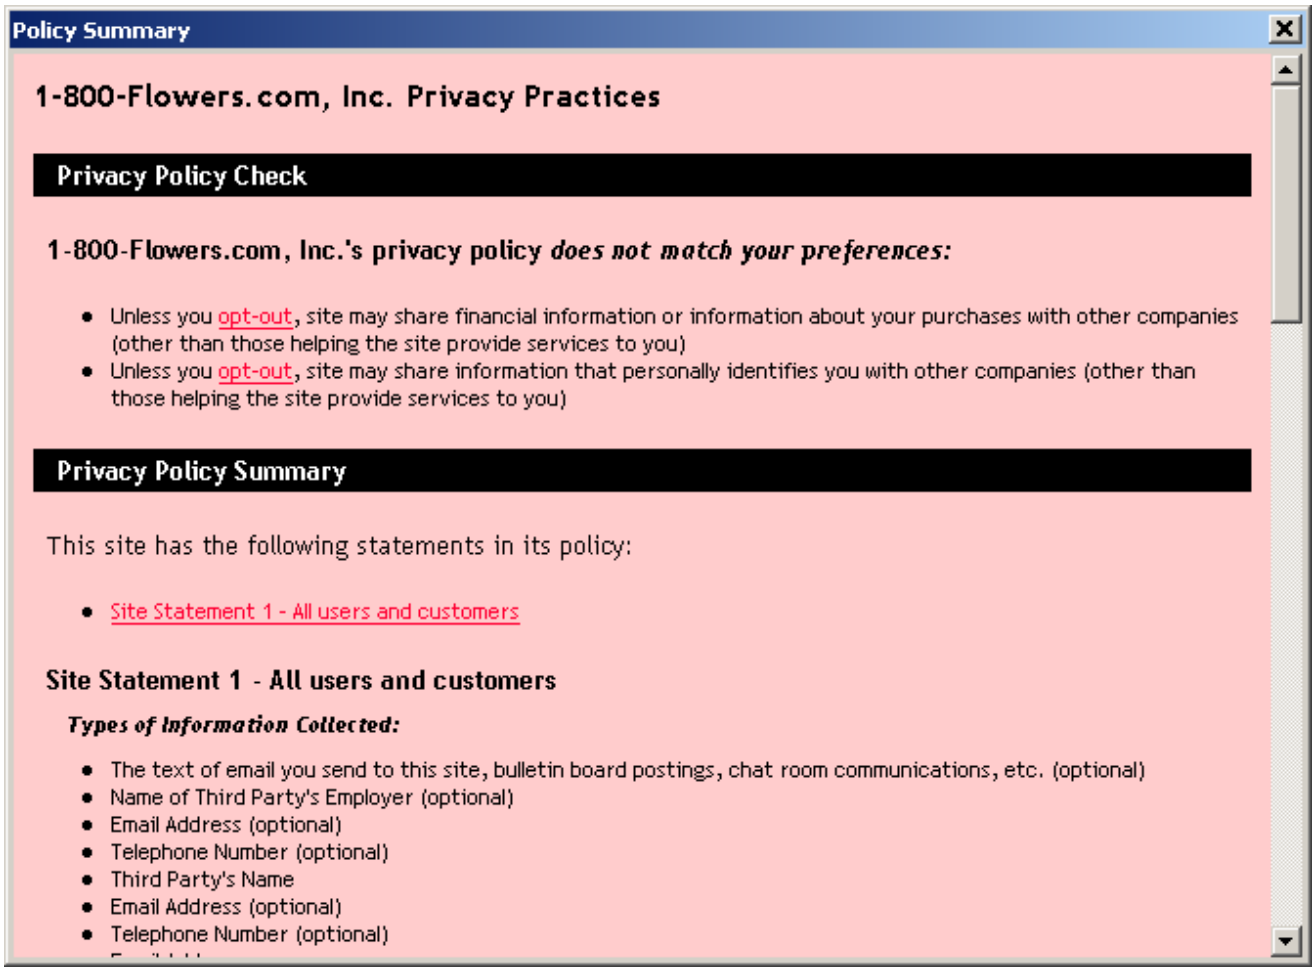

Figure 6. A policy summary for a site that does not match the user's preferences

Our goal was to make the policy summary window into an automatically generated privacy "nutrition label." Our experience with previous P3P user agent prototypes gave us some insights into how to make the policy summary accessible to end users [3]. Additional work is needed to better understand what aspects of the policy summary users are most interested in and what terminology is most readily understandable. The user study discussed in this paper is a step towards gaining this understanding.

3.2 Preference Configuration

Another challenge in designing the AT&T Privacy Bird software was developing a configuration window for setting user privacy preferences. As already discussed, the complexity of the P3P vocabulary and the expert language used in the definitions of vocabulary elements, makes it difficult to convey information about P3P policies to end users. A graphical user interface that allowed users to specify preferences over every possible combination of vocabulary elements would be overwhelming. Therefore, we used survey data [2] and our experience with P3P user agent prototypes [3] to focus on a subset of the P3P vocabulary that we believe American Internet users are most interested in. Our GUI bundles together related vocabulary terms, reducing the number of choices available to the user. In addition we have developed terminology designed to be more accessible to end users than the terminology used in the P3P specification.

As shown in Figure 7, the AT&T Privacy Bird preference settings window offers users twelve possible conditions that may trigger warnings. Users can read through the twelve descriptions and select any or all of them, or they can select from the pre-configured high, medium, or low settings. Selecting one of the pre-configured settings causes the corresponding check boxes to be checked automatically. This gives users easily-accessible information about what each of these settings actually means.

The preference setting window also gives users the option of importing settings files written in a limited version of the APPEL language [5]. This allows expert users to take advantage of the full flexibility of the P3P vocabulary when specifying their preferences. It also makes it possible for individuals or organizations to create and distribute pre-configured "recommended settings" files that reflect their views [1].

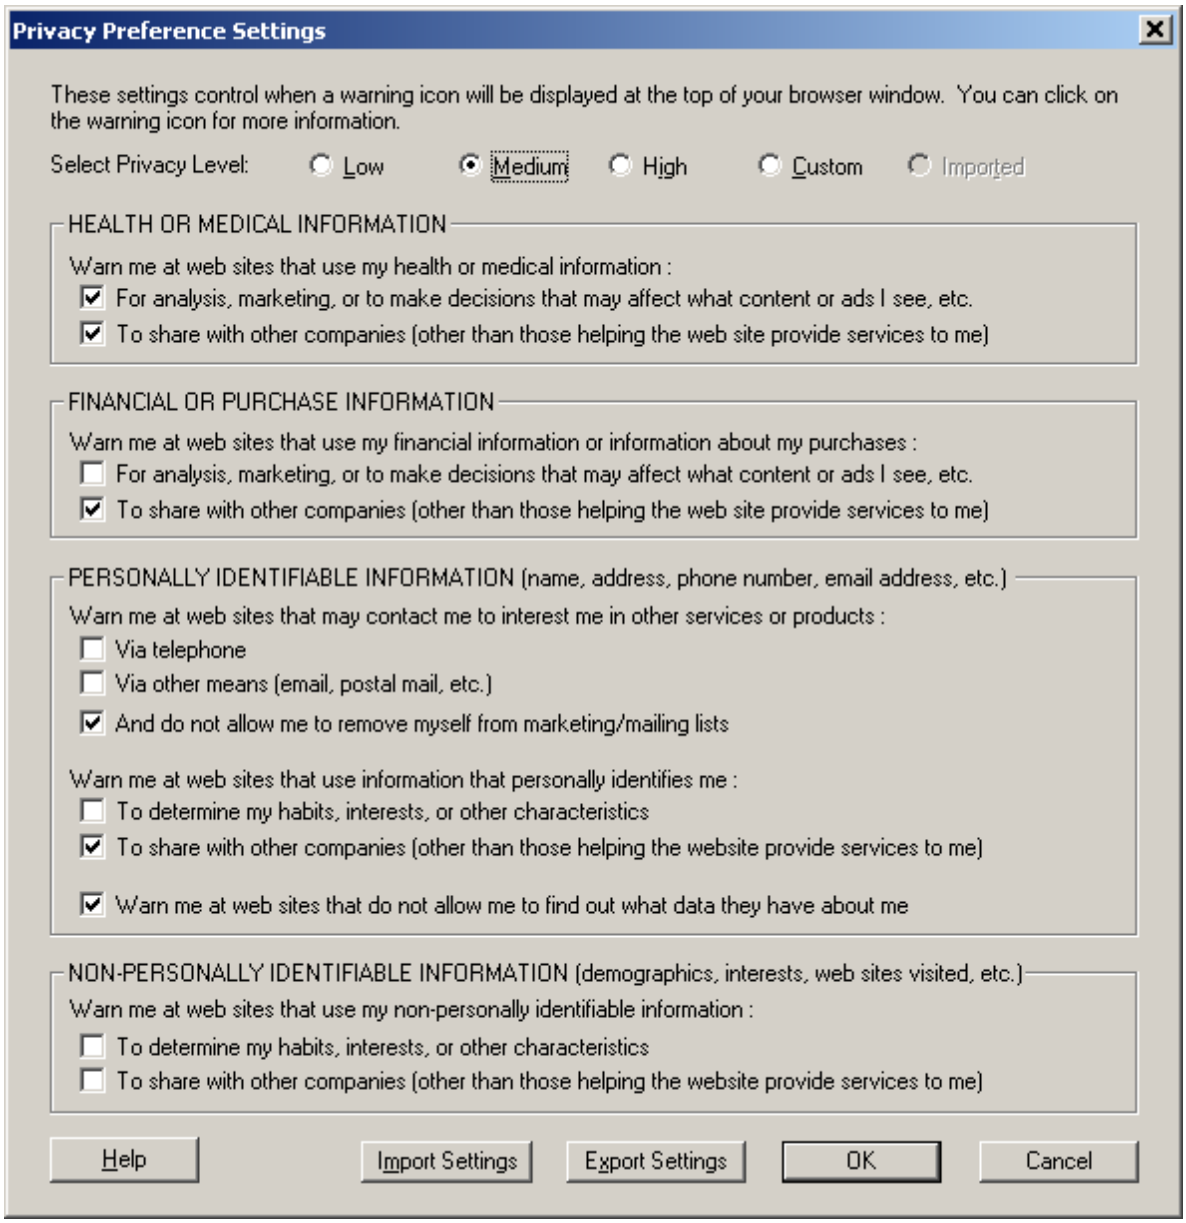

Figure 7. Privacy preference settings window

4. USER STUDY

We released a beta version of the AT&T Privacy Bird to a group of AT&T WorldNet customers in October 2001. We then released a free public beta from the *http://privacybird.com/* web site in February 2002. The software was advertised on several AT&T web sites and included in lists of P3P implementations on the W3C web site and at *http://p3ptoolbox.org/*. One of the authors spoke about the software at several conferences and workshops, and was interviewed for a variety of publications. Probably as a result of mentions in news reports, links to the AT&T Privacy Bird were created on many other sites around the world. This publicity resulted in approximately 20,000 downloads of the software in the first six months of the public beta trial.

In August 2002 we conducted a user study in order to evaluate and improve the AT&T Privacy Bird software and to gain insights into how it is being used.

4.1 Survey Methodology

The AT&T Privacy Bird download site asks individuals to provide an email address and check a box if they are willing to participate in AT&T Privacy Bird user studies. In the first six months of the trial, approximately 4,000 individuals indicated that they would be willing to participate in user studies. In August 2002 we randomly selected 2000 of their email addresses and sent them an invitation to fill out a 35-question, web-based questionnaire. Each email contained a unique URL to ensure that each individual filled out the questionnaire only once.

We received 331 completed surveys (a 16.55% response rate). In addition, about three-dozen individuals emailed us or submitted incomplete surveys with comments indicating that they had never installed the software. Some wrote that they did not realize until after they downloaded the software that it would not run on their operating system or browser. Others were unable to get the software to install correctly on their system.[2]

4.2 Demographics and Internet Use

Our survey asked several demographic questions in order to help us better understand the types of individuals who are using the AT&T Privacy Bird software. The demographics of the users who completed our survey are somewhat different from the demographics of Internet users in general; our respondents are older, better educated, and more predominantly male than the general Internet population. We found that 84% of our respondents are male. The average age of our respondents is 46, and 45% of our respondents are 50 or over. A recent survey published by Consumer WebWatch [10][3] estimated that 51% of American Internet users are male and that 22% of American Internet users are 50 or over. 91% of our respondents have had at least some college education, and 31% have had post-graduate education as well, while Consumer WebWatch estimated that 62% of American Internet users have had at least some college education and 14% have had post-graduate education.

Our respondents tend to be experienced Internet users. Only 16% of our respondents began using the Internet in the past three years, while Consumer WebWatch [10] estimated that 55% of all American Internet users have begun using the Internet in the past three years. 90% of our respondents report using the Internet at least once a day, with 78% reporting that they use it several times a day or more. In addition, 92% of respondents said they have purchased products or information from a web site.

About 70% of respondents live in the United States, 14% in Australia, 6% in Canada, and 2% in the United Kingdom, Holland, and France. Respondents came from several other countries as well, with no other country representing more than 2% of respondents. Given that the software is designed for English speaking users and the download web site and survey were written in English, it is not surprising that most of the respondents reported living in English speaking countries. Overall, demographics were similar across countries; however, US respondents tended to have more Internet experience and be more frequent Internet users. We found that 82% of US respondents reported having used the Internet for at least five years, as compared to 68% of respondents outside the US. In addition, 97% of US respondents had made purchases at a web site, as compared with 77% of respondents from outside the US.

[2] Prior to conducting the survey we identified a problem with the software that prevents it from installing properly on some systems; we believe that this may have impacted about 2% of individuals who tried to install the software.

[3] Note that the Consumer WebWatch study was limited to American Internet users. However, the significant demographic differences between our respondents and the Consumer WebWatch estimate still appear (and in some cases increase) if we examine only the responses we received from respondents who live in the US.

Table 1 compares the demographics of our respondents with the Consumer WebWatch [10] estimates of American Internet user demographics in 2002.

Table 1. Respondents' demographics compared to Consumer WebWatch estimate of American Internet users

	AT&T Privacy Bird Survey Respondents	**Consumer WebWatch**
Male	84%	51%
Over 50	45%	22%
At least some college education	91%	62%
Post-graduate eduction	31%	14%
Began using Internet in past three years	16%	55%
Go online several times a day	78%	35%

The skewed demographics of our survey respondents are noteworthy. While there may be some bias due to the fact that people interested in privacy software may be less likely than other people to respond to surveys, the magnitude of the difference between our demographics and those found in representative samples of Internet users is quite large. It would be interesting to determine whether our demographics reflect the demographics of the users of other privacy software, for example cookie cutter software or anonymity tools. We would expect users of other privacy software to have similar demographics; however, no user studies have yet reported demographics for users of other privacy software. It would also be interesting to determine why our demographics are skewed towards older, educated, male, experienced Internet users. Do these users have more interest in privacy software, more knowledge about privacy software, or more awareness of P3P than other users?

4.3 Attitudes about Privacy

We asked a number of questions to determine the attitudes of AT&T Privacy Bird users about privacy and their familiarity with P3P and with privacy policies. 34% of respondents said they had never heard of P3P, and 44% said they had heard of P3P but didn't know much about it. We identified 21% of our respondents as "P3P experts" because they said they knew a lot about P3P or had actually created a P3P policy. We expect this percentage is significantly higher than the percentage of P3P experts among Internet users in general, as P3P experts tend to be particularly interested in trying out new P3P software. We also asked respondents how often they read web site privacy policies before installing our software. 29% of respondents said they never read privacy policies, 49% said they occasionally read privacy policies, 20% said they read privacy policies at most web sites where they were considering providing personal information, and 2% said they read privacy policies at most or all web sites they visit. Similarly, a November 2001 survey estimated that 31% of

American Internet users spend little or no time looking at web site privacy policies [11]. The Consumer WebWatch study estimated that 10% of American Internet users read privacy policies at all the cites they visit, 25% read privacy policies at "most" of the sties they visit, 48% read privacy policies at "only some" of the sites they visit and 17% never read online privacy policies [10]. Throughout this paper we use the term "privacy policy readers" to refer to those who read privacy policies at most web sites or at most web sites where they are considering providing personal information. Women were nearly twice as likely as men to be classified as privacy policy readers.

We asked respondents how concerned they were about threats to their personal privacy when using the Internet. 53% said they were very concerned, 37% said they were somewhat concerned, 9% said they were not very concerned, and less than 1% said they were not concerned at all. We also asked respondents how concerned they were about three specific web site data practices. 98% of respondents said they were very or somewhat concerned about web sites sharing the information they collected from them with other companies. 96% of respondents said they were very or somewhat concerned about web sites collecting personally identifiable information and combining it with data about their browsing activities to determine their habits, interests or other characteristics. 65% of respondents said they were very or somewhat concerned about web sites collecting information about their web browsing activities and storing it in their records without connecting their name, email address, or other personally identifiable information to it. Women and respondents from outside the US were more concerned about all of these activities and about online privacy in general than other respondents. P3P experts were the least concerned about these activities and about online privacy in general. The level of concern expressed by our respondents is similar to the level of concern found in other studies. For example 87% of American respondents in a 1998 panel study of attitudes about online privacy reported being very concerned or somewhat concerned about threats to their personal privacy when using the Internet [2] and in 1998 Westin estimated that 81% of all Americans were very concerned or somewhat concerned [9].

We asked respondents how concerned they were about web cookies. 36% of respondents said they were very concerned, 43% said they were somewhat concerned, 17% said they were not very concerned, and 3% said they were not concerned at all. Only 1 respondent reported not knowing what a cookie is. We also asked respondents about their familiarity with third-party cookies. 18% of respondents said they had never heard of third-party cookies, 41% said they had heard of them but didn't really know what they are, 37% said they knew a lot about third-party cookies, and 4% said they had helped setup a web site or create policies for a web site that uses third-party cookies. P3P experts were less concerned about cookies and much more knowledgeable about third-party cookies than other respondents.

Overall our respondents appear to be much more knowledgeable about cookies and more concerned about them than most Internet users. The Consumer WebWatch study [10] estimated that only 49% of American Internet users know what a cookie is, while a 2001 Wall Street Journal study [12] estimated that 71% of American Internet users know what a cookie is and 46% believe cookies pose a threat to privacy.[4] The 1998 study of attitudes about online privacy (a panel study, not a random sample) found 52% of respondents were concerned about cookies and 12% said they did not know what a cookie is [2]. It is not surprising that individuals with the knowledge and motivation to download privacy software also have substantial knowledge and concern about cookies. However, it is noteworthy that most of them report having little or no knowledge about third-party cookies, as much policy discussion and cookie filtering efforts have focused on third-party cookies.

4.4 Use and Evaluation of AT&T Privacy Bird

More than half of the respondents (52%) indicated that they were still using the AT&T Privacy Bird software and about a third (34%) said they had used it for more than three months. 14% reported trying it and uninstalling it within a few days. Many of those who were no longer using the software reported that it caused their browser or system to crash.[5] We also received many comments on the survey and in our support email from individuals who wanted to run AT&T Privacy Bird on Unix or Macintosh computers, or with the Netscape or Opera web browsers (which are not supported currently).

We asked respondents to rate AT&T Privacy Bird's usefulness and the likelihood that they would recommend the software to a friend. Women and respondents from outside the US found AT&T Privacy Bird most useful, and were most likely to recommend it to a friend. Indeed these same groups were more likely than others to have learned about it from a friend themselves. Not surprisingly, those who were still using AT&T Privacy Bird after at least one month of use were more likely to recommend it than those who had uninstalled it or who had been using it for less than a month.

A frequent criticism of AT&T Privacy Bird was that a yellow bird appeared at most web sites (because most web sites are not yet P3P-enabled[6]). One user wrote "Biggest concern is the bird in most cases stayed yellow, neither red or green. You have a good idea but until every one plays the same game maybe your idea will have only limited use." Another wrote, "Privacy Bird is a very

---

[4] It is not clear why the Wall Street Journal estimate is so much higher than the Consumer WebWatch estimate a year later. The Consumer WebWatch report cited earlier data indicating a slight *increase* in knowledge about cookies over time, so the difference in estimates is probably not attributable to an actual change in Internet user's awareness of cookies. Rather, the differences between these studies may be due to sampling differences or differences in the way the question was phrased. In any case, the percentage of AT&T Privacy Bird users who reported knowing what a cookie is is significantly higher than either estimate.

[5] Based on our own testing and from the support email we received, we believe that while Privacy Bird appears to be stable on most systems, it is unstable on a small percentage of systems. We need to do further work to identify the cause of the instability and correct it.

[6] In August 2002, Ernst & Young reported that 24% of the top 100 domains and 16% of the top 500 domains visited by US Internet users had been P3P enabled [6].

well designed, easy to use piece of software. Its weakest point is that few sites provide privacy polices for it to evaluate." And another wrote "Great idea, just need the industry to catch up. If there was critical mass, and you could choose to ignore non-secure sites, then the pressure would be there – but today, there's not. But keep going!"

We asked respondents to predict the usefulness of AT&T Privacy Bird if most web sites became P3P enabled. The average usefulness rating on a 5 point scale (where 5 is very useful and 1 is completely useless) jumped from 2.9 for today's web to 4.0 if most web sites were P3P-enabled. Respondents also felt the software would be more useful (4.1) if it was able to block cookies at web sites where the red bird was displayed. P3P experts were less likely than other respondents to find a cookie blocking feature useful, probably reflecting their lower level of concern about cookies. Some respondents commented that they would also like to see AT&T Privacy Bird block spyware and pop-up ads.

We asked respondents to evaluate how easy or difficult it was to use several aspects of AT&T Privacy Bird. On a 5-point scale (where 1 is very difficult and 5 is very easy) the average rating was 4.6 for installation, 3.9 for changing privacy settings, and 3.3 for understanding the policy summary. This indicates that overall respondents found the software easy to use, but suggests there may be room for improving the readability of the policy summary.

4.4.1 Policy summary

We asked respondents to rate the amount of information contained in the policy summary. 64% thought the policy summary contained the right amount of information, while 15% said it contained too much information and 20% said it did not contain enough information. Women tended to be more likely than men to want the policy summary to contain more information (and women also found the policy summary easier to understand than men did). Privacy policy readers and P3P experts were the most likely to prefer more information, while respondents from outside the US were most likely to prefer less information.

We asked respondents if there was anything specific they would like future versions of AT&T Privacy Bird to look for when examining web site privacy policies. None of the respondents suggested any specific aspects of privacy policies to consider or highlight in the policy summary. Some suggested that AT&T Privacy Bird have the ability to check the reputation of web sites or block cookies. One respondent suggested "an option to log which sites I gave info to in case I wanted to exercise rights to enquire on what data held under national Data Protection legislations." The lack of specific suggestions about aspects of privacy policies to consider suggests that the type of information that currently appears in the policy summary is satisfactory and that future efforts should focus on presenting this information in a way that is easier to understand.

We asked respondents how often they reviewed a policy summary. 15% of respondents said they had never done it, 34% said they had done it once or twice, 36% said they had done it several times, and 15% said they had done it 10 or more times. Privacy policy readers and P3P experts reviewed policy summaries more frequently than other respondents. Respondents who used the Internet less than once a day also reviewed policy summaries less frequently.

4.4.2 Privacy settings

While those who changed their privacy settings reported it was relatively easy to do so, 25% of users reported that they never changed their settings after the initial installation, 52% reported changing them once or twice, 21% reported changing them several times, and 2% reported changing them 10 or more times. P3P experts were twice as likely as non-experts to have changed their settings several times, probably because they were experimenting with the software to see what it would do. A few respondents commented that they did not fully understand the privacy setting options, and one suggested "some sort of wizard or dialog that, upon installation, walked you through the setting of your privacy concerns. This could be an opportunity to educate the user as to what the various elements of a privacy policy entail."

4.4.3 Icon and sounds

Many AT&T Privacy Bird users had strong feelings about the optional sound effects. Nearly half of the respondents (45%) reported turning the sounds off completely, while 19% configured AT&T Privacy Bird to play sounds at all web sites and 37% configured the software to play sounds only when a certain color bird appeared. Women and users from outside the US were most likely to turn the sounds on, and P3P experts were most likely to turn them off. One user complained "maybe a little sound would be ok, but than damned crow caw really grates on you after a while" while others praised the sounds: "[I] Like the bird sounds, and animation," "I think he is so cute and I love it when he talks to me," and "Oh, how we love the squawking red crow when we cross paths with Microsoft or MSN!" Other sound haters wrote "One thing that did drive me nuts was the bird chirping which was promptly shut off." and "I was driven almost to a state of collapse, I used to jump when I heard the same bird call in my yard…." Some users suggested a configuration option in which the bird sound would be played only on the first visit to a particular web site rather than every time a page is loaded.

Another annoyance to some users was the fixed position of the AT&T Privacy Bird icon in the browser tool bar. Some suggested placing it in the system tray, while others suggested allowing users to move it to whatever location they preferred. One suggested that the bird flutter across the screen as an alternative to the sound effects.

4.5 Impact on Online Behavior

We asked respondents how often they read privacy policies before and after installing AT&T Privacy Bird. The percentage of people who never or occasionally read privacy policies decreased from 78% before installing the software to 51% after installing the software. While this indicates that even individuals interested enough in privacy to install privacy software rarely read privacy policies, it also suggests that AT&T Privacy Bird has had some effect on raising the awareness of privacy policies among its users. After installing the software, 24% of respondents said they read privacy policies at most web sites where they saw a red bird, 11% said they read privacy policies at most web sites where they were considering providing personal information (down from 20% before installing AT&T Privacy Bird), and 12% said they read privacy policies at most web sites where they saw a red bird and they were considering providing personal information. This suggests that the presence of the red bird is a more significant motivator to read a site's privacy policy then the presence of a

form requesting personal information. Indeed, our intention was that AT&T Privacy Bird users who are concerned about web site privacy policies should not have to read these policies at every site where they are considering providing personal information; rather they should be able to visit green-bird sites without worrying about privacy concerns, but seek more information at red bird sites before providing personal information. The fact that more people seem to be reading privacy policies at some sites, but fewer report reading privacy policies at every site where they are considering providing personal information suggests that many users are using the AT&T Privacy Bird as we intended.

We asked users whether they had learned anything about web site privacy policies as they used AT&T Privacy Bird that caused them to change their online behavior. 88% indicated that their use of AT&T Privacy Bird had resulted in some change in behavior. About 37% of respondents reported that they fill out fewer forms online, 37% reported taking advantage of opt-out opportunities, 29% reported that they stopped visiting some web sites, and 18% reported comparing privacy policies at similar sites and trying to frequent the sites with the better privacy policies. Women and non-US respondents were most likely to stop visiting some web sites or fill out fewer forms online, while US respondents and privacy policy readers were most likely to take advantage of opt-out opportunities. Men, non-US respondents, privacy policy readers, and P3P experts were most likely to compare privacy policies. One respondent commented "Basically, I use Privacy Bird like a warning light. Whenever it's red I treat the website as hostile and am extra careful about the information I provide and activities I perform there."

Some users wrote in some additional ways AT&T Privacy Bird has impacted their behavior. One wrote about sending email to sites where the red bird appeared, and another claimed "I told one mutual fund web site about Privacy Bird's findings, and they improved their pages because of it!" Another wrote about doing a "regular clean up of all cookies," and another wrote about changing browser cookie settings to block most cookies. Others wrote about using AT&T Privacy Bird to educate themselves about online privacy. Several respondents indicated that they were implementing P3P policies or studying P3P for a company or as part of an academic project. One such respondent wrote that AT&T Privacy Bird was helpful for learning "about the technical problems involved in maintaining a useful privacy policy."

We asked respondents whether they found that most web site privacy policies matched their expectations. 5% of respondents said that most privacy polices were better than they had expected, 40% said they were worse than they had expected, 33% said they were about what they had expected, and 22% said they had not been sure what to expect. Not surprisingly, privacy policy readers and P3P experts were most likely to find that privacy policies matched their expectations. Respondents from outside the US were also somewhat more likely to find that privacy policies matched their expectations, despite the fact that they were slightly less likely to read privacy policies.

Some respondents commented that they did not find web site privacy policies trustworthy. One respondent wrote "the information provided in the site privacy policies is like listening to a corporate board all double talk at the same time…. For that reason, the privacy policy format is not trustworthy." Several respondents suggested adding a mechanism for filing complaints against web sites and some recommended that AT&T Privacy Bird look each web site up in a database of privacy complaints automatically. One respondent suggested "It would be helpful to be able to link to a site/database that would contain information about whether a site actually follows their stated privacy policy or is known to ignore it and sell their gathered information." We did not ask any questions about privacy seals from TRUSTe, BBBOnline, or other organizations (which AT&T Privacy Bird displays in the policy summary at sites that have them) and none of our respondents mentioned them in their comments. These seals were developed primarily to improve the trustworthiness of online privacy policies.

We asked respondents whether they would be likely to purchase an item from the web site that had the best privacy policy if it was possible to identify which of the web sites that offered that item had the best policy. 33% of respondents said they would always purchase the item from the site with the best privacy policy; 54% said they would probably purchase the item from the site with the best privacy policy as long as the price and services offered by that site were about the same as at other sites; 6% said they would always purchase the item from the site with the best price, regardless of its privacy policy; and 7% said they do not plan to make online purchases. Women and respondents from outside the US were most likely to purchase the item from the site with the best privacy policy.

5. DISCUSSION

The AT&T Privacy Bird user study provided many interesting insights that will not only help us improve the AT&T Privacy Bird software, but also shed some light on the use of P3P user agents and other privacy software. Very little work has been done to study how people use privacy software, and little is known about how to make privacy concepts accessible to end users.

From the responses to several of our survey questions it appears that women liked the AT&T Privacy Bird software more than men did. This is particularly interesting because women appear to represent a minority of AT&T Privacy Bird users. It would be interesting to learn more about why women like the software better, as well as why they are less likely to download it than men. It is also noteworthy that people outside the US seemed to like AT&T Privacy Bird better than US users did. This is especially interesting because the software was designed with US users in mind.

This study indicates that the aspect of the AT&T Privacy Bird user interface that is probably in most need of improvement is the policy summary. Other surveys have highlighted consumer concerns about the readability of privacy policies and desire for standardized privacy policy formats [11]. Getting this standardized format "right" will take some additional work. Allowing users to choose from shorter and longer formats might also be helpful.

The results of this study highlight the potential that privacy software has as a tool to educate consumers about privacy. A large proportion of AT&T Privacy Bird users began reading privacy policies more often and being more proactive about protecting their privacy as a result of using this software. Unfortunately, the usefulness of P3P user agents is severely limited by the number of web sites that have implemented P3P. In addition, as long as yellow-bird sites are in the majority, AT&T Privacy Bird in and of

itself does not act as an incentive to sites to become P3P-enabled unless they have privacy policies that will generally result in green birds. Our survey results also suggest that if it becomes easier to compare privacy policy across e-commerce web sites, a significant group of consumers would likely use this information in their purchase decisions. As more web sites become P3P-enabled, P3P would be a useful feature to add to search engines and online price comparison services.

6. REFERENCES

[1] Cranor, L. and Reagle, J. Designing a Social Protocol: Lessons Learned from the Platform for Privacy Preferences Project, in J.K. MacKie-Mason and D. Waterman (eds.) *Telephony, the Internet, and the Media*. Lawrence Erlbaum Associates, Mahwah, NJ, 1998. http://www.w3.org/People/Reagle/papers/tprc97/tprc-f2m3.html

[2] Ackerman M.S., Cranor, L.F., and Reagle, J. Privacy in E-Commerce: Examining User Scenarios and Privacy Preferences, in *Proceedings of EC'99* (Denver CO, November 1999), ACM Press, 1-8.

[3] Cranor, L. *Web Privacy with P3P*. O'Reilly & Associates, Sebastopol CA, 2002.

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

[5] Cranor, L., Langheinrich, M., and Marchiori, M. *A P3P Preference Exchange Language 1.0 (APPEL1.0).* World Wide Web Consortium Working Draft, April 2002. http://www.w3.org/TR/WD-P3P-Preferences.

[6] Ernst & Young. P3P Dashboard Report, August 2002. http://www.ey.com/global/download.nsf/US/P3P_Dashboard_-__August_2002/$file/P3PDashboardAugust2002.pdf

[7] Esposito, D. Browser Helper Objects: The Browser the Way You Want It, MSDN Library, January 1999. http://msdn.microsoft.com/library/default.asp?url=/library/en-us/dnwebgen/html/bho.asp.

[8] Goldfeder, A. and Leibfried, L. Privacy in Internet Explorer 6. MSDN Library, October 2001. http://msdn.microsoft.com/library/default.asp?url=/library/en-us/dnpriv/html/ie6privacyfeature.asp.

[9] Harris, Louis and Associates and Westin, A.F. E-commerce & Privacy: What Net Users Want. Privacy & American Business, Hackensack NJ, 1998.

[10] Princeton Survey Research Associates. A Matter of Trust: What Users Want from Web Sites. Consumer WebWatch, April 2002. http://www.consumerwebwatch.org/news/1_abstract.htm.

[11] Privacy Leadership Initiative. Privacy Notices Research Final Results. Conducted by Harris Intereactive, December 2001. http://www.ftc.gov/bcp/workshops/glb/supporting/harris%20results.pdf.

[12] Wall Street Journal. Exposure in Cyberspace. Conducted by Harris Interactive, 21 March 2001, p. B1.