

Can user agents accurately represent privacy notices?

Lorrie Faith Cranor^{*} and Joel R. Reidenberg[†]

Abstract

The Platform for Privacy Preferences (P3P) is a W3C specification that provides a standard computer-readable language for web sites to encode their privacy policies. This standardization allows for the creation of web browsers and other user agents that can display privacy warnings and signals that are meaningful to users or that automate actions in accordance with user instructions. This paper shows that P3P user agents will necessarily include judgmental design decisions and that the accuracy of the P3P user agent interactions becomes a critical matter. The accuracy of P3P user agents raises significant legal concerns about privacy agreements, inadvertent deception, and liability. The technological mediation designed to make it easier for users to understand the privacy practices of web sites risks adding ambiguity, confusion and legal uncertainty. This paper argues that one way to avoid having privacy practices represented inaccurately by P3P user agents is to certify P3P user agents for the accuracy of their representations of web site P3P policies. While there are some things that P3P user agents might do that would be readily identified as inaccurate or misleading, there is a large gray area in which user agents might present factual information side-by-side with subjective judgments. These judgments may be deemed misleading by some people but not others. Many of the issues raised here are new, but this is not likely to be the last time that these issues arise. As work progresses on computer-mediated search and negotiation technologies with a wide variety of applications, these issues are likely to surface repeatedly. In this paper we explore these issues and suggest some possible solutions, as well as a number of open questions.

^{*} Lorrie Faith Cranor <<http://lorrie.cranor.org/>> is a Principal Technical Staff Member at AT&T Labs-Research. She is chair of the P3P Specification Working Group at the World Wide Web Consortium and one of the creators of the AT&T Privacy Bird software. This paper represents her personal views.

[†] Joel R. Reidenberg <<http://reidenberg.home.sprynet.com>> is a Professor of Law at Fordham University School of Law. He participated as an invited expert in some of the deliberations of the P3P Working Groups.

Introduction

An important role of many software programs is to present voluminous or complicated information in an accessible format. Computer users frequently interact with systems that aggregate large volumes of information or convert computer-readable information or expert notation into more generally understandable language. For example, grocery store scanners convert bar codes into item descriptions and prices and display those descriptions and prices to the cashier and customer. Scanners produce aggregate purchase reports that are used by the store manager when restocking the shelves and creating financial reports. These systems are usually accurate, but they do occasionally make mistakes. Scanning errors are generally fairly easy for cashiers or customers to detect. Reporting errors too, can often be detected when balance sheets produce unexpected results or stocking problems persist. But, other types of computer systems may display inaccurate information that is much harder to detect.¹

When individuals delegate decision-making to a computer, they not only trust the computer to communicate information accurately, but they also trust the computer to assess correctly the information and to make decisions correctly, or at least in the same way as they would. Individuals treat the computer as their “agent” and thus many computer programs that make decisions on behalf of users or assist users in accessing information are referred to as “user agents.”² In practice, users often have only a vague understanding of how their user agents make decisions and users trust, somewhat naively, that their agents will make the same decisions the users would. Thus, parents who purchase software that promises to protect their children from adult content on the Internet often are surprised when that software blocks access to inoffensive content or fails to block access to offensive content. Parents typically do not research thoroughly the factors that determine a particular filter’s blocking decisions, and detailed information about filtering criteria are not always readily available to consumers.³

¹ Peter G. Neumann catalogs hundreds of examples of computer failures, some of which have had catastrophic results, in *Computer Related Risks* (Addison Wesley, 1995).

² The terms “agent” and “user agent” have many different meanings. In the computer science research literature, the terms generally refer to software programs that are able to perform tasks on behalf of their users often involving interaction with other agents, observing and responding to user behavior, or adapting to changing situations. These agents are often referred to as “intelligent” because they include computer code allowing them to make decisions based on a dynamic set of inputs, sometimes with unpredictable results. In contrast, the P3P user agent examples we discuss in this paper are static and predictable: they are programs that display information or take defined actions on the basis of a well-specified set of possible inputs. What these P3P user agents and intelligent agents have in common is that users delegate tasks to the agents and users believe they can trust the agents. In the future, we may see P3P user agents that also take on features of intelligent agents. Some of the issues raised in this paper will be further complicated when the agents involved operate dynamically.

³ In 2000 the COPA Commission concluded (III.3) “The current lack of information about how well [filtering] technologies work, and lack of transparency about what they might block, is a major hurdle for their adoption by families or caregivers. “ in *Final Report of the COPA Commission* (October 20, 2000), at <http://www.copacommission.org/report/>. For a discussion of bias in content filtering, see also Jonathon

Traditionally, most computer systems that individuals relied on for information or decision-making were closed systems; a single entity controlled both the information that went into the system and the user interface that presented this information to end users. For example, companies that maintained large databases of information (journal article indexes, financial data, legal documents, telephone directories, etc.) typically provided the user interface for the system and thus controlled the way the information would be presented to users. Any user agent software for interacting with these systems was provided by the companies that maintained the databases. For example, Sabre provided travel agents with the software to access the airline reservation system and Lexis and Westlaw each provided proprietary software for clients to access their information databases.

More recently, open standards have been developed in a variety of areas to allow for the interoperability of components developed by different companies.⁴ For example, metadata⁵ standards for journal articles (which include standard formats for elements such as title, author, year of publication, and abstract) allow publishers to open up their databases (perhaps for a fee) to other companies that, in turn, may wish to mine this data and present it in an alternate format to end users.⁶ Open standards often result in the development of competing products and services that use the same underlying data but with different interfaces. The World Wide Web itself is built on open standards. Individuals create content using HTML and other languages, and companies build web browsers that are able to display this content. There are differences in the way web browsers render the same content (some due to browser bugs but others due to implementers making different design decisions).⁷ If not explicitly specified in the content, for example, implementers can decide what colors to use to display page backgrounds and text, and whether or not to underline links. Most browsers allow users to specify their own preferences about the font size, and some also allow users to specify other preferences such as background and text colors, and whether or not to display graphics. This flexibility allows users to adapt content to meet their own needs, which may include compensating for poor vision, a small display, or some unusual circumstance unanticipated by the content provider.⁸ However, the flexibility means that the content

Weinberg, *Rating the Net*, 19 *Hastings Comm/Ent LJ* 453 (1997); Lawrence Lessig, *What Things Regulate Speech*, 38 *Jurimetrics* 629 (Summer 1998).

⁴ See Bruce Perens, *Open Standards: Principles and Practice*, at <http://perens.com/OpenStandards/Definition.html>.

⁵ The term “metadata” is used to refer to data about data. For example, title, author, publisher, and number of pages are metadata about a book.

⁶ See the Open Archives Initiative web site, at <http://www.openarchives.org/>.

⁷ A number of web sites offer comparisons of the way different web browsers render content. See for example <http://browsercomparison.esmartweb.com/>.

provider no longer controls and often has no way of knowing exactly how the content will be displayed to users.

While companies typically like to retain as much control as possible over how their information is presented, the adverse implications of web browsers changing the size and color of their text are usually fairly limited. However, more serious concerns arise when content is displayed out of context or manipulated so that the content of the text may actually be modified. Disputes about the practice of one web site “framing”⁹ another site’s content and about services that edit web site content (for example to remove adult language)¹⁰ have become common.

As new standards are developed to encode a wide variety of information as computer-readable meta data, the potential exists for disputes over how that information is converted into a human-readable format and displayed to end users. User interface implementers will have to make decisions about the precise language to use, the organization of the information, the information to include or exclude in a display, and whether or how to summarize the information. When the display of metadata may have legal implications or may influence financial decisions, or when the metadata may be used as part of an agent-mediated negotiation, the display is likely to be especially controversial. Questions arise about the legal implications of metadata displayed inaccurately or in a misleading way, as well as what constitutes an accurate display in the first place.

In this paper, we discuss the Platform form Privacy Preferences (P3P) as an example of a specific metadata standard for which several different user agents are already available. We describe the differences among these user agents and some of the questions that have arisen about them. We explore the legal implications of user agents that do not accurately represent privacy disclosure statements. We then examine some approaches to addressing the problems posed by such user agents. Finally, we discuss the implications for technologies beyond P3P.

⁸ The World Wide Web Consortium’s Web Accessibility Initiative has developed guidelines for developing web browsers and web content that will be accessible to users with a variety of needs. See <http://www.w3.org/WAI/>.

⁹ See e.g. Bruce P. Keller, *Condemned to Repeat the Past: The Re-emergence of Misappropriation and Other Common Law Theories of Protection for Intellectual Property*, 11 *Harv. J. Law & Tec* 401 (1998)(analyzing framing)

¹⁰ See e.g. Chris Oakes, *The Web’s New Graffiti?*, *Wired*, June 9, 1999, available at <http://www.wired.com/news/technology/0,1282,20101,00.html>; *The Proxomitron—Universal Web Filter*, <http://www.webattack.com/download/dlproxo.shtml>. In an unusual situation, Beaver College reportedly considered changing the college’s name because of problems with content filtering software. See Craig Bicknell, *Beaver College Not a Filter Fave*, *Wired*, Mar. 27, 2000 available at <http://www.wired.com/news/politics/0,1283,35091,00.html>

The Platform for Privacy Preferences

P3P is a specification of the World Wide Web Consortium (“W3C”) and provides a standard computer-readable language for web sites to encode their privacy policies.¹¹ This standardization allows for the creation of web browsers and other user agents that can display meaningful privacy warnings and signals or that can automate actions in accordance with user instructions. For example, P3P user agents may block cookies¹² based on user preferences and web site P3P policies, may provide an indication whether a site’s policy matches a user’s stated preferences, may provide a privacy rating for a web site, or may display a translation of the P3P policy into a human-readable language. Perhaps, most importantly, users may choose whether or not to provide personal information to a web site, or even whether to interact with a particular web site at all, on the basis of the information provided by their P3P agent.

The accuracy of the P3P user agent, thus, becomes a critical matter. The P3P specification places few requirements on how user agents should display information about web site privacy policies.¹³ This flexibility is intended, in part, to promote creative approaches from user agent implementers that will meet the needs of users. As a result, P3P user agents will necessarily include judgmental design decisions.¹⁴ While computer mediated interactions often introduce bias¹⁵ in the way an interaction is represented, P3P is one of the first such interactions in which companies may risk having contractual terms represented inaccurately by third-party user agents. When user agents present inaccurate or extremely biased information (accidentally or even intentionally) the user agent misrepresents a web site’s privacy practices. Users then proceed with actions such as disclosing or withholding personal information based on a false understanding of the web site’s practices. These computer-mediated agreements concerning the use of personal information may then be invalid because of the misrepresentations or may bind a party to unintended commitments. Thus, the technological mediation designed to make it easier

¹¹ See *The Platform for Privacy Preferences 1.0 (P3P1.0) Specification*, W3C Recommendation 16 April 2002, <http://www.w3.org/TR/P3P/>.

¹² Cookies are bits of text transferred between web sites and web browsers that are used to help sites keep track of a user’s preferences or activities at a particular web site. Cookies can provide useful functions for users, such as remembering the contents of an electronic shopping cart. Cookies can also be used to build profiles of individual users against their wishes. Many web browsers include configuration options that allow all cookies to be blocked, however, users often find that complete blocking prevents them from accessing some web sites and services. Browsers and other tools that can selectively block cookies allow users to benefit from the cookies they find useful and block the cookies they find privacy invasive.

¹³ *Supra* note 11, at 1.1.4.

¹⁴ See e.g. Joel R. Reidenberg, *Lex Informatica: The Formulation of Information Policy Rules through Technology*, 76 *Texas L. Rev.* 553 (1998); Lawrence Lessig, *Code and Other Laws of Cyberspace* (1999)

¹⁵ See Batya Friedman and Helen Nissenbaum, “Bias in Computer Systems” in Batya Friedman, ed. *Human Values and the Design of Computer Technology*, Cambridge University Press, 1997, p. 21-40.

for users to understand the privacy practices of web sites risks adding ambiguity, confusion and legal uncertainty.

P3P Policies

P3P policies are encoded in a computer-readable language called XML.¹⁶ The P3P specification defines dozens of terms called *XML elements* that can be composed to form a P3P policy. The specification includes a definition for each XML element and a set of rules that govern how the elements can be combined to form a syntactically valid P3P policy.

The following are examples of some information about web site privacy policies that can be described by XML elements in a P3P policy:

- The legal entity making the representation of the privacy practices contained in the policy
- Whether the site provides access to various kinds of information
- Dispute resolution procedures that may be followed for disputes about a services' privacy practices or a protocol violation
- Purposes for data processing
- The kind of data to be transferred or inferred
- The legal entity or domain beyond the service provider and its agents where data may be distributed
- The kind of data retention policy that applies to the data

While some of these elements contain fields that allow web sites to fill in free form information (for example the postal address of a web site), most of the information will be represented by one or more sub-elements. Web site administrators must select the sub-elements that best represent their practices. This is typically accomplished by answering a series of multiple-choice questions. For example, when describing the purpose of data collection, web site administrators must select from 12 possible choices. Each choice is defined in the P3P specification. The definitions are generally one to three sentences long, and designed primarily for use by web site administrators rather than end users. For example, the definition of one type of purpose called "pseudonymous decision," is represented by the XML element <pseudo-decision/>:¹⁷

¹⁶ XML stands for eXtensible Markup Language. See *XML in 10 Points* at <http://www.w3.org/XML/1999/XML-in-10-points>.

¹⁷ *Supra* note 11, at 3.3.4.

Information may be used to create or build a record of a particular individual or computer that is tied to a pseudonymous identifier, without tying identified data (such as name, address, phone number, or email address) to the record. This profile will be used to determine the habits, interests, or other characteristics of individuals *to make a decision that directly affects that individual*, but it will not be used to attempt to identify specific individuals. For example, a marketer may tailor or modify content displayed to the browser based on pages viewed during previous visits.

This definition is precise enough that an individual familiar with a web site's data practices should be able to determine whether or not that site engages in the specific type of data activity. However, the term "pseudonymous decision" is an invention of the drafters of the P3P specification, and thus is not a term that will have obvious meaning if used without explanation. The definition also contains terms such as "profile" and "pseudonymous" that are not likely to be meaningful to many web site visitors. Thus, the P3P specification advises "user agent implementers may each make their own choices about what words and symbols to present to users to provide information about a Web site's privacy policy. Implementers need not use the definitions found in this specification verbatim in their user interfaces."¹⁸

The P3P specification allows web site administrators to choose from among three approaches when describing the type of information their sites collect. Administrators may use one of these approaches or any combination of them. The simplest approach is to include elements in the P3P policy that represent one or more of seventeen categories of information defined in the specification. These categories include among others: physical contact information, online contact information, unique identifiers, and preference data. By using these data categories, administrators can avoid having to enumerate every piece of data they might request with every form on their sites. However, for sites that collect only a limited amount of data, this approach may leave users with the impression that the sites collect far more data than they really do. For example, if a site declared that it collected physical contact information, visitors would correctly assume that the site might collect name, address, phone number, and other information. A site that collects only visitors' zip codes, might wish to make a more specific statement. Thus, P3P also allows sites to enumerate the specific data they collect using elements defined in a "P3P base data schema." The schema includes many of the types of data typically collected by web sites. In addition, P3P allows sites to declare their own data schema in order to enumerate data elements not included in the "base data schema."

P3P User Agents

P3P user agents are software tools designed to fetch P3P policies, interpret them, and display them or perform actions based on those policies on behalf of an Internet user. P3P user agents can be built into web browsers, implemented as software that users can

¹⁸ *Supra* note 11, at 1.1.4.

download and add onto their browsers, implemented as stand-alone applications, or built into services that users can access over the Internet.

Microsoft Internet Explorer 6

Microsoft implemented a P3P user agent as part of its Internet Explorer 6 (IE6) web browser, publicly released in August 2001. This P3P user agent can be configured to make cookie-blocking decisions on the basis of a web site's P3P policy.¹⁹ Microsoft offers six cookie-blocking settings, plus a special language that sophisticated users might use to create their own customized settings.²⁰ When users configure their cookie settings they can move a slider bar to see a few lines of explanation as to what each of the six standard settings does. These explanations include phrases like "cookies that use personally identifiable information without your implicit consent" to represent a set of possible element combinations from a P3P policy. The IE6 help files do not offer further explanation, but the Microsoft web site does explain the full details.²¹ The web site explains that the phrase "cookies that use personally identifiable information" refers to cookies that contain or link to one or more of four possible "unsatisfactory" categories of data and are also either used for one or more of five possible "unsatisfactory" purposes or shared with one or more of four possible categories of "unsatisfactory" recipients. The phrase "without your implicit consent" modifies the previous phrase to exclude any purposes or recipients for which an opt-out is provided. The web site offers several tables that make clear the precise meaning of each of the six settings.

The IE6 cookie settings would be improved by the use of phrases that might be more transparent to end users, and by integrating the definitions of each setting into the help files provided with the software. However, for sophisticated users who do some research, these precise definitions for each setting are available and can eventually be found.

IE6 also features a "View Privacy Report" option that produces a human-readable translation of a site's XML P3P policy. When a user selects this option, the browser fetches the site's P3P policy and copies selected XML elements from the policy into an HTML template. For each P3P XML element in the template, the browser substitutes the English-language definition from the P3P specification, with only minor editorial changes. The resulting HTML file is then displayed to the user. This strategy, in theory, results in a very accurate representation of a site's P3P policy because the translation uses almost the exact definitions from the P3P specification. However, the translation also

¹⁹ More precisely, IE6 cookie blocking decisions are based on a compact representation of a site's P3P policy called a "compact policy." For details of the cookie blocking options offered by IE6 see *Privacy in Internet Explorer 6* at <http://msdn.microsoft.com/library/default.asp?url=/library/en-us/dnpriv/html/ie6privacyfeature.asp>.

²⁰ See *How to Create a Customized Privacy Import File* at <http://msdn.microsoft.com/library/default.asp?url=/workshop/security/privacy/overview/privacyimportxml.asp>

²¹ *Supra* note 19

results in a representation that is very lengthy and not necessarily all that meaningful to end-users. Furthermore, because IE6 includes only selected XML elements from the P3P policy in the HTML template, some statements made in a P3P policy will be omitted from the human-readable version presented to the user.

Some of the omissions in the IE6 Privacy Report are arguably benign. At worst, omissions may leave users with some unanswered questions. At best, the omissions may improve the readability of the privacy policy by leaving out details that many users may not care about anyway. For example, the IE6 Privacy Report omits contact information for the web site and omits the text of human-readable explanations encoded in P3P “consequence” elements. However, the IE6 Privacy Report includes a link to the site’s full human-readable privacy policy, which would typically contain this information.

Some omissions, however, may not be benign. For example, in one section of the template, IE6 lists the types of data collected by a web site under the heading “What kind of information does this Web site collect?” When creating a privacy report, IE6 selects all of the data category elements from the P3P policy and places those elements under this heading. However, the report omits any data elements referred to explicitly by name from the P3P base data schema or from data schemas defined by a web site. As a result, users who request privacy reports at a web site that declares explicit data elements rather than data categories will be left with the impression that the site collects no data. As sites are permitted to use a combination of approaches when declaring the types of data collected, a site that declares some data by category and some data explicitly will appear in its IE6 privacy report to collect only a subset of the data actually collected. A user making decisions about whether or not to interact with a web site on the basis of this privacy report would be misled.²²

Netscape Navigator 7

In May 2002, Netscape released a preview of its Navigator 7 web browser, which includes P3P functionality similar to that found in IE6.²³ Netscape followed a similar approach to Microsoft and Navigator 7 allows users to make cookie blocking decisions on the basis of P3P policies. Navigator 7 also provides the ability to translate XML P3P policies into an English “policy summary.” Netscape uses similar terminology to describe the cookie blocking settings, but provides no information in help files and does

²² Because users are generally aware of the information they input at web sites, the risk that users will make faulty decisions as a result of this problem may be lower than other types of misleading reports such as the omission of purpose elements. Nevertheless, there are situations where information is collected without users typing it in (for example, cross-site cooking tracking and clickstream data), and thus users might not be aware that data is collected if the practice is omitted from the privacy report. Furthermore, astute users may wonder why data a site requests in its web forms is not listed in the privacy report and may conclude, incorrectly, that the site has omitted critical information from its P3P policy. This problem was reported to Microsoft by members of the P3P Specification Working Group shortly after IE6 was released; however, as of August 2002 the problem has not been corrected.

²³ See Netscape 7.0 Preview Release 1 Release Notes at <http://wp.netscape.com/eng/mozilla/ns7/relnotes/7pv1.html>.

not provide precise definitions of the terms on the Netscape website. In addition, the Netscape policy summary also omits data elements listed explicitly by name.

Unlike the IE6 privacy report, the Netscape policy summary replaces the definitions from the P3P specification with abbreviated and more user-friendly phrases. For example, the “pseudonymous decision” element is shown in the Netscape policy summary as “Determine your habits and interests to make a decision that affects your online experience.” This simplified phrase does not capture the pseudonymous aspect of this purpose. Overall, this approach makes the resulting policy summary shorter and more readable, but at the risk that the Netscape terminology may not capture accurately all the nuances of the P3P specification definitions. For example, Netscape’s translation of the pseudonymous decision elements fails to capture the pseudonymous nature of this purpose.²⁴

AT&T Privacy Bird

In February 2002, AT&T released a public beta of Internet Explorer add-on software called the AT&T Privacy Bird. Unlike the Microsoft and Netscape P3P implementations, the Privacy Bird beta does not block cookies. Instead, this user agent fetches P3P policies automatically from every site a user visits and displays a visual indication (with optional sound effects) of the match between a web site’s privacy policy and the user’s preferences. A green bird icon indicates a match, a red bird icon indicates a mismatch, and a yellow bird icon indicates that the site does not have a P3P policy. Users can click on the bird to retrieve a human-readable translation of the site’s XML policy called the policy summary.

The Privacy Bird configuration window allows users to indicate which of 12 potential privacy concerns should trigger a warning. While each trigger is described in the configuration window with a short description such as “Warn me at web sites that use information that personally identifies me to determine my habits, interests, or other characteristics,” the accompanying help files provide a more detailed explanation of each warning trigger and a list of specific XML P3P elements that can trigger each warning. The description in the configuration window is too brief to precisely identify the conditions that will trigger each warning, but the help files provide unambiguous information for those users who want a more precise explanation. Users who wish to configure warning triggers beyond the 12 offered choices can import their own settings files.

²⁴ Netscape does provide a link “for more information” about this purpose. A user who clicks on this link receives the message: “We may use the information that you provide to us to determine your habits, interests or characteristics, without linking them to you personally, and based on those habits, interests, or characteristics make a decision that directly affects you. For example, this web site may decide to show the user a modified web page depending on whether you have accessed the page before.” This explanation is a curious paraphrase of the definition from the P3P specification, and seems to have trouble deciding whether to refer to users in the second or third person and whether to refer to web sites in the first or third person.

The Privacy Bird policy summary is similar to the IE6 privacy report and Netscape policy summary in that the Privacy Bird includes only a subset of the information available in a full P3P policy. Unlike IE6 and Netscape, however, the Privacy Bird does not omit specific data elements from its list of collected data. Similar to Netscape, Privacy Bird replaces the full definitions of XML elements with shorter, user-friendlier phrases at the risk that these phrases may provide an incomplete representation of the elements they describe. For example, Privacy Bird represents the pseudonymous decision element with the phrase “To make decisions about what content or ads you see at the web site, etc. without identifying you.” The Privacy Bird policy summary also includes an explanation of any mismatches that occur between a user’s preferences and a site’s policy.

Potential for Confusion

User agent implementers who want their products to be useable by non-expert end users need to find ways of simplifying the terms and definitions in the P3P specification. As a result implementers often bundle together multiple elements and, then, describe those elements with a single term, with paraphrased definitions, or with jargon replaced by more readily accessible terms.²⁵ This simplification process is essential to usable product design, but reduces the precision of the P3P terms and may introduce some confusion.

As noted earlier, P3P user agents may provide confusing and potentially misleading information when describing settings with insufficient detail or presenting summary information about P3P policies that is inaccurate, imprecise, or incomplete. The companies that developed the three user agents described here (presumably) made good faith efforts to represent accurately P3P policies. As all three of these companies had participated in the development of the P3P specification, it is also likely that any bias in their implementations is likely to be similar to any bias present in the P3P specification itself.²⁶ However, future P3P user agent implementations may be designed purposely to misrepresent privacy policies, or to represent them with a particular bias that arguably results in an inaccurate representation.

Some user agent implementers may wish to convey information that supports their particular agenda. For example, a marketing company might develop a P3P user agent that described telemarketing with positive sounding terms such as “personalized home shopping opportunity,” while a privacy watch-dog group might label any site that does telemarketing as an “evil privacy invader.” Hopefully such overt bias would be obvious to users, who could decide whether the implementer’s agenda matches their own. More troublesome, perhaps, would be a user agent in which the bias is more difficult to detect and would more likely result in users misunderstanding a web site’s privacy policy. Imagine, for example, a user agent that provided a numeric rating—perhaps on a scale of

²⁵ See Lorrie Faith Cranor, *Web Privacy with P3P* (O’Reilly and Associates, 2002), chapter 14.

²⁶ While P3P was designed to present factual information about privacy practices without judging these practices, the fact that W3C decided to pursue P3P as a technical solution represents a bias towards a particular approach to addressing privacy concerns. See Lorrie Faith Cranor, “Bias and Responsibility in ‘Neutral’ Social Protocols” *Computers and Society*, September 1998, p. 17-19.

1 to 10—to describe a web site’s P3P policy. The agent would necessarily make subjective judgments when determining the rating for each web site. Unless the strategy for assigning ratings was clearly explained to users, the user’s assumptions about the meaning of a 10 rating versus a 1 rating may not match the user agent’s actual behavior.²⁷

Some companies have expressed concern about legal issues surrounding P3P policies that will be displayed to a user in a format not under corporate control.²⁸ To date, however, we have found few documented objections to the way implementers have chosen to display P3P policies in their currently available implementations. Beyond the problems with omitted data already discussed, and concerns about the readability of some of the phrases, only a few other specific concerns have been raised. One such concern relates to user agent wording. In particular, the issue is whether wording might suggest to users that web sites will routinely collect all of the types of data and use the collected data for all of the purposes mentioned in their P3P policy, or whether the wording suggests only that the data might be collected and might sometimes be used for these purposes.²⁹ Many of the definitions in the P3P specification contain the word “may” to express the later idea.³⁰ However, the word “may” is omitted in some of the information displayed in user agent policy summaries. For example, the beta release of the Privacy Bird includes the heading “How your information will be used.” Some companies have suggested that this heading be changed to “How your information may be used.”

²⁷ Indeed in 2000 a company called Enonymous hired people to rate thousands of web site privacy policies on a scale of 1 to 4. They provided user agent software that informed users about the ratings of the web sites they visited. This rating system proved controversial because the criteria Enonymous used for assigning a rating was not well documented and not applied uniformly by their human raters. A well-known privacy advocate praised the system, and then retracted his statement when he discovered that his own web site mysteriously had been assigned a poor rating. See Declan McCullagh, *Odd Privacy Ratings Exposed*, *Wired News* (12 April 2000) at <http://www.wired.com/news/print/0,1294,35587,00.html>.

²⁸ In a 15 October 2001 memo to the P3P Specification Working Group, Cheryl Charles, Senior Director of BITS (The Technology Group for The Financial Services Roundtable), expressed the concern “There are potential conflicts between how a P3P implementation characterizes site behavior and a company’s own plain language privacy policy, which could appear to lead to charges of bad faith,” at <http://lists.w3.org/Archives/Public/www-p3p-public-comments/2001Oct/0015.html>.

²⁹ Representatives of several companies have raised this concern in personal conversations with one of the authors.

³⁰ The word “may” was added to many of the definitions in the P3P specification after one company complained about this problem. They actually proposed that the P3P specification include separate “will” and “may” elements so that companies could distinguish between routine and occasional practices. However, the consensus of the working group was that such a distinction would not be useful to end users. See section 3.3.4 of 18 October 2000 P3P 1.0 Working Draft at <http://www.w3.org/TR/2000/WD-P3P-20001018/> (“Note, that the working group discussed at length the possibility of allowing sites to distinguish between purposes they may engage in and purposes they will engage in. The consensus of the working group was that such a distinction is not necessary. However, some members disagreed with this conclusion stating: Yes, no and may all need to be response options in the vocabulary. If no and may are the only options, then the meaning of may is corrupted to equal yes....”)

Legal Implications

The accuracy of P3P implementations raises a number of significant legal issues for users, web sites and implementers. At the outset, inaccurate representations by user agents of a web site's privacy policy can undermine the main purpose of the P3P standard, namely the enabling of automated decisions based on notice of information practices and consent by the user. The validity of automated actions or agreements between web sites and users for the use of personal information is jeopardized by inaccurate implementations. Yet, even if agreements and automated actions are considered valid, an inaccurate translation of a web site's policy creates ambiguity as to the applicable privacy protections for the personal information. Are the applicable protections those understood by the user based on the user agent's unfaithful translation or those actually disclosed by the web site? From the user's perspective, the inaccurate portrayal of a site's policy might appear as fraud or deception. Should the web site be responsible for this problem or might the implementer bear liability? Web sites are also sensitive that their privacy policies not be portrayed to users in a false light. If an implementation unfairly tarnishes a web site, could the implementer be held liable for defamation?

Although established legal principles offer some guidance for these questions, the novelty of P3P and its user implementations do not fit precisely into the traditional doctrine.

Validity of Privacy Agreements

The P3P standard contemplates an agreement predicated on notice and consent between the web site and the user over the use of personal information. American law recognizes that parties may form binding contracts through electronic means and recent legislation specifically supports electronic contracting.³¹ U.S. courts have enforced agreements formed electronically including "clickwrap" and "browsewrap" agreements when users have had the opportunity to review the contractual terms and were required to signify in some fashion their agreement with the terms.³² These agreements can be binding even though the users have not actually read the terms. But, in one prominent case, *Specht v. Netscape Communications*,³³ the court refused to enforce a "browsewrap" agreement because the contractual terms were only available through a link that users did not have to follow or otherwise signify that they had an opportunity to view the terms. In essence,

³¹ See e.g. Electronic Signatures in Global and National Commerce Act, 15 U.S.C. §§ 7001-7031 available at http://www.ecommerce.gov/ecomnews/ElectronicSignatures_s761.pdf ; N.C.C.U.S.L., Uniform Electronic Transactions Act (1999) available at <http://www.law.upenn.edu/bll/ulc/fnact99/1990s/ueta99.htm>

³² See e.g. *Moore v. Microsoft*, 741 N.Y.S.2d 91 (app. Div., Apr. 15, 2002); *I.Lan Systems, Inc. v. Netscout Service Level Corp.*, 183 F. Supp. 2nd 328 (D. Mass., Jan. 2, 2002); *ProCD v. Zeidenberg*, 86 F.3d 1447 (7th Cir., 1996). See also Robert A. Hillman & Jeffrey J. Rachlinski, *Standard-Form Contracting in the Electronic Age*, 77 N.Y.U. Law Review 429, 486-494 (2002)

³³ *Specht v. Netscape Communications Corp.*, 150 F. Supp. 2d 285 (S.D.N.Y., July 5, 2001);

the validity of these agreements appears to turn on whether the users had effective notice of the terms and a real choice to accept the terms prior to the purported conclusion of the agreement.

An inaccurate display of the terms of a web site's privacy policy by a software agent may jeopardize the validity of any agreement between the web site and the user. Although under traditional contract principles, a willfully ignorant buyer assumes the risk of the transaction, inaccurate translations of web site policies by user agents fall outside the bounds of willful ignorance and the typical clickwrap and shrinkwrap challenges. With an inaccurate translation, the user does not truly have an opportunity to review the actual terms of the web site's policy and does not have an opportunity to signify assent to those terms. Unless the user has knowledge of the way a software agent summarizes and, potentially distorts, a web site's policy, the actual terms proposed by the web site for the use of personal information will not be available or readily identifiable for the user. Courts might be more willing to find an invalid consent and, thus, no agreement for the protection of the user's privacy. But, for users, the invalidity of a privacy agreement may be more detrimental than several undesirable terms. Invalidity would mean that the web site is not bound contractually to respect the privacy of users' information.

If, however, an inaccurate user agent does not invalidate consent, a user might still object to the privacy agreement and seek to avoid some of the terms on the basis of mistake. The Restatement (Second) of Contracts summarizes well the circumstances that enable a party to avoid an agreement on the basis of mistake.³⁴ The mistake must go to a basic assumption upon which the contract was formed, have a material effect on the agreed upon exchange, and may not relate to a risk that the party bears. Poor translations of web site privacy policies by software agents mean that users and sites understand different terms for the use of personal information. But, this misunderstanding would relate to a 'basic assumption' only with some difficulty. The inaccuracy relates to the terms of use rather than the existence itself of the data. Similarly, for a user to object successfully to terms of use, the user would have to demonstrate that the mistake was material. The user would need to show that the interaction with the web site would not otherwise have taken place. And, lastly, the user would need to demonstrate that the web site bore the risk of an inaccurate translation of its privacy policy. Since the user chooses the agent, this argument will be difficult.³⁵

Nevertheless, to the extent that a web site is aware of the way in which a prominent software agent such as the IE6 implementation or the AT&T Privacy Bird simplifies P3P privacy statements, there is no mutual mistake. The web site's knowledge would vitiate such a claim, but open up the problem of misrepresentation and deception.

³⁴ Restatement (Second) Contracts, § 152

³⁵ See e.g. Suzanne Smed, *Intelligent Software Agents and Agency Law*, 14 Santa Clara Computer and High Technology Law Journal 111 (1998)(arguing that agency law should apply to software agents.)

In short, the enforceability of an agreement based on a P3P agent's simplified presentation of a web site's privacy policy is uncertain.

Wrong Terms Bind

If a privacy agreement is nonetheless considered formed between a user and web site through a software agent that distorts the web site's privacy statement, the applicable terms for the use of personal information are unclear. Either the web site will be bound by the terms as presented to the user or the consumer will be bound by the unseen or different terms actually posted by the web site in the P3P statement. In either case, one of the parties will face disappointment.

On a superficial level, this situation resembles cases of transmission error. In these cases, courts tend to enforce transmission errors against the seller.³⁶ This jurisprudence suggests that the user agent's simplification of terms would be binding against the web site since the web site proposes the privacy policy. However, the situations are not truly analogous. In the transmission cases, the seller typically chooses the means of transmission and, hence, can more appropriately be assigned the risk of transmission. By contrast, in the P3P implementation, the user rather than the web site chooses the user agent. This choice would ordinarily imply that the user should bear the risk of flaws in the agent. However, the complexity and non-transparency of user agents reduce the capacity of non-expert users to evaluate the risks and make informed choices. In fact, to the extent that several user agents dominate the consumer market, web sites will be in a better position to know how those agents simplify and distort their P3P policy statements. Thus, the applicability of the transmission error cases becomes problematic.

In the event that the web site is aware of the distortion that prominent software agents might make to the site's privacy policy, then the terms appearing to the user are more likely to be applicable to the collected personal information. The web site's knowledge of the software agent's translation ought to shift the risk of error to the web site.

Consumer Deception

Beyond these uncertainties for privacy agreements, the confusion between actual P3P policy statements and the user agent simplifications may lead to claims of consumer fraud. Federal and state laws prohibit "unfair and deceptive practices"³⁷ in order to protect the public from unscrupulous business practices. A company violates this prohibition by making a representation or promise to the public that is either false or that the company does not respect. The Federal Trade Commission and a number of states have pursued 'unfair and deceptive practices' claims against companies that did not honor privacy policies posted on their web sites.³⁸

³⁶ See E. Alan Farnsworth, Contracts § 3.9 (1990)

³⁷ See e.g. 15 U.S.C. §45(a); Mass..Gen..Law Ch. 93A, §11

³⁸ See e.g. FTC, Internet Site Agrees to Settle FTC Charges of Deceptively Collecting Personal Information in Agency's First Internet Privacy Case, <http://www.ftc.gov/opa/1998/9808/geocitie.htm> (Aug. 18, 1998);

For consumers, web sites and implementers, the risk of consumer fraud is significant. The user agent that inaccurately portrays the privacy statement of web sites to consumers causes a deception of those consumers who then interact with the web site. For example, the way IE6 and Netscape 7 simplify the P3P statement of a web site's data collection policy and omit data elements may mislead consumers. In effect, the user agent that does not provide a fully accurate translation of the P3P statement misrepresents the privacy policy of the web site to the consumer.

Nevertheless, these misrepresentations will not ordinarily be volitional on the part of the web site. The 'deception' is caused by the implementer's lack of fidelity to an accurate translation of the P3P policy. There would, thus, not be the requisite 'unfair and deceptive practice' by the web site. However, to the extent that a web site is aware of the user agent's misrepresentations and does nothing to alert consumers, then the misrepresentations might be imputed to the web site. For example, if a web site describes a privacy policy using explicit data elements from the P3P Specification rather than the data categories, then the IE6 Privacy Report will represent to users that the web site collects no personal information. Major web sites with professional development teams will certainly be aware of this translation problem in IE6.

Because the misrepresentations are generated by the P3P implementations, implementers may also face liability for 'unfair and deceptive practices.' By putting a product—the agent—on the market that inaccurately translates web site privacy policies, implementers cause harm to consumers who rely on those inaccurate translations. The 'unfair and deceptive practice' comes from providing the public with a software agent that secretly distorts privacy statements and causes consumers to release personal information under false pretenses. Implementers may, however, avoid this risk by documenting and disclosing how the implementation translates P3P statements.

Implementation Torts

In addition to the problem of potential consumer deception, inaccurate implementations may give rise to two types of other harms: defamation and negligence. For businesses, the misrepresentation of their privacy policies by an inaccurate user agent may be defamatory. For users, the harm caused by an inaccurate agent may be the result of an implementer's negligence.

Web Site Defamation

Some web site developers are already concerned that software agents will inaccurately portray their information practices and corresponding P3P policy statements in a negative fashion. Depending on the nature of the inaccuracy, the software agent may defame a business and the developer might be held liable. As one recent court explained:

Agreement between Attorneys General of Arizona, California, Connecticut, Massachusetts, Michigan, New Jersey, New Mexico, New York, Vermont and Washington and DoubleClick, Aug. 26, 2002, available at http://www.oag.state.ny.us/press/2002/aug/aug26a_02_attach.pdf

“Business defamation is committed when a false and defamatory statement is communicated which ‘prejudice[s] [the plaintiff] in the conduct of its business and deter[s] others from dealing with it.’”³⁹

The way user agents translate P3P privacy policies may, thus, qualify as business defamation. Suppose, for example, that a user agent displays a warning label “Privacy Violator” for any site that uses cookies for any purpose or, like IE6, indicates to users that the site collects less information than is actually the case and that might otherwise be observed. The web site’s reputation may be compromised. For a web site to prevail in the case of business defamation against the developer of a software agent, the web site will still need to show that the designer knew of the falsity produced by the agent, acted with reckless disregard to the truth or acted with negligence.

While the harm to the reputation of web sites from inaccurate translations of their P3P statements will be easy to show, designers of software agents acting in good faith are unlikely to know of the requisite falsity with which specific web sites might be portrayed and are even less likely to have intended any particular harm to a specific web site. Web sites might still contend that the designer acted with reckless disregard for the consequences of the design choices. But, the argument would be hard to sustain. If the designer intended to harm a class of web sites, such as all sites using cookies, then the volitional features of the implementation present a stronger case for business defamation. This aspect, and the potential for crippling lawsuits, might hinder the development of P3P agents by privacy advocacy groups.

Implementer’s Negligence

Although designers of software agents might not face liability for defamation if their products portray web sites inaccurately, negligence law still applies. Both web sites and users might seek to impose liability on the developer of a software agent that poorly translates P3P statements into human-readable descriptions. A negligence claim requires that the victim was owed a duty of care, that the duty of care was breached, and that the breach caused injury. While designers might owe a duty of care to the users who deploy a software agent, the duty is more tenuous for web sites. Even if there is a duty of care, without a customary quality standard for programming, there is no way to establish breach of the duty. For example, the design of IE6 tells a user that the web site collects no personal information if the web site indicates the collection of data by specific element.⁴⁰ To a layman, this design is negligently defective. Yet, under the legal standard, breach of the duty of care is likely to be difficult to establish. The open standard contemplates by its very concept that software agents will be of varying quality. No software is perfect and the designer’s inaccurate translation cannot breach the duty without a clear industry quality standard.

³⁹ Media3 Techs., LLC v. Mail Abuse Prevention Sys., LLC, 2001 U.S. Dist. LEXIS 1310 (MA, 2001)

⁴⁰ See section “Microsoft Internet Explorer 6”

Solutions

In essence, the technological mediation by software agents that is designed to ease the ability of users to understand the privacy practices of web sites risks adding ambiguity, confusion and legal uncertainty. The potential confusion and legal uncertainty demonstrates that the trustworthiness of agents is critical for the success of technical solutions for the protection of privacy. Trustworthiness of agents is in the mutual interests of both consumers and web sites. Several approaches may assure consumers and web sites of the trustworthiness of agents as well as provide implementers with greater legal certainty against potential liability claims.

User agent documentation

The first step in promoting accuracy is comprehensive user agent documentation. User agents without any documentation conceal how the user agent represents web site policies.⁴¹ Transparency will reveal the simplification and choices made by the implementer. Comprehensive documentation enables users and web sites to interpret the meaning of descriptions generated by the agent and to understand how the agent reacts. Transparency also provides an incentive for the implementer to make reasonable choices and creates an opportunity for feedback from sites and users who believe that the choices distort P3P statements.

Transparency will not, however, assure that the simplification and choices faithfully translate P3P statements. Transparency will only expose issues to those sophisticated users and web sites who review the documentation and who examine how particular design choices interact with specific sites, P3P statements and agent configurations.

Certification

A more complete and direct way to assure that user agents accurately represent web site P3P policies is through certification. Certification benefits users and web sites by assuring that the parties interact on the basis of a shared understanding of the terms for the use of personal information. Certification further benefits web sites by the assurance that their policies will not be misrepresented to users. Implementers can benefit by the assurance that the design choices do not potentially cause the defamation of web sites.

Certification, though, raises its own set of issues. The establishment of criteria to evaluate accuracy is likely to be difficult. While there are some things that P3P user agents might do that would be readily identified as inaccurate or misleading, there is a large gray area in which user agents might present factual information side-by-side with subjective judgments that may be deemed misleading by some people, but not others. The fact that consumers and businesses often have different views about what constitutes a

⁴¹ Software agent researchers acknowledge the importance of agent documentation. See Batya Friedman and Helen Nissenbaum, Software Agents and User Autonomy, *Proceedings of Autonomous Agents* 97 at <http://doi.acm.org/10.1145/267658.267772> (“Sometimes, in order to use the services of an agent as desired, a user must know how the agent goes about its task. When the designer of a software agent does not make information accessible to the user, then the user’s autonomy can be undermined.”)

good privacy policy increases the likelihood that developing a standard for user agents would be a controversial endeavor.

The choice of one or more certifying authorities presents further difficult practical and political decisions. One option is to rely on the market and self-regulatory organizations to provide certification. This option necessitates that the organizations providing the certification also be sufficiently trustworthy and that the measurement criteria be valid. An opposing option is to rely on a government agency, like the Federal Trade Commission, to certify the accuracy of P3P products. This alternative may be cumbersome and also requires a valid set of measurement criteria. A combination of these two options may be the most promising way to proceed. A potential model already exists for this combined approach under the Children’s Online Privacy Protection Act.⁴² COPPA allows groups to submit privacy implementing guidelines for approval by the Federal Trade Commission. The FTC has established criteria for approval of these guidelines.⁴³ Once approved, the guidelines provide a “safe harbor” establishing compliance with COPPA. For P3P implementations, the FTC may through a public proceeding set out criteria for the accuracy of software agents.⁴⁴ Independent organizations would then have a trustworthy benchmark or alternatively, the FTC itself might issue safe harbor notices for compliance with the criteria of accuracy.

Finally, once the criteria and the certifying authority are settled, the certification of a user agent may still be difficult to assess in an objective manner. The divergent perspectives of businesses and consumers may affect the measurement of how a user agent behaves.

User Agent Guidelines

If the establishment of a certifying authority proves infeasible, the development of a set of requirements, guidelines, or best practices for user agents might go a long way towards improving user agent reliability. Such guidelines might even be used as part of a self-certification system. For example, the W3C has developed accessibility guidelines that web site developers and user agent implementers often use for self-assessment.⁴⁵ Companies sometimes refer to their compliance with these guidelines when advertising their products.

As discussed earlier, the definitions in the P3P specification were not written with end users in mind. Thus, user agent implementers are faced with the choice of presenting users with information that will be difficult for them to understand, or making their own decisions about the translation of the definitions into user-friendly terminology with the

⁴² 15 U.S.C. §§ 6501-6506

⁴³ See 16 CFR Part 312

⁴⁴ The FTC’s jurisdictional authority would be based on its “unfair and deceptive practices” authority under 15 U.S.C. §45(a).

⁴⁵ *Supra* note 8

risk that the translations will not be completely accurate. P3P user agent guidelines might offer a standard set of user-friendly terms for representing the P3P vocabulary in English (or perhaps several natural languages). User agent implementers would likely welcome the opportunity to adopt standard user-friendly terms and avoid the risks associated with inventing their own. However, without a certification authority, implementers may choose to implement user agent guidelines only partially.

Selecting the most appropriate organization and process to develop P3P user agent guidelines is very important for the eventual credibility of any resulting guidelines. As the organization that developed the P3P specification, W3C seems like a logical candidate. Although W3C is not setup to take on the role of a certifying authority, W3C does have experience developing guidelines. However, to develop good guidelines for P3P user agents will likely require participation by usability experts, as well as require conducting user studies, for example to ensure that proposed user-friendly terminology is actually understandable by end users. This would be a new type of undertaking for the W3C and might require considerable resources to complete.

Outside the P3P development efforts, US industry groups have begun talking about the desirability of offering a “privacy nutrition label” or a standard format “short notice” or “layered notice.” Much of this discussion has been in response to criticism that financial privacy notices are lengthy and difficult to read.⁴⁶ One group is also investigating ways of using P3P as part of this effort.⁴⁷ This or a similar effort could potentially result in the development of a standard mapping of P3P vocabulary elements into a user-friendly format. Such a mapping could serve as a set of guidelines for P3P user agent implementers. However, the often narrow special interest focus will limit the appeal of resulting guidelines. Similarly, it is also uncertain whether industry alliances will have sufficient representative expertise and the commitment to undertake such an effort.

Some Implications for Computer-Mediated Negotiations

Like P3P implementations, transaction negotiations that are mediated by computer will be predicated on a set of machine standards and corresponding implementations that seek to simplify those standards for human interaction. A variety of (mostly experimental) agents are already available that can negotiate to buy or sell items on behalf of their users, make airline reservations, schedule meetings, and perform other tasks.⁴⁸ Open

⁴⁶ For example, The Center for Information Policy Leadership, a division of the Hunton & Williams law firm, has created an industry group to develop a proposed short notices “template.” The Center has conducted focus groups to vet their proposed template with consumers.

⁴⁷ See *Privacy Regulation Report*, With Industry Divided Over Layered-Notice Approach, Privacy Group Readies “Template”, 5 August 2002.

⁴⁸ The MIT Media Lab’s Agents Group lists dozens of agent-related projects on its web site at <http://agents.www.media.mit.edu/groups/agents/projects/>.

standards are under development for specifying rights to electronic works,⁴⁹ and researchers are constructing the building blocks of a “semantic web” in which electronic agents can seek information from any appropriately annotated web page, combine the data with information from other sources, reason with the data, and process the data to assist their users.⁵⁰

Most of the agent and semantic web research to date has focused on developing the standards for computer-to-computer interaction. Researchers and standards groups strive to develop common languages so that two computers can exchange information or engage in a transaction with a shared common “understanding” of the meaning of their dialogue.⁵¹ The languages are developed in consultation with experts in the particular domain area of the intended application (e.g. privacy or digital rights management), and language specifications include definitions that are precise and meaningful to these experts. However, little attention is usually paid to how an agent will communicate with end users, especially with those end users who are not domain experts. Even when the difficulties of this problem are acknowledged, implementers outside the standards processes generally address these issues. Thus, while users may be able to delegate tasks to their computer agents,⁵² users may be unable to specify properly the task to be performed or may not fully understand the information presented by the agent after it has

⁴⁹ See Clint Boulton , OASIS Sets Sights on XML for DRM, April 2, 2002, at http://www.internetnews.com/dev-news/article.php/10_1002301 (“[The OASIS committee] will devise a universal method for specifying and managing rights for a wide variety of business models....”)

⁵⁰ See Tim Berners-Lee, James Hendler and Ora Lassila, The Semantic Web: A new form of Web content that is meaningful to computers will unleash a revolution of new possibilities, *Scientific American*, May 2001, at <http://www.scientificamerican.com/article.cfm?articleID=00048144-10D2-1C70-84A9809EC588EF21&catID=2> (“The Semantic Web will bring structure to the meaningful content of Web pages, creating an environment where software agents roaming from page to page can readily carry out sophisticated tasks for users.”)

⁵¹ See Michael Strobel, Communication Design for Electronic Negotiations on the Basis of XML Schema, *Proceedings of WWW10*, 2001, at <http://doi.acm.org/10.1145/371920.371924> (“A critical factor for the efficiency of the future negotiation processes on this market and the success of the potential settlements is an a-priori agreement among the negotiating parties about how the issues of a negotiation (item attributes, transaction terms and conditions) are represented as abstract objects in the negotiation and what this representation means to each of the negotiating parties. If, for instance, party X offers a delivery date of ‘12/10/2000’ for a workstation to party Y, one potential conflict arises if this syntax is misinterpreted by Y as ‘October 12’ whereas X intended to offer ‘December 10’. A semantical problem could occur if the meaning of this date to X is the point in time where the product will leave the premises of X, whereas Y assumes this is the day the workstation will arrive on the premises of Y. This problem is referred to as the ontology problem of electronic negotiations.”)

⁵² Indeed delegation of tasks is one of the primary motivations for software agents. See Ben Shneiderman and Pattie Maes, Direct manipulation vs. interface agents, *Interactions* 4(6) November/December 1997, p. 42 – 61, at <http://doi.acm.org/10.1145/267505.267514> (“whenever workload or information load gets too high, there is a point where a person has to delegate....because there will be tasks that we just cannot deal with because of our limited attention span or limited time, and we need other entities to be able to represent us and act on our behalf.”)

accomplished the task.⁵³ As illustrated by our experience with P3P user agent implementations, when developers try to simplify the agent interface to make it easier for end users to understand, developers risk introducing simplifications that will be inherently less precise and less accurate than the standard itself. This creates a problem of agent trustworthiness,⁵⁴ which is particularly problematic when agents are used to automate negotiations and conclude agreements.⁵⁵

Conclusion

As work progresses on computer-mediated search and negotiation technologies with a wide variety of applications, the problematic issues associated with the accuracy of user agents are likely to surface repeatedly. The trustworthiness of user agents and their fidelity to the translation of machine standards to plain language are critical for the successful deployment of user agents. Legal uncertainty surrounds many of the conflicts that arise from inaccurate user agents ranging from the enforceability of computer-mediated contracts to business defamation. The law can eventually settle these issues, but the cost and time horizon may impede the development of robust user agents. A more promising and faster resolution to the accuracy dilemma might be the emergence of certification mechanisms. Certification can provide greater confidence and greater certainty to all parties in the development of computer-mediated negotiations.

Acknowledgements

The authors would like to thank Eleanor Lackman for her able legal research assistance.

⁵³ Agent researcher Pattie Maes has stated that understanding and control are two major challenges in designing user agent interfaces. *Ibid* (“Understanding means that the agent-user collaboration can only be successful if the user can understand and trust the agent, and control means that users must be able to turn over control of tasks to agents but users must never feel out of control.”); Usability expert Donald Norman also views human-agent interaction as a major challenge, *see* Donald Norman, How might people interact with agents, *Communications of the ACM* 37, 7 (July 1994), 68-71, at <http://doi.acm.org/10.1145/176789.176796>.

⁵⁴ Microsoft Chairman Bill Gates has cited P3P as a component of his company’s “trustworthy computing” initiative. *See* Kevin Maney, Gates e-mail touts ‘trustworthy’ progress, *USA TODAY*, July 19, 2002, p. B8 at http://www.usatoday.com/tech/news/2002-07-18-gates-email-_x.htm.

⁵⁵ *See e.g.* Margaret Jane Radin, Online Standardization and the Integration of Text and Machine, 70 *Fordham Law Review* 1125 (2002); Robert A. Hillman & Jeffrey J. Rachlinski, Standard-Form Contracting in the Electronic Age, 77 *N.Y.U. Law Review* 429, 486-494 (2002)