Yale University
Department of Computer Science

Usability of Browser-Based Tools for Web-Search Privacy

Felipe Saint-Jean Joan Feigenbaum

YALEU/DCS/TR-1424
March 2010

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED
MAR 2010		**00-00-2010 to 00-00-2010**

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
Usability of Browser-Based Tools for Web-Search Privacy	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Yale University ,Department of Computer Science,New Haven,CT,06520**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES

14. ABSTRACT

Web search is currently a source of growing concern about personal privacy. It is an essential and central part of most users' activity online and therefore one through which a signi cant amount of personal information may be revealed. In an earlier paper [8], we showed that there are security weaknesses in some popular general-purpose Web-privacy tools. We went on to show that, if one is willing to use a tool aimed speci cally at search privacy, it is possible to avoid these weaknesses and we presented PWS (for Private Web search"), a Firefox extension that avoids them. In [8], we also claimed that PWS is easier to use than Firefox extensions aimed at general Web privacy. This paper presents the results of a user study that supports that claim. Speci cally, subjects had signi cantly more di culty using the TPTV bundle (Tor, Privoxy, Torbutton, Vidalia") for web search than they did simply using Google with no privacy enhancements. Users of PWS did much better. In an attempt to understand the reasons for adoption of web-privacy technology (or the lack thereof) we also surveyed the study participants about their level of concern about web privacy and their reasons for using or not using brower- based privacy tools. Most users expressed concern about privacy and willingness to take action to address it, but they also said that they would do not use Firefox extensions such as TPTV or PWS because of the latency that these extensions introduce to the search process.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**Same as Report (SAR)**	**24**	

Usability of Browser-Based Tools for Web-Search Privacy

Felipe Saint-Jean* Joan Feigenbaum†

Abstract

Web search is currently a source of growing concern about personal privacy. It is an essential and central part of most users' activity online and therefore one through which a significant amount of personal information may be revealed. In an earlier paper [8], we showed that there are security weaknesses in some popular general-purpose Web-privacy tools. We went on to show that, if one is willing to use a tool aimed specifically at *search* privacy, it is possible to avoid these weaknesses, and we presented PWS (for "Private Web search"), a Firefox extension that avoids them. In [8], we also claimed that PWS is easier to use than Firefox extensions aimed at general Web privacy.

This paper presents the results of a user study that supports that claim. Specifically, subjects had significantly more difficulty using the TPTV bundle ("Tor, Privoxy, Torbutton, Vidalia") for web search than they did simply using Google with no privacy enhancements. Users of PWS did much better. In an attempt to understand the reasons for adoption of web-privacy technology (or the lack thereof), we also surveyed the study participants about their level of concern about web privacy and their reasons for using or not using brower-based privacy tools. Most users expressed concern about privacy and willingness to take action to address it, but they also said that they would do not use Firefox extensions such as TPTV or PWS because of the latency that these extensions introduce to the search process.

1 Introduction

The penetration of computers and networks into almost every aspect of daily life has conferred tremendous benefits, including easy access to unprecedented amounts of information via the World Wide Web. Unfortunately, it

*Supported by NSF grants 0331548 and 0534052. Email: felipe.saint-jean@yale.edu.

†Supported in part by NSF grants 0331548 and 0716223 and IARPA grant FA8750-07-2-0031. Email: joan.feigenbaum@yale.edu.

has also created new privacy threats. Web search is an essential and central part of users' activities online, and search-engine companies such as Google and Yahoo! may be able to build very accurate profiles of their users by analyzing the users' search behavior. Motivated by users' desire to submit queries to search engines without revealing identifying information, we designed and built a Firefox extension called Private Web Search (PWS), which we reported on in our previous paper [8].

PWS and other web-privacy tools that use similar techniques exhibit network effects; that is, the benefit that any single user derives from using them grows with the total number of users. Thus, it is important to ask whether there is widespread interest among users in adopting such tools and, if there is, whether widespread adoption is feasible. For it to be feasible, typical users would have to find the tools easy to use, and they would have to be satisfied with the tools' performance. These are the questions we address in this paper.

Our main contribution is the presentation and analysis of the results of a study of 41 users at Yale University. The study was designed to discover how well users were able to install PWS, how well they were able to install TPTV (the well known "Tor, Privoxy, Torbutton, Vidalia" bundle for general Web privacy), and how effectively they were able to search using PWS or TPTV (compared to their effectiveness when using Google without any privacy tools). Users were also surveyed about their level of concern about privacy and their reasons for adopting a browser-based web-privacy tool such as PWS or TPTV or for choosing not to adopt such a tool.

Users were divided into three groups. The control group used Google without any additional privacy-enhancing technology; Google was chosen because it is the most widely used search engine and thus presumably easy for study subjects to use. A second group used PWS, the search-privacy tool that we designed and whose usability we wanted to test. The third group used TPTV; we chose TPTV, because it is regarded as the simplest and easiest-to-use of the official Tor distributions, and Tor [9] is the most widely used anonymity network. Our premise was that, if TPTV is the easiest general web-privacy tool for typical users to install and use, then a finding that PWS is easier to use than TPTV would be significant.

Users in the second and third groups had to install the assigned Firefox extension (PWS or TPTV, respectively). All users had to solve as many *search tasks* as they could in 45 minutes. A search task is a trivia question with an unambiguous, correct answer, together with a search method (Google, PWS, or TPTV). (See the Appendix below for the list of trivia questions that users were given.) From the data collected, we were able

to compare both PWS users and TPTV users with the control group with respect to accuracy, speed, and unrecoverable errors that they encountered. At the end of each session, we asked subjects to fill out a short survey about their privacy concerns and their reasons for adopting search-privacy tools (or choosing not to adopt them).

The results support our hypothesis that PWS perfoms better than TPTV. PWS users encountered fewer unrecoverable errors and were able to solve search tasks more effectively than TPTV users. Moreover, PWS users did not encounter as much of a performace degradation as expected (in comparison with Google users). However, the survey results indicate that users have very little tolerance for increased latency in Web search – they do not want to wait longer than they are accustomed to for their search results. This lack of tolerance for delay will be a significant challenge in the design of search-privacy tools that can be widely adopted.

The reminder of this paper is organized as follows. Section 2 describes related prior work. Sections 3 and 4 give brief overviews of PWS and TPTV, respectively. Section 5 presents the study design, and Section 6 presents the results. We draw conclusions in Section 7.

2 Related Work

Clark, van Oorschot, and Adams [1] presented a set of meaningful criteria with respect to which one should evaluate the usability of Tor-based Web-privacy technology; we used these criteria in the design of our study. Clark *et al.* also used a cognitive walkthrough to evaluate the usability of several Web-privacy tools. By contrast, we evaluate PWS and TPTV with a user study that yields empirical results.

Dingledine and Mathewson [2] studied the relationship between anonymity networks (such as Tor) and usability. They make the point that usability has a positive effect on the overall security of an anonymity network. This motivates our study of the usability of PWS and TPTV, because both are Tor-based.

By contrast, TrackMeNot [12] is a Firefox extension that seeks to enhance privacy by adding "cover traffic" to each user's query stream. In addition to sending the user's real queries to the search engine, a TrackMeNot-enabled browser also sends a random stream of "fake" queries; the claim is that this can be done in a such a way that the search engine cannot distinguish real queries from fake and hence cannot accurately profile the user. The cover-traffic approach to search privacy is incomparable to the Tor-based

approach used by PWS; therefore, we did not include TrackMeNot in our study.

3 Overview of PWS

In this section, we briefly review the design of PWS and the major difference between it an earlier web-privacy tools, including the TPTV bundle. Details can be found in [8].

PWS is comprised of several modules that collaborate to handle sensitive information on each "level" of the interaction of browser and search engine (see Figure 1). When the user executes a query using PWS, the browser connects to the local HTTP proxy. The proxy filters the HTTP request, then sends it to the search engine over the Tor network. Later, the proxy receives the response from Google through Tor, filters the HTML to remove all "active components," and sends the answer back to Firefox for display. Active components are programs (written in Javascript, Flash, Java, and many other popular languages) that run in Web pages; they are essential aspects of the functionality provided by many websites, and thus general web-privacy tools such as TPTV *cannot* remove them without destroying users' ability to use the Web. Unfortunately, active components can send large amounts of sensitive information (including personally identifying information) to a server from the client in whose browser they are running, and so they can destroy the privacy gained by routing the client-server interaction over Tor. A major contribution of our earlier work [8] was the observation that, while it is infeasible for a *general* web-privacy tool to remove active components from web pages and maintain functionality, it is feasible to remove them and maintain *search* functionality. This is precisely what PWS does.

Thus far, PWS can only be used with Google [5]; it would be straightforward to extend it to let users select from multiple search engines.

4 Overview of TPTV

The TPTV bundle is the standard Tor distribution for Windows users who want to navigate the Web anonymously. As its name suggests, TPTV is not a single program but rather a combination of applications each of which does part of the job:

Tor is a widely used anonymity network (more precisely, a network that defends users against traffic analysis [9, 3]).

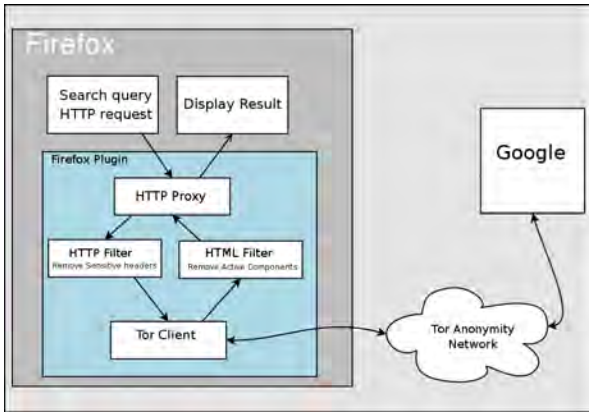

Figure 1: PWS Architecture

Privoxy is a local Web proxy that controls the information "leaked" by browsers [7].

TorButton is a Firefox extension that manages and simplifies the configuration of Tor and Firefox and the interaction between them [10, 11]. In addition, TorButton prevents the browser from leaking several types of sensitive information while attempting to maintain functionality of most websites. In particular, it blocks known Javascript leaks, including window size, timezone, and user agent, but does not remove the Javascript from the displayed page; it disables caches that may be exploited by timing attacks; and it prevents Firefox from storing information on the client machine in several ways that are known to cause leaks (*e.g.*, saving passwords and other forms of autocompletion).

Vidalia is the user interface for Tor. It allows the user to monitor and configure Tor in a user-friendly manner [14].

The Windows installer is used to install and configure all pieces of the TPTV bundle. We chose this setup for our study, because it requires almost no configuration and thus is likely to be the easiest way for most people to get started on using TPTV.

5 User-Study Design

Clark, van Oorschot, and Adams [1] identify the following *Core Tasks* that a user should be able to accomplish in a Tor-based system:

CT-1 Successfully install Tor and the components in question.

CT-2 Successfully configure the browser to work with Tor and the components.

CT-3 Confirm that the web traffic is being anonymised.

CT-4 Successfully disable Tor and return to a direct connection.

All of these core tasks informed the design of our study. We also included a fifth task that is specific to Web search: Successfully search the Web using the assigned Firefox extension (PWS or TPTV). To accomplish this, we presented users with a set of simple but realistic search tasks.

Clark *et al.* [1] also identify the following usability guidelines for Tor-based systems:

G1 Users should be aware of the steps they have to perform to complete a core task.

G2 Users should be able to determine how to perform these steps.

G3 Users should know when they have successfully completed a core task.

G4 Users should be able to recognize, diagnose, and recover from non-critical errors.

G5 Users should not make dangerous errors from which they cannot recover.

G6 Users should be comfortable with the terminology used in any interface dialogues or documentation.

G7 Users should be sufficiently comfortable with the interface to continue using it.

G8 Users should be aware of the application status at all times.

In addition to G1, G2, and G3, which refer to the user's ability to perform Core Tasks, we consider G4 and G5 particularly important. This is because users' encountering errors from which they can't recover seems like an obvious barrier to adoption; furthermore, it is something that we can measure in our study.

Preparation. Before beginning our experiment, we submitted our study design to Yale's human-subjects-research authorities for Institutional Review Board (IRB) approval. Our study was promptly classified as "minimal

risk" and qualified for expedited approval. The only problem we encountered concerned remuneration of users. We originally proposed to pay users $20 per session, but we were directed to reduce it to Yale's standard $10 per session. The fact that $20 per session would not have blown our budget was considered irrelevant by the people in the IRB office; they did not want us to destabilize the campus-wide market for experimental subjects.

Sessions. Sessions were divided into three parts. In the first, we measured users' ability to complete installation and setup successfully; control-group users did not have to do this part. In the second, users completed as many search tasks as they could in 45 minutes. The third part was devoted to the survey described in Subsection 5.4 below.

Sessions took place in a medium-sized office that accommodated up to three subjects at a time, together with the experiment monitor. (The first author of this paper served as experiment monitor for all of the sessions, but anyone familiar with all of the software involved could easily have performed this duty.) The office used for these sessions was not used by anyone else for any purpose throughout the duration of the study.

Subject Recruiting. Users were be recruited from the population of Yale University students who use computers and Google daily. Recruiting was accomplished by posting signs in popular spots that offered a $10 payment for participation in a session.

Data Collection. Users interacted with a Web data-collection application that fed them search tasks and received their answers; this allowed us to determine how many of the trivia questions each user answered correctly. The Web data-collection application also enforced the 45-minute time limit. Network traffic was monitored on the laptops that subjects used by means of a non-intrusive Firefox extension, and each URL visited as a result of each query issued was recorded; this allowed us to determine whether the subject actually found the answer to the trivia question on the web or already knew it before he or she started searching. It also allowed us to determine whether the subject actually used the requested search method.

In addition to the automatic data collection, the experiment monitor kept track of every time the subject asked for help in restarting the experiment after an unrecoverable error.

Starting point. Users were assigned a laptop computer with standard specifications (Win XP, Firefox 2.0). The initial configuration was saved as a virtual machine in order to ensure that every user started from exactly the same configuration.

User Training. Once users were familiar with the environment, they were given brief instructions to ensure that all users could find the Web

data-collection application when they needed it.

Search Tasks. The core of the experiment was a set of search tasks, each of which was defined by a trivia question and a search method (Google, PWS, or TPTV). As a response to a search task, the user was required to give an answer to the question and the resource (URL) at which the answer was found. In order to be able to evaluate users' success objectively, the answers to the questions in search tasks had to be verifiable facts. To ensure that this would be the case, we used a commercially available trivia database as our source of questions [13]. Trivia questions are well suited to this experiment, because they have unambiguous, verifiable answers, and some of them seem to be hard enough to require more than one query. Moreover, the use of objective trivia questions allowed us to minimize the amount of personal information about the subjects that would be revealed in the search process: Trivia questions are broad and general; neither the search queries nor the answers are likely to contain private information about users.

In the form in which one downloads them from [13], the database questions are grouped by topic and sorted in increasing order of difficulty within topic. In order to avoid bias that could result from users' widely varying familiarity with each topic and to provide users with questions of all levels of difficulty, the database questions were randomly permuted before being fed to the study subjects. Only one permutation was performed, however; all study subjects received the questions in the same order.

5.1 Installation

To get started, users in the PWS and TPTV groups were asked to install the assigned software. We pointed them to the relevant download site and corresponding instructions. We verified whether the installation was done correctly and helped them to finish correctly if they had failed to do so on their own.

5.2 Switching

The goal of this step was to test whether users in the PWS and TPTV groups could successfully switch between their assigned privacy tool and Google. (Once again, users in the control group skipped this step.) In this step, the content of the questions was not relevant. Users were assigned a set of easy search tasks some of which were supposed to be performed with the assigned privacy tool and some with plain Google. Success was measured by counting the number of queries performed in the correct mode.

5.3 Searching

The goal of this step was to measure the performance penalty attributable to the use of PWS or TPTV. Performance was measured in terms of both speed and accuracy. Users in all three groups were given the same set of trivia questions and instructed to answer as many as possible, as accurately as possible, in the 45 minutes allotted. Both the total number of questions answered (*i.e.*, the total number of search tasks completed) and the total number of correct answers found on the web were recorded.

5.4 Survey

The final thing we did was to survey the subjects about their level of concern with web privacy and their reasons for using or not using brower-based privacy tools. The questions asked were:

SQ1A: When using Google to search the Web, do you avoid certain topics (select all relevant answers):

a) I never do.

b) At public places

c) On a computer shared with other people

d) At my workplace

e) At home

SQ1B: I avoid certain topics when using Google because (select all relevant answers):

a) I never do.

b) I'm concerned that other people with access to the computer will have access to my search history.

c) I'm concerned that someone may intercept the traffic between me and Google.

d) I'm concerned that Google will learn certain things about me.

e) My employer has a corporate policy governing personal use of company-owned resources.

f) Other:

SQ1C: If you never refrain from searching (*i.e.*, if you selected (a) in the previous question), why is this?

a) I don't consider my search history to be private data.

b) I do consider my search history to be private data, but I trust that it is well protected.

c) Other:

SQ2A: How much do you agree with the following statement: "When I use Google to search the Web, Google has a good chance of associating my identity with each of my queries if it chooses to do so."

a) Strongly disagree

b) Weakly disagree

c) Weakly agree

d) Strongly agree

SQ2B: How much do you agree with the following statement: "Google keeps a fairly complete search history associated with my identity."

a) Strongly disagree

b) Weakly disagree

c) Weakly agree

d) Strongly agree

SQ3: Suppose that Google were able to associate each query you issue with you, and you had an equally accurate alternative method for searching that protected your identity but performed more slowly. You would consider using it if getting the answer to a query (underline one answer per row):

a) Took an additional 1 second or less: Never, Sometimes, Always

b) Took about 5 additional seconds: Never, Sometimes, Always

c) Took about 10 additional seconds: Never, Sometimes, Always

d) Took about 30 additional seconds: Never, Sometimes, Always

Search Method	Installation Success Rate
TPTV	59.9%
PWS	84.6%

Table 1: Percentage of users who successfully completed installation

e) Took about 60 additional seconds: Never, Sometimes, Always

SQ4: If Google were able to associate each query you issue with you, and you had an equally accurate alternative method for searching that protected your identity, you would consider it using it for queries about (select all relevant answers):

a) Health

b) Sex

c) Politics

d) Illegal activities

e) Yourself

f) People you know

g) Your job or your employer

h) Would never use it

6 Results

6.1 Installation

Google users (the control group) did not have to install any software, and hence there are no results to report on their success with installation. PWS and TPTV users were told to follow the installation instructions on the respective websites; when a user said that he or she was finished with this step, we checked whether the Firefox extension had been installed successfully. The results appear in Table 1. Because PWS uses the standard installation procedure for Firefox extensions, the significantly higher success rate for PWS users may be attributed to Firefox's high usability. It is worth noting, however, that we designed PWS so that Firefox users would have a single

Search Method	Switching Success Rate
TPTV	73%
PWS	77%

Table 2: Percentage of users who successfully performed at least 4 out of 5 potential switches.

extension-distribution package precisely in order to make the setup process manageable for most users.

Of the 15.4% of PWS users who did not complete the installation process successfully, most failed by not confirming the installation of an untrusted extension. TPTV users experienced a more diverse set of failure modes. Although the instructions pointed to a single Windows package, some users downloaded the wrong package. Others failed to restart the browser after the installation process was done.

6.2 Switching

Users in the second and third groups were instructed to perform each of the first 5 search tasks using a randomly assigned search method – either plain Google or their assigned extension. A user "succeeded" in this experiment if he or she performed 4 of the 5 tasks using the correct search method. Results are presented in Table 2.

Switching performance for the two extensions was not as similar as these numbers suggest. 46.6% of the users in the TPTV group, when instructed to use TPTV for the first search task, failed to activate Tor. They knew what they were supposed to do but could not find the user-interface element that they needed and asked the experiment monitor for help. Once they received that help, many of them were able to perform subsequent switches correctly. The large number of users who needed help activating Tor in the first task is reflected in Table 5 below.

6.3 Accuracy and Speed in Searching

Having chosen search tasks that are realistic and have factual, verifiable answers, we set out to measure the performance degradation that PWS and TPTV users suffered in terms of accuracy and speed. The results appear in Tables 3 and 4.

As can be seen in Table 3, neither PWS nor TPTV imposed a very large

Percentage of Correct Answers in Completed Search Tasks		
Search Method	Average	Median
Google	92.1%	95.5 %
TPTV	86.3%	93.3 %
PWS	81.5%	92.0 %

Table 3:

Number of Search Tasks Completed in 45 minutes		
Search Method	Average	Median
Google	35	36
TPTV	18.2	18
PWS	30.9	27

Table 4:

penalty on users writh respect to accuracy, but both imposed some penalty. Given that TPTV provides users with exactly the same web-search interface as Google, why should TPTV impose any accuracy penalty at all? Plausible explanations include unrecoverable errors (as explained below) and the Tor-induced slow down of the search process (which may have caused some users to give an answer based on the results of one query when several queries were required to find the correct answer).

It is not clear exactly why PWS imposes a higher accuracy penalty than TPTV, but the following observations seem relevant. The PWS user interface is different from the Google-TPTV interface, and it is likely that users are not familiar enough with it to search with peak accuracy; this would not be an insurmountable barrier to adoption, because they would become familiar with it over time if they continued to use it, and the initial accuracy penalty is not prohibitive. Some of Google's helpful features, like query suggestion and spell checking, are not present in this early version of PWS, and their absence probably hurts accuracy. It is also worth mentioning that PWS only provides the top 20 results in a single page.

In terms of speed, measured by how many search tasks users could answer in a fixed amount of time, Google set the benchmark at an average of 35 tasks in 45 minutes. For detailed results, see Table 4. The poor performance of the TPTV group can be explained by two basic facts. First, TPTV uses Tor for all Web requests, and Tor increases latency considerably; PWS is faster in

Search Method	Percentage of Users who Made No Unrecoverable Errors
Google	100.0%
TPTV	20.0%
PWS	76.9 %

Table 5:

part because it uses Tor only for the interaction with Google. This difference is a natural consequence of different adversary models and is explained in our earlier paper on the design and implementation of PWS [8]. Second, TPTV users encountered several problems that slowed them down. The relevance of a user's encountering this type of frustrating problem during his or her first hour of experience with a search-privacy tool goes beyond accuracy; it can have a direct and severely negative impact on adoption.

6.4 Unrecoverable Errors in Searching

As discussed in Subsection 6.1 above, some users were unable to complete the installation process and needed help before they could proceed with the rest of the session. After installation, some users encountered other errors from which they could not recover on their own. The figures in Table 5 include both classes of unrecoverable error.

It is interesting to consider the type of unrecoverable errors that TPTV users encountered. Besides not being able to install the software correctly, they faced three significant problems.

33.0% were faced with a Google page in a language they cold not understand.

33.0% were told by Google that it could not answer queries because the user's machine was infected by spyware.

46.6% were not able to figure out how to activate TPTV after installing.

As a result, a significant fraction (80%) of TPTV users could not finish the study without help from the experiment monitor.

6.5 Survey

We now report the users' answers to the survey about their privacy concerns and their willingness (or the lack thereof) to use certain types of privacy

technology.

SQ1A: When using Google to search the Web, do you avoid certain topics (select all relevant answers):

21.95% I never do.

60.98% At public places

68.29% On a computer shared with other people

56.10% At my workplace

14.63% At home

SQ1B: I avoid certain topics when using Google because (select all relevant answers):

17.07% I never do.

65.85% I'm concerned that other people with access to the computer will have access to my search history.

19.51% I'm concerned that someone may intercept the traffic between me and Google.

26.83% I'm concerned that Google will learn certain things about me.

34.15% My employer has a corporate policy governing personal use of company-owned resources.

00.00% Other

SQ1C: If you never refrain from searching (*i.e.*, if you selected (a) in the previous question), why is this?

50.00% I don't consider my search history to be private data.

50.00% I do consider my search history private data, but I trust it is well protected.

00.00% Other

SQ2A: How much do you agree with the following statement: "When I use Google to search the Web, Google has a good chance of associating my identity with each of my queries if it chooses to do so."

Seconds	Aggregate			Google			PWS			TPTV		
	Always	Sometimes	Never	Always	Sometimes	Never	Always	Sometimes	Never	Always	Sometimes	Never
at most 1	97.56	2.44	0.00	100.00	0.00	0.00	100.00	0.00	0.00	93.33	6.67	0.00
about 5	90.24	7.32	2.44	100.00	0.00	0.00	92.31	7.69	0.00	80.00	13.33	6.67
about 10	56.10	34.15	9.76	69.23	23.08	7.69	53.85	38.46	7.69	46.67	40.00	13.33
about 30	17.07	36.59	46.34	46.15	30.77	23.08	7.69	38.46	53.85	0.00	40.00	60.00
about 60	7.32	34.15	58.54	23.08	38.46	38.46	0.00	38.46	61.54	0.00	26.67	73.33

Table 6: Percentage of users who would trade N seconds of delay for identity protection, aggregated and by group, for $N \in \{1, 5, 10, 30, 60\}$

12.20% Strongly disagree

29.27% Weakly disagree

36.59% Weakly agree

21.95% Strongly agree

SQ2B: How much do you agree with the following statement: "Google keeps a fairly complete search history associated with my identity."

7.3% Strongly disagree

14.6% Weakly disagree

53.6% Weakly agree

21.9% Strongly agree

The answers to SQ3 are given in Table 6 and Figure 2. It is quite interesting that some of these results, especially those in Figure 2, indicate that users who experienced more privacy-related delay when trying to complete the search tasks expressed less willingness to trade increased latency for increased privacy. In general, control group users, who searched with plain Google and thus experienced no privacy-related delay, expressed much more willingness to trade latency for identify protection than users in groups 2 (PWS) and 3 (TPTV). TPTV users, who experienced the most delay, were also the most likely to say that they would *never* trade N seconds of delay for identity protection, for all values of $N \in \{5, 10, 30, 60\}$.

Study subjects were given no tangible incentive to excel in their assigned tasks. Nonetheless, most of them strove hard to perform the search tasks as quickly as possible. Anecdotally, we can report that several PWS and TPTV users found their experiences quite frustrating and explicitly expressed anger

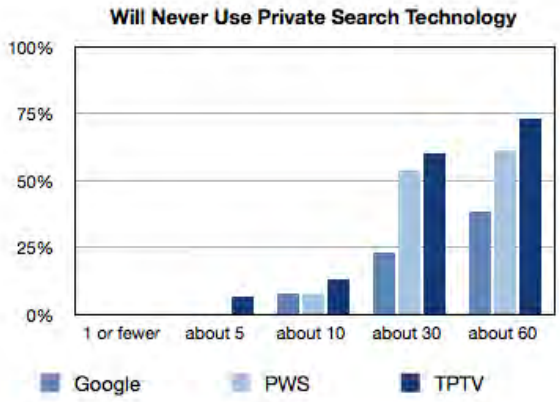

Figure 2: Percentage of users in each group who said they would **never** trade N seconds of delay for identity protection, for $N \in \{1, 5, 10, 30, 60\}$

(even fury) by the end of the session. According to them, this frustration was caused by delay: The Firefox extensions they were using did not allow them to complete search tasks as fast as they wanted to.

SQ4: If Google were able to associate each query you issue with you, and you had an equally accurate alternative method for searching that protected your identity, you would consider it using it for queries about (select all relevant answers):

75.61% Health related queries

92.68% Sexual related content

53.66% Political related queries

87.80% Illegal activities

21.95% Things about yourself

7.32% Things about people you know

12.20% Job related queries

7 Conclusions

From one point of view, the results of this study make PWS look good. At the expense of a small degradation in accuracy, which might disappear over

time as they gained experience with it, non-expert users found PWS to be easier to install and faster to use for search than TPTV. This is evidence that it is good to provide users with a tool aimed specifically at *search* privacy: Such a tool can be made both more secure (because pages can be stripped of active components) and more usable than a general web-privacy tool.

From another point of view, the results cast doubt on the worth of this well established approach to search privacy: Users said that delays caused by the Tor network make any Tor-based system, including both PWS and TPTV, highly undesirable for search. One rule of thumb in the study of Web applications is that delays of more than 10 seconds cause users to lose focus [6]. In our study, the average time to completion for a search query was approximately 30 seconds, for both PWS and TPTV. (Note that the fact that it was 30 for both does not contradict the results in Table 4, because many search tasks require multiple queries.) Despite the fact that Tor is the best widely available anonymity network, it is too slow for use in search. Tor designers and coders are well aware that latency is a barrier to building usable Web applications on top of Tor [4] and are working to improve the situation. However, because design, implementation, and deployment of a very low-latency anonymity network is a difficult open problem, it is unrealistic to expect improvement in the near future.

Finally, we note some obvious limitations of our study design. Searching for answers to trivia questions is only one special case of Web search. It is possible that privacy technology would affect users' overall search activity in ways not capturable in a study of this special case. For example, Google enhances the results it displays to a user based on that user's long-term search history, but that history is likely to be irrelevant to searches for the answers to the questions in the Appendix below and was in any case unavailable to Google in this study. Despite this limitation, our comparison of three well defined groups' performance on trivia questions does shed some light on the effects of privacy technology in this special case. Another limitation of our study concerns the definition of "unrecoverable error." We considered an error "unrecoverable" if the subject gave up and asked the experiment monitor for help; it is possible that, without the implicit pressure imposed by the fact that this was an experimental session of limited duration, a determined user who had as much time as he or she was willing to devote to it could "recover" from one or more of these errors.

References

[1] Jeremy Clark, P. C. van Oorschot, and Carlisle Adams. Usability of anonymous web browsing: an examination of tor interfaces and deployability. In *SOUPS '07: Proceedings of the 3rd Symposium on Usable Privacy and Security*, pages 41–51, New York, NY, USA, 2007. ACM.

[2] Roger Dingledine and Nick Mathewson. Anonymity loves company: Usability and the network effect. In *Fifth Workshop on the Economics of Information Security*, 2006.

[3] Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The second-generation onion router. In *Proceedings of the 13th USENIX Security Symposium*, August 2004. `http://tor.eff.org/tor-design.pdf`.

[4] Roger Dingledine and Steven J. Murdoch. Performance improvements in tor or, why tor is slow and what we're going to do about it.

[5] Google. `http://www.google.com/`.

[6] Jakob Nielsen. *Usability Engineering*, chapter 5. Morgan Kaufmann, 1994.

[7] Privoxy. `http://www.privoxy.org`.

[8] Felipe Saint-Jean, Aaron Johnson, Dan Boneh, and Joan Feigenbaum. Private web search. In *WPES '07: Proceedings of the 2007 ACM Workshop on Privacy in Electronic Society*, pages 84–90, New York, NY, USA, 2007. ACM.

[9] Tor. `http://tor.eff.org`.

[10] Torbutton-download. `http://freehaven.net/~squires/torbutton/`.

[11] Torbutton-options. `http://www.torproject.org/torbutton/options.html.en`.

[12] Trackmenot. `http://mrl.nyu.edu/~dhowe/TrackMeNot/`.

[13] Triviadatabase. `http://www.triviadatabase.com/`.

[14] Vidalia. `http://www.torproject.org/vidalia/`.

Appendix

What follows is the list of trivia questions that users were given, together with the answers that they were supposed to find on the web. Items are presented in the form Q — A, where Q is the question, and A is the answer.
In "A Clockwork Orange", who was Alex's favorite composer? — Ludwig Von Beethoven
Who wrote "The Crucible"? — Arthur Miller
Who wrote the book, "Dr. Jekyll and Mr. Hyde"? — Robert Louis Stevenson
What artist painted "Les Demoiselles d'Avignon"? — Pablo Picasso
In what city does Louis get interviewed in, in "Interview With The Vampire"? — San Francisco
In "Flowers for Algernon", who is Algernon? — The Mouse
Who wrote the book, "Pride and Prejudice"? — Jane Austen
Who wrote the book "Fowl Tips"? — Wade Boggs
What artist was on the cover of "Time" and "Newsweek" in 1975? — Bruce Springsteen
"The Lord of The Rings" was written by whom? — J.R.R. Tolkien
A person who poses for a painter is usually called a what? — Model
Famous painter Picasso went by the first name of? — Pablo
Finish the title of this book: "Of Mice And ..."? — Men
What is the name of the snowy owl that Hagrid bought for Harry Potter? — Hedwig
What was in the Trojan Horse? — Soldiers
Who wrote the "Cat in the Hat"? — Dr. Seuss
Who writes the "A Series of Unfortunate Events" book series? — Lemony Snicket
Who wrote the book "Mommie Dearest"? — Christina Crawford
Who wrote the book "The Rainmaker"? — John Grisham
Who wrote the book "The Partner"? — John Grisham
Who wrote the book "Tuesdays with Morrie"? —
Mitch Albom
Who did Rosie O'Donnell play in the Broadway show "Grease"? — Rizzo
Which Broadway musical featured the songs of Billy Joel? — Movin' Out
Which Steve Martin film was turned into a musical in 2005? — Dirty Rotten Scoundrels
Revived in 2002, "Man of La Mancha" is based on which novel? — Don Quixote
Which play won both the Pulitzer and Tony for Best Play in 2001? — Proof

Which musical legend is Liza Minnelli's mother? — Judy Garland
What city was Edgar A. Poe in when he died? — Baltimore
What does NHL stand for? — National Hockey League
What does M.V.P. stand for? — Most Valuable Player
Who wrote the 1959 book "Hawaii"? — James A. Michener
In 1956, J. R. R. Tolkien wrote which classic? — Lord Of The Rings
In 1937, John Steinbeck wrote which classic? — Of Mice And Men
Famous author Steinbeck has a full name of? — John Steinbeck
Famous author Updike has a full name of? — John Updike
What town is named after the author of "Last of the Mohicans"? — Cooperstown
Which of the following movies was not adapted from a James Clavell novel? — Gunga Din
What does the T.S. stand for in T.S. Eliot's name? — Thomas Stearns
Who killed Macbeth in the play "Macbeth"? — Macduff
What novel introduced the noun "Droogies"? — A Clockwork Orange
Who wrote "The Time Machine"? — H.G. Wells
Which poet wrote the poem "The Road not Taken"? — Robert Frost
Who wrote the book, "The Greene Murder Case"? — S. S. Van Dine
Who wrote the book "Gone with the Wind"? — Margaret Mitchell
Who is the author of "Angels & Demons"? — Dan Brown
Where does James have his adventure in "James and the Giant Peach"? — A Giant Peach
In Greek Mythology, who preceded the rising of the sun each morning? — Eos
Which Goddess was the Patron Saint of Athens? — Athena
A milk punch is made up of milk, sugar and which of the following? — Rum
What gives a "Brain" its blood vessels? — Grenadine
What is in the shot "Liquid Cocaine"? — Goldschlager
Which beer is brewed in the "land of sky blue waters"? — Hamm's Beer
Adolphus Busch started Anheuser-Busch in which U.S. city? — St. Louis
Kirin Brewery was founded in what country? — Japan
Pabst Brewing Company was headquartered in which U.S. city? — Milwaukee
Anheuser-Busch Brewery had its headquarters in which U.S. city? — St. Louis
Which beer is from the "land of sky blue waters"? — Hamm's
Which beer is advertised as "beechwood aged"? — Budweiser
Which beer is nicknamed "America's fire brewed beer"? — Stroh's
Miller Brewing Company featured what bird in its trademark? — Eagle

Kirin is a brewing company found mainly on what continent? — Asia
What country in the world is known for XXXX beer? — Australia
What is the most popular exported New Zealand beer? — Steinlager
From which country does Carlberg's beer originate? — Denmark
In what U.S. city is Weinhard Brewing Company located? — Portland, OR
Which beer claims to have "bottled beer taste in a can"? — Keystone
Where is Red Stripe beer brewed? — Kingston, Jamaica
Where is the Coors Brewing Company located? — Golden, Colorado
Which brand of beer has its brewery located in Latrobe, Pennsylvania? — Rolling Rock
How many cans are in a Lone Star "Texas 6-Pack"? — Eight
A gallon is equal to how many ounces? — 128 Ounces
Curacao is made from what dried peel? — Orange
What type of Schnapps is Ice 101? — Peppermint
Which soft drink is the oldest in the U.S.? — Dr. Pepper
What is Stilchester? — A Cheese
What might Italians call maize? — Polenta
From which country do french fries originate? — Belgium
In Peru, which color potatoes are grown, in addition to white? — Purple And Black
What are the two ingredients in a roux? — Flour And Fat (Butter)
What is also known as Liberty Cabbage? — Sauerkraut
What is kartofflen? — Potato Dumplings
What is the name of the flatbread eaten with most Indian cuisine? — Naan
What type of cuisine offers Dim Sum? — Chinese
What is Cioppino? — A Seafood Stew
What does Etoufee mean? — Smothered
What famous dish uses arborio rice? — Risotto
With which vegetable are Norwegian Lefse made? — Potatoes
What is the name of the bar where "Buffalo Wings" originated? — The Anchor Bar
What type of cheese is an ingredient in Tiramisu? — Mascarpone Cheese
What is a "sabayon"? — A Custard Dessert
What temperature should you not exceed when melting chocolate? — 120*F
What is the featured flavor in Mexican mole sauce? — Chocolate
What does the Italian term "Al Dente" mean in regards to pasta? — To The Teeth
What is the main flavoring agent in a Mornay Sauce? — Gruyere Cheese
What gives the drink known as a "Black Cow" its color? — Coffee
Trader Vic claims credit for creating what drink? — Mai Tai

What nationwide U.S. fast food chain opened the first drive-in? — A&W

In 2005, how much beef and steak did the average American eat? — 67 Pounds

What is the first sign of the western zodiac? — Aries

What is the name of the currency used in Finland? — Markkaa