

Personal knowledge questions for fallback authentication: Security questions in the era of Facebook

Ariel Rabkin
UC Berkeley
asrabkin@cs.berkeley.edu

ABSTRACT

Security questions (or challenge questions) are commonly used to authenticate users who have lost their passwords. We examined the password retrieval mechanisms for a number of personal banking websites, and found that many of them rely in part on security questions with serious usability and security weaknesses. We discuss patterns in the security questions we observed. We argue that today's personal security questions owe their strength to the hardness of an information-retrieval problem. However, as personal information becomes ubiquitously available online, the hardness of this problem, and security provided by such questions, will likely diminish over time. We supplement our survey of bank security questions with a small user study that supplies some context for how such questions are used in practice.

Categories and Subject Descriptors

K.4.4 [Computers and Society]: Electronic Commerce—*Security*

General Terms

Security, Human Factors

Keywords

Security, questions, authentication

1. INTRODUCTION

Online banking is becoming a widely used way of controlling personal finances. Many users find the convenience offered by electronic access from personal computers irresistible, despite the possible security risks. By the same token, criminals have found online banking an irresistible target. A recent study of online criminal markets has found that stolen bank login credentials are among the most frequently offered and sought contraband goods. Further, there are large criminal networks organized to turn these compromised credentials into actual crime [6].

To meet this threat, banks have deployed increasingly sophisticated authentication mechanisms. Most banks exhort, or require, users to pick “strong” passwords, not easily guessed by an attacker. Strong passwords, however, are hard for many users to remember. For usability reasons, banks often couple their password authentication mechanism with some sort of “lost password” mechanism, which users can fall back on if they have forgotten their passwords. While the security and usability of passwords in practice has been studied extensively, these “fallback authentication” mechanisms have been much less studied by the academic community.

1.1 Security Questions

One solution, used by many banking sites, has been to rely on security questions. These come in two varieties. One sort of security questions asks about sensitive (though not necessarily private) information such as social security and bank account numbers, and ATM PIN codes. We refer to these as sensitive security questions. Another set of security questions, which we term personal security questions, ask about personal history, and family background, such as one's mother's maiden name. These have also been referred to in the literature as “Personal Verification Questions”. Personal security questions, in turn, can be divided into those selected by the user, perhaps from a menu of choices, and those specified entirely by the institution, such as ZIP code, mother's maiden name, or date of birth.

Both sorts of security questions differ in a crucial respect from passwords. The ideal password is a high entropy string of characters, and is chosen entirely by the user, and then memorized. However, users are not expected to memorize answers to security questions. Instead, the answer should already be part of a user's long-term memory (or, in the case of a bank account number, be written down and readily accessible).

In the context of fallback authentication, the user is assumed to be unable to remember arbitrary strings — otherwise they would have been able to remember their password. Thus, the ideal security question should have an answer that is completely determined by the question, so that the user need not memorize or guess. As a result, the security questions posed essentially determine the answers given. This shifts much of the responsibility for secure authentication away from users and onto the designers of the authentication system. In a security question scheme, the only choice users should be conceived of having is that of which questions to answer; all the other decisions are made by the mechanism designer.

In principle, mechanism designers could allow users to

write their own questions, thus returning a measure of control to the user. In practice, this approach appears uncommon. Only one of the 20 sites that we examined allowed users to write their own security questions.

1.2 This paper

We believe that survey of current authentication practices in online banking will be a useful contribution to the broader topic of security questions. Banks are well motivated to reduce identity fraud, and have ample data on user behavior with which to improve their mechanisms. They therefore represent the commercial state-of-the-art in security question-based authentication. Common flaws in online banking authentication will likely be found elsewhere, as well.

We have conducted a survey of fallback authentication mechanisms used by a population of 20 diverse online banking sites, listed in Appendix II. We focus on personal security questions: they are more diverse than sensitive security questions, and their security properties are less well understood. Further, we argue that their security properties have shifted recently for the worse, in a way not true of sensitive security questions.

We caution that our results should not be used to compare the security of various online banking websites. Fallback authentication is only one part of a site's security, and successful fallback authentication may trigger intensive auditing and profiling of user behavior. We present results organized by institution, rather, in order to allow future work to make direct comparisons with our data, and to give a sense of the variety of authentication mechanisms we encountered.

We begin with a discussion of prior work. We then discuss the mechanisms we saw in use, and analyze the personal security questions in detail. In addition to our examination of authentication mechanisms, we have conducted a modest survey of online banking users, in order to validate our assumptions about how security questions are used. Last, we propose some approaches to improving personal security questions.

2. PRIOR WORK

Authentication is a well-studied topic, and we limit this discussion of prior work to personal-knowledge based techniques, and to studies of industrial practice.

2.1 Academic research

There is a significant literature on various sorts of authentication questions. O’Gorman, Bagga, and Bentley propose a family of question-based protocol called Query-Directed Passwords (QDP). The scheme imposes restrictions on the questions and answers, and specifies how QDP should be joined with other techniques (such as PIN, address of physical devices, and client-side storage device or wallet card). The intent is to hide the questions from attackers, by equipping users with copies in advance. In a sense, QDP is a fusion of knowledge-based with token-based authentication [16].

Another variant form of authentication questions is the preference-based technique proposed by Jakobsson et al. [11]. In this scheme, users are asked to make a series of preference judgments, and if their answers are close enough to the user's previously-established preferences, they are authenticated. The scheme is motivated by the fact that preferences of the form “do I like cats?” are more durable than memory for

facts, and are often harder to guess. A working prototype is available online.¹

Most security question systems require users to specify the correct answers in advance. The Adaptive Challenge Question scheme avoids this, by asking users about their browsing history in the recent past [2]. Unfortunately, this scheme has limited applicability. It requires that the authenticating site have access to this history, and that the user's browsing sessions can be identified. It also presumes that adversaries do not have access to the browsing history, either from a centralized site, or from observing the network sessions directly.

Security questions as used today have been studied in the past, largely from a prescriptive point of view. Just has proposed a set of criteria for evaluating personal security questions, and has sketched a number of possible design alternatives [12]. Outside of academia, the government of Canada has published guidelines for security questions in authentication [15]. Unlike these publications, we seek to analyze the questions we see in use today, rather than set guidelines for future authentication system designers.

The most notable prior empirical study of security questions is that of Haga and Zviran, published in 1991 [10]. They asked users to answer a set of personal security questions, and then measured the successful answer rate for the users, and also that of the users' friends, family, and significant others. The list of twenty personal security questions used in their study corresponds reasonably well with current practice. Half are asked essentially word-for-word by banks in our survey, and there are approximate matches for several other questions. Unfortunately, their results may be of limited applicability today, since the rise of the World Wide Web has made a vast quantity of personal information available online, thus significantly aiding an attacker.

More recent studies of commercial authentication techniques are rare in the published literature. The only such study we are aware of is that of Mannan and van Oorschot, which examined the websites of six financial institutions, not overlapping with those examined in this paper [13]. The websites they observed had surprisingly weak security models, and fairly loose length and complexity requirements for both passwords and personal security questions. Password-reset techniques were not a primary focus of their study, and so their results are not directly comparable with those of this study. Insofar as they show that financial institutions content themselves with fairly lax security requirements, their results confirm those presented here.

Griffith and Jakobsson have demonstrated that mother's maiden name, perhaps the canonical example of a personal security question, can be deduced with significant probability via public records [9]. They suggest a number of techniques for deducing mothers' maiden names from public records, and attempt to quantify the fraction of individuals at risk. Techniques very similar to theirs would likely work to answer a number of other security questions that ask about names of family members.

2.2 Industrial practice

A security question technique that we are aware of, but have not considered in this paper, is the use of data already held by an institution to authenticate users. For instance, several of the major American credit bureaus ask users mul-

¹At <http://www.i-forgot-my-password.com/>.

multiple choice questions about their past financial activity in order to authenticate them. This technique exploits the privileged access that credit bureaus have to sensitive information about past financial transactions. A similar technique, sold commercially by RSA Security, asks questions based on public record databases [19]. In both cases, the security guarantee depends on there being a significant difference in the ability of the attacker and the authentication system to extract information from public records. This assumption is hard to validate in practice, and may become less true over time, if attackers improve their techniques, or if user privacy demands limit institutional access to public records.

A commercial authentication technique that **has** been analyzed in the academic literature is email based authentication. In this approach, the ability to receive mail at a prearranged email address is used as proof of identity. A security analysis of this approach was offered by Garfinkel in [8]. Garfinkel observes that emailing users a link to retrieve new password is often a secure technique. Few attackers are able to reliably snoop on user email, and if the mechanism is designed intelligently, it is only vulnerable to adversaries with near real-time access. This in practice means adversaries with control of either the network path, or of the user's computer, both of which are trusted by most users in practice.

Email-based authentication is used by a significant minority of financial institutions. One institution in our study used this technique as its sole fallback authentication technique. A few others used email-based authentication in conjunction with security questions.

3. SURVEY OF BANK MECHANISMS

3.1 Methodology

As mentioned above, we attempted to obtain “forgotten” passwords at 20 financial websites. Our sample included brokerages, deposit banks, and credit card issuers, and the banks sampled included a mix of national, regional, and online-only banks. Four of these sites had no online password-recovery mechanism, and were excluded from further study.

For each remaining website, we performed and recorded the fallback authentication procedure, and compiled a list of all available security questions. We also checked whether or not the recovery mechanisms changed depending on whether the request originated from a host that had previously been used to access that account. All the banks were examined between November 2007 and February 2008. The accounts used were regular accounts, not recently opened, held by volunteers known to the author; the data was collected by those account-holders. This imposed certain limits on our data collection: we did not want to risk the account-holders being locked out of their accounts, and therefore did not attempt to thoroughly explore the mechanisms in use.

We seek only to evaluate the security questions that we encountered. We do not purport to do a full security analysis of online banking, or even of the fallback authentication mechanisms in use. Institutions may well adapt the strength of their authentication mechanisms to a variety of cues, such as frequency of access, source address, and so forth, that are difficult for us to control for. Further, such mechanisms change over time, and our study should not be taken to re-

fect the current state of such mechanisms. ²

3.2 Results

As mentioned, four sites did not have any sort of web-based fallback authentication system. A fifth relied purely on email, sending a new password to users upon request. The remaining fifteen institutions all relied on security questions of some sort, either secret or personal. We describe the questions here; a detailed chart of our results is presented in appendix II.

The authentication mechanisms of these fifteen institutions fall into three rough categories. One set, of six institutions relies on personal security questions, coupled with a username or Social Security number. Of these, one bank also verified that a user could receive email at a prespecified address. A second set of four institutions requires both personal security questions and account details, such as bank account or credit card number and sometimes PIN as well. A fifth required account details plus date of birth. Two institutions relied solely on account numbers and PINs, in one case coupled with an unprompted recall challenge (“enter your secret word here”). A third institution relied on account numbers, plus proof of email address. The last institution (a bank) allowed users to choose at authentication time whether to answer personal security questions, or supply the account number.

There were a few apparent patterns in institutional choice of authentication mechanism. The largest banks tended to rely more on sensitive security questions. The credit cards largely avoided personal security questions, presumably on the grounds that an attacker with a purloined credit card number and CVV code can use the card directly, and has little reason to authenticate via the online management interface. The brokerages, in contrast, tended to have lenient authentication procedures. The small number of institutions in our sample and the diversity of mechanisms employed precludes our placing any real confidence in these generalizations, however.

There was significant variation in how personal security questions were chosen, both at setup and at authentication time. At setup time, users were generally presented with one or more subsets of the total pool of questions, and asked to supply an answer to one question per pool. These sets would often have substantial overlap, but sometimes would be disjoint. We do not understand why institutions did not content themselves with a single pool of questions. Sometimes, the pool of questions from which users were obliged to choose was a strict (and seemingly random) subset of the total available question pool; we do not understand the purpose of this mechanism, either. At authentication time, most institutions that employed user-selectable personal security questions only asked a single such question. Only two out of ten asked more than two such questions.

The total number of questions in each institution's question pool varied widely. One bank had a pool of over one hundred questions, of which users needed only answer one

²In passing, we note that we encountered surprisingly few overt defenses against automatic attack. None of the institutions employed CAPTCHAs during password recovery to prevent automatic attacks. Only two institutions appeared to vary their mechanism depending on whether the user was connecting from a known host. At least two sites exposed whether a given username exists, before a client succeeds in authenticating. This is generally a poor security practice.

at authentication time. (That institution required users to supply answers to three questions, of which only a single one was asked at authentication time.) At the other extreme, another institution asked users to supply answers to five questions, half the total pool, and asked all five at authentication time. There did not appear to be a clear correlation between the number of questions users were asked at authentication time, and the size of the question pool. We do not know how many answers must be correct before a user is authenticated, nor whether approximate matching is in use.

There were several striking negative results in our sample. No institutions offered users the opportunity to write their own questions. Only four institutions combined security questions with email-based authentication. Only three of these institutions required email-based authentication. None utilized SMS messages to cell phones for fallback authentication. (A surprising omission, given that several institutions in our sample use SMS messages in other authentication contexts.)

3.3 Security Analysis

Sensitive security questions are reasonably easy to reason about. The answers are generally infeasible to guess, and so attackers must somehow learn them. Users generally know who has access to their PIN number or bank account number, and can change them relatively easily if they suspect compromise. In contrast, Social Security numbers have only limited utility for authentication. They are frequently compromised in institutional data losses, and are frequently sold in bulk on the black market [6]. The Social Security Administration has a policy of not assigning individuals a new number unless supplied with proof of fraud [18]. As a result, Social Security numbers cannot be reset to a “secret” value between disclosure and attack. Therefore, the pool of Social Security numbers available to attackers is likely to grow over time. However, Social Security numbers are by no means useless for security: they are likely able to defeat casual attackers, such as curious acquaintances, who might be able to guess or learn answers to personal security questions. Further, requiring them will raise the bar somewhat on identity theft, at fairly modest cost to the user and to the bank.

In contrast, personal security questions are comparatively difficult to analyze. The questions themselves are far more varied than sensitive security questions, and attackers can learn or guess the answers in a variety of ways. Such questions, though, are commonly used in practice, and thus demand careful consideration and analysis.

4. PERSONAL SECURITY QUESTIONS

In this section, we describe our procedure for analyzing personal security questions, and the results of our analysis. We make some observations about trends in the question pool

4.1 Special cases and user choice

Most institutions that rely on personal security questions allow their users to choose the questions for which they will supply answers. However, at a few institutions, one or more security questions are mandated by the mechanism designer. In particular, four require users to specify their date of birth to authenticate, three require a ZIP code, and one requires

the user’s mother’s maiden name. These questions have the benefit of having unambiguous answers for most users. Unfortunately they have the drawback of being comparatively insecure, as will be discussed in a subsequent section of this paper.

4.2 Topics

We extracted a total of somewhat over 200 personal security questions from the sites in our study. Though these question varied widely, there were a number of general topics and specific questions that came up repeatedly.

Names of friends and family were a common topic. There were 34 questions about first names, 13 about middle names, and ten each about last names and nicknames. In total nearly a third of all security questions were about names of individuals. “Favorites” were also popular. Banks asked about such varied topics as “favorite culinary ingredient”, “the last name of your favorite president”, and “your favorite restaurant in college”. Such questions accounted for roughly a sixth of all questions in our sample.

In addition to the broad agreement about topics, there were also specific questions that were used by many different banks. Four banks, out of 11, ask about grandmothers’ first names. Six ask “What was the name of your first pet?”; a seventh bank asks about current pets. “Favorite sports team” and “high school mascot” each come up four times; “mother’s middle name” come up three.

In some cases, this matching appears to be more than coincidence: Roughly a quarter of the questions used by one bank were also found at one other bank. This correlation may reflect copying, or that both institutions acquired questions from the same source, or perhaps some other possibility. Whatever the cause, such copying is of some security importance, since it means that an attacker can operate more efficiently by attacking accounts at both institutions simultaneously.

4.3 Classification

We defined six possible weaknesses in personal security questions. Three of the weaknesses we looked for — inapplicability, ambiguity, and lack of memorability — diminish the usability of the question; these usability concerns are essentially the same as those pointed out by Just in [12]. In all three cases, users can likely determine at the outset whether such questions are useful for them. A moderate number of unusable questions are not a significant usability concern. However, such questions do reduce a user’s choices, and thus increase the odds that a user will be forced to choose a weak question. Further, they risk confusing users, or overloading them with useless options, and thus may make it harder for users to select good questions.

Another three weaknesses, guessability, attackability, and automatic attackability — reduce the security provided by a question. We define these terms briefly below, and include a detailed discussion of our classification approach in Appendix I. Our tagged data is available online ³.

Inapplicable Some security questions are simply inapplicable to a large fraction of the public. For instance, “Which high school did your spouse attend?” is inapplicable to unmarried individuals. “In what city is your

³from <http://www.eecs.berkeley.edu/~asrabkin/securityquestions.tgz>

vacation home?” is also widely inapplicable. We classified a question as inapplicable if a recognizable demographic, equal to at least 15% of the public, would be unable to use them. This threshold was chosen somewhat arbitrarily, however there seemed to be only a few edge cases. Only children exceed the threshold, orphans do not.

Not Memorable Some security questions have answers that comparatively few individuals would reliably recall. Membership in this category is of course hard to judge, and we only labeled the most egregious offenders this way. One representative example was “last name of your kindergarten teacher?”

Ambiguous While inapplicable questions sometimes have no truthful answers, ambiguous questions have too many for a significant fraction (at least 20%) of the public. We classified a question as ambiguous if much of the public could truthfully give more than one answer to the question, or if answers shift rapidly over time. We ignore the possibility of small lexical variation, and restrict this category to questions that admit several semantically distinct true answers.

Guessable Many questions have answers that can be guessed with significant probability (in excess of 1%) even without any knowledge about the user. For instance, 30% of Americans marry between 25 and 30, and so a random number in that range would be a good guess for the question “How old were you when you were married?” We classified a question as guessable if we could identify an answer that was likely to be correct in excess of 1% of the time for a random member of the American public. The 1% threshold is often used in assessing the security of authentication techniques. Note, however, that we are using it to measure the strength of particular questions, not the security of a complete system.

Attackable For other questions, an attacker who knows the victim’s identity can learn an answer with substantial probability. For instance, a victim’s resume would reveal the answer to “with which company did you hold your first job?” Employment histories are by no means secret; they are commonly listed on web pages, biographic descriptions, resumes, and the like.

Automatically attackable Sometimes, the above process can be automated. For instance, “what year did you graduate from college” has an answer that can be automatically mined from Facebook profiles. We classified a question as automatically attackable if it had an answer that would be visible in the structured portion of a user’s profile page on Facebook or similar social networking websites. Date of birth and ZIP code, which are often mandatory questions, fall into this category.

4.4 Security of personal security questions

As discussed above, personal security questions are often a key part of bank authentication mechanisms. While few banks rely solely on such questions, it is valuable to analyze the security of questions in isolation, since such an analysis is necessary in order to understand the security of mechanisms overall.

We examine the security of the mechanisms, assuming they are used precisely as directed, and assuming users choose which question to answer uniformly at random from the space of offered questions. Many users no doubt are careful about choosing their security questions; however, the frequency of poorly-chosen passwords suggests that users seldom go to extra effort to improve security. Interestingly, no bank we examined gave users explicit guidance on choosing a security question: users were not encouraged to pick hard-to-answer questions. In addition, our user survey suggests that many users treat memorability, rather than security, as the dominant factor in choosing security questions.

We assume throughout that an attacker has guessed or obtained the real name of the targeted user, the user’s bank, and also the username and/or Social Security number (as needed). These assumptions are not unduly pessimistic: Social Security numbers are frequently compromised in institutional data losses. Our user survey (discussed below) suggests strongly that users often pick guessable usernames for online banking. And while an attacker may not know a user’s bank a priori, there are often few enough banks in a given geographic region that an attacker who knows the user’s approximate location can try every likely institution.

4.4.1 Random guessing

A number of banks offer users the choice of personal security questions which, if answered honestly, have easily guessed answers. As discussed above, we arbitrarily chose a 1% chance of guessing the right answer as our threshold of guessability. For most of the institutions we examined, the fraction of guessable questions was at least 33% (See Appendix II for details). Thus, an automated attack, not using any personal information, might be expected to succeed at least 0.3% of the time on the first guess for these institutions assuming the questions asked during authentication are a random sample of the available questions.

Subsequent guesses will raise this chance, though the authentication mechanism will likely lock out repeated attempts, or trigger alarms, beyond a certain threshold. In our experience, banks universally allow at least second and third tries; therefore, a success probability in excess of 1% might be expected. This probability must be put into perspective: The attack in question, since it requires no personal information about the target, can be conducted automatically, and on a large scale. Armed with a set of user names and Social Security numbers (or whatever other identification information is needed), an attacker can attempt to compromise many thousands of accounts in parallel, with negligible cost per target. Even a low success probability against any particular user could support an economically viable attack. Worryingly, such an attack could be difficult to detect against the steady background of legitimate fallback authentication attempts.

4.4.2 Automatically using online information

An overwhelming majority of today’s college students and recent college graduates maintain an account at some social networking site, such as Facebook, MySpace, or LiveJournal [4]. These sites allow users to expose structured information about themselves, such as their educational background, age, birthday, and friends, via their personal profiles. This information can significantly help attackers seeking to fraudulently authenticate. Roughly 12% of our question sample

was automatically attackable, meaning that answers to those questions could be found on a social networking site. The common ZIP code and date of birth questions fall into this category. While it is not found on most social networking sites, one’s mother’s maiden name can often be discovered from public records [9].

While this information is sometimes restricted by privacy policies, an attacker might very easily control enough compromised Facebook accounts to get around this barrier. Social networking sites are not typically viewed as needing strong protection, and few of them use SSL extensively to prevent password interception. Perhaps even more seriously, applications built on top of social networking sites have access to user profiles. A malicious or compromised application could readily leak large volumes of user data to attackers. While recent work has proposed a more secure model for such embedded applications, they represent a serious risk to use privacy at present [5].

An attacker likely will obtain a set of names and social security numbers, or names and bank account numbers, and separately have a collection of personal information, also indexed by name. These tables cannot simply be joined, because many individuals share the same name. (Though note that there may be enough rare names to constitute a large vulnerable population.) Several remedies are open to an attacker, including random guessing. Social Security numbers are not assigned randomly, but instead are assigned by blocks to particular regions at particular times [20]. An attacker may be able to use the additional information conveyed by a Social Security number to join it with a purloined user profile. Any additional information attackers have associated with a social security number — for instance, geographic location, email address, birth date, and so forth — will also aid the attacker.

4.4.3 *Dedicated Human Attackers*

A still more potent attack is that directed against some particular known user by a reasonably dedicated human. A great deal of personal information is available online in unstructured or loosely structured documents. Archival copies of old personal web pages, short newspaper profiles, club membership rosters and the like are all potent sources of personal information to a human attacker, and are all growing in volume and coverage. While reliably answering personal security questions using these sources is beyond the reach of today’s commodity information retrieval techniques, human adversaries are able to make use of them.

Questions such as “what is your home town” are comparatively easy for humans to answer, if the result is indicated by a document in the first few pages of search engine results. Names of pets or family members are not viewed as private, and are often made public via personal web pages and the like. Insidiously, users may have little awareness or control over online information about themselves either published by others, or published and archived, making it difficult for users to assess and minimize their risk. Genealogical information, for instance, is often published without the subjects being informed; old personal webpages or discussion list emails may be available through archival websites.

We do not attempt to compute a probability in this context — while we had definitions for “guessable” and “automatically attackable” questions that allowed us to assign a probability to attacker guesses, we have no such probabilistic

standard for “attackable” questions. Such a definition would be difficult in any event; the degree to which a question is attackable depends not merely on the user, but also on how persistent, and how clever, the attacker is.

4.4.4 *Personal acquaintance*

Personal security questions of the sort used today appear unlikely to keep out an attacker with intimate knowledge of the target, such as a friend or former spouse. The two most common topics for security questions were personal preferences (favorite sports teams, restaurants, etc) and the names of family members such as uncles, siblings, and cousins. In both cases, this is information that a personal intimate (and especially a family member) would very likely know.

In 1991, Haga and Zviran measured the ability of romantic partners to guess security question answers [10]. They succeeded 38% of the time in correctly answering personal security questions about their partners, half as often as those partners themselves. We note, however, that their methodology only allowed one guess, and that subsequent guesses will raise this success rate. Further, online sources of personal information may be of help even to attackers who are well acquainted with their target; thus, we suspect a repeat of their experiment would show an improvement in the success rate for fraudulent authentication by friends and family.

4.5 Usability-security tensions

There is a tradeoff between security and usability in question-based authentication. Quite often, users will not recall or type their answer perfectly, and as a result users must be allowed to make multiple attempts before being locked out. All else equal, the more ambiguous the question, the more attempts users must be granted, in order for the system to have a given success rate. A generous lockout threshold, in turn, impairs the security of the system. Thus, the presence of ambiguous questions exacerbates the tension between security and usability of the authentication mechanism. Since a large fraction (roughly 30%) of the questions we encountered were ambiguous, this tension is likely significant in question design.

A similar usability-security trade-off applies to the memorability of answers to security questions. Information that individuals recall easily is more likely to be recorded online than information that even the individual in question cannot recall. In our sample, slightly over 70% of memorable questions were attackable or automatically attackable, while only 25% of the non-memorable questions were attackable. (Most of the non-memorable but attackable questions concerned genealogical facts that a user might not remember, but that might be recorded in a public database.)

The usability-security tension in authentication is not confined to personal security questions. Account numbers are much harder to guess than usernames, but less convenient for users who are unlikely to memorize nine-digit numbers but will remember a frequently-used string.

5. USER SURVEY RESULTS

User behavior significantly influences the security of a system. By picking easily guessed answers or predictable usernames, users expose themselves to attack. Conversely, users can sometimes gain security by picking answers that are not literally true, and that are therefore hard to learn. This can increase the security of an authentication system be-

yond that which it would have if users precisely followed the system's instructions.

We are unaware of any recently published work on how users answer personal security questions, and attempted to find some preliminary answers by conducting a small survey. We prepared a questionnaire addressing several aspects of users' experience with online banking, and in particular, with personal security questions. We asked about their choice of user names, their habits in selecting and answering security questions, their difficulty in remembering answers, and their overall feelings towards such mechanisms.

Of our sample of 46 users, around 85% reported using online banking "often". Three-quarters of our sample were college graduates. Most described themselves as having "extensive" experience with computers; a fifth had taken courses in computer security. This sample is of course not representative of the general public. While this undoubtedly biases our results, they remain useful: a population heavily weighted towards computer science and computer security students is almost certainly at least as cautious about security as the general public. Our results thus represent a crude upper bound on the public's security consciousness. Our results are suggestive, but not conclusive. Nonetheless they are sufficiently relevant, and sufficiently persuasive, that we present them here.

First, the users in our sample seemed remarkably casual about the security of their online banking. Only 7% of users claim to worry "a lot" about security; half worry "some" and 43% worry "very little". Only 20% claimed to put "a lot" of thought into their choice of questions. Only 44% said security is a very important factor when choosing user-selected security questions.

Second, our study suggests that usernames are easily guessed. Roughly two-thirds of respondents admitted to basing their usernames on their real names; 40% admitted to using the same username both for online banking and for public services such as email and instant messaging. (Usernames for these public services, in turn, can be automatically harvested from online directories and social networking sites.)

Third, our results suggest that concerns about usability of security questions are well-founded. 70% of respondents claimed memorability is a very important factor when choosing user-selected security questions, in contrast to 44% who claimed that security was a very important factor.

Last, users are typically honest when answering security questions. 38% claimed they always give truthful answers, 18% "seldom lie", 31% "sometime lie", and only 13% "usually" falsify answers to such questions. Designers of security questions, then, cannot rely on users to pick secure answers to weak questions.

6. EVALUATING AND MITIGATING THE THREAT

Current security question schemes are vulnerable, to a greater or lesser extent, to a variety of adversaries. Even an adversary with no knowledge except general information about answer frequencies can succeed non-negligibly often. Adversaries able to mine social networking websites can do much better, and dedicated human adversaries can answer a large fraction of currently used security questions with significant probability.

6.1 Severity of the threat

Despite the apparent weakness of the fallback authentication mechanisms used by many banks, there have been no publicized cases to date of large-scale attacks on personal Internet banking via this route. We offer two possible explanations of this. First, that other attacks, such as phishing, are more profitable or easier for attackers to mount. Second, it appears that attackers have difficulty removing money from compromised personal bank accounts, and thus prefer alternate targets and attack techniques. Financial institutions utilize a number of techniques to detect suspicious transactions, which may pose a larger barrier than the mechanisms for authenticating users. Consequently, the risk of direct economic loss from attacks on online banking authentication mechanisms may be fairly modest. Even so, compromised bank accounts are frequently bought and sold by the criminal underworld, suggesting attackers do indeed find value in them [6].

Even without being able to readily transfer money from compromised accounts, attackers have other avenues to exploit such security breaches. An attacker with enough compromised online brokerage accounts could induce changes in the price of a chosen stock, and profit from this price shift. Spammers profit in a similar way by manipulating stock prices via stock tout-driven trading [7]. Some users may be vulnerable to blackmail based on the past transaction history recorded by their online banking site. Others may be intimidated by the threat of disruption to their bank account, even though they are indemnified by their bank against fraud. Many online banking sites make check images available, which include bank account and routing numbers. This information may be sufficient for attackers to subsequently arrange fraudulent money transfers.

There are fundamental problems with conventional personal security question schemes. Users will seldom share personal information that is truly secret with a bank. As a result, it is hard to imagine banks posing security questions about a patron's medical or sexual history. Personal security questions must therefore ask about information that is not truly private, but that has not yet been made publicly accessible. This means that the answer must either not have been shared despite being in principle sharable, or that finding the answer requires solving a hard information retrieval problem.

One of the hallmarks of the Internet age is that users are willing to share a great deal of personal information online, some of it quite intimate. Another is that search technology improves rapidly and unpredictably. Even today, tools such as Maltego [1] aim to do just the sort of cross-database join that an attacker would perform in order to identify a victim uniquely. As a result, the long-term prospects for question-based security of the form seen today appear dim.

6.2 Possible Defenses

In a sense, much of the research in this field seeks to evade the objection posed above by asking questions whose answers cannot be readily discovered by an attacker. One way to do this is to rely on questions with ephemeral answers. An example of this is the use of adaptive challenge questions [2] based on recent browsing history. Another is to ask users for durable information they may not consciously remember and are unlikely to have recorded in a public machine-readable form. The preference-based questions suggested in [11] are

an example of this approach.

There are several more modest steps that could improve user experiences with security questions. One step that banks should take is to explain to users what the security consequences of their answers are, and that answers should be private and unpredictable. Another is to avoid asking questions that lend themselves to predictable answers. There is little justification for questions such as “Last name of favorite president?” that invite users to give easily guessed answers. The authentication system can, however, detect these weak questions by examining the distribution of answers. Banks often ask users to supply answers for several security questions, and therefore can choose which to ask at authentication time. It should be possible to preferentially issue users a challenge question with an unusual correct answer. This would reduce the success rate of random guessing.

Information sharing between financial institutions would help detect certain attacks. Absent such sharing, an attacker can use a compromised host or an illicitly learned question answer against many different institutions, hoping to score a hit. Exchanging lists of suspicious addresses, or suspicious login attempts, would help damp this attack. Some commercial products, notably RSA’s Identity Verification, already do these sorts of comparison [19].

A fairly modest step to prevent automated guessing would be to require CAPTCHAs, in order to increase the cost of low-probability guessing attacks. However, CAPTCHAs may offer only modest benefits, since they are vulnerable to both replay attacks and to attack by low-paid humans.

A more radical approach could effectively embed CAPTCHAs inside security questions. Institutions could ask a question about an image or audio file that had been previously specified by the user. Rather than asking users the name of their first pet, users could be asked to upload an image of their first pet (or child, or grandparent), and associate with it an answer to the question “what is the name of the pictured individual?” With the proliferation of digital cameras, microphones, and broadband connections, the hardware requirements for such questions are steadily becoming pervasive.

This is not a perfect solution. Image-hosting sites such as Flickr and Facebook expose large volumes of images, tagged with user IDs and captions. If similar or identical images appear on these sites and in security questions, attackers, either human or software, may be able to guess the answers to the associated security questions. This sort of question, by integrating multimedia content, makes the attacker’s information retrieval problem much harder, since searching by image or voiceprint is beyond current technology.

7. CONCLUSION

This paper has had two goals. First, to document the current state of commercial identity-based authentication, and in particular, the sorts of personal security questions in current use. Second, to analyze those questions in the light of today’s information-rich Internet.

We believe that personal security questions, as currently used in fallback authentication in online banking, are surprisingly weak. Optimistically, this suggests that even simple security questions are useful in practice in authenticating users. Less positively, it suggests that even institutions with money on the line have difficulty designing high-quality au-

thentication questions. If current trends continue, questions of the form used today may become dangerously insecure.

That said, a good deal more work is needed to fully assess the security of personal security questions. User behavior and preference remains unclear. A substantial study would be needed to truly demonstrate how security questions are used in practice. Ideally, such a survey would be conducted on live data, and observed user behavior would be correlated with demographic information to determine how different sub-populations behave. Unfortunately, such studies cannot be readily conducted with public resources, limiting academic work on the topic. A limitation of current research into user behavior that most studies, including this one, have used highly atypical populations, generally students and staff at major universities. Ideally, researchers would study the behavior of “typical” user populations. Alas, there is a paucity of data describing the “typical” user, limiting work in this area.

The quality of personally sensitive data available via the web is also uncertain. No large-scale study has been done of the accuracy and volume of such data, nor of the ease of correlating a preexisting list of names and attributes with social network data. As information retrieval improves, it may be possible to mine unstructured text on the web for answers to questions. However, the efficacy of such approaches for answering personal security questions about arbitrary individuals remains uncertain. One approach may be to establish some standard benchmarks to assess the hardness of the information-retrieval problem posed by security questions. The two key metrics applicable to the topic of this paper are the fraction of correct answers to personal security questions, averaged over some set of users, either by algorithmic or human agents.

Last, an analysis of security questions, alone, may result in a distorted view of authentication systems. Personal security questions are only a part of most fallback authentication schemes. Activity profiling, secret security questions, and out of band verification via email or SMS have important roles in safeguarding user data and personal finances. The relative strengths of these techniques have not been explored in the public literature, and it remains to be seen whether these techniques will be sufficient in the future.

Fortunately, a number of promising alternatives to today’s security questions are in development. These techniques may significantly strengthen future authentication mechanisms. Thus, research in this area has the potential to mitigate vulnerabilities before they become serious threats.

Acknowledgments

We wish to thank David Wagner, Vern Paxson, and Doug Tygar for helpful discussions about authentication. Steve Houston and Mao Ye contributed to an earlier version of this paper and collected a significant portion of the data. Markus Jakobsson supplied invaluable advice and encouragement.

8. REFERENCES

- [1] Maltego. Available online: <http://www.paterva.com/maltego/>.
- [2] F. Asgharpour and M. Jakobsson. Adaptive Challenge Questions Algorithm in Password Reset/Recovery. In *First International Workshop on Security for Spontaneous Interaction: IWISSI '07*, September 2007.
- [3] J. Chao. Trend in canine names reflects the times. *San Francisco Examiner*, page B1, Oct 12 1997.
- [4] Facebook. Statistics. <http://www.facebook.com/press/info.php?statistics>.
- [5] A. Felt and D. Evans. Privacy Protection for Social Networking Platforms. In *Web 2.0 Security and Privacy 2008*, May 2008.
- [6] J. Franklin, V. Paxson, A. Perrig, and S. Savage. An inquiry into the nature and causes of the wealth of internet miscreants. In *Proceedings of the 14th ACM conference on Computer and Communications Security*, 2007.
- [7] L. Frieder and J. Zittrain. Spam Works: Evidence from Stock Touts and Corresponding Market Activity. *Berkman Center Research Publication*, 2006.
- [8] S. Garfinkel. Email-based identification and authentication: an alternative to PKI? *IEEE Security & Privacy Magazine*, 1(6):20–26, 2003.
- [9] V. Griffith and M. Jakobsson. Messin' with Texas: Deriving mothers maiden names using public records. In *Applied Cryptography and Network Security (ACNS)*. Springer, 2005.
- [10] W. Haga and M. Zviran. Question-and-answer passwords: an empirical evaluation. *Information Systems*, 16(3):335–343, 1991.
- [11] M. Jakobsson, E. Stolterman, S. Wetzels, and L. Yang. Love and authentication. In *CHI '08: Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, pages 197–200, New York, NY, USA, 2008. ACM.
- [12] M. Just. Designing and evaluating challenge-question systems. *IEEE Security and Privacy Magazine*, 2(5):32–39, Sept.-Oct. 2004.
- [13] M. Mannan and P. van Oorschot. Security and usability: The gap in real-world online banking. In *New Security Paradigms Workshop (NSPW'07)*, September 2007.
- [14] NYC Department of Health. "Health Department Announces Most Popular Dog Names and Breeds of 2005". Available online: <http://www.nyc.gov/html/doh/html/pr2006/pr122-06.shtml>, December 27 2006.
- [15] Office of the Privacy Commissioner of Canada. "Guidelines for Identification and Authentication". Available Online. http://www.privcom.gc.ca/information/guide/auth_061013_e.asp, October 2006.
- [16] L. O'Gorman, A. Bagga, and J. Bentley. Call Center Customer Verification by Query-Directed Passwords. In *Financial Cryptography: 8th International Conference, FC 2004, Key West, FL, USA, February 9-12, 2004: Revised Papers*. Springer, 2004.
- [17] L. Saad. Lincoln Resumes Position as Americans' Top-Rated President. Available online: <http://www.gallup.com/poll/26608/Lincoln-Resumes-Position-Americans-TopRated-President.aspx>, February 19 2007.
- [18] Social Security Administration. "Identity Theft and Your Social Security Number". SSA Publication No. 05-10064. Available online: <http://www.socialsecurity.gov/pubs/10064.html>, October 2007.
- [19] Social Security Administration. RSA Identity Verification. Available online. <http://www.rsa.com/node.aspx?id=3347>, 2008.
- [20] Social Security Administration. The SSN Numbering Scheme. Available online. <http://www.socialsecurity.gov/history/ssn/geocard.html>, Not Dated.

Appendix: Our classification scheme and some notable questions

Evaluating questions is a partially subjective process. We did go to some lengths, however, to make our judgments consistent and reasonable. We were aided by the fact that many questions are substantially similar to one another, and so only a comparatively modest number of judgments needed to be made. This appendix documents our classification in more detail, and describes significant clusters of questions that we observed in our sample and why they were classified as they were.

Guessable

Sometimes, the problem is blatant and the degree of vulnerability easily evaluated. One online bank asked "What is the last name of your favorite president?" Polls indicate that Lincoln is viewed as the greatest American president, with around 20% of the public ranking him first. [17] An attacker with the optimal guessing strategy would thus likely succeed one try in five. Even random guessing of presidents would succeed in excess of 2% of the time, since there have only been 38 presidential last names. Likewise "What was the make of your first car?" suffers from an overly small space of answers, and clustering of results around a few popular answers.

Sometimes the problem is more subtle. Many questions ask for the name of a pet, or a first pet. As it happens, the most common pet name in America is Max, and in the cities where the statistics are available, 1% of registered pets are named Max [3, 14]. Likewise, first names are drawn from a known distribution, and if asked for a child's name, Michael is a good guess: three percent of newborn boys are named Michael, and so around 1.5% of recently-born children will have that name. Consequently, we classified questions asking for first names and pet names as guessable.

We offer some caveats, though, about the guessability of first names. While American first-name frequency data is readily available from the Census Bureau and elsewhere, these frequencies shift over time. Questions, by asking about

particular relatives, give only rough guidance to the age of the individual in question. The census data may also be misleading in another respect: many online banking users have grandparents who not live in the United States, and whose first names will not obey American first-name frequencies from the era.

Preferences are often guessable. Several sites, including Facebook, track which books, movies, and the like are most commonly listed as favorites. This information would aid attackers in making optimal guesses, particularly if they have information about the age or location of their targets. We found 24 guessable “favorite X” questions.

Inapplicable

We classified 105 questions as “often inapplicable”. Most of these questions made strong assumptions about the family or lifestyle of the user. A large fraction asked about spouses, weddings, siblings, and children. Questions referring to spouses and marriage accounted for 30% of the total pool of often inapplicable questions. Other inapplicable questions take for granted the existence of middle names, or nicknames. Still others assume that users can identify a favorite sports team or athlete. One bank asked “In what city is your vacation home?” — surely a question with limited applicability.

Ambiguous

Some personal security questions simply fail to have unique answers. For many individuals, “what is name of a college you applied to but did not attend?” describes several institutions, as does “what is the name of a school you attended?”

As an important special case, questions about preference (such as favorite actor, restaurant and so forth) were marked as ambiguous, and accounted for just over half of the ambiguous questions in our sample. Personal preferences of this sort often shift from month to month, and therefore users may require many attempts in order to hit upon the favorite that they had in mind when they initially answered the questions.

Attackable

A question was marked as attackable if it relied on information that is commonly revealed online, in some searchable way. By far the largest subset of attackable questions asked for names of family members, such as grandparents, in-laws, children, and the like. We considered these to be attackable because this information is commonly available in genealogical references, biographical summaries, and the like. Previous work has shown that genealogical information can often be culled, or automatically deduced, from public records [9].

Another large subset of attackable questions asked about educational history: schools attended, major during college, and so forth. Still others were questions for which a good guess could be made, given enough information about the user’s address or birthplace. For instance, “the hospital your youngest child was born in” can often be predicted if the child’s city of birth is known. This information may also be available from public records.

A last set of attackable questions required a two-step process to answer. At present, humans are much more effective than purely software-based agents in discovering answers to “two-phase” questions, that require both discovering a fact, and then performing some sort of inference on it. While

“what was your high school mascot” might be beyond the reach of current automated techniques, a human could learn the answer by first discovering the victim’s high school, and then looking up the school’s mascot.

Automatically Attackable

In some cases, the attack procedure discussed above can in fact be automated, thanks to the increasing prevalence of structured data about personal history on social networking websites. Perhaps the cleanest example of an automatically attackable question is “In what year (YYYY) did you graduate from high school?” This information is commonly found on Facebook pages, LiveJournal personal profile pages, and the like.⁴

We considered preferences to be automatically attackable if they were solicited by social networking sites, and in particular by Facebook. Most social networking sites encourage users to list their favorite TV shows, songs, and so forth. Since personal preferences often shift in difficult-to-reconstruct ways, even legitimate users will need to guess several times. We suspect that trying each of a user’s “favorite songs” listed on Facebook is an effective attack strategy for questions of this sort.

While questions requiring significant inference or multi-step research to answer are not automatically attackable, we did consider questions automatically attackable if an algorithmic process starting from public information would reveal the answer with high probability. The answer to “What college was your college rival” can be worked out with a fairly modest lookup table, coupled with the (public) information of which college the target user attended.

We did not classify names of family members as automatically attackable, although in some cases they may be, depending on the public records made available online in a particular jurisdiction, and the degree to which they can be automatically processed.

⁴This question is also guessable by our terms: almost all answers will be a year within the last 60, and most will be within the last 20 or so.

Appendix II: Summary Tables

Key:

DOB = Date of birth

MOB = Month of birth

CVV = three digit credit card verification number

MMN = Mother's maiden name

Email = Check for ability to receive email

Table 1: Summary of Authentication Requirements

Site	Category	SSN	User name	Security questions	acct number	PIN	email
Ameritrade	brokerage		yes	1+ MMN			
AmTrust	online bank	last 4 digits	yes	1+ ZIP			
Bank of America	major bank	or uname	or SSN	1	yes	yes	
Chase	credit card	yes			yes		yes
Citi Cards	credit card			unprompted recall	plus CVV		
Discover	credit card	last 4 digits		DOB	plus exp. + CVV		
Emigrant Direct	online bank	last 4 digits	yes	5 and MOB			
Fidelity	brokerage	or uname	or SSN	1+ DOB	or known host		
FNBO Direct	online bank		yes	1	yes		
GMAC Direct	online bank	or email	yes	1			or SSN
ING Direct	online bank		last 4 digits	2 + ZIP + DOB +phone #			yes
M and T	regional bank	yes	yes	2 or acct+pin	or sec. Qs	or sec. Qs	
Presidential	online bank			3	yes		
USAA	regional bank	last 4 digits	yes	ZIP+country+DOB			
Washington Mutual	major bank		yes				yes
Wells Fargo	major bank	or uname	or SSN		yes	yes	

Table 2: Institutions without automatic password reset mechanisms

Institution	Category
Cal State 9	Credit Union
Advantis	Credit Union
UFBDirect	Online Bank
ScotTrade	Brokerage

Key:

Secure = Neither guessable, nor attackable

A.A. or G. = either automatically attackable or guessable

Table 3: Question Statistics

Bank	Ambig.	Not Memorable	Inapplic.	Guessable	Attackable	Auto. Attackable	A.A. or G.	Secure	Total
Ameritrade	4	0	1	0	1	1	0	3	4
Amtrust Direct	23	24	65	38	73	12	47	35	107
Bank Of America	1	0	11	10	16	4	13	8	22
Citi cards online	11	4	7	7	7	2	7	12	21
Emigrant Direct	4	0	1	3	5	3	4	0	5
FNBO Direct	6	0	6	8	6	1	8	1	10
Fidelity	3	0	3	2	3	0	2	1	5
GMAC Direct	1	1	5	3	13	4	7	8	14
ING Direct	2	0	2	4	10	3	6	2	9
M and T	1	0	3	2	4	1	3	3	6
Presidential	12	0	1	8	6	4	8	3	12
Total	68	29	105	85	144	35	106	75	215

Note: "Attackable" numbers include "automatically attackable" questions.

Note 2: Citi Cards did not use these security questions for routine fallback authentication, however they are used to authenticate "suspicious" logins or transactions, and we include them in this study, as they are representative of security questions used in financial institutions.