"Wait, Do I Know This Person?": Understanding Misdirected Email

Emilee Rader
Department of Media and Information
Michigan State University
emilee@msu.edu

Anjali Munasinghe
Department of Computer Science and Engineering
Michigan State University
munasin2@msu.edu

ABSTRACT

Email is an essential tool for communication and social interaction. It also functions as a broadcast medium connecting businesses with their customers, as an authentication mechanism, and as a vector for scams and security threats. These uses are enabled by the fact that the only barrier to reaching someone by email is knowing his or her email address. This feature has given rise to the spam email industry but also has another side-effect that is becoming increasingly common: misdirected email, or legitimate emails that are intended for somebody else but are sent to the wrong recipient. In this paper we present findings from an interview study and survey focusing on characteristics of misdirected email messages, possible reasons why they happen, and how people manage these messages when they receive them. Misdirected email arises as a result of signifiers (usernames) which were selected by people for social and self-representation purposes, that are also used by machines for addressing. Because there is no mechanism for dealing with misdirected emails in a systematic way, individual recipients must choose whether to take action and how much effort to put forth to prevent potential negative consequences for themselves and others.

CCS CONCEPTS

Human-centered computing **Empirical studies in collaborative and social computing**.

KEYWORDS

misdirected email, authentication, usernames, mixed-method

1 INTRODUCTION

Despite the popularity of social media and messaging apps, email remains a tool and an infrastructure that is used widely and for a variety of purposes. Email is used to communicate with others and receive information from organizations [2]; it functions as a personal archive and a to-do list [1]; and it is used as a personal identifier and authentication mechanism [9]. These uses are enabled by the fact that the only barrier to sending an email message to someone is knowing his or her email address. This feature has given rise to the spam email industry but also has another side-effect that is becoming increasingly common. Imagine opening your email to find multiple password reset notifications for an account that you do not remember creating, messages from an online dating service you do not remember signing up for, billing reminders for a credit card you did not open, or photographs of other people's children that you do not recognize. These are all examples of misdirected email: legitimate email messages that were intended for someone else, but were sent to you instead.

Misdirected email happens when the sender, whether it is an individual or a system, sends an email message to an email address that does not belong to the intended recipient, but in fact belongs to a totally different person. This creates a problem for all parties involved, but in different ways. The sender of the misdirected email is unable to reach the correct recipient and might not even know that this has happened; the correct recipient does not receive the email; and, the person who actually receives the email has access to information and communications not intended for them. This can expose the sender and intended recipient to unwanted disclosure and misuse of private information, and creates an additional burden on the actual recipient of figuring out what, if anything, to do about the email.

Email addresses perform three different kinds of functions. The first is a social function: email addresses, and usernames in particular, are a representation of a person's social identity online that is used for communicating with other people and organizations, at least in cases where the email user is allowed some discretion in choosing his or her username [13]. The second function is a technical addressing function, that allows messages to be routed and delivered to the correct mailbox [20]. And the third is an authentication function [24]: people use email addresses as usernames for other online accounts, and access to a valid email account is frequently

ACM ISBN 978-1-4503-5970-2/19/05.
https://doi.org/10.1145/3290605.3300520

used as an authentication factor ("something you have [11]). Misdirected email presents a unique opportunity to explore and describe unintended consequences that occur at the boundaries between the social and the technical.

We contribute to the literature on sociotechnical systems an investigation focusing on recipients of misdirected email. Using data collected via a semi-structured interview study, a survey, and a single-user corpus of misdirected email, we define and characterize misdirected email from the point of view of people who receive it. We also describe their reactions to it—how they feel about it, what they do about it, and the problems it causes them, as well as the opportunities it provides to help and feel connected to others. We conclude by discussing our participants' speculations about why misdirected email happens, and connecting these observations back to broader issues that arise at the interaction between technology and people in sociotechnical systems.

2 RELATED WORK

There are few published papers about misdirected email. Most of this literature focuses on the likelihood and consequences of the threat of unwanted disclosure due to an email message being sent to the wrong person by mistake, and how to prevent this from happening. Misdirected email has consequences for both senders and recipients, and McLaughlin [19], writing in a law journal, described the risk that misdirected email could cause a professional ethics problem for both parties if it violates attorney-client privilege. Senders are therefore obligated to remedy the mistake as quickly as possible; recipients are obligated to notify the sender if they received information by email that was not intended for them. McLaughlin focuses on the potential consequences of these kind of unwanted disclosures, and on how both senders and recipients should handle these consequences.

Poole et al. [23] conducted a semi-structured interview study in which participants were asked about unintended disclosures, or "misclosures", they recalled committing. The most common type of misclosure their participants reported was sending misdirected email, and the participants attributed their mistakes to not paying enough attention to the recipient field in the email message. Sender carelessness when composing an email message was also the focus of several papers proposing automated ways to help senders be more accurate in specifying the recipients of email messages they were composing [3, 16, 26]. These projects took the approach that misdirected email can be detected before it happens assuming that unintended recipients are outliers when compared with the people one typically communicates with by email about certain topics. By developing a model of the content and contacts of each user of an email system, these papers argue it would be possible to notify the user about outlier recipients before sending an email. Most of these papers are theoretical; however, one described the implementation of this type of solution as an extension for Mozilla's Thunderbird email client. Based on a 4-week user study, the researchers concluded that system had been instrumental in helping four different participants (out of 26 in the study) identify misdirected email messages before they were sent [4].

A user interface focused intervention related to misdirected email was developed by Lieberman and Miller [15] who created the system "Facemail" to display recipient photos along with email addresses near the "To:" field of a message. This was intended to provide visual cues to the user about who the people are that the message will be sent to. In a lab experiment, they determined that this system would be helpful especially for situations in which the sender did not intend to "Reply-to-All", and that displaying faces helped participants notice potential misdirected emails before they were sent.

These papers treat misdirected email as a potential threat that should be mitigated by preventing senders from committing addressing mistakes that lead to unwanted disclosures. However, they do not provide any information about what the everyday experience of receiving misdirected email (rather than sending) might be like for people, or discuss how prevalent misdirected email currently is. As we will show in this paper, this problem is bigger and more complicated than just helping people pay more attention to avoiding typos when specifying recipients. The increase in email volume [8], and the shift of personal communication to platforms other than email (e.g., texting, messaging apps, and social media) [2], suggests that misdirected email may be changing in similar ways, becoming more common and less personal. Therefore, we conducted this research to describe the experience of receiving misdirected email, and better understand why it happens and how recipients manage it.

3 METHOD

The findings of this paper are based on three datasets: interviews, a survey, and a single-user corpus of misdirected email messages collected over more than a decade by one of the authors of this paper. The findings present descriptive evidence from all three datasets, although most of the Findings section focuses on the interviews. The interview study enabled us to collect detailed examples of misdirected email participants had received, but the small sample size and sampling frame focused on recipients did not allow us to draw conclusions about prevalence. We conducted the survey to learn more about how common misdirected email is over a broader sample. The corpus of misdirected emails is a single-unit case study [10], useful for this exploratory research primarily to describe a longitudinal perspective on misdirected email that the interview and survey methods

could not provide [7]. However, it should be emphasized that the corpus represents misdirected email messages received by only one person. A supplementary file is available with this paper and provides more information about the interviews and survey.

Interviews

We conducted 22 semi-structured interviews during October 2017. Participants were recruited in two ways. Twelve participants (7 women and 5 men) were recruited via snowball sampling starting from a study advertisement on Facebook. Ten (5 women and 5 men) came from a paid subject pool consisting of members of the community surrounding a large public university located in the Midwest region of the USA. Participants worked in a variety of professions including photographer, night auditor, clergy and programmer. Three local participants were affiliated with the university as a staff member, an instructor, and a graduate research assistant. The average age of participants was 36.8 years (*Range*: 23–63).

Eligible participants used email regularly, and were over 18 years old. Friends, family members, and colleagues of members of the research team as well as undergraduate students were not eligible to participate. In addition, the screening questionnaire asked about 5 different types of experiences related to misdirected email: "I have received an email message that seems like it was intended for someone else; I have sent an email message to the wrong person by mistake; I have been asked in an email to confirm an account that I don't remember creating; I have given out a 'throwaway' email address that was fake or did not belong to me to a website; and I have used a 'throwaway' email address that was fake or did not belong to me to sign up for an online account. We were seeking participants who had received misdirected email, as well as those who may have caused misdirected emails to be sent to others, because we did not know how common receiving misdirected email messages would be. Therefore, any response except "None of the above was considered eligible.

The interviews ranged from about 14 to 49 minutes (M=28 min), excluding the consent process and demographic post-questionnaire, and took place over the phone so that we could recruit a more geographically diverse sample. Questions focused on email use in general, how participants define and recognize spam email, and experiences with misdirected email. (The recruiting message, screening questionnaire, and interview protocol all described misdirected email as "an email message that seems like it was intended for [or meant for] someone else .) Most of the interview focused on what participants recalled about specific misdirected email messages they had received, and what happened next. Interview length varied based on how many examples of misdirected email participants mentioned. Many participants had several examples readily available and were eager to talk about them, because they receive misdirected email frequently. For the participants who did not, we began by providing a high-level description of misdirected email, similar to how it was described in our recruiting message, and proceeded with a sequence of 2-3 specific prompts until participants recalled an example they felt was relevant. Participants who needed more prompting typically provided an example of a misdirected email they had sent, or a spam message they had received that they thought could be a misdirected email. Each participant received a $15 Amazon.com gift card for participating.

Interviews were audio recorded and transcribed, and identifying information was removed. This meant that we needed to remove specific references to participants' email usernames, which were frequently discussed during the interviews. In the remainder of this paper, any [text in brackets] is something the participant said that has been altered or redacted to protect their identity.

We used an iterative, inductive coding approach in which all members of the project team conducted initial structural and thematic coding on the same two interviews. The analysis took place over 5-6 months during which the research team met at least weekly to discuss the coding in progress. The process we followed was based on MacQueen et al. [17]. We first did structural coding for broad topics. We then developed an initial list of codes inductively in the first round of analysis where all three members of the research team coded the same two interviews. This produced an initial set of codes. These were then refined by two members of the research team as they coded the same additional five interviews, during which time the entire team met several times to iterate on the set of themes we were coding for. After this, the set of themes had stabilized and was well understood by both coders, and they subsequently divided up the remaining fifteen interviews and each analyzed half of them. The themes focused on things like participants' approaches and strategies for dealing with misdirected email messages, their feelings and justifications for how they dealt with them, and their speculations about why they received the misdirected emails they discussed.

Survey

Survey data collection took place from March 20 to April 2, 2018. Respondents were recruited by Qualtrics.com using their panel service, with quotas for gender (50% men and 50% women) and age (18-29: 25%, 30-49: 38%, 50-64: 21%, 65+: 15%), in order to recruit a sample that resembles the population of US adults who use the internet on those two characteristics. The age quota was based on information from the Pew Research Center's Internet/Broadband Fact Sheet

from Feb. 5, 2018 about the age distribution of US adults who use the internet [22], and data from the US Census Bureau's 2016 American Community Survey about the age distribution of the United States population [25]. Eligible respondents were at least 18 years old and indicated that they had at least one email account on a popular free email service. Respondents who completed the survey received points from the online panel service worth approximately US$1-$2 that could be combined with the incentives from other surveys and redeemed for items like gift cards, frequent flyer miles, credit for online games, etc. The specific incentive amount was determined by the Qualtrics panel service.

The final dataset for analysis includes 380 respondents. Their average age was 45 (*SD*=17, *Range*=18–85). There were 196 women and 181 men in the final dataset; 3 reported "Other". 81% of respondents reported "white as one of the ethnicity categories that described them, and 80% reported that their income was less than $75,000 per year. The survey asked questions that were based on patterns we were seeing in the interview analysis, including reasons for creating new email accounts, how people choose usernames, and experiences with misdirected email. The median survey completion time was 10 min (*M*=13), including consent and screening.

Single-Case Email Corpus

We analyzed a corpus of misdirected email that had been collected for over a decade by the first author of this paper. The emails were received at the primary Gmail address of the author, who was an early adopter of Gmail and whose address is of the form firstname@gmail.com. As the messages were received, they were tagged with a label, and saved. The main criterion for an email to be collected was that it seemed like a real or legitimate email message, but was from a sender or organization that the account holder did not have an existing business or personal relationship with, or otherwise seemed to be directed at an intended recipient other than the account holder. The earliest message in the corpus is from September 9, 2005, and the latest from July 21, 2018.

We used Google Takeout[1], the interface for downloading an archive of one's data from Google products, to download an MBOX file[2] containing all messages from the email account that had been tagged as misdirected email. We then parsed and cleaned the MBOX file to identify each "To:" and "From:" address associated with each email, and ltered the dataset so that only the first message in each thread—the initial misdirected email, excluding any replies—was retained for analysis. The final dataset includes 2932 initial misdirected email messages from 1788 distinct senders that were received between Jan. 1, 2007 and Dec. 31, 2017. There were only three messages received before this time window, and emails from 2018 were excluded to make year-over-year comparisons more clear.

Limitations

These methods and datasets have several limitations. The interview study uses a convenience sample, and this means that no generalizations should be made about the amount and type of misdirected email our interview participants received. The survey sampling frame was focused on age and gender quotas to ensure diversity on those aspects, not representativeness for additional demographic categories. All of the data collected using both interview and survey methods is self-report, and as such represents the perceptions, attitudes and beliefs of the people who participated. Social desirability bias may have affected their willingness to self report some behaviors, such as whether they had used password reset links to take over the accounts others had created using their email address. The survey asked for respondents' first and last names and an email username on a free email service (one of the eligibility criteria) but not the full email address. They were informed that this information would be deleted once it had been used to compare the username with the name, but this may have caused some people to stop participating at that point in the survey. Finally, the misdirected email corpus, while containing messages similar to those described by interview participants, is a single-user case study and therefore not representative of all recipients of misdirected email.

4 FINDINGS

Is it Spam or Misdirected Email?

Definitions in the research literature describe spam email as unwanted or unsolicited messages sent in large numbers intended to sell a product or service [5, 12, 18]. Misdirected email could be considered a form of spam in that it is unsolicited; therefore, in our interviews we discussed with participants both how they define spam, and how they recognize when an email message they receive is misdirected. Contrary to typical definitions of spam, 18 of our 22 interview participants described spam email as messages having malicious or harmful intent, such as phishing, spreading viruses, or soliciting victims for a money-related or other type of scam. P06 described spam email messages this way: ...they're either like a phishing scam, or they're being used to kind of try to pretend to peddle something when in fact they're trying to get my personal information. So for most of our participants, spam emails were messages that made them feel suspicious

[1] https://takeout.google.com/

[2] https://www.loc.gov/preservation/digital/formats/fdd/fdd000383.shtml

or uncomfortable, and that they felt they needed to protect themselves from.

Misdirected email messages, while also unwanted, were described very differently from spam. Fifteen out of 22 interview participants (68%) said something about how misdirected email messages seemed like they were meant for someone else. For example, P14 said, "I don't know what's happening [to cause this] but it's just, I know they don't mean to write this to me. P03 described how he differentiates misdirected email from spam:

> "So I can't tell you a sort of a hard and fast rule. [Misdirected emails] are usually not too hard to spot because they're not asking me for anything, usually. Or they just feel like... a thing a person would really write as opposed to a thing that a person trying to get past the spam filter would write." (P03)

To figure out whether a misdirected email was intended for them or not, participants described focusing on the sender of the message, and the content of the message itself. For example, P01 described thinking, "Wait, do I know this person?" about the sender of one of the misdirected emails she received. P05 talked about receiving a message "that just felt very out of context, and I didn't really know what they were talking about"—the out-of-context aspect of it signaled to her it was a misdirected email. P02 said he had received "PDFs of contracts and all sorts of weird stuff that he said was obviously was not mine , but he also said that it sometimes takes him a minute after he starts reading a message to realize, "Oh, this... no, this isn't for me. As P16 described, "I just read them as if they were sent to me and then by the time I finish reading them I can tell they are not sent to me because I have no idea what they're talking about.

Misdirected Email Content and Frequency

Eight of our interview participants described some of the misdirected email messages they had received as including personal correspondence from complete strangers, such as an email scheduling a job interview, a reminder to show up in court for a deposition, and email threads about someone else's family reunion or birthday party. Some of these messages contained sensitive personal information that could be used to steal someone's identity or otherwise cause them problems, such as mortgage loan documents, credit card info, airline tickets, contracts, mobile phone bills, bank statements, and actual PayPal payments. A few said that they had been asked in an email for parental permission for someone else's child to create an account on a gaming website, and one had received a report card for someone else's child. Several also said they had been signed up for a university recruiting email list, political candidate email list, or neighborhood-related email list, none of which these participants remembered having any relationship with.

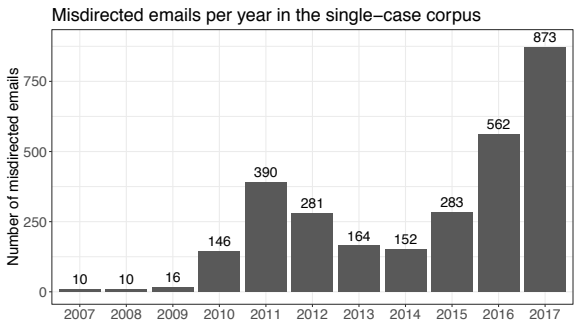

Figure 1: The number of misdirected emails in our single-case corpus per year. Note that this corpus contains misdirected emails received by a single user over 10 years, 2007-2017.

While not all of these eight participants mentioned how frequently they received misdirected emails, those who did provided estimates ranging from a few per month, to eleven misdirected email messages in the two days before the interview took place (P14). These participants' estimates are similar to the frequency of misdirected emails in our single-user email corpus, which most closely resembles the high end of participants' statements about frequency. In 2017, for example, the fewest misdirected emails in the misdirected email corpus in a given week is 5, and the most is 33. During 27 of the weeks in 2017, the corpus contains between 10 and 20 new misdirected emails each week. In addition, the quantity of misdirected emails in the corpus increased over time; Figure 1 illustrates this.

Eight other interview participants had only received a small number of misdirected email messages exclusively from someone who works for the same employer as they do, and very infrequently. These participants said that it was usually easy for them to recognize these mails as intended for someone else. For example, P07 worked at the same organization as a relative who had a very similar name. She said she could differentiate the occasional email intended for her relative, because the senders "were not normally people that I talked to before, and they were asking for information that I don't handle.

The remaining 6 interview participants each could recall only one or two examples of occurrences related to misdirected email, such as receiving a password reset notification for an account they did not remember creating, or sending a misdirected email message. None of them spoke about receiving email containing a stranger's sensitive personal information or personal correspondence, or receiving a work-related misdirected email message.

This wide range of experiences across the interview participants was surprising, and made it difficult to estimate whether receiving misdirected email is broadly common or not. To find out about how prevalent misdirected email might be from a larger, more diverse sample of people, we asked our survey respondents about whether they had ever experienced it. The survey included a choose all that apply' question that started by asking "Which of the following statements describe you?" and then listed the same set of statements that were used as screening questions for the interview study. 164 survey respondents (43%) said they had experienced either receiving an email message that seemed like it was meant for someone else, or had been asked to confirm an account that they did not remember creating.

Sympathy for Senders and Intended Recipients

Receiving misdirected email was sometimes an annoying and bewildering experience for our interview participants. P03 described, in an exasperated tone, how he had received 172 misdirected emails from the same woman starting in 2007 and ending in 2015. He had replied to these emails to ask the sender to stop, but felt like that didn't turn out to be effective. He said,

> "So here's [an example misdirected email]. 'Hey, [P03's name], I just finished taping a lesson of [redacted] but I need to convert the DVC to a DVD. Yikes. Can you help me this weekend to convert it?' So what's weird about that was... This is a person who is interacting with her. And I just kept thinking that at some point they're gonna talk about how he never responds to her emails, and at some point she's gonna realize that he's not getting them. And it just never happened!" (P03)

P13 also recounted how she repeatedly received misdirected email messages from "a grandmother or a grandfather type . She said that these emails were amusing,

> "...because they kept emailing to invite me to family dinners and stuff. And so I would reply and say, 'This is the wrong email address. I'm sure you would have a great time at Thanksgiving dinner but you're emailing the wrong person. And they clearly didn't really know how email worked, and so they would respond a lot." (P13)

Others talked about the misdirected emails they received as something that made them feel like they were part of something bigger than themselves. For example, P01, who received an email that someone else's child who had the same name as she did had received a prize in an art competition, said, And I was like, these parents have to find out, their kid is a winner! So I emailed them back and let them know that they have the wrong address... there was a sense of camaraderie. And even P03 talked about receiving misdirected email as if he were making the best of a situation that he wouldn't seek out—"If I had the ability to turn these off entirely I probably would"—but felt like "It's a good opportunity to just interact with people I otherwise would never encounter.

Many participants reported feeling sympathy for both the senders and intended recipients of the misdirected email messages they'd received, and a desire in at least some instances to help them by responding to let them know about their mistake. They were also motivated to do this by the hope that the sender would not do this again in the future. Fifteen out of 22 participants talked about responding to a misdirected email message. Messages that seemed more like personal correspondence were more likely to be replied to by participants, as well as emails containing someone else's sensitive personal information.

For example, P05 described realizing, "Oh, these are real people. They're actually trying to find a person. They're trying to talk to another individual. They've mistaken me for somebody else. She said she wanted to help them connect with the person that they were actually looking for. And P15 replied to a misdirected email from the director of a child's band camp concerning payment for their services, to let them know they had reached the wrong person. She described her motivation to do this: ...he had followed up a couple of times and finally I was like, oh, this is a real human who's expecting someone to pay her elementary school kid's band bill. Nearly all of the work-related misdirected email messages our interview participants talked about also involved some kind of follow-up from the participant. For example, when P10 received a work-related email meant for someone with her same name, she responded to it because "since it was work related, I knew that it needed to be passed on to the correct person cause it was important.

About 20% of our interview participants talked about actually tracking down or communicating with the intended recipient of a misdirected email message, instead of the sender. P13 talked about how after she received a cell phone bill that contained personal information including the intended recipient's phone number, she "texted them a screenshot of the email" to let them know someone was receiving their personal information and to encourage them to change the email address on the account. And P22 received a work email from a manager meant for a different employee. He said that he "looked to see who was handling that project, and then I sent the email to the individual I believed it was intended for, and CC'd the manager who had sent it. So in this case

he was able to track down the intended recipient because he had access to directory information through his employer.

Participants also talked about being careful with the personal information they had unwittingly received. P02 described receiving bank statements, transaction information, and account codes from a bank in Brazil that were intended for someone else. He said, "so I wound up talking to the bank. It was like, this is not the right email address, you guys need to not send out this stuff. P14 talked at some length about how it was creepy" to receive photos of other people's children: And then sometimes I feel creepy about getting photos of children, so I will email them to tell them that, you have the wrong person, please stop emailing me. P14 mentioned that she feels badly about having access to this information. She said this is because if they were her children, she wouldn't want strangers looking at photos of them. She felt like she was invading the privacy of the children and parents by receiving them, even though this was out of her control. She said,

> "Like, if I had a child I wouldn't really want photos of my children floating around to unknown people, so there's that too. But some of it is more like, can you please be better about your privacy with your own children's photos." (P14)

Auto-Complete and Sender Carelessness

Interview participants placed the blame for misdirected email messages mostly on the senders of the emails. This was partially because the participants themselves were guilty of having sent an occasional misdirected email, and they reflected on this when speculating about why they might have received the misdirected email messages they talked about during the interview. In fact, 18 out of 22 interview participants said they'd sent a misdirected email, and 89 out of 380 survey respondents (23%) also said that they'd done this. (These survey respondents checked the statement, "I have sent an email message to the wrong person by mistake.")

When interview participants described sending misdirected email themselves, they nearly all said that they just hadn't noticed that the auto-complete had chosen the email address of someone else that started with the same combination of letters as the intended recipient's address. P19 described it like this: "That's where I start typing in a name and it auto fills, and it auto fills with the wrong person, same first name, wrong last name, but I don't notice and it gets sent. Participants reported that sending a misdirected email message was a rare occurrence for them. They also universally believed that sender mistakes like this are unavoidable, and just part of using email. As P04 said, "I don't think you can change human behavior with software. You're gonna have to say to people, be a little bit more careful... But of course, people aren't interested in that, they're in a hurry, right?"

Name Similarity and Username Collisions

Auto-complete would not cause as much misdirected email if people had usernames that were sufficiently distinct from one another. Participants recognized this, and speculated that similarity between their own names and usernames and what they presumed the intended recipients' to be is part of why they end up receiving misdirected email. Most had specific and sometimes numerous examples that informed this hypothesis. For example, P13 talked about becoming familiar with specific intended recipients with whom she shared a name through all of the misdirected email she had received that was intended for them. She said,

> "The email address I use is my first name with the date of my birthday, so they're just like, they've mistyped one letter or something like that, so they're all people with my same name. So I know that there's one that lives in Florida and there's one that lives in Texas and there's one that lives in Vermont and there's one that lives in the UK, and so I know if I get someone, the email from Florida, I know that, 'Oh, that's so-and-so who lives in Florida, and I know that if I get an email confirming an order from The Gap and it's going to Texas, I know it's, 'Oh, that's [P13's first name] in Texas. So yeah, I definitely know that there's specific people." (P13)

Several other participants believed they had a common name, and that this combined with the username they had chosen for their email address—based on their name—was why they received so many misdirected email messages. For example, P03 said his name was "the most popular for babies born in [P03's birth year]" and that he has the canonical Gmail address" (firstname.lastname@gmail.com) for his name. He continued, And so as a result, every week to two weeks I get an email that's for a different person with my name. And P20, whose email address is also of the so-called canonical form, kept returning throughout the interview to the idea that before he started receiving misdirected email, he thought his name was uncommon. He said, "I've received enormous amounts of email meant for what is quite to my surprise, a rather large group of [people with the same first and last name as P20] around the world. P14 talked about how she has an Apple Me.com email account where the username is her first name, and she believes she receives misdirected email messages at that account because of other people with the same first name: ...it's like every other [P14 first name] in the world that has a Me.com account forgets to put in the numbers". Additionally, the username of the Gmail

account our single-user misdirected email corpus was collected from is a first name that was one of the most popular names for babies born in the USA in the 1990's and 2000's.

Even if one's name is not common, username collisions can still occur. Many of the work-related misdirected email examples were a result of a name mixup. For example, P04 only received work-related misdirected email, but talked about at least two other people with similar names to hers that she used to get email for—one with a different first name and similar last name, and another with the same last name. P10 gave a similar example of a work-related misdirected email: "I think in the very specific case, both of our first names and then our last names kind of look similar, hers ends with [redacted] as well...

How People Choose Usernames

Username collisions were identified by participants as a source of misdirected email, based on the messages they had received. To explore whether broader patterns in username choice might be partially responsible for this, we asked survey respondents about what characteristics are important to them when choosing usernames. If people prefer to have usernames that resemble their real names and don't prioritize uniqueness when creating them, it is more likely that multiple people—and in some cases many, many people—would share very similar usernames, making misdirected email more likely for those people than others.

To help us describe preferences for username choice, we asked survey respondents to list up to five email accounts that they used regularly. For each email account they listed, they were asked: "When you created the email account [account listed by respondent], how did you choose the username for that account? (The username is the part of the email address before the @ sign.)" The response options were: "I created the exact username I wanted; A username was assigned to me, and I kept it; A username was assigned to me, but I changed it later; The system suggested a username to me when I created the account, and I accepted it; Most of the usernames I tried to create were already taken, so I ended up with this one; I don't remember; and None of the above. Seventy-four percent (280 out of 380) respondents reported that they had chosen their username for the first account they listed, and 71% did so for the second account.

Then respondents were asked, "Please imagine that you are creating a new email account on a brand new email service that currently has very few users, and you can have any email username that you want. What username would you choose?" The next question followed up with, "Please explain why you would choose the username [username they entered]" (These usernames were discarded; we were only interested in the explanations.) The first author made an initial pass through the responses and developed a set

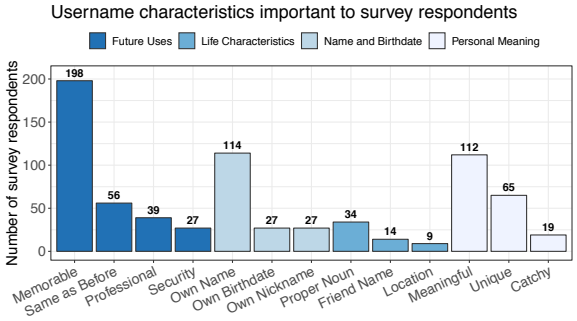

Figure 2: The number of survey respondents who said these username characteristics were important to them. "Professional": wanted a more professional-sounding username; "Proper Noun": the name of a team, band, etc.; "Location": a place, zip code, or phone number.

of categories of reasons why a particular username would be chosen. These categories reflected concerns respondents had, expectations about future uses of their email account, and characteristics they wanted their usernames to have.

Then, two members of the research team coded the responses. The graph in Figure 2 shows counts for each of the codes which had Fleiss' inter-rater reliability above 0.5, and on which either one of the coders said that code was present. The most commonly mentioned characteristic was that respondents wanted their usernames to be memorable (52%), to have personal meaning (30%), and to contain their own name (30%). 307 out of 380 survey respondents said one of these three things (81%). Only 65 respondents (17%) said they wanted their username to be unique or uncommon when compared with others. This indicates that most respondents wanted their usernames to be personally meaningful and to contain their own name, but rarely considered whether their preferred username would be different from others' usernames. This means that avoiding username collisions does not seem to be a priority for people when choosing usernames.

When Misdirected Email (Maybe) Isn't an Accident

Another observation that 9 out of 22 interview participants made about some of the misdirected emails they had received was that someone else might be giving out the participant's email address intentionally, even though it did not belong to them. P13, who lives in the Midwest region of the USA, gave an example of an email she received from "a clothing store that exists only in the United Kingdom" that she had never been to. She said, "it was like, Welcome to X store, and, thanks for signing up for emails from the store. P13 suspected that someone must be giving out an email address—hers—that they know isn't theirs because "they're trying to evade, they are trying to tell the person at the store that they're signing

up for the email but they don't want to receive the email or something like that.

And in fact, 7 out of 22 interview participants (32%) and 69 out of 380 survey respondents (18%) said they had given out an email address that wasn't theirs at some point in the past. The survey respondents selected the statements, "I have given out a 'throwaway' email address that was fake or did not belong to me to a website , or "I have used a 'throwaway' email address that was fake or did not belong to me to sign up for an online account". These percentages may underestimate how common this is, due to social desirability bias, although the interview participants who mentioned doing this did not seem to feel badly about it.

For example, P09 said he had given out a throwaway email address "Mostly out of convenience, I guess. I didn't really want to ever get contacted, I just wanted to get through a gatekeeper... it was a random fake [email address]. I just wanted to see what was behind the gate online for a website. And P10 said, "It's some clothing store in the mall... sometimes I change the last letter of my email when I enter it in, because it's awkward to respond back with no, thank you, I don't want to give you my email. P11 described, "I was signing up for a forum and I didn't imagine myself using it more than once, so I typed in a somewhat real email address to sign up... I used my initials and some numbers to make it look like a authentic one. The 7 interview participants who reported doing this did not voice concern that the emails they were avoiding might end up being sent to a real person. They were also not among the 9 participants who thought someone might be giving out their email address to avoid spam, so it may not have occurred to them that the email address they thought they made up could actually belong to someone else.

Misdirected Account Creations and Password Resets

One type of misdirected email interview participants received that was difficult for them to deal with was automated confirmation messages or password reset emails from accounts that had been set up by someone else, using their email address. P02 described one example:

> "I got this funny one a few weeks ago... there was a different [first name with P02's first initial, P02's last name], this guy was in [city] and he was creating, he is going to train to be an Uber driver. So I got his Uber account information and the qualifications and training for being an Uber driver..." (P02)

There is evidence from our misdirected email corpus that this behavior—using an email address that does not belong to you to sign up for an account—can be quite common. Figure 3 shows the number of misdirected emails in the corpus

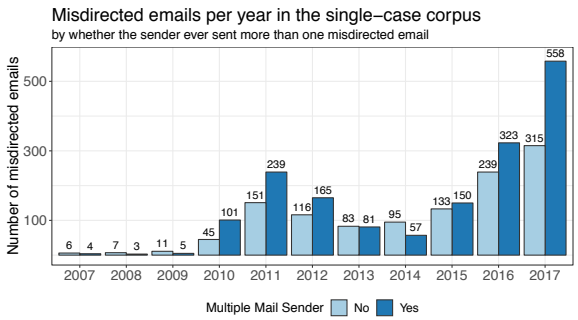

Figure 3: The number of misdirected emails in the single-case corpus from unique versus repeat senders. Note that this corpus contains misdirected emails received by a single user over 10 years, 2007-2017.

from senders who sent multiple misdirected email messages. There are clearly more messages from repeat senders than from unique senders in recent years, and the size of the gap between those two categories has been increasing as well. In 2017 especially, the number of misdirected emails from multiple email senders is much higher than single email senders. The five most prolific multiple email senders in the corpus are systems, not people, and are all platforms on which users can sign up for accounts: esubscriptions@kiplinger.com, a personal finance website which sends a regular newsletter to subscribers (116 mails); password reset emails from Facebook (security@facebookmail.com, 101 mails) and Epic Games, the creator of the popular multiplayer video game Fortnite (help@epicgames.com, 76 mails); account creation emails from NBA 2K, another popular video game series (fromnoreply@nba2k.com, 67 mails); and notifications about the creation of a new Gmail address linked to this one (mail-noreply@google.com, 40 mails).

When interview participants received messages like this, they said that it can sometimes be difficult to remove their email address from the account that has been created. Some confirmation emails include a link that enables the recipient to disavow the account, as described by P14: "Instagram has an option where you can click the link to say that isn't your email address and they'll remove your email address from that account". However, to participants it seemed like most account creation emails do not provide an easy way to disavow. P11 said that he had been signed up for an instant messaging application by someone else using his email address, without a way to disavow. He said that he has tried to contact "the administrator or someone I could speak to about the account being false or the email being incorrect", but that "it was hard to talk to anybody". P14 also had a similar experience; she said her email address was associated with someone else's AT&T account, so she was receiving their

mobile phone bills. She had quite a difficult time convincing AT&T customer service to do anything about this:

> "I had to call AT&T and argue with them because I didn't want to be liable for the bill because it was coming into my email address... it was just very stupid but I had to escalate it and talk to them for hours to make sure I got my email address removed from that account." (P14)

Several participants expressed a desire for more websites to include links that would allow them to remove their email address from an account that had been created without their consent. P20 summarized this as wanting a way to respond and indicate "This is not meant for me. This is the wrong address. I did not sign up for this. Please cancel it.

Three participants described coping with unwanted accounts that had been created using their email address by using a password reset link to log in and delete the account or change something so they would no longer receive any emails because of it. P02 "stole back" a Snapchat account that had been created with his email address, and continued to use it. About this, he said: "I hate to admit that I'm okay with it, because for me, it's like, well, you should've known your email address. How do you not know your email address? There's no reason why in this day and age you don't know what your email account is. P20 also reset a password, so he could log in and cancel an account that someone used his email address to sign up for. It was on a "a website which allowed people with trucks to get in touch with people who needed stuff moved". Here is how P20 described what happened:

> "And I guess he signed up and said that he wants to get an email when a job was available in his area. And then a couple of days later, he hadn't received anything so he logged back into the website and told them to give him any jobs in the entirety of London. And then, he still didn't receive anything, because they were coming to my email address... My phone was just downloading hundreds of emails. And then finally, he just told the website to give him every single job that became available in all of England. And so, it sent me... It was sending me close to 1,000 emails an hour. And that was obnoxious, that was actually one of two times where I used the site's password recovery feature..." (P20)

This example from P20 is extreme. However, all of the participants who received misdirected emails related to accounts they did not create and didn't want found it to be a source of concern. Some participants worried about the potential for negative consequences to themselves, due to others' use of their email address. For example, P03 talked about receiving email confirmations for orders through Ebay and PayPal that looked like they were being shipped to a "legitimate address"; he said he was concerned about his email address "accruing some weird reputation" because someone else was using it to buy things. He talked about worrying about whether "it's something that seems like it's gonna cause me trouble, like a credit score or something. And P14 had received email messages as a result of an online dating profile created with her email address without her knowledge, which made her uncomfortable. She said she contacted the site asking, "Please ban [my email address] so that no one could ever use this email address [on the dating site]" but did not know if that was effective.

If misdirected email is increasing, taking time to manage, and also causing concern, why don't people just stop using the account which received the misdirected messages and get a new email account? We asked interview participants about whether they had ever considered this as a way to stop problems like this from happening. Despite the frustrations, all felt like it would be more effort to change their email address than to cope with the day-to-day issues that arose for them related to misdirected email. P15 described it like this: ...[misdirected email] is just not burdensome enough that I feel like I need to change my address which I've had for a really long time now to get away from it. P13 talked about having the email address where she receives misdirected email "for probably over 15 years", and has used it in a lot of places around the web for various accounts, so it would be difficult to change to a new one for all those things. And P02 said he would never even consider changing his email address to avoid receiving misdirected email. He feels a strong connection to his online usernames and places the blame on the other people involved: "To me, like I said, it's my identity. It's just how I relate to people and how people relate to me, especially online... So that's all on them. I'm not gonna change who I am because they keep screwing up.

P01 also described how her email address has become a kind of stable identifier for her, persisting over time—even more so than her physical address. She said, "Especially the older family members, they've only ever had one email address for me... it used to be when people moved, their address changed and their phone number changed, and now your phone number doesn't necessarily change when you move and neither does your email address. In all, seven of our interview participants talked about their email addresses this way.

To find out whether these comments from our interview participants might signal a broader pattern, we asked our survey respondents about how long they'd had their primary email account (survey question text: "How long have you had the email account [account listed by respondent]?"). They

reported that the first account they listed was the one they had the longest, with 74% of 380 respondents reporting they'd had the first email account they listed for 4 or more years. 134 respondents (35%) said that they'd had that account for "More than 10 years". This means that people are making a long-term commitment when they initially create an email account and choose a username, whether they realize it or not, and that username choices have lasting consequences.

5 DISCUSSION

The experience of receiving misdirected email is surprisingly common, and this paper provides evidence that it may be increasing in frequency. Our survey findings show that 43% of a sample of 380 US adult internet users had experienced receiving an email message not intended for them, which was much higher than we had anticipated. Misdirected email is therefore a problem that could affect a large number of people. However, not everyone receives misdirected email, and some experience it far more than others. While it can be irritating, and take time to manage, our interview participants reported that misdirected email also sometimes makes them feel connected to others, and like they are a part of something bigger than themselves.

Careless senders are undoubtedly at least partially responsible for causing misdirected email, as pointed out in previous research [23]. We identified two additional sources for misdirected email messages that had not been previously described. First, we provided evidence that some people give out "fake email addresses in order to avoid spam email that they believe they made up on the spot, but might actually be valid addresses for real people. And second, examples from our interview participants illustrate how they receive, and struggle to deal with, account confirmation and password reset messages for accounts that do not belong to them. We also provided evidence that recipients of misdirected email believe the volume of these messages they experience may be related to how common their name is and what characteristics they prioritized when creating their email account username. And, because creating an email account is a long-term commitment for people, these problems with receiving misdirected email cannot easily be solved by abandoning the email address that one has had for years and is a representation of their identity.

Our findings suggest ways that human choices and behavior interact with features of email as a tool and as infrastructure to create the conditions where misdirected email arises. We focus below on implications of two of these points of interaction between the human and the technical that studying misdirected email has allowed us to identify: email addressing and username choice, and the use of email as a secondary authentication mechanism for other accounts [24].

The Email Addressing Problem

Popular email systems like Google's Gmail that offer people the ability to create their own email username are different from most other widespread addressing systems. Many systems have a central authority that controls addressing, like the Internet Assigned Numbers Authority (IANA)[3] which manages domain names and coordinates IP address number spaces. Likewise, people are not usually allowed to select their own postal address or full telephone number; these are assigned by municipalities or telecom companies. Addresses usually exist as part of a larger system designed to facilitate the transmission of information and communication from senders to recipients, and conform to requirements that help make delivery more reliable.

However, when creating an email account, people choose usernames for social reasons, not to ensure effective information transmission by the system overall. They want email addresses that signal things about them as people, such as making them look more professional, or telling others that they were an early adopter of a popular platform (e.g., because they don't have numbers at the end of their username). They prefer usernames that are more memorable, like their own name, and that are easier for others to type. But email addresses serve a dual purpose: they are technical identifiers as well as social identifiers. This creates name collisions due to the social constraints placed on a technical addressing system, that result in the misdirection of email messages to the wrong recipients.

Popular free email platforms like Gmail complicate this further by having few namespace restrictions on usernames. Gmail crossed 1 billion monthly users in 2016 [14]; that's a lot of email accounts that must have unique usernames. Individuals, though, have no cause to suspect that name collisions are possible when they create a new email account. The "Birthday Paradox" describes how people greatly underestimate the likelihood that they share a birthday with another person (in a group of about 25 people, the likelihood that some pair of them will share a birthday is about 50%) [21]. Likewise, it seems unlikely that people would be able to accurately estimate how many others have email usernames similar to theirs, and therefore may also underestimate the likelihood of name collisions at this scale. In other words, email users have no visibility into larger-scale patterns in username similarity that result from patterns in people's preferences for email address characteristics.

Email systems could address this breakdown by calculating edit distance metrics between existing and proposed usernames at account creation time. This would help people learn about the likelihood of potential name collisions and facilitate greater username differentiation, while still allowing

[3] https://www.iana.org/

people to have some flexibility in username choice. However, given the preferences stated by our survey respondents for creating usernames that are memorable and personally meaningful, people may not be satisfied with more stringent username differentiation constraints, and systems providing free email accounts likely want the barrier to account creation to be as low as possible [6].

Also, as many of our interview participants described, people are always going to make mistakes when specifying email recipients that result in messages being sent to the wrong person. Algorithms that help to reduce sender mistakes due to typos or auto-complete inattention can be used to prevent this type of misdirected email from ever being sent [3, 16, 26]. However, our research shows that this is just one of multiple causes of misdirected email, and this type of misdirected email is the easiest for recipients to deal with by simply responding to the sender to let them know the message reached the wrong person. Another type of misdirected email, described below, involves three parties: an account holder on another online service, the online service itself, and the recipient of misdirected email, and is much harder for recipients to deal with.

The Email as Authentication Problem

Our research has uncovered a perplexing phenomenon that has not been previously reported: unintended recipients who have received an authentication request or other notification about an account that they didn't create. This presumably occurs when someone creates an account using an email address that does not belong to them. This behavior cannot be explained by recipient auto-complete inattention in an email client. We were unable to discern the root cause of this problem from our data, because we did not speak with anyone who admitted to doing this. In fact, it may be difficult for the third-party account creator to even know this has happened; if they never receive the authentication email, they may assume the fault lies with a failure of the system and not something they did themselves.

The use of email as "something you have by platforms for authentication purposes seems to have arisen as a convenient identify verification mechanism, and a simple fix to misdirected authentication requests seems possible. Every account confirmation and password reset email message should contain a disavow link that actually works and doesn't generate more unwanted email. This was explicitly requested by multiple interview participants, and would go a long way to alleviating recipients' fears about being held accountable or liable for the actions, and in some cases the debts, of others.

This fix is so simple, in fact, that one wonders why it is not more common. It may be that mistaken account creations and password resets seem like a small problem from a platform perspective, which they deal with by taking no further action if there is no action on the new account within a certain period of time. But platform operators have no way of knowing that someone who received tens of account creation emails from their platform in the last year also received similar mails in similar quantities from other platforms, too. Unfortunately, the true cost to the recipient in evaluating misdirected email messages, especially if they are password reset messages from platforms where the recipient already has their own account, is higher than the time it takes them to click a disavow or unsubscribe link. They must evaluate whether the message means their account has been compromised or not, and then figure out what steps, if any, they should take in response. This is an identity collision rather than a name collision, with stakes that are higher for recipients than for platforms.

Misdirected email is, in some cases, an indication of a breakdown that has arisen because a system designed for interpersonal communication is now being used for other purposes. There were other signs of this type of breakdown in our data as well. For example, some participants spoke about systems that collect email addresses in exchange for access to information, or about being asked for their email address when they made a purchase in a brick-and-mortar store. They said that they did not want to give out their real email address in these contexts, because they did not want to receive unwanted email related to the transaction. In these situations, rather than giving out their "real" email address, an individual may trade an email address they do not own" instead. And, the recipient of the subsequent misdirected emails bears some of the consequences of this.

Future Work

Misdirected email can result in unintended disclosures, and can be a burden for people who feel that creating a new email account is not a viable option. In this paper, we have described the conditions that cause misdirected email, and how it affects recipients. We have also described two types of problems that give rise to misdirected email. While there are a few simple individual-level fixes that may be able to reduce the amount of misdirected email people receive, the technical properties of the email infrastructure and the incredibly large number of people who use email on a daily basis could make it hard to prevent at scale.

Our research focused mostly on recipients of misdirected email; senders proved much harder to locate, especially when the misdirected messages were sent automatically by a third party system. In addition, characteristics of certain canonical" usernames may make it more likely for some people to receive misdirected email than others, making it easier to recruit high-volume recipients than senders. This points to a clear opportunity for future work to understand more about senders' perspectives. Another area for future research

related to senders is the privacy-related problems and complications misdirected email may cause. Many of our interview participants expressed concern about the personal information they had received by mistake, and future work should address consequences and harms that may arise for senders from these unintended disclosures.

Finally, rather than focusing on recipients or senders, future work should also focus in more depth on the different contexts in which misdirected email arises. For example, our interviews showed that misdirected email occurs in the context of interpersonal communication, organizational communication, and also automated communication from systems to individuals. Focusing on misdirected email in different contexts would shed more light on situational differences in motivations, dynamics and constraints, and offer context-specific solutions to the sociotechnical breakdowns that cause misdirected email.

ACKNOWLEDGMENTS

We thank Cindy Ochoa and Nina Capuzzi for their assistance with piloting and early analysis efforts, and the BITLab @ MSU research group for helpful discussions and feedback.

REFERENCES

[1] Tarfah Alrashed, Ahmed Hassan Awadallah, and Susan Dumais. 2018. The Lifetime of Email Messages: A Large-Scale Analysis of Email Revisitation. In *Proceedings of the Conference on Human Information Interaction & Retrieval*. 120–129. https://doi.org/10.1145/3176349.3176398

[2] Frank Bentley, Nediyana Daskalova, and Nazanin Andalibi. 2017. "If a person is emailing you, it just doesn't make sense": Exploring Changing Consumer Behaviors in Email. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 85–95. https://doi.org/10.1145/3025453.3025613

[3] Vitor R Carvalho. 2011. Email Information Leaks. In *Modeling Intention in Email*. Studies in Computational Intelligence, Vol. 349. Springer Berlin Heidelberg, Chapter 3, 35–51. https://doi.org/10.1007/978-3-642-19956-1

[4] Vitor R Carvalho, Ramnath Balasubramanyan, and William W Cohen. 2009. Information Leaks and Suggestions: A Case Study using Mozilla Thunderbird. In *Conference on Email and Anti-Spam*.

[5] Lorrie Faith Cranor and Brian A LaMacchia. 1998. Spam! *Commun. ACM* 41, 8 (1998), 74–83. https://doi.org/10.1145/280324.280336

[6] Dinei Florêncio and Cormac Herley. 2010. Where Do Security Policies Come From?. In *Proceedings of the Sixth Symposium on Usable Privacy and Security*. https://doi.org/10.1145/1837110.1837124

[7] Bent Flyvbjerg. 2006. Five Misunderstandings About Case-Study Research. *Qualitative Inquiry* 12, 2 (2006), 219–245. https://doi.org/10.1177/1077800405284363

[8] Iftah Gamzu, Liane Lewin-Eytan, and Natalia Silberstein. 2018. Unsubscription: A Simple Way to Ease Overload in Email. In *Proceedings of the ACM International Conference on Web Search and Data Mining*. 189–197. https://doi.org/10.1145/3159652.3159698

[9] Simson L. Garfinkel. 2003. Email-Based Identification and Authentication: An Alternative to PKI? *IEEE Security and Privacy* November/December (2003), 20–26. https://doi.org/10.1109/MSECP.2003.1253564

[10] John Gerring. 2004. What Is a Case Study and What Is It Good for? *American Political Science Review* 98, 2 (2004), 341–354. https://doi.org/10.1017/S0003055404001182

[11] Paul A Grassi, James L Fenton, Elaine M Newton, Ray A Perlner, Andrew R Regenscheid, William E Burr, Justin P Richer, Naomi B Lefkovitz, Jamie M Danker, Yee-Yin Choong, Kristen K Greene, and Mary F Theofanos. 2017. *Digital identity guidelines: authentication and lifecycle management*. Technical Report. National Institute of Standards and Technology, Information Technology Laboratory, Gaithersburg, MD. https://doi.org/10.6028/NIST.SP.800-63b

[12] Galen A Grimes, Michelle G Hough, and Margaret L Signorella. 2007. Email end users and spam: relations of gender and age group to attitudes and actions. *Computers in Human Behavior* 23, 1 (2007), 318–332. https://doi.org/10.1016/j.chb.2004.10.015

[13] Benjamin M Gross. 2009. Names of Our Lives. In *International Conference on Computational Science and Engineering*. 747–752. https://doi.org/10.1109/CSE.2009.474

[14] Frederic Lardinois. 2016. Gmail Now Has More Than 1B Monthly Active Users. https://techcrunch.com/2016/02/01/gmail-now-has-more-than-1b-monthly-active-users/

[15] Eric Lieberman and Robert C Miller. 2007. Facemail: showing faces of recipients to prevent misdirected email. *Proceedings of the Symposium on Usable Privacy and Security* (2007), 122–131. https://doi.org/10.1145/1280680.1280696

[16] Tingwen Liu, Yiguo Pu, Jinqiao Shi, Quangang Li, and Xiaojun Chen 0004. 2014. Towards misdirected email detection for preventing information leakage. *IEEE Symposium on Computers and Communications* (2014), 1–6. https://doi.org/10.1109/ISCC.2014.6912554

[17] Kathleen M. MacQueen, Eleanor McLellan, Kelly Kay, and Bobby Milstein. 1998. Codebook Development for Team-Based Qualitative Analysis. *CAM Journal* 10, 2 (1998), 31–36. https://doi.org/10.1177/1525822X980100020301

[18] Cristobal Martinez and Christina Thorpe. 2017. Analysis of Spam: Honeypot Experiment. In *European Conference on Cyber Warfare and Security*. 692–701.

[19] Chris McLaughlin. 2016. Legal Ethics and Social Media . *NC St. BJ* 21 (2016), 24–26.

[20] Adam Moskowitz. 2003. On Choosing Usernames. *;login: The Magazine of Usenix and Sage* 28, 4 (2003), 5–6.

[21] Joseph I. Naus. 1968. An Extension of the Birthday Problem. *The American Statistician* 22, 1 (1968), 27–29. https://doi.org/10.1080/00031305.1968.10480438

[22] Pew Research Center. 2018. Internet/Broadband Fact Sheet. http://www.pewinternet.org/fact-sheet/internet-broadband/

[23] Alan B Poole, Kelly E Caine, Arthur D Fisk, and Wendy A Rogers. 2010. Errors of Disclosure in Computer Mediated Systems. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*. 512–516. https://doi.org/10.1518/107118110X12829369603326

[24] Rob Reeder and Stuart Schechter. 2011. When the Password Doesn't Work: Secondary Authentication for Websites. *IEEE Security and Privacy* 9, 2 (2011), 43–49. https://doi.org/10.1109/MSP.2011.1

[25] United States Census Bureau. 2018. American Community Survey 2016. https://www.census.gov/acs/www/data/data-tables-and-tools/data-profiles/2016/

[26] Polina Zilberman, Gilad Katz, Asaf Shabtai, and Yuval Elovici. 2013. Analyzing group E-mail exchange to detect data leakage. *Journal of the American Society for Information Science and Technology* 64, 9 (2013), 1780–1790. https://doi.org/10.1002/asi.22886