What's in a Name?
Understanding Profile Name Reuse on Twitter

Enrico Mariconti[⋆], Jeremiah Onaolapo[⋆], Syed Sharique Ahmad[‡], Nicolas Nikiforou[⋆],
Manuel Egele[†], Nick Nikiforakis[‡], and Gianluca Stringhini[⋆]
[⋆]University College London, [‡]Stony Brook University, [†]Boston University
{e.mariconti,j.onaolapo,n.nikiforou,g.stringhini}@cs.ucl.ac.uk
{syahmad,nick}@cs.stonybrook.edu [†]megele@bu.edu

ABSTRACT

Users on Twitter are commonly identified by their profile names. These names are used when directly addressing users on Twitter, are part of their profile page URLs, and can become a trademark for popular accounts, with people referring to celebrities by their real name and their profile name, interchangeably. Twitter, however, has chosen to not permanently link profile names to their corresponding user accounts. In fact, Twitter allows users to change their profile name, and afterwards makes the old profile names available for other users to take.

In this paper, we provide a large-scale study of the phenomenon of profile name reuse on Twitter. We show that this phenomenon is not uncommon, investigate the dynamics of profile name reuse, and characterize the accounts that are involved in it. We find that many of these accounts adopt abandoned profile names for questionable purposes, such as spreading malicious content, and using the profile name's popularity for search engine optimization. Finally, we show that this problem is not unique to Twitter (as other popular online social networks also release profile names) and argue that the risks involved with profile-name reuse outnumber the advantages provided by this feature.

Categories and Subject Descriptors

K.6.5 [**Management of Computing and Information Systems**]: Security and Protection

General Terms

Security, Measurement

Keywords

Social network; OSN; Profile name; Impersonation

WWW 2017, April 3–7, 2017, Perth, Australia.
ACM 978-1-4503-4913-0/17/04.
http://dx.doi.org/10.1145/3038912.3052589

1. INTRODUCTION

Users on Twitter are identified by their profile name, such as *@taylorswift13*. A user's profile name is also used to directly *mention* accounts on Twitter, as well as to identify their profile page's URL.[1] However, Twitter provides profile names as a mere convenience to its users. Internally, the social network identifies accounts with unique numerical identifiers, so-called *user IDs* (e.g., the number 17919972 for Taylor Swift's Twitter account). While user IDs are globally unique and persistent, they are usually not observed by end-users. With these robust identifiers in place, Twitter allows users to change the profile names of their accounts over time. As the concepts of stable user IDs as well as changeable profile names are crucial to this paper, we will use the terms user ID and account interchangeably to refer to the persistent notion of an account as identified by its user ID. Furthermore, we use the term profile name to refer to the changeable name the user chose for a given account. There are multiple reasons why a user might want to change their profile name, including changes of jobs and affiliations, changes of names due to marriage, or the selection of a different nickname. Interestingly, once a user changes her profile name, Twitter releases the old name for other users to adopt. The same happens if a user decides to delete her account. Only in the case of an account that gets suspended due to a violation of the terms of service,[2] does Twitter make the profile name unavailable for other users to adopt.

There are many reasons why a Twitter user might take a profile name that was previously in use and then abandoned, both legitimate and malicious. It is possible that a person innocently selects a profile name that was previously in use, perhaps because it corresponds to a particularly common first and last name (e.g., *JohnSmith*). It is also possible, however, that a malicious entity will select abandoned profile names on purpose, in an attempt to leverage the residual reputation that this profile name might have. Miscreants trying to spread spam on Twitter might use this reputation in the hope of attracting more followers. Another use of abandoned profile names is Blackhat Search Engine Optimization (SEO) [19]. Since profile names identify the URL of the profile pages of accounts on Twitter, once a profile name is freed there will be multiple links on the web pointing to a non-existent page. By selecting an abandoned profile name with many URLs already pointing to it, a malicious

[1] https://twitter.com/taylorswift13
[2] https://twitter.com/tos

user can start promoting his content and influence search results. Note that this is not a mere hypothetical scenario; the practice of harvesting abandoned Twitter profile names and using them for SEO purposes has been observed in the wild.[3]

We first introduced the concept of profile name reuse on Twitter in our preliminary work [20]. In that work, we showed that profile names are reused in the wild, and we identified a number of accounts that adopted abandoned profile names and used them to send spam. In this paper we perform a more in-depth, larger-scale measurement of the phenomenon of profile name reuse on Twitter. We show that over a period of one year, 1% of popular accounts with more than one million followers that appear in our datasets changed their profile name, and this name was later taken by another account. We provide a number of case studies in which a popular profile name was used to ridicule the original owner, or to post malicious content. We also show that Twitter users often prevent profile name hijacking by creating placeholder accounts that immediately adopt the abandoned profile name and point users to the new one. These case studies show that there is a concrete threat linked to freeing abandoned profile names on Twitter.

To understand how profile name reuse manifests at a large scale, we collected a 1% sample of all public tweets posted over a period of six months, between October 10, 2015 and April 12, 2016. In total, we identified 106,935 profile names that have been shared by 196,200 unique accounts. In doing so, we identify a set of profile names that were taken over by multiple accounts during the observation period. We identify different categories of accounts, ranging from those that seem to be acting for legitimate reasons to those that repeatedly iterate among different profile names with the purpose of spreading malicious content. We analyze the general characteristics of these accounts, together with the topics that they discuss on Twitter and the URLs that they post, finding that accounts that take abandoned profile names are more likely to post malicious content than regular Twitter accounts, as well as more likely to be suspended by Twitter.

Although this paper focuses on Twitter because of its size and popularity, the phenomenon of profile-name reuse is not necessarily unique to this social network. In fact, we discovered that two other major social networks (Tumblr and Pinterest) also allow profile name reuse. We conclude that freeing profile names after they are not used anymore is not a good design choice for a social network, as it exposes its users to security risks. We advocate for social networks to avoid this practice, or, at least, to monitor with particular attention accounts that adopt a previously-freed profile name.

In summary, this paper makes the following contributions:

- We show that 1% of popular Twitter accounts abandoned their profile name between 2015 and 2016 and had it taken over by a third party; we then provide a number of case studies in which abandoned profile names are used to spread malicious content or to ridicule the original owner.

- We identify 196,200 accounts partaking in this practice, which shared 106,935 profile names over a period of six months. We show that accounts that take over abandoned profile names are more likely than regular Twitter accounts to post malicious content and get suspended by Twitter.

- We show that Twitter is not the only social network allowing profile name reuse, but that Tumblr and Pinterest allow this practice too. We argue that this practice should not be permitted because it enables malicious users to perform reputation hijacking and impersonation attacks.

2. DATA COLLECTION

Our dataset consists of a 1% random sample of all public tweets posted on Twitter over a period of six months, between October 10, 2015 and April 12, 2016. In total, the dataset contains 667,294,613 tweets posted by 70,803,606 distinct users (user IDs). Twitter's streaming API provides access to this data as a stream of JSON dictionaries that include information about the message and about the account that posted it. We call the dataset of 70,803,606 users **U**. We used the dataset **U** for two purposes: to extract a set of popular profile names that were abandoned and taken by a third party, and to identify a set of regular accounts whose profile name was taken over. In the following, we describe these two datasets in detail.

Popular account dataset. To extract popular profile names that were freed and taken over by a third party, we proceeded as follows. First, we define a popular account as an account having at least one million followers. Therefore, we extracted all accounts in **U** that had more than one million followers at the time the data was collected. We identified 4,263 accounts that fit that criterion. Second, we queried the Twitter API six month after data collection was completed to assess whether the corresponding user IDs were still associated with the same profile names that we observed during the data collection. After this procedure, we obtained 42 popular profile names that were freed and taken by another account. This is approximately 1% of all accounts with more than one million followers. We call this set of users, **C**.

For 146 popular accounts, the Twitter API did not respond with valid information. In these cases, we visited the profile URL associated with the profile name, and we also queried the Twitter API for the user ID originally associated to that profile name. If the profile name was marked as suspended by Twitter and the user ID was not active anymore, we concluded that the original account was simply suspended for violating Twitter's terms of service. If, however, the user ID was still active but the profile name that it had freed was blocked, we concluded that some third party entity took it over and used it for malicious purposes (e.g., to spread malicious content), resulting in suspension by Twitter.

General account dataset. Besides assessing popular accounts, we also analyzed the behavior of regular accounts. To this end, we analyzed **U** to identify accounts that engaged in profile name reuse during our observation period. We identified 196,200 accounts that shared 106,935 profile names during the measurement period, posting 3,290,286 tweets. These accounts constitute 0.27% of all accounts observed in our sample. We call this dataset **G**.

[3] http://www.inetsolutions.org/quickly-easily-find-high-authority-expired-twitter-accounts/

Auxiliary dataset. To better characterize accounts that partake in profile name reuse, we use a dataset of random Twitter accounts. This dataset provides a baseline of typical behavior of accounts on Twitter and helps to highlight differences in behavior shown by accounts involved in profile name reuse, compared to a regular account. To this end, we randomly sampled one million accounts from **U**. We call this dataset **R**.

Limitations. Although our dataset allowed us to measure the practice of profile name reuse on Twitter, it has some limitations. First, we can only detect two accounts as sharing the same profile name if they both posted tweets during the observation period, and if these tweets were both captured in the 1% sample that we collected. This means that our dataset is likely to make us underestimate the phenomenon of profile name reuse.

The limited visibility provided by our dataset can also impact the conclusions that we draw from our analysis. While accounts that are observed switching many profile names during the observation period are likely to be participating in profile name reuse schemes and are potentially malicious, the fact that an account is present in our dataset with a single profile name is not a guarantee that the account never changed its profile name. As we will describe in detail in Section 4, we took extra care in ensuring that we can infer the profile name change history of an account, but some of these limitations persist.

Ethics. Dealing with social network data raises ethical concerns. In this paper, we only used publicly-available Twitter data, and we successfully obtained ethics approval from the University College London ethics committee (Project ID 6521/004). To treat data ethically, we followed the guidelines outlined by Rivers et al. [22]. In particular, we ensured not to link multiple datasets together with the goal of further deanonymizing the users contained in them, and we stored our data according to the UCL Data Protection Officer guidelines.

3. PROFILE NAME REUSE FOR POPULAR TWITTER ACCOUNTS

One of the main reasons for someone to take an abandoned profile name on Twitter is to hijack the "name recognition" held by that account. This opens up the opportunity for malicious actors to reach a larger audience and mount impersonation attacks. To understand the reasons why people take over abandoned Twitter profile names, we analyzed the accounts in the dataset **C** in detail. To recap, the profile names of these accounts belonged to popular accounts with more than one million followers. These accounts then changed their profile name, which was taken over by other accounts. In total, we found that 42 profile names were abandoned and taken over by another user. This accounts for approximately 1% of all accounts that appear in our sample and have more that one million followers. Comparing this number with the profile names that are reused in the dataset **G** presented in the previous section, we can conclude that popular accounts are four times more likely to have their profile name reused than generic Twitter accounts. We then analyzed these accounts in more detail to find out the reasons behind profile name reuse for popular accounts. Broadly speaking, we identified four trends.

Profile names taken over by a third party with no malicious activity. As mentioned earlier, profile name reuse is not necessarily malicious. A Twitter user could select a profile name that was previously in use by mere chance (e.g., in the case of people with the same first and last names), or they could do that on purpose, but without a malicious intention (e.g., a fan of a celebrity who has the chance to take that person's old profile name). We observed nine cases in which popular accounts changed their name and were taken over by someone else without a clear malicious intention. These examples include the old profile name of singer Lorde (`@lordemusic`), which is now owned by an account whose tweets are private, and the old profile name of the TV show "The Big Bang Theory" (`@BigBang_CBS`), which is now owned by the official fan club of the series.

Profile names taken over by a third party to set up a parody account. The reputation that a Twitter profile name gains over the years makes it a useful asset for people who want to discredit or just ridicule public figures. We observed two cases were popular accounts changed their profile name and their old names were taken over by parody accounts: Brazilian footballer Alex De Souza (`@Alex10Combr`) and Colombian radio host Vicky Davila (`@vickydavilafm`).

While parody is not a malicious activity and falls under free speech, chances are that the original owners of those accounts did not intend to give visibility to accounts portraying them in a satirical fashion. Allowing anyone to adopt a freed profile name, however, can make it easier for such parody accounts to gain popularity. Moreover, not all parody accounts are harmless. An example is what happened to Annaliese Nielsen, an activist who recorded a video of her arguing with a minicab driver and threatening to ruin his reputation. This video generated outrage in certain online communities [11], causing the activist to be harassed on Twitter, to the point that she eventually deleted her account. The freed profile name (`@tornadoliese`) was then taken over by trolls, who then set up a parody account in which they ridicule the former owner. This example shows how allowing profile name reuse on Twitter can introduce a new attack vector for online harassment.

Profile names taken over by a third party with clear malicious intentions. We identified that twenty profile names that were freed by their owners, were subsequently used in violation of Twitter's terms of service, and were eventually suspended. Examples of these profile names include highly visible television outfits, such as, BBC Science News (`@bbcscitech`) and the Entertainment Channel (`@eonline`). These high-profile incidents show that attackers are actively using the reputation gained by popular profile names to perform malicious activity. Unfortunately, we were not able to collect evidence of the specific type of malicious activity performed by these accounts before they were suspended by Twitter since the offending behavior was not part of our 1% sample.

Profile names "protected" by a placeholder account. The aforementioned case studies show examples of how malicious actors are misusing abandoned profile names to their advantage. Unfortunately, Twitter does not provide an easy countermeasure for users who want to change their profile name, while not allowing anyone else to take the old one. We observe that some people managing popular Twitter accounts understand the risks of profile-name reuse, and overcome this problem by creating *placeholder accounts* that take the old profile name. These accounts usually do not post messages, but have a pointer in their profile description to the new profile name of the account. This strategy has been adopted by high-profile accounts such as Manchester City Football Club (`@MCFC`) and singer Enrique Iglesias (`@enrique305`).

4. PROFILE NAME REUSE IN GENERAL TWITTER ACCOUNTS

In this section we analyze the accounts in our dataset **G** in detail, with the goal of understanding the reasons why profile names are reused on Twitter, and measure the modus operandi of accounts that are switching between multiple profile names. We start by defining different types of accounts involved in profile-name reuse, and we continue with a detailed measurement of their characteristics and activity. We then look for the presence of links pointing to abandoned profile names, and investigate the possibility of using these links for SEO purposes, and to inflate the popularity of Twitter accounts. Finally, we go through some interesting case studies, exposing complex ecosystems of accounts sharing profile names and posting about common topics.

4.1 Types of accounts involved in profile name reuse

Accounts can be involved in profile-name reuse from various perspectives. There are accounts changing their profile name and taking another one, which was never held by another account, accounts that take an abandoned profile name upon creation without knowing it, and accounts that systematically take abandoned profile names for their gain (for example to hijack the reputation linked to those profile names). We group these activities into three types of accounts as follows:

First unique account. This type of account represents the typical behavior of a user changing their profile name and taking another one that was never used before, or deleting their account and consequently freeing its profile name. We consider an account as belonging to this group if it is the first one in our dataset holding the original profile name, and if it is the first one holding the new profile name of their choice as well. As we will explain later, if an account changes its profile name to a name that was previously used, or changes its profile name more than once during the measurement period, we consider this account as a "multi account." Note that some misclassifications are still possible for this category, in particular if the account used multiple profile names in the past and this was not captured by our dataset. We observed 101,244 accounts (over 196,200 total accounts) belonging to this category (that is, first unique).

Second unique account. This type of account is one that holds a single profile name in our dataset, and that profile name was freed by another account. These accounts represent the cases in which an account takes a profile name that was previously used, either by chance or on purpose. One possibility is that the abandoned profile name is a popular first and last name (e.g., *@johnsmith*) and someone happens to have that same first and last name, or that an account owner decides to delete their account and start over again, using the previous profile name. Due to the limitations in our dataset, it is possible that a profile name switched between other profile names in the past and we did not record that. To mitigate this problem, we look at the creation date of each account. If the creation date is before the last time in which the first unique user corresponding to that profile name tweeted, it means that the account was already active back then, holding a different profile name that we did not observe. In this case, we consider the account as a "multi account." We observed 15,911 second unique accounts in our dataset.

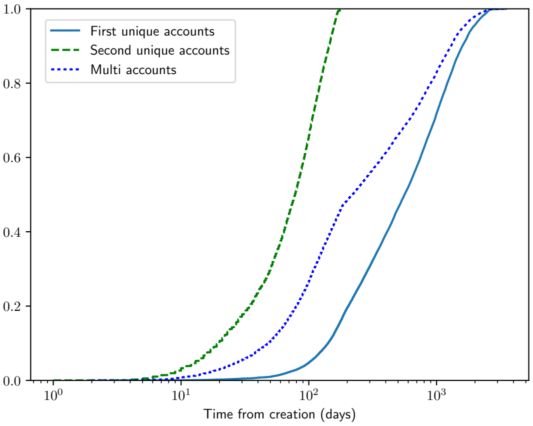

Figure 1: Cumulative distribution function of the time since the opening of first, second, and multi user accounts.

Multi accounts. These are the accounts that switched among more than two profile names. They are either observed holding three or more profile names in our dataset (two profile names that were previously used by someone else), or they are identified by following the procedures explained in the previous two categories. Multi accounts represent a systematic behavior in which an account changed many profile names, potentially to hijack the reputation of those names. We observed 79,045 accounts of this type in our dataset.

4.2 General characteristics of accounts reusing profile names

The first analysis that we performed was looking at the general characteristics of accounts reusing profile names. For each type of account described in the previous section, we obtained its creation time, its number of followers, and its number of tweets. We first looked at the life span of accounts reusing profile names. A Cumulative Distribution Function (CDF) of this metric for the different types of accounts is reported in Figure 1. The figure shows that second unique accounts are younger because, by definition, they are using only one profile name that has already been used by someone else during the measurement period. The earliest account creation time for them is therefore the beginning of the measurement period. Multi accounts are also generally younger than first unique accounts.

We then wanted to quantify the activity levels of the different types of accounts on Twitter. To this end, we studied the number of tweets posted on Twitter by the various types of accounts involved in profile name reuse compared to the

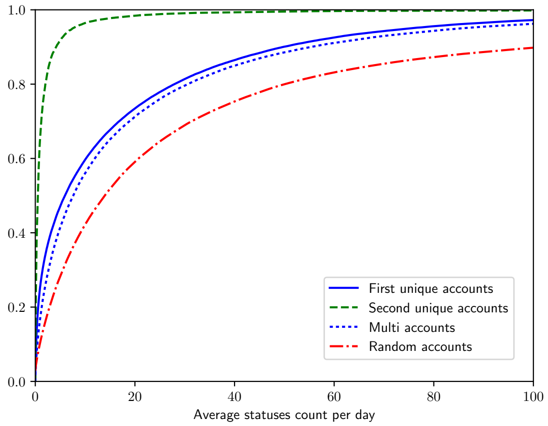

Figure 2: Cumulative distribution function of the number of tweets per day of first, second, and multi accounts.

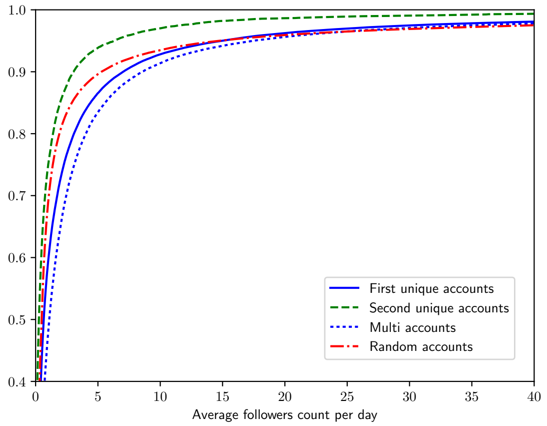

Figure 3: Cumulative distribution function of the number of followers per day of first, second, and multi accounts.

activity of the random sample of one million accounts in **R**. Figure 2 shows the average number of tweets posted by the accounts in our dataset per day. Accounts that are involved in profile name reuse generally post less tweets than the general Twitter population. First unique and multi accounts do not seem to show a significantly different posting behavior, while second unique accounts are much less active than the others, with 90% of them posting less than 10 tweets per day. As we mentioned in Section 2, this large discrepancy between the accounts under consideration and the random set of Twitter accounts could be due to our collection methodology for **R** accounts.

As another general characteristic of Twitter accounts, we decided to investigate the number of followers of the accounts under consideration. The number of followers is commonly regarded as a measure of reputation on Twitter [26], and it ultimately regulates how many users will see the statuses posted by an account. To exclude biases due to the age of accounts, we plotted the CDF of the average number of followers gained per day (Figure 3). Interestingly, multi accounts are able to attract more followers on average than general accounts, showing that switching profile name to popular ones could help in gaining more followers. First unique accounts also attract more followers than random accounts, while the number of followers obtained by second unique accounts is still less than for the other types.

Themes discussed by accounts reusing profile names. Next to obtaining general statistics about the activity associated with accounts involved in profile-name reuse, we wanted to understand the popular themes that these accounts tweet about. We also wanted to understand if the different accounts reusing profile names were tweeting about different themes or not. To achieve this, we conducted Term Frequency-Inverse Document Frequency (TF-IDF) analysis on the text of the tweets. TF-IDF extracts important words within a given text corpus, and the important words in turn provide information about the general theme(s) of the text corpus.

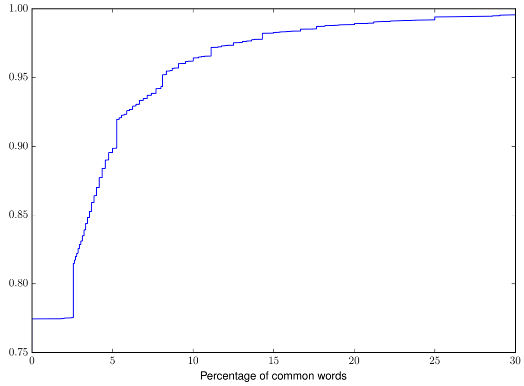

Figure 4: CDF of the percentage of common words across the tweets of the different accounts that shared the same profile name.

Before carrying out TF-IDF analysis, we preprocessed the text corpus as follows. First, we removed all non-English tweets; to identify such tweets, we just need to check the language from the tweet metadata contained in the JSON objects returned by the Twitter API. We could have used online automatic translation tools to translate the non-English text to English, but we found that they fail to translate some of the tweets properly. Besides, automatic translation would result in some loss of context which might bias our results. Second, we filtered out all words with less than five characters, and also removed all non-printable characters. We then carried out the TF-IDF analysis on the resulting text. In total, we processed tweets in English from 54,238 accounts associated with 26,678 profile names during the period of observation.

Figure 4 shows the CDF of the percentage of common words across the tweets of the different accounts that shared the same profile name. The words were extracted through TF-IDF analysis as earlier described. As it can be seen, 80% of the accounts that shared profile names had less than 3% words in common. This can be an indicator that accounts reusing a profile name are often not controlled by the previous owner, hence they tweet about completely different topics. We also grouped the important words we obtained from TF-IDF analysis into different themes. The main themes we identified included *Blog, Music, Porn, Video, Follow, Celebrities, Business, Retweet*, and *Football*. For example,

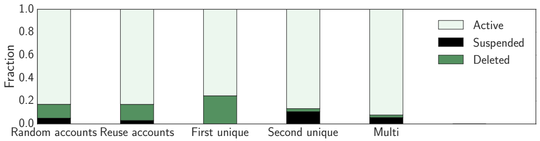

Figure 5: Ratio of user IDs sharing profile names that were suspended or deleted, compared to the ratio of random accounts. "Reuse accounts" comprises all accounts that shared profile names, regardless of their typology.

the *Video* theme comprises the following related words: *video*, *youtube*, *live*, *stream*, and *vine*. We discovered that second unique accounts are more likely to tweet about follow-back schemes, blogs, football, and business, than other themes. As we show in Section 4.6, many second unique accounts engage in follow-back schemes. It is worth noting that, among other patterns, we discovered that a significant fraction of the tweets of multi accounts is about "business."

4.3 IDs and profile-name-sharing behaviors

We wanted to understand the sharing behavior of profile names by accounts in **G**. In other words, we wanted to look into how many accounts share the same profile name, and understand the complex dynamics of accounts which change their profile name multiple times. Most of the accounts in **G** (90%) had only one profile name. Similarly, most of the profile names that were involved in profile name sharing (94%) were using only two different profile names. As we discussed in Section 4.1 these measurements are likely to be a lower bound of the actual phenomenon, due to limitations in our dataset. In Section 4.6 we will provide evidence of elaborate profile-name-sharing schemes, in which groups of accounts share a set of profile names among themselves.

4.4 Suspension and deletion dynamics of accounts reusing profile names

Taking another account's abandoned profile name is not an indicator of malicious intentions per se. As we have discussed, however, the possibility of taking abandoned profile names can facilitate malicious activity due to reputation hijacking and impersonation opportunities. As an indicator of malicious activity, we wanted to understand to what extent accounts taking part in profile name reuse get suspended by Twitter. There are many reasons why an account can get suspended, including sending spam or impersonating another person.[4]

To assess the current status of a Twitter account, we proceed as follows. First, we take advantage of the Twitter API to see if an account is still active at the time of check. If it is not, it might be the case that the account was deleted by the owner, or that it was suspended by Twitter because of a violation of the terms of service. As we mentioned in Section 1, when an account gets suspended by Twitter its profile name does not get freed. We can take advantage of this fact to determine if the account was suspended (in case its associated profile name is not available anymore on Twitter) or whether it was deleted. This technique is not flawless, because it could happen that an account changes its profile name and gets suspended later on — in this case, the original profile name with which we observed the account would be freed, and we would erroneously consider the account to be deleted instead of suspended. We acknowledge this problem but also believe that this happens rarely and therefore does not affect the nature of the observed trends.

We applied the aforementioned technique to our dataset of accounts involved in profile reuse, as well as to the dataset of one million random Twitter accounts (**R**). Figure 5 reports a summary of these results. For random Twitter accounts, the ratio of suspension is 5%, while the ratio of accounts being deleted is 12%. If we look at the overall situation of accounts reusing profile names, regardless of their type (first unique, second unique, multi), the fraction of suspended accounts is slightly lower (3%) while the one of deleted accounts is slightly higher (14%). These numbers are however influenced by the fact that the set of first unique accounts does not contain any suspended account. This makes sense since, by design, first unique accounts released their profile name for somebody else to take, and this profile name would have not been made available if the accounts were suspended. 24.5% of the first unique accounts deleted their profile during the measurement period, while the remaining ones simply changed their profile name, thereby releasing the old one.

Interestingly, second unique accounts have a ratio for suspension that is much higher than random accounts (10%). This could be due to the fact that many of these accounts are aggressively sending spam, as we will show in Section 4.6. Multi accounts have a slightly higher rate of suspension than random accounts (5.6%), but this small increase suggests that many of these accounts are involved in schemes that are not clearly malicious, such as schemes for reputation boosting (see Section 4.6).

4.5 Analysis of web links towards reused profile names

In this section, we describe our analysis of the backlinks of reused Twitter user names, in an effort to quantify the misplaced trust that websites have bestowed upon these names. To this end, we started with the set **G** of 106,935 Twitter profile names that were reused. We carried out this analysis immediately after the end of the data collection phase. We used Moz-Open Site Explorer[5] to gather inbound links for each Twitter profile name contained in our set. We discovered that out of the 106,935 profile names, 12,037 (11.3%) of them had active inbound links on the web. The total number of discovered inbound links for these 12,037 profile names was 20,457. Next, for each page hosting a link to a reused Twitter account, we gathered three additional attributes, namely, its Rank, Safety, and Category. The Rank attribute is the Alexa rank of a linking page's main domain (TLD + 1) and was sourced from lookups in Alexa's top one million dataset, as well as scraping of Alexa's API[6] for the links having ranks greater than one million. In this manner, we were able to collect Rank information for 19,178 pages, leaving 1,279 pages without rank information. The categories of Safety and Category were provided by Trendmicro's public website categorization engine[7]. We found that the vast majority of domains linking to the hijacked Twitter user names are safe sites and not exploit-ridden sites that

[4] https://twitter.com/tos

[5] https://moz.com/researchtools/ose

[6] http://data.alexa.com/data?cli=10&url=alexa.com

[7] http://global.sitesafety.trendmicro.com

link to Twitter accounts for SEO purposes. We also found that the vast majority of websites linking to reused Twitter user names are popular benign websites with ranks that imply thousands of visitors on a daily basis.

The discovered websites belong to a wide range of categories. The top three categories are Social Networks, Computers and Internet, and Sports, in that order. The remaining categories include Entertainment, News/Media, and Politics, among others.

We wanted to see if multi accounts have a preference for reusing profile names with more links pointing to them. We counted the number of links pointing to profile names that were reused at some point during our observation, grouping the accounts as earlier described (that is, second unique and multi account groupings). We found that users in the multi accounts category have a higher preference for reusing profile names that have more links pointing to them, than second unique accounts, likely for SEO purposes. Figure 6 highlights this observation, which we confirmed by applying chi-square tests to the data. Specifically, we tested the hypothesis that multi accounts have more websites linking to them with respect to second accounts. We tested the two categories by dividing the accounts in those to which at least two links in the wild were pointing at or not. The p-value of the statistical test is less than 0.0001, stating that there is evidence of difference between multi and second accounts.

Multi accounts have more websites pointing to them with respect to second unique ones, therefore this means that multi accounts are more likely to use profile names that have multiple links on the web pointing to them. We further observed that the majority of Twitter accounts in our dataset that reuse profile names and have inbound links have less than ten such links, with around 90% of accounts having five or less inbound links.

Overall, if we combine the information of ranking, category, and links pointing to reused profile names, we can safely conclude that the majority of websites linking to these hijacked Twitter user names are popular benign sites which have no knowledge of the fact that they are no longer linking to the accounts of popular users and celebrities, but rather to accounts that are now under the control of potentially malicious users. This confirms that profile name reuse in social networks is far from a theoretical danger since it is already happening in the wild.

4.6 Case studies

In this section, we analyze two case studies of accounts reusing profile names. These examples show different ways in which abandoned profile names are used in the wild.

Activity of second unique accounts. In Section 4.4 we showed that second unique accounts have a higher chance to get suspended by Twitter (10% compared to 5% for a random population of accounts). We investigated possible reasons why these accounts could have been suspended. We could not identify a common theme, since the activity of most accounts seemed unrelated to each other. We, however, identified evidence of accounts sharing links to YouTube videos that are now deleted for violation of their terms of service, as well as links pointing to malware and pornography. We also identified accounts from this type engaging in follow-back schemes [26]. Our hypothesis is that these accounts obtained abandoned profile names and started posting malicious content, hoping to leverage the residual popularity of these accounts to attract more victims.

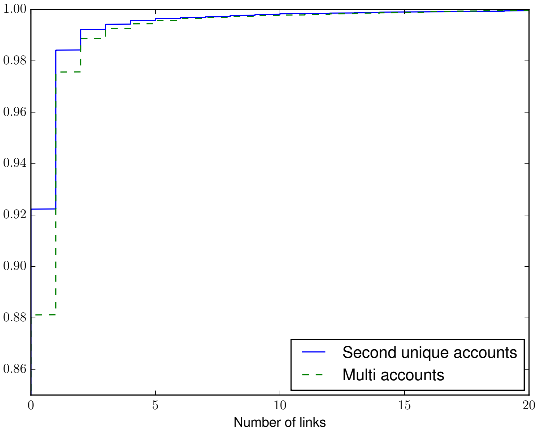

Figure 6: CDFs of link distribution per group of accounts associated with reused Twitter user name.

A group of accounts sharing profile names. We identified an interesting group of 48 accounts that shared 187 distinct profile names. Every time one of the accounts released a profile name, another one took it. Figure 7 shows the timeline of this case study. Overall, the accounts involved in it changed profile name 246 times, often taking names that were previously abandoned. These accounts appeared to portray models from Asia, and the accounts were tweeting in Indonesian. One of the popular words that appeared in the tweets sent by these accounts was “follow,” indicating that the accounts that reused those profile names were likely involved in follow-back schemes, similar to the ones presented in previous work [26].

5. DISCUSSION

In this section we first aim to understand whether the problem of profile name reuse is unique to Twitter. We then reason about our findings and offer some suggestions that social networks should put in place to avoid the security problems highlighted in this paper.

5.1 Profile name reuse on other social networks

In this paper, we analyzed the phenomenon of profile name reuse on Twitter. As a further step, we wanted to understand to what extent other popular social networks allow their users to change their user name. To this end, we examined the terms of service of eight popular social networks other than Twitter: Facebook, Google+, LinkedIn, Pinterest, Reddit, Snapchat, Tumblr, and Youtube. We looked for clauses specifying whether a user is allowed to change their user name and if the released user name can be used by other users in the future. In case the terms of service did not contain a definite answer to these questions, we created two accounts on the platform, a_1 and a_2. We then changed a_1's user name (if that was allowed), and set up a_2 to pick a_1's old user name, checking if this was allowed. The results of our investigation are summarized in Table 1. We identi-

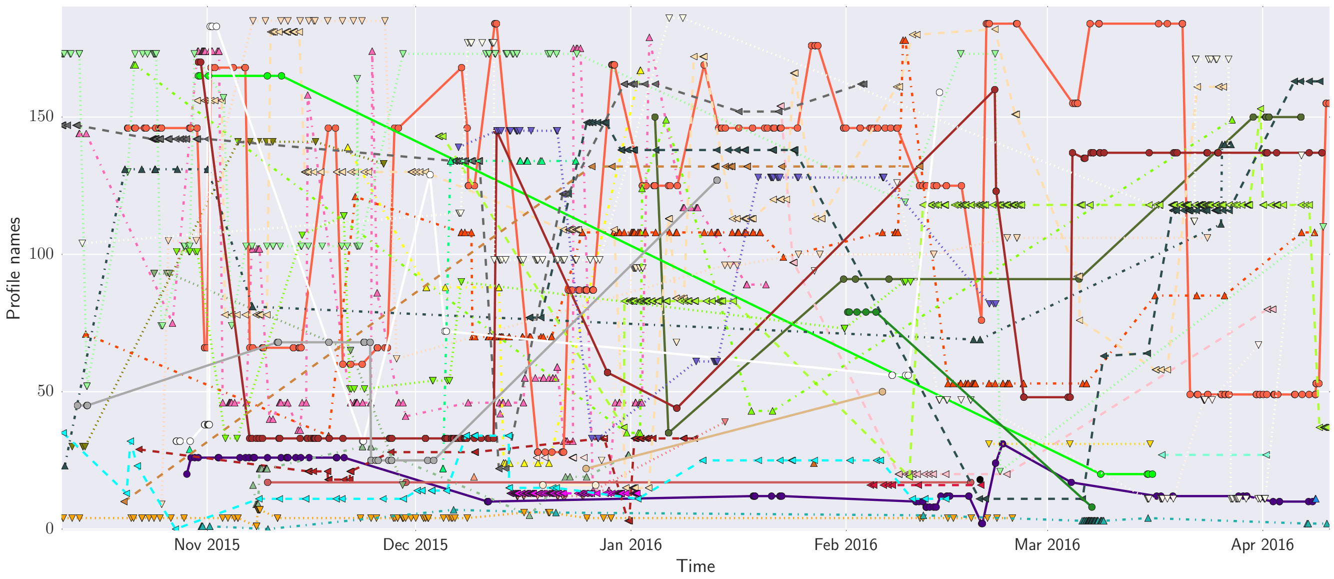

Figure 7: Timeline of a group of 48 accounts that shared 187 distinct profile names during the measurement period. Each line shows a different account, while the Y axis shows different profile names. Each dot represents a tweet sent by the account with a certain profile name. As it can be seen, accounts in this group typically changed profile name multiple times over the measurement period. It also happened in multiple occasions that a previously freed profile name was taken by another account in the group.

Service	Allows change	Allows reuse
Facebook	✓	✗
Google+	✗	✗
LinkedIn	✓	✗
Pinterest	✓	✓
Reddit	✗	✗
Snapchat	✗	✗
Tumblr	✓	✓
Twitter	✓	✓
Youtube	✗	✗

Table 1: Possibility of changing user name on popular online social networks. With the exception of Google+, Youtube, Reddit, and Snapchat, all other popular social networks allow users to change their user name. Moreover, Twitter, Pinterest, and Tumblr allow users to take a user name that used to belong to another account. This policy makes these social networks vulnerable to user name squatting.

fied three categories of social networks, illustrated in detail in the following.

Social networks not allowing a change of user name. These social networks present the strictest settings, not allowing users to change their user name at all. Reddit[8] and Snapchat[9] belong to this group. Google+ and Youtube allow users to change the name of their profile up to three times in 90 days. This name is the one listed on the profile page of a user and used to directly mention her, but is not part of the URL of the profile page. At the same time, users are also allowed to set a *handle* for their page or channel, which allows people to more easily remember the URL associated to it (e.g., *https://youtube.com/taylorswiftVEVO*). After this name has been set, however, it is not possible to change it anymore.[10] We mark both Google+ and Youtube as social networks that do not allow a profile name change.

Social networks allowing a change of user name, but not its reuse. Similar to Google+ and Youtube, Facebook and LinkedIn also allow users to change their name — a legitimate use of this being, for example, a user changing their last name after their spouse's. Moreover, Facebook[11] and LinkedIn[12] allow users to change the user name that is associated to the URL of their public profile page. Facebook, however, limits their users to a single change of user name in the account's lifetime. More importantly, after a user name is changed on these networks, it is not made available for others to use. Facebook and LinkedIn, therefore, are not

[8] https://www.reddit.com/r/help/wiki/faq

[9] https://support.snapchat.com/en-US/a/change-username

[10] https://support.google.com/plus/answer/2676340

[11] https://www.facebook.com/help/203523569682738

[12] https://www.linkedin.com/pulse/20140424124611-12064186-how-to-customize-your-\linkedin-public-profile-url

vulnerable to the type of impersonation attacks described in this paper.

Social networks allowing both a change of user name and its reuse. Finally, Pinterest[13] and Tumblr[14] match Twitter's capability in profile name reuse, both allowing users to change their user name *and* return the old user name to the pool of available ones. A user name on these networks identifies both the user, and the URL of their profile page (or their blog in the case of Tumblr). As a partial mitigation, Tumblr releases the old user name only after 24 hours from the change. These findings show that the security issues and the phenomena highlighted in this paper are not unique to Twitter, but can be found on other social networks too.

5.2 Recommendations to social networks

By using large-scale measurements and individual case studies, in this paper we described the problem of profile-name reuse on Twitter and how other social networks that follow similar name-reuse choices will likely suffer from the same type of abuse. We acknowledge that allowing profile name reuse can have some benefits and it is a user friendly policy that gives a high degree of freedom to users. We showed, however, that this policy also has security implications.

The best solution to these security issues is, in our opinion, not to allow accounts to reuse abandoned profile names. Some social networks such as Facebook and LinkedIn allow the change of user name, but they do not allow the reuse of profile names that have already been used by someone; we suggest that Twitter, Tumblr, and Pinterest should adopt a similar policy to easily and effectively tackle the problem. As a less strict approach, social networks could allow accounts to take abandoned profile names, but should then start monitoring them for signs of malicious activity.

6. RELATED WORK

Being one of the most popular online social networks, Twitter attracted significant interest from the research community, who studied its general characteristics [15], how reputation on the network works [5], investigated peculiar traits of the service, in particular its microblogging features [16], and looked at the unfollow patterns of Twitter users [12].

Particular focus was given to the security issues around Twitter. Grier et al. performed the first large-scale study of abuse on the platform [10], while Thomas et al. studied marketplaces where one can buy fake Twitter accounts [29]. Stringhini et al. studied services that sell compromised accounts as followers to customers that are willing to pay for them [23, 26]. De Cristofaro et al. studied the ecosystem of services that deliver likes to the Facebook pages of their customers [6]. Based on the insights from this research, a number of systems have been proposed to detect malicious activity on Twitter, such as fake accounts [3, 17, 18, 24, 28], compromised accounts [8] or malicious accounts controlled by botnets [4, 25, 31]. Goga et al. studied the problem of impersonation on Twitter [9]. In this work, we showed how reusing abandoned profile names could facilitate the impersonation problem on Twitter.

The problem of profile name reuse on Twitter was originally presented in our preliminary work [20]. In that previous paper, we identified the security issues linked to profile name reuse on Twitter, identified 19,000 profile names that had been reused over a period of one month, and provided general statistics about them. In this paper, we took the study much further, analyzing more accounts for a much longer period of time, and performing a deeper analysis on the modus operandi and characteristics of accounts that reuse abandoned Twitter profile names. Jain et al. also published a study in which they show that users on Twitter temporarily change their profile names [13].

The problem of profile name reuse shares certain similarities with the phenomenon of *cybersquatting* since attackers essentially squat profile names that they do not own in an attempt to confuse visitors about the nature of a Twitter account. In a similar fashion, in domain squatting, attackers register domains that are confusingly similar to popular authoritative domain names, and abuse this similarity for various types of advertising fraud, phishing, and malware delivery [1, 2, 7, 14, 21, 27, 30].

7. CONCLUSION

In this paper we studied the phenomenon of profile name reuse on Twitter. We identified a number of interesting ways in which profile names are reused, some of which are malicious. We also showed that Twitter is not the only social network vulnerable to the issues highlighted in this paper. We hope that this work will help to raise awareness of the issues with freeing profile names after they have been abandoned.

Acknowledgments

We wish to thank the anonymous reviewers for their comments. This work was funded by the H2020 RISE Marie Sklodowska Curie action (MSCA) grant number 691925, by the EPSRC grant number EP/N008448/1, by the NSF under grant CNS-1527086, and by a Boston University Hariri Institute for Computing Research Award. Enrico Mariconti was funded by the EPSRC under grant 1490017, while Jeremiah Onaolapo was supported by the Petroleum Technology Development Fund (PTDF), Nigeria.

8. REFERENCES

[1] P. Agten, W. Joosen, F. Piessens, and N. Nikiforakis. Seven months' worth of mistakes: A longitudinal study of typosquatting abuse. In *Symposium on Network and Distributed System Security (NDSS)*, 2015.

[2] A. Banerjee, D. Barman, M. Faloutsos, and L. Bhuyan. Cyber-fraud is one typo away. In *IEEE Conference on Computer Communications (INFOCOM)*, 2008.

[3] F. Benevenuto, G. Magno, T. Rodrigues, and V. Almeida. Detecting Spammers on Twitter. In *Conference on Email and Anti-Spam (CEAS)*, 2010.

[4] Q. Cao, X. Yang, J. Yu, and C. Palow. Uncovering Large Groups of Active Malicious Accounts in Online Social Networks. In *ACM Conference on Computer and Communications Security (CCS)*, 2014.

[5] M. Cha, H. Haddadi, F. Benvenuto, and K. Gummadi. Measuring User Influence in Twitter: The Million

[13] https://help.pinterest.com/en/articles/edit-your-profile

[14] https://www.tumblr.com/docs/en/blog_management

Follower Fallacy. In *International AAAI Conference on Weblogs and Social Media (ICWSM)*, 2010.

[6] E. De Cristofaro, A. Friedman, G. Jourjon, M. A. Kaafar, and M. Z. Shafiq. Paying for likes?: Understanding facebook like fraud using honeypots. In *Internet Measurement Conference (IMC)*, 2014.

[7] B. Edelman. Large-scale registration of domains with typographical errors. *Harvard University*, 2003.

[8] M. Egele, G. Stringhini, C. Kruegel, and G. Vigna. COMPA: Detecting Compromised Accounts on Social Networks. In *Proceedings of the Network and Distributed System Security Symposium*, San Diego, CA, February 2013.

[9] O. Goga, G. Venkatadri, and K. P. Gummadi. The doppelgänger bot attack: Exploring identity impersonation in online social networks. In *Internet Measurement Conference (IMC)*, 2015.

[10] C. Grier, K. Thomas, V. Paxson, and M. Zhang. @spam: the underground on 140 characters or less. In *ACM Conference on Computer and Communications Security (CCS)*, 2010.

[11] G. E. Hine, J. Onaolapo, E. De Cristofaro, N. Kourtellis, I. Leontiadis, R. Samaras, G. Stringhini, and J. Blackburn. A longitudinal measurement study of 4chan's politically incorrect forum and its effect on the web. *arXiv preprint arXiv:1610.03452*, 2016.

[12] H.Kwak, C. Lee, H. Park, and S. Moon. What is Twitter, a social network or a news media? In *World Wide Web Conference (WWW)*, 2010.

[13] P. Jain and P. Kumaraguru. @i to @me: An anatomy of username changing behavior on twitter. *Arxiv Preprint*, 2015.

[14] M. T. Khan, X. Huo, Z. Li, and C. Kanich. Every second counts: Quantifying the negative externalities of cybercrime via typosquatting. 2015.

[15] B. Krishnamurthy, P. Gill, , and M. Aritt. A Few Chirps About Twitter. In *USENIX Workshop on Online Social Networks*, 2008.

[16] H. Kwak, H. Chun, and S. Moon. Fragile online relationship: a first look at unfollow dynamics in twitter. In *SIGCHI Conference on Human Factors in Computing Systems (CHI)*, 2011.

[17] K. Lee, J. Caverlee, and S. Webb. Uncovering social spammers: social honeypots + machine learning. In *International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2010.

[18] S. Lee and J. Kim. WarningBird: Detecting Suspicious URLs in Twitter Stream. In *Symposium on Network and Distributed System Security (NDSS)*, 2012.

[19] R. A. Malaga. Worst practices in search engine optimization. *Communications of the ACM*, 2008.

[20] E. Mariconti, J. Onaolapo, S. S. Ahmad, N. Nikiforou, M. Egele, N. Nikiforakis, and G. Stringhini. Why Allowing Profile Name Reuse Is A Bad Idea. In *European Workshop on System Security (EUROSEC)*, 2016.

[21] T. Moore and B. Edelman. Measuring the perpetrators and funders of typosquatting. In *Financial Cryptography and Data Security*, 2010.

[22] C. M. Rivers and B. L. Lewis. Ethical research standards in a world of big data. *F1000Research*, 2014.

[23] G. Stringhini, M. Egele, C. Kruegel, and G. Vigna. Poultry Markets: On the Underground Economy of Twitter Followers. In *SIGCOMM Workshop on Online Social Networks*, 2012.

[24] G. Stringhini, C. Kruegel, and G. Vigna. Detecting Spammers on Social Networks. In *Annual Computer Security Applications Conference (ACSAC)*, 2010.

[25] G. Stringhini, P. Mourlanne, G. Jacob, M. Egele, C. Kruegel, and G. Vigna. EvilCohort: Detecting Communities of Malicious Accounts on Online Services. In *USENIX Security Symposium*, 2015.

[26] G. Stringhini, G. Wang, M. Egele, C. Kruegel, G. Vigna, H. Zheng, and B. Y. Zhao. Follow the Green: Growth and Dynamics in Twitter Follower Markets. In *ACM SIGCOMM Conference on Internet Measurement*, 2013.

[27] J. Szurdi, B. Kocso, G. Cseh, J. Spring, M. Felegyhazi, and C. Kanich. The long "taile" of typosquatting domain names. In *USENIX Security Symposium*, 2014.

[28] K. Thomas, C. Grier, J. Ma, V. Paxson, and D. Song. Design and Evaluation of a Real-Time URL Spam Filtering Service. In *IEEE Symposium on Security and Privacy*, 2011.

[29] K. Thomas, D. McCoy, C. Grier, A. Kolcz, and V. Paxson. Trafficking fraudulent accounts: The role of the underground market in twitter spam and abuse. In *USENIX Security Symposium*, 2013.

[30] Y. Wang, D. Beck, J. Wang, C. V., and B. Daniels. Strider typo-patrol: discovery and analysis of systematic typo-squatting. In *USENIX Workshop on Steps to Reducing Unwanted Traffic on the Internet (SRUTI)*, 2006.

[31] Y. Zhao, Y. Xie, F. Yu, Q. Ke, Y. Yu, Y. Chen, and E. Gillum. Botgraph: Large scale spamming botnet detection. In *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2009.