(How) Do People Change Their Passwords After a Breach?

Sruti Bhagavatula
Carnegie Mellon University

Lujo Bauer
Carnegie Mellon University

Apu Kapadia
Indiana University Bloomington

Large-scale password breaches are a common occurrence. To protect their users' accounts and data after these breaches, companies often attempt to convince their users to change their passwords on the affected sites. In this article, through our unique infrastructure that monitored the security behaviors of 249 participants over almost two years, we report on the extent to which people change their passwords after breaches. We also analyze the strength and quality of their new passwords. In particular, we examined the real-world passwords chosen by participants with accounts on one of nine breached domains. Our results show that users typically did not change their passwords after a breach, and when they did, the new and old passwords were similar enough that an attacker could easily guess their new password [4]. Additionally, users continued to use their old stolen passwords across other online accounts, making those accounts also vulnerable. Our findings suggest that it is difficult to ensure users' accounts are safe in the wake of breaches if users themselves are the ones responsible for changing their passwords. For example, sites could require password resets and multi-factor authentication. These findings highlight the need for designing systems that keep users safe without putting the onus on them to take correct remedial action.

Password breaches have been on the rise, affecting companies from Yahoo! to popular gaming sites such as League of Legends and Neopets [1]. Passwords were often stored insecurely before being stolen (e.g., in plain text or using cryptographic mechanisms known to be vulnerable, such as the use of 'unsalted hash functions'). Even when stored with stronger cryptographic protections (i.e., when hash functions are used appropriately), many passwords can be easily reverse-engineered through dictionary-based password guessing attacks. Overall, password breaches leave users vulnerable unless they change their passwords on the affected sites [1]. Previous work has also shown that, on average, a user exactly or partially reuses their passwords on over 50% of their accounts [7, 10, 23]. In such cases, when a person's password on one site is compromised, they incur the risk that an attacker will be able to gain access to their other accounts that use the same or similar passwords (i.e., through an attack known as 'credential-stuffing'). To mitigate and prevent this from occurring, passwords on those other sites also need to be changed. In order to make informed recommendations to companies on best risk-mitigation practices after a breach, it is instructive to examine people's current password-changing behavior following a breach.

Password breaches we studied

We studied nine of the most significant password breaches that became public in 2017 and 2018: Imgur (breach announced Nov. 2017) [18], Deloitte (Sep. 2017) [15], Disqus (Oct. 2017) [31], and Yahoo! (Feb. and Oct. 2017) [16, 17], MyFitnessPal (Mar. 2018) [9], Chegg (Sep. 2018) [5], CashCrate (Jun. 2017) [22], FLVS (Mar. 2018) [14], and Ancestry (Dec. 2017) [20]. We selected these breaches because they were part of Identity Force's list of biggest breaches in 2017 [6], Digital Information World's list of biggest breaches in 2018 [24], or the password breaches reported on HaveIBeenPwned [1] (a service that alerts users if their passwords were found in breached password lists).

(How) Did people change their passwords after a breach?

We analyzed the password-changing behavior for participants who had accounts on at least one of these nine breached domains. For these 63 participants, we examined who changed their passwords on breached domains and whether they changed their passwords on other sites that were similar to the breached ones. We also measured whether the quality of their passwords improved. We measured this quality along the aforementioned dimensions of similarity, reuse, and strength. A constructive change implies less similarity, less reuse, and an increase in strength of passwords.

We identified 63 participants who entered passwords on at least one of the nine breached domains we study, implying that they had an account on the domain and therefore were *potentially* affected when the domain was breached. In particular, these participants entered a password on at least one of the breached domains before the breach announcement date and were active in the behavior-monitoring study for at least 90 days after the announcement. We checked whether the identified participants changed their password on the affected domain. If they did, we checked how similar the new and old passwords were and whether the new password was stronger than the older one.

Overall, only 21 of the 63 participants changed a password on a breached domain following the breach announcement. However, only two of the 21 changed their password within a month of the breach announcement, only five total within two months, and only eight within three months. Even in cases when their passwords were *definitely* stolen, only a small fraction of participants changed their passwords: 49 of the 63 participants had Yahoo! passwords, and only 18 of them changed their Yahoo password even though all were affected [17].

Users often use the same passwords across multiple online accounts [2]. Password reuse is dangerous because if one of a user's passwords is breached, an attacker is likely to be able to gain access to their other accounts that use the same or similar passwords. Following previous work [23], we considered two passwords to be *partially reused* if they shared at least a three-character substring and *exactly reused* if the same password was used for multiple accounts. We computed the extent of reuse of a password as the fraction of that participant's *other* passwords that exactly or partially reuse the password in question. For example, if a user had "iluvDONUTS90", "ih8bagels20", "uluvDONOTS200", and "iluvDONUTS90" as their passwords for four different accounts at a given time, the amount of reuse of the first password is 2/3. We examined the change in password reuse for each participant who changed a password on a breached domain. We did this by comparing the reuse before the password change and the reuse a month after it. For nine participants, the new passwords on the breached domains were more reused; for ten it was less reused; and for two it was equally reused, implying that the majority of participants' password changes did not result in higher-quality passwords.

Most of the time, participants either made their passwords stronger or kept them equally strong [19]. Nine of the 21 participants created stronger passwords while 11 created equally strong passwords. Only one created weaker passwords. On average, participants created stronger new passwords that would be only slightly more difficult for an adversary to crack when compared to their old passwords. However, participants' new passwords were similar enough to their old passwords, on average, that an adversary with access to someone's old password may be able to guess their new password with less effort.

In summary, participants' new passwords were slightly stronger but often similar to their old passwords. Additionally, the new passwords on these breached sites were still often similar to passwords they used on other domains, continuing to leave the users vulnerable.

(How) Did people change passwords beyond the affected sites?

Even if they change their passwords on breached domains, people may still be at risk of compromise on other sites through credential-stuffing attacks. In credential-stuffing attacks, the attackers use stolen username-password combinations to breach users' other accounts that use the same or similar username and password combinations. Therefore, it is important for users not only to change passwords on the breached site but also on other sites where they used the same or similar passwords. As part of our study, we measured how many participants changed similar passwords (i.e., similar to the breached passwords) on other sites and the quality of these changes.

If a participant changed their password on a breached domain, we examined whether they changed any of their similar passwords on domains beyond the breached domain in the month that followed. We considered two passwords *similar* if they shared a substring that is at least as long as half the length of the longer password. For example, the passwords "iluvDONUTS90" and "ih8DONUTS90" are similar since they share the substring "DONUTS90" that is at least half as long as the longer password, "iluvDONUTS90". We measured similarity by examining passwords similar to the most recent passwords entered on any domain before the breach announcement.

The 21 participants who changed a password on a breached domain had, on average, 30 passwords similar to their older, breached password. More than half of these participants (14 participants) changed a similar password. However, on average, they changed only four of these 30 similar passwords within a month of changing their password on the breached site. Nine of the participants changed to a password that shared a substring of three or more characters with their old password; these nine participants' new passwords on average shared a substring 44% the length of the longer password with their older counterparts.

The 14 participants changed their similar passwords to be on average 10% stronger (based on accepted password guessing metrics) than their original passwords on the breached domains and 18% stronger than the passwords being changed. However, the majority (63%) of the changes resulted in weaker or equal-strength passwords.

Overall, participants changed very few passwords on breached domains and even fewer similar passwords on other domains. Even when they did change a password, the change was often to more reused passwords, passwords similar to the old ones, and not to stronger passwords.

How much better were people's passwords on breached domains than on other domains?

The main findings of our study show that people's password change behavior after breaches is less than ideal. While this finding could be bad news for people's cyber security, it is important to study whether this non-ideal behavior applies only to changes to breached passwords (where the impact of password change is the most crucial) or to their overall password changes. To provide a baseline against which to compare breach-related password changes, we computed password-change statistics for all password changes by all 249 participants over the two years spanned by the dataset.

Overall, the participants' password updates resulted in relatively similar changes in strength, regardless of whether they were on breached domains. However, breach-related password changes resulted in more dissimilar new passwords (i.e., their new passwords were reused less than before across their other accounts). Of the 249 participants, 223 individuals changed a password and made at least one password change that involved carrying over a substring of at least three characters. In such cases, old and new passwords shared a substring, on average, 85% the length of the longer of the two. Looking at *all* password changes made by those 223 participants over the two year period, 70% of these password changes resulted in weaker or equally strong passwords. However, new passwords were on average 23% stronger than older passwords and the median change in password strength was neutral (i.e., the old and new passwords were equally strong). In other words, changes related to breaches resulted in less strong but more dissimilar new passwords compared to participants' overall password changes.

Finally, to obtain an overall baseline for how often participants reused passwords, we computed the average amount of reuse

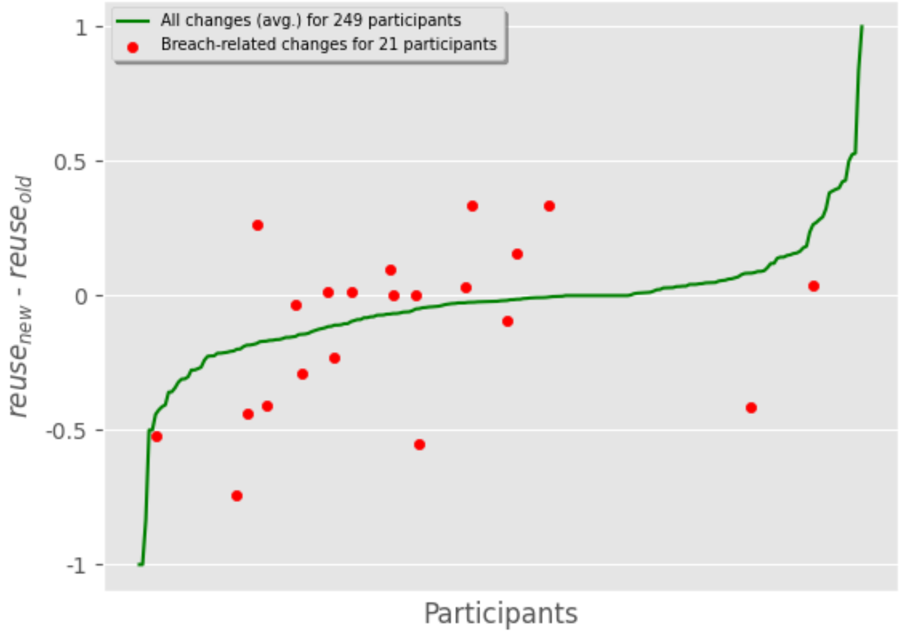

Figure 1: Change in password reuse across each password change per participant. Each point on the X-axis represents a participant and they are sorted by how much more reused the participants' changed passwords were on average than their old passwords. Y-axis values below zero indicate that passwords became less reused, which is more desirable.

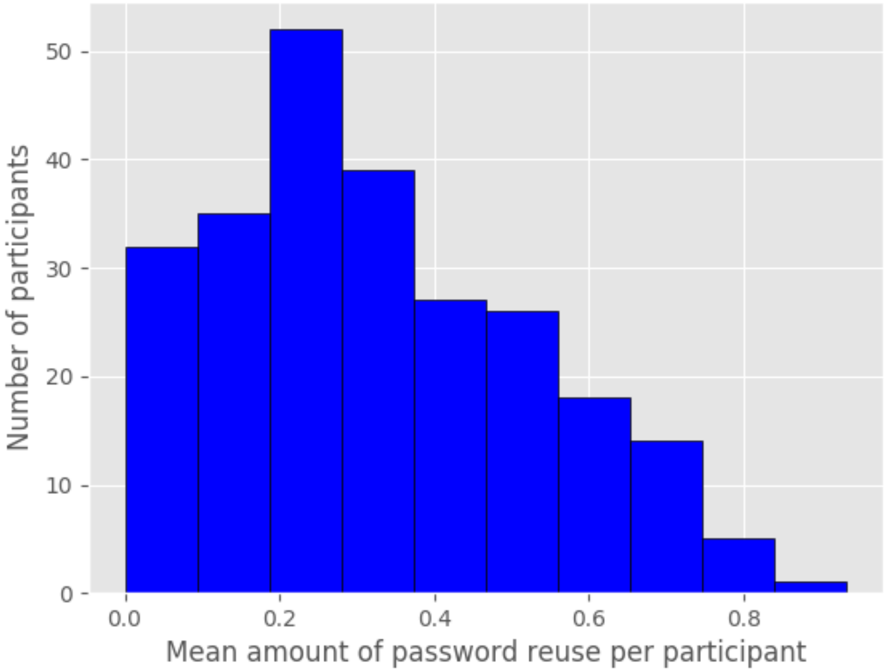

Figure 2: The average amount of reuse of each participant's unique passwords. The amount of reuse of one password is represented as the fraction of a participant?s passwords that exactly or partially reuse (i.e., share a three-character or longer substring with [23]) that password; the X-axis describes the average of these fractions for each participant.

of each participant's unique passwords entered per domain during the time period spanned by the dataset. In particular, if a participant had three unique passwords on `google.com` and five on `yahoo.com`, we computed the average reuse of those eight passwords even if some of the `yahoo.com` passwords were exactly reused on `google.com`. We computed the reuse of each password at the time of that password's first entry.

Figure 1 shows, per participant, how often participants' new passwords were reused across their internet accounts compared to their old passwords. In this figure, red dots above the green line indicate that a participant's breached password changes resulted in lower quality passwords (i.e., more reuse). More than half of the participants changed their passwords on a breached domain to be more reused across their other accounts than their old breached domain password, compared to their baseline changes in password reuse. Figure 2 shows the average amount of reuse of all of each participant's unique passwords entered per domain.

Users can be more protected after breaches

Our results show that users typically do not change their passwords after password breaches even when their old passwords are definitely compromised. The low number of participants who changed a password on a breached site and the even lower number of those whose changes were to better (i.e., less reused, less similar, or stronger) passwords suggest that password breach notifications are failing dramatically, both at causing users to take action and even more at causing them to take *constructive* action.

Government regulators should take note of the ineffectiveness or absence of breach notifications and consider imposing stricter requirements on companies to implement better practices [11,27,30,32,33]. In particular, they should require companies to repeatedly send notifications until they have positive confirmation that the notifications have been understood and that any instructions have been followed. Existing breach notification laws should be altered such that affected companies must force password resets and verify that new passwords are "better" in terms of them being dissimilar from users' old passwords before allowing users to successfully access their accounts again.

Websites can also do more to protect their users after a breach. To ensure old and new passwords are dissimilar, as a user is entering a new password, websites can display which parts of the newly entered password match the old one and allow the user to modify their new password until the amount of overlap is below a threshold (e.g., less than 50% the length of the longer password, as we used in this study). Similarly, to provide real-time feedback on how to make a stronger password than the old one, websites can make use of in-browser password strength meters [19] and display prompts on what properties of the entered password could be changed so that its strength becomes satisfactory [26]. The other indicator of password quality that we studied in this article is how often a user's password is reused on their other online accounts. Given that websites typically won't share their users' passwords with other websites, it may be difficult for websites to ensure their users' passwords are not reused on other sites. However, at the time of password creation and updating, websites can still display a warning against password reuse and prompt users to verify that they are not reusing their new password on other sites, even if the websites cannot verify this.

From a preventative standpoint, companies should use an authentication method other than just passwords, for example, multi-factor authentication, or they should allow the use of single sign-on (e.g., sign on via Google). Companies should be required to store their passwords securely by hashing and salting their passwords to avoid credential-stuffing and rainbow-table attacks on plaintext or weakly hashed passwords [21, 25]. Companies can also make use of external services such as HaveIBeenPwned and force users to change their passwords when they encounter a matching hash in this database. They can additionally encourage the use of password managers to help their users with password management.

Some facets of good password maintenance behavior may be difficult for users to grasp [3, 12, 13, 28, 29]. Users may find managing passwords to be simply overwhelming. For instance, the affinity towards changing to weaker or equal-strength passwords could be because when people feel compelled to choose new passwords they have poor awareness of password strength or the additional memory burden leads them to pick weaker passwords [8, 28] (e.g., they might change just enough characters to satisfy system requirements). Related to partial password reuse, people may find it difficult to understand how their "different" password is still similar to other passwords (i.e., they might be unintentionally partially reusing passwords). Therefore, it is vital that password education focuses on encouraging people not to re-use passwords as well as to use password-managing tools. Potential efforts in this space could include integrating password-reuse trackers within tools that people may already use and trust to store their passwords. Some password managers, such as 1Password, already warn users if one of their saved passwords is reused. Password managers, including those built into web browsers, could go further and more actively discourage password reuse. Users may additionally benefit from the above kinds of tools automatically changing users' passwords to mitigate reuse and the presence of low quality passwords on their behalf.

Conclusions

Password breaches are now a common occurrence and pose a major security and privacy threat to a large number of users. Yet, we find that after a breach users are unlikely to change their passwords, and even if they do, they do not pick suitably secure passwords. Our findings suggest that companies need to improve their communication with users as well as mechanisms to enforce password changes after a breach. Even so, deeper challenges remain with users' rampant reuse of passwords across websites. Tools such as password managers should be improved to prevent or reduce such reuse as well as provide efficient means to update passwords following a breach.

References

[1] Have I Been Pwned: Pwned websites, 2019. URL: https://haveibeenpwned.com/PwnedWebsites.

[2] Jacob Abbott, Daniel Calarco, and L. Jean Camp. Factors influencing password reuse: A case study. In *Research Conference on Communications, Information and Internet Policy*, 2018.

[3] Farzaneh Asgharpour, Debin Liu, and L. Jean Camp. Mental models of security risks. In *International Conference on Financial Cryptography and Data Security*, 2007.

[4] Sruti Bhagavatula, Lujo Bauer, and Apu Kapadia. (How) Do people change their passwords after a breach? In *Workshop on Technology and Consumer Protection*, 2020.

[5] Sean Cavanagh. Education company Chegg acknowledges data breach, puts 40 million users on notice. *Market Brief*, 2018. URL: https://marketbrief.edweek.org/marketplace-k-12/tutoring-company-chegg-acknowledges-data-breach-puts-40-million-users-notice/.

[6] Heidi Daitch. 2017 data breaches - the worst breaches, so far. *IdentityForce*, 2017. URL: https://www.identityforce.com/blog/2017-data-breaches.

[7] Anupam Das, Joseph Bonneau, Matthew Caesar, Nikita Borisov, and XiaoFeng Wang. The tangled web of password reuse. In *Network and Distributed System Security Symposium (NDSS)*, 2014.

[8] Geoffrey B. Duggan, Hilary Johnson, and Beate Grawemeyer. Rational security: Modelling everyday password use. *International journal of human-computer studies*, 2012.

[9] Paul Fipps. Important message regarding MyFitnessPal account security. *MyFitnessPal*, 2018. URL: https://content.myfitnesspal.com/security-information/notice.html.

[10] Dinei Florencio and Cormac Herley. A large-scale study of web password habits. In *International conference on World Wide Web (WWW)*, 2007.

[11] Maximilian Golla, Miranda Wei, Juliette Hainline, Lydia Filipe, Markus Dürmuth, Elissa M. Redmiles, and Blase Ur. What was that site doing with my facebook password?: Designing password-reuse notifications. In *ACM Conference on Computer and Communications Security (CCS)*, 2018.

[12] Hana Habib, Jessica Colnago, Vidya Gopalakrishnan, Sarah Pearman, Jeremy Thomas, Alessandro Acquisti, Nicolas Christin, and Lorrie Faith Cranor. Away from prying eyes: Analyzing usage and understanding of private browsing. In *Symposium on Usable Privacy and Security (SOUPS)*, 2018.

[13] Ameya Hanamsagar, Simon S. Woo, Chris Kanich, and Jelena Mirkovic. Leveraging semantic transformation to investigate password habits and their causes. In *Conference on Human Factors in Computing Systems (CHI)*, 2018.

[14] Benjamin Herold. Florida virtual school reveals huge data breaches. *Education Week - Digital Education*, 2018. URL: http://blogs.edweek.org/edweek/DigitalEducation/2018/03/florida_virtual_school_data_breaches.html.

[15] Brian Krebs. Krebs on security: Deloitte breach affected all company email, admin accounts. *Krebs on Security*, 2017. URL: https://krebsonsecurity.com/2017/09/source-deloitte-breach-affected-all-company-email-admin-accounts/.

[16] Mohit Kumar. Yahoo hacked once again! Quietly warns affected users about new attack. *The Hacker News*, 2017. URL: https://thehackernews.com/2017/02/yahoo-hack.html.

[17] Selena Larson. Every single Yahoo account was hacked. *CNNMoney*, 2017. URL: https://money.cnn.com/2017/10/03/technology/business/yahoo-breach-3-billion-accounts/index.html.

[18] Natasha Lomas. Imgur says 1.7m emails and passwords were breached in 2014 hack. *TechCrunch*, 2017. URL: https://techcrunch.com/2017/11/27/imgur-says-1-7m-emails-and-passwords-were-breached-in-2014-hack/.

[19] William Melicher, Blase Ur, Sean M. Segreti, Saranga Komanduri, Lujo Bauer, Nicolas Christin, and Lorrie Faith Cranor. Fast, lean, and accurate: Modeling password guessability using neural networks. In *USENIX Security Symposium*, 2016.

[20] Francis Navarro. Ancestry.com suffers big data leak - 300,000 user credentials exposed. *Komando.com*, 2017. URL: https://www.komando.com/happening-now/435921/ancestry-com-suffers-big-data-leak-300000-user-credentials-exposed.

[21] Philippe Oechslin. Making a faster cryptanalytic time-memory trade-off. In *Annual International Cryptology Conference*, 2003.

[22] Pierluigi Paganini. Cashcrate cash-for-surveys site breached, 6 million accounts stolen. *Security Affairs*, 2017. URL: https://securityaffairs.co/wordpress/60083/data-breach/cashcrate-data-breach.html.

[23] Sarah Pearman, Jeremy Thomas, Pardis Emami Naeini, Hana Habib, Lujo Bauer, Nicolas Christin, Lorrie Faith Cranor, Serge Egelman, and Alain Forget. Let's go in for a closer look: Observing passwords in their natural habitat. In *ACM Conference on Computer and Communications Security (CCS)*, 2017.

[24] Saima Salim. Revealed: The 21 biggest data breaches of 2018. *Digital Information World*, 2018. URL: https://www.digitalinformationworld.com/2018/12/biggest-data-breaches-of-2018.html.

[25] Kurt Thomas, Jennifer Pullman, Kevin Yeo, Ananth Raghunathan, Patrick Gage Kelley, Luca Invernizzi, Borbala Benko, Tadek Pietraszek, Sarvar Patel, Dan Boneh, and Elie Bursztein. Protecting accounts from credential stuffing with password breach alerting. In *USENIX Security Symposium*, 2019.

[26] Blase Ur, Felicia Alfieri, Maung Aung, Lujo Bauer, Nicolas Christin, Jessica Colnago, Lorrie Faith Cranor, Henry Dixon, Pardis Emami Naeini, Hana Habib, Noah Johnson, and William Melicher. Design and evaluation of a data-driven password meter. In *Conference on Human Factors in Computing Systems (CHI)*, 2017.

[27] Blase Ur, Felicia Alfieri, Maung Aung, Lujo Bauer, Nicolas Christin, Jessica Colnago, Lorrie Faith Cranor, Henry Dixon, Pardis Emami Naeini, Hana Habib, Noah Johnson, and William Melicher. Design and evaluation of a data-driven password meter. In *Conference on Human Factors in Computing Systems (CHI)*, 2017.

[28] Blase Ur, Jonathan Bees, Sean M. Segreti, Lujo Bauer, Nicolas Christin, and Lorrie Faith Cranor. Do users' perceptions of password security match reality? In *Conference on Human Factors in Computing Systems (CHI)*, 2016.

[29] Rick Wash, Emilee Rader, and Chris Fennell. Can people self-report security accurately?: Agreement between self-report and behavioral measures. In *Conference on Human Factors in Computing Systems (CHI)*, 2017.

[30] Jane K. Winn. Are better security breach notification laws possible. *Berkeley tech. LJ*, 2009.

[31] Jason Yan. Security alert: User info breach. *Disqus Blog*, 2017. URL: https://blog.disqus.com/security-alert-user-info-breach.

[32] Yixin Zou, Shawn Danino, Kaiwen Sun, and Florian Schaub. You 'might' be affected: An empirical analysis of readability and usability issues in data breach notifications. In *Conference on Human Factors in Computing Systems (CHI)*, 2019.

[33] Yixin Zou and Florian Schaub. Beyond mandatory: Making data breach notifications useful for consumers. *IEEE Security & Privacy*, 2019.