More identities, more problems: Managing multiple identities in digital self-presentations*

Samantha C. Phillips[1] and Kathleen M. Carley[1]

CASOS Center, Software and Societal Systems Department
Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA 15213
{samanthp, carley}@andrew.cmu.edu

Abstract. Presenting multiple identities with conflicting expectations can prompt revisions to self-presentation, with potential consequences for future social interactions and psychological well-being. Social media provides a unique context to examine how individuals navigate multiple identities amid real-world social and structural constraints. This study investigates how the number and diversity of identities in a profile description predict subsequent updates by analyzing 7.3 million descriptions written by over 370,000 users on X (formerly Twitter). We first classified identities extracted from descriptions into topical categories, allowing us to assess the relative stability of different types of identities and to approximate the diversity of identities within each description. Our analysis reveals that users who present more numerous and diverse identities are significantly less likely to add new identities, and substantially more likely to revise or remove existing ones. These results suggest that simplification is the dominant strategy for managing multiple identities in the context of social media, where audiences are diverse and potentially unknown. We conclude by outlining implications for platform design and policy aimed at supporting flexible and inclusive identity expression online.

Keywords: Impression management · Identities · Social media.

1 Introduction

People constantly refine the identities in their self-presentation in response to feedback from social interactions and context, a process referred to as impression management [16,7]. Each identity (e.g., personal trait, group membership,

* This work was supported in part by the Knight Foundation and the Office of Naval Research grant MURI: Persuasion, Identity, & Morality in Social-Cyber Environments, N00014-21-12749. Additional support was provided by the Center for Computational Analysis of Social and Organizational Systems (CASOS) and Center for Informed Democracy and Social-cybersecurity (IDeaS) at Carnegie Mellon University. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Knight Foundation, Office of Naval Research, or the U.S. Government.

relational role) is associated with a set of social norms, shaping behavior and attitudes to align with perceived expectations [17]. However, when multiple presented identities impose conflicting demands, individuals experience identity conflict [25,1,15]. To maintain self-consistency and social credibility, people may simplify their self-presentation by foregrounding certain identities while suppressing others [3,12]. This simplification of identity performance can have negative psychological and social consequences, including increased intergroup bias and diminished well-being [25,20,21]. Therefore, we are motivated to understand the contexts in which people manage identity conflict by decreasing the complexity of their self-presentation. While prior work has demonstrated strategies to manage multiple identities [12,11,3,1,9], there is little analysis of management practices in real-world settings like social media.

On social media platforms like X (formerly Twitter), users can write a short, text-based description that appears alongside their profile [19,18]. These descriptions allow users to explicitly present a subset of their identities and researchers to observe the evolution of organically-occurring self-presentations. Unlike offline contexts where communication is synchronous and the audience is largely known, social media users have to contend with invisible audiences and context collapse. Context collapse is the flattening of social contexts associated with different social norms and responses, driving the consideration of multiple, distinct audiences in a single self-presentation [14,2]. The pressure to present a single, verifiable self-presentation is especially challenging in the face of diverse audiences. To navigate these competing pressures, we hypothesize that users with more numerous and diverse identities will be less likely to add one or more identities, more likely to remove one or more identities, and more likely to simultaneously add and remove one or more identities - effectively simplifying their self-presentation.

In this work, we analyzed 7.37 million profile descriptions from over 370,000 X users who tweeted about COVID-19 vaccines at least twice between September 2020 and August 2021. We first identified which types of identities are most frequently added to or removed from profile descriptions. To do so, we extracted discrete identity phrases from each profile and manually categorized approximately 10,000 of the most common identities into 26 topical identity categories (e.g., political, gender, sports, relationships). This classification enabled us to assess which categories are especially ephemeral or likely to carry higher signaling costs. Next, we examined whether the number of identities or the number of identity types in a description predict the likelihood of subsequent updates. Specifically, we estimated fixed-effects logistic regression models predicting whether users modified their profile descriptions - by adding, removing, or simultaneously adding and removing one or more identities - between consecutive tweets. This analysis directly tests our hypothesis that more numerous and diverse identity performances constrain future signaling flexibility and increase the likelihood of simplification through revision.

Understanding how presenting multiple identities influence profile description changes offers insight into how people navigate context collapse, audience

ambiguity, and reputational risk online. These dynamics not only shape individual identity expression and behavior, but also carry broader implications for designing healthier, more inclusive digital spaces.

2 Related works

A substantial body of research has examined how and why people update their self-presentation both on and offline. Goffman's theory of impression management emphasizes the contextual and dynamic nature of identity performance [7,13]. Shifts in perceived audience and social context, therefore, drives changes in self-presentation. Updates may also reflect incorporation of social feedback, such as suppressing identities associated with hostile reactions [22,28]. In addition, people may increase public conformity to certain identities to gain social acceptance or positively differentiate themselves from relevant groups [24].

Furthermore, people tend to present identities that align with their ideal self, updating their self-presentation to align with desirable characteristics [29,8]. There is evidence this holds online as well [23]. Self-presentation may also change as individuals refine their self-image, pursue aspirational roles, or experience major life events [6]. Individual differences also play a role - those high in self-monitoring are more likely to adapt their identity performance [22]. Note that we control for individual differences by including fixed effects in our models predicting updating behavior.

Several researchers have proposed strategies for managing conflicting multiple identities [12]. For example, Burke identifies three approaches to management identity conflict: shift the meanings of the identities, withdraw from one of the identities, avoid situations where conflicting identities are salient [3]. Roccas and Brewer suggest individuals make sense of multiple identities through one of the following four approaches: define ingroup as the intersection of multiple groups, adopt one identity as dominant, compartmentalize identities depending on the context, merge identities into larger, inclusive ingroup [20]. Amiot and co-authors developed a dynamic model of identity development and integration, including compartmentalization (i.e., present different identities in different contexts) and integration (i.e., create higher order categorizations to resolve conflict and perform identities simultaneously) [1]. More recently, Jones and Hynie synthesized literature on multiple identity management to propose three techniques [11]: reconciliation involves integrating aspects of different identities, realignment involves prioritization of one identity over the others, and retreat indicates avoidance of conflicting identities.

Affordances of social media likely influence the propensity to use and effectiveness of multiple identity management techniques. For example, maintaining different identities in separate contexts is achievable through the creation of multiple profiles [31]. In other cases, people attempt to present multiple identities in a single profile. Integration and similar identity management strategies that involve enacting multiple identities concurrently may be especially challenging on social media due to the platform's structural constraints like context collapse,

audience ambiguity, and the persistence and visibility of identity signals [2,27]. In the presence of multiple, distinct audiences, users have to contend with different expectations and normative standards. Audience ambiguity adds further uncertainty. When the audience is unknown or fluid, users lack the contextual cues needed to tailor their self-presentation, increasing the perceived risk of misinterpretation or reputational harm. In addition, identity performances such as profile descriptions, usernames, or posts are often permanently archived and broadly accessible, making inconsistencies across identity claims more noticeable and potentially more consequential. Users are thus incentivized to simplify their self-presentation by foregrounding one identity that aligns with perceived audience expectations while suppressing others that may be misunderstood, stigmatized, or contextually inappropriate. Additionally, features such as character-limited descriptions and algorithmic filtering of information exposure further constrain opportunities for complex identity expression, reinforcing a tendency toward realignment or retreat rather than reconciliation. Our observational analysis of profile description changes deviates from prior work on identity conflict management that uses surveys and interviews.

3 Data and Methods

3.1 Data

To obtain a sample of profile descriptions by the same users over time, we used a streaming keyword search via Twitter v1 API to collect tweets about COVID-19 vaccines between September 1, 2020 through August 31, 2021. In total, we extracted over 29.5M tweets and 12.1M unique descriptions written by 5.8M unique users. We only analyze the 8.8M descriptions written by 2.5M users that tweeted at least twice, so we could observe if they changed their description.

3.2 Identity extraction and categorization

To divide descriptions into discrete phrases, we use a method developed by Pathak et al. [18] that involves splitting the text on common delimiters (e.g., ",", "-"). This method returned 8.9M phrases, 4.8M of which only occurred once. Next, an author manually reviewed the 14,253 phrases that occurred in at least 0.00001% of descriptions. This initial review resulted in the removal of approx. 4,000 phrases determined to be nonsensical, not applicable to the user, or referring to their online behavior (e.g., "100% follow back", "anytime", "terms of use"). Of 12.1M unique descriptions collected, 5.5M do not contain any identities, 2.2M contain one identity, 1.4M contain two identities, and the remaining 3.0M contain three or more identities. Of the 6.6M descriptions containing at least one identity, the average number of identities is 2.93 (std. dev. 2.3) and median number is 2.

Buliding upon prior analyses of identities in profile descriptions [30,18], we next categorized the 10,633 identity phrases into 26 categories (in order of number of identities per category): astrology, COVID-19, gender/sexuality, political, activism, nature/animals, location, relationships, business/finance, media,

travel, art, religion, profession, race/ethnicity/nationality, technology/science, location type, medicine, social science, disability/illness, military, games, school, sports, personal attribute, age. To validate the categories, we asked gpt-4o to assign one or more of the categories to a sample of 200 identities, resulting in a Krippendorf's alpha of 0.91.

Of the descriptions containing at least one identity, the average number of categories is 2.58 (std. dev. 1.7) and median number is 2. Note that although a single identity can be in multiple categories (e.g., "teacher" is in both profession and school). The number of identities and categories in a given description is highly correlated with a Pearson's correlation of 0.91 (0.84 when excluding descriptions with 0 identities that necessarily contain 0 categories).

3.3 Identity stability

We identified whether or not a change occurred (addition, removal, or both addition and removal of one or more identities) in the profile description between each pair of consecutive tweets by the same user. Of the 15,705,392 pairs of consecutive tweets by the same user in our dataset, 14,437,354 displayed no change in description and 1,268,038 displayed any change. Specifically, there were 425,311 instances of one or more identities added, 392,812 instances of identities removed, and 449,915 of both occurring.

3.4 Analysis

We estimated high-dimensional fixed-effects logistic regression models using the feglm() function from the fixest package in R to predict whether users updated (i.e., added, removed, both, any change) their profile descriptions during the time period between consecutive tweets. We run models using the number of identities or number of types of identities as predictors and control for time between observations. The model included user fixed effects to control for all time-invariant individual characteristics, enabling estimation of within-user effects while still computing absolute predicted probabilities. This approach scales efficiently to large datasets and allows us to identify threshold values of identity complexity where users become more likely to update than not. Note that we avoid scaling the predictors for interpretability. When we use scaled predictors, the direction and significance of estimates are the same and model fits are nearly identical.

4 Results

4.1 Types of identities vary in stability

Figure 1 provides the average and 95% CI of the number of additions or removals per identity in each identity category. We averaged across identities in each category instead of using counts due to discrepancies in the number of identities assigned to each category. We find variance in the degree of stability across

categories. Some types of identities are notably ephemeral, likely due to being event driven (e.g., COVID-19) or high costs of signaling (e.g., political). Others are very stable which may reflect positive feedback on the platform due to broad appeal (e.g., sports, games) and generally low costs of signaling (e.g., age).

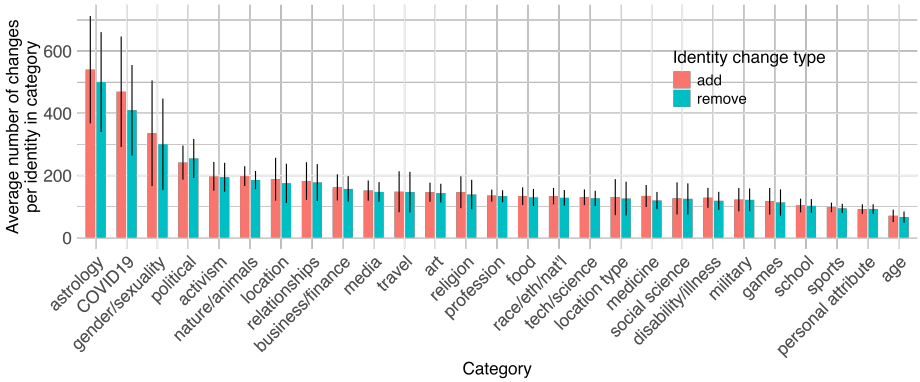

Fig. 1: Average and 95% CI of the number of additions or removals per identity in each identity category

4.2 Multiple identities encourage description updating

As represented in Figure 2, each additional identity or category represented in a description decreases likelihood to add one or more identities and increases likelihood to remove or both add and remove one or more identities. Specifically, each additional identity or category is associated with a 90.3% or 91.8% decrease in likelihood to add another identity or set of identities, respectively. The likelihood to add an identity averaged across users is only above 0.5 when the number of identities, and therefore categories, is 0.

Users are nearly 6 or 8 times more likely to remove one or more identities for each additional identity or category in their description, respectively. Notably, the absolute likelihood of removing an identity averaged across users is only greater than 0.5 when descriptions contain four or more identities or discrete categories. The influence of multiple identities on both adding and removing one or more identities between observations is much weaker: each additional identity or category increases the likelihood of updating their description through both actions by 18.1% or 26.4%. Overall, more identities or types of identities in a user's description is associated with 1.4% or 2.3% higher likelihood of updating, respectively.

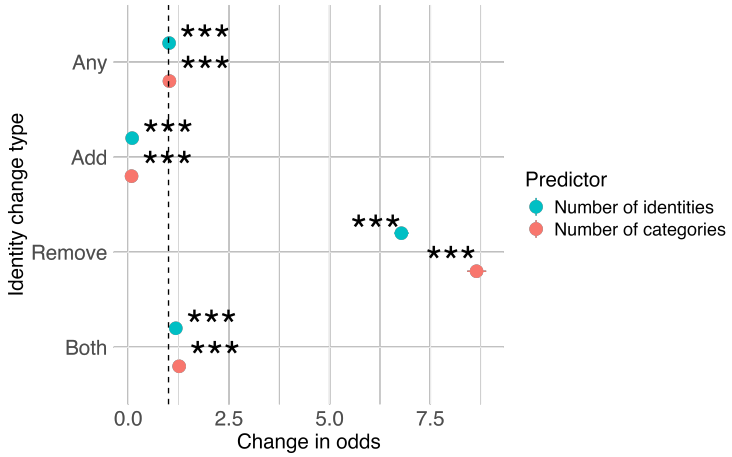

Fig. 2: Change in odds of adding (Add) , removing (Remove), both adding and removing (Both), or any updating (Any) one or more identities in a description given the number of identities (blue) or number of categories (red). Error bars represent 95% CI (though error bars are small). $^{***}p < 0.0001$

5 Discussion

This study provides empirical evidence that the identity count and complexity shape the dynamics of digital self-presentation. Analyzing over 7 million profile descriptions from X (formerly Twitter), we find that the types of identities people choose to display - and how many they include - strongly influence the likelihood and nature of subsequent updates. While some identity types are especially ephemeral (e.g., political, COVID-19), others remain stable over time (e.g., sports, age), suggesting that identity signaling is sensitive to perceived social risks and contextual volatility. Types of identities that are relatively stable may provide low cost opportunities for presenting an individualized, complex self online.

We observe that users with more numerous and diverse identities are significantly less likely to add additional identities, and far more likely to remove or revise existing ones. Our results demonstrate the asymmetrical pressures of managing identity complexity in digital self-presentation. Descriptions that contain any identity at all are less than 50% likely to add additional identities, suggesting that once identity signaling begins, users become more selective and cautious about expanding their description. In contrast, when descriptions contain four or more identities, the likelihood of removal or revision exceeds 50%, indicating a tipping point at which identity overload may provoke streamlining or editing. As the set of performed identities grows, there is more likely to be tension between their associated norms and audience expectations. In response, users may engage in self-presentation simplification, streamlining or editing their description. We hypothesize that these actions are taken in order to maintain coherence, manage

audience ambiguity, and reduce the reputational risk associated with signaling potentially conflicting roles.

This is in contrast to identity conflict mitigation strategies that involve enacting multiple, socially distinct identities concurrently [3,20,1]. There is survey evidence that people most often try to balance their conflicting identities [11]. However, integration becomes harder when the identities are publicly presented to multiple, possibly invisible audiences, where individuals are likely to experience more fear of misinterpretation or rejection by one group due to signals aimed at another [14,2,27]. These dynamics highlight the psychological and social toll of navigating audience ambiguity in digitally persistent environments.

Self-presentation updating is not merely cosmetic, it strengthens engagement, fosters emotional connection, and helps users find and belong to communities [10]. Our results underscore the need for platform design and moderation policies that offer more responsive and empowering online spaces and acknowledge the challenges of impression management under identity constraint. First, platforms could investigate the benefits to building affordances that support audience segmentation, such as customizable profile views or identity-specific visibility settings. These features allow users to present different facets of themselves to different groups. However, this could result in adverse effects like increased social fragmentation, so a measured approach is necessary. Second, enabling temporal flexibility, such as ephemeral or situational identity markers, could reduce the burden of curating a fixed identity across evolving contexts [5]. For example, Facebook has integrated the ability to set temporary profile photos[1]. Third, platform policies and moderation practices should account for the pressures of identity management, especially for users navigating multiple marginalized or politicized identities. Tools that help users experiment with identity expression in lower-stakes ways (e.g., drafts, preview features, or private modes) may encourage more authentic and adaptive self-presentation.

Our results open several directions for future research. First, we do not have access to users' motivations, perceived audience, or emotional experiences - factors that likely shape decisions to add or remove identities. Integrating survey data or interviews could enrich our understanding of the psychological processes behind public-facing identity performance evolution. Second, this work uses researcher-designated identity categories, biasing results through the cultural lens of the authors. Using the number of categories as a proxy for identity complexity relies on these researcher-designated categories and does not consider the relationships (e.g., similarity, overlap) between identities directly. Also, identities that represent discrete domains can still represent culturally coherent combinations (e.g., "democrat" and "she/her" [26]), which is not captured by our measure. Future work should develop more theoretically-grounded identity complexity measures measuring the average semantic similarity between (i.e., conceptual heterogeneity) and overlap of holders (i.e., cultural atypicality) of each pair of identities presented [20]. Third, our analysis does not incorporate the content of users' posts or engagement behavior. Prior work has shown that

[1] https://www.facebook.com/help/816899741763786

identity disclosures often co-occur with shifts in language use and network composition [4], suggesting that identity updating is part of a broader behavioral realignment. Future research could extend our findings by examining how identity constraint and updating relate to changes in content features and social connections. Finally, this study focuses on a specific platform during a time of heightened sociopolitical salience (i.e., the COVID-19 pandemic) and uses a non-random set of users (i.e., users that tweeted about COVID-19 vaccines). Extending this approach across platforms, time periods, and cultural contexts could reveal how identity constraint and signaling costs vary with platform norms and audience expectations.

In sum, this work indicates that the dominant strategy for managing the risks of presenting multiple identities to broad and potentially incompatible audiences is not integration, but simplification. We contribute to research on multiple identities and impression management by providing real-world evidence of how self-presentation is shaped by structural and social constraints. These findings offer theoretical insights into identity performance and disclosure in digitally mediated contexts and practical guidance for designing more inclusive and flexible platforms.

References

1. Amiot, C.E., De la Sablonniere, R., Terry, D.J., Smith, J.R.: Integration of social identities in the self: Toward a cognitive-developmental model. Personality and social psychology review **11**(4), 364–388 (2007)
2. Boyd, D.: Social network sites as networked publics: Affordances, dynamics, and implications. In: A networked self, pp. 47–66. Routledge (2010)
3. Burke, P.J.: Relationships among multiple identities. In: Advances in identity theory and research, pp. 195–214. Springer (2003)
4. Choi, M., Romero, D.M., Jurgens, D.: Profile update: the effects of identity disclosure on network connections and language. EPJ Data Science **13**(1), 45 (2024)
5. Choi, S., Williams, D., Kim, H.: A snap of your true self: How self-presentation and temporal affordance influence self-concept on social media. New Media & Society p. 1461444820977199 (2020)
6. Corvite, S., Zhang, B.Z., Haimson, O.L.: Social media's role during identity changes related to major life events. Proceedings of the ACM on Human-Computer Interaction **6**(CSCW2), 1–22 (2022)
7. Goffman, E.: The presentation of self in everyday life. In: Social theory re-wired, pp. 450–459. Routledge (2023)
8. Higgins, E.T.: Self-discrepancy: a theory relating self and affect. Psychological review **94**(3), 319 (1987)
9. Hirsh, J.B., Kang, S.K.: Mechanisms of identity conflict: Uncertainty, anxiety, and the behavioral inhibition system. Personality and social psychology review **20**(3), 223–244 (2016)
10. Hollenbaugh, E.E.H.E.E.: Self-presentation in social media: Review and research opportunities. Review of communication research **9** (2021)
11. Jones, J.M., Hynie, M.: Similarly torn, differentially shorn? the experience and management of conflict between multiple roles, relationships, and social categories. Frontiers in psychology **8**, 1732 (2017)

12. Kulich, C., De Lemus, S., Kosakowska-Berezecka, N., Lorenzi-Cioldi, F.: Multiple identities management: Effects on (of) identification, attitudes, behavior and well-being (2017)
13. Leary, M.R., Kowalski, R.M.: Impression management: A literature review and two-component model. Psychological bulletin **107**(1), 34 (1990)
14. Marwick, A.E., Boyd, D.: I tweet honestly, i tweet passionately: Twitter users, context collapse, and the imagined audience. New media & society **13**(1), 114–133 (2011)
15. McHardy, J., Ratcliffe, A.: Identity conflict: A framework and empirical investigation. Sheffield Economics Research Papers (SERPS) **2019019**(2019019) (2019)
16. Mead, G.H.: Mind, self, and society from the standpoint of a social behaviorist. (1934)
17. Owens, T.J., Robinson, D.T., Smith-Lovin, L.: Three faces of identity. Annual Review of Sociology **36**(1), 477–499 (2010)
18. Pathak, A., Madani, N., Joseph, K.: A method to analyze multiple social identities in twitter bios. Proceedings of the ACM on human-computer interaction **5**(CSCW2), 1–35 (2021)
19. Priante, A., Hiemstra, D., Van Den Broek, T., Saeed, A., Ehrenhard, M., Need, A.: # whoami in 160 characters? classifying social identities based on twitter profile descriptions. In: Proceedings of the first workshop on NLP and computational social science. pp. 55–65 (2016)
20. Roccas, S., Brewer, M.B.: Social identity complexity. Personality and social psychology review **6**(2), 88–106 (2002)
21. Settles, I.H.: When multiple identities interfere: The role of identity centrality. Personality and Social Psychology Bulletin **30**(4), 487–500 (2004)
22. Snyder, M.: Self-monitoring of expressive behavior. Journal of personality and social psychology **30**(4), 526 (1974)
23. Strimbu, N., O'Connell, M.: The relationship between self-concept and online self-presentation in adults. Cyberpsychology, Behavior, and Social Networking **22**(12), 804–807 (2019)
24. Tajfel, H., Turner, J.C., Austin, W.G., Worchel, S.: An integrative theory of intergroup conflict. Organizational identity: A reader **56**(65), 9780203505984–16 (1979)
25. Thoits, P.A.: Multiple identities and psychological well-being: A reformulation and test of the social isolation hypothesis. American sociological review pp. 174–187 (1983)
26. Tucker, L., Jones, J.: Pronoun lists in profile bios display increased prevalence, systematic co-presence with other keywords and network tie clustering among us twitter users 2015-2022. Journal of Quantitative Description: Digital Media **3** (2023)
27. Vitak, J.: The impact of context collapse and privacy on social network site disclosures. Journal of broadcasting & electronic media **56**(4), 451–470 (2012)
28. White, H.C.: Identity and control. Princeton NJ **3** (1992)
29. Wicklund, R.A., Gollwitzer, P.M.: Symbolic self completion. Routledge (2013)
30. Yoder, M.M., Shen, Q., Wang, Y., Coda, A., Jang, Y., Song, Y., Thadani, K., Rosé, C.P.: Phans, stans and cishets: self-presentation effects on content propagation in tumblr. In: Proceedings of the 12th ACM Conference on Web Science. pp. 39–48 (2020)
31. Zhong, C., Chang, H.w., Karamshuk, D., Lee, D., Sastry, N.: Wearing many (social) hats: How different are your different social network personae? In: Proceedings of the International AAAI Conference on Web and Social Media. vol. 11, pp. 397–406 (2017)