

Someone Like Me: How Does Peer Parity Influence Participation of Women on Stack Overflow?

Denae Ford
NC State University
Raleigh, NC, USA
dford3@ncsu.edu

Alisse Harkins
NC State University
Raleigh, NC, USA
amharkin@ncsu.edu

Chris Parnin
NC State University
Raleigh, NC, USA
cjparnin@ncsu.edu

Abstract—Stack Overflow is a learning community for software developers to share and solve programming problems with each other. However, women are often deterred from contributing questions or answers. Research external to programming communities suggest the presence of peers can increase activity from underrepresented users in unfamiliar spaces. To investigate the concept of peer parity, we studied how women participate on Stack Overflow and if the presence of more women on a thread enhanced their activity. We found that women who encountered other women were more likely to engage sooner than those who did not. We discuss how these findings can support women in programming communities through peer mentorship and increase engagement.

Keywords—Stack Overflow, online Q&A, gender, peer parity

I. INTRODUCTION

Exchanging knowledge on the internet can be *intense* [1]; especially when online users do not notice anyone like themselves posting [2]. Stack Overflow, a popular programming Q&A web site, reported in a recent survey that only 19.2% of respondents have ever provided a new answer to a question while less than 8% identified as women [3]. Based on these results, there may be less than 1% of women who are answering programming questions online. Researchers found women to be disproportionately underrepresented in the community [4].

Several challenges exist that can explain low participation rates. In a recent study on Stack Overflow participation of women, subjects mentioned that one reason they not post on Stack Overflow is that “*They are just not even on the same race track [5].*” This notion of observing people on the same “race track” or having similar individuals to compare oneself to is known as *peer parity*. Other challenges include asking a question as users may be uncomfortable revealing that they do not know something. Similarly, asking can give the perception that a person may seem pretentious as they display their knowledge on a topic that others lack [6]. How users phrase a question matters—emotional tones of technical questions can influence how they get answered [7]. Often times questions lacking the affect expected from community users may not get answered. These are all factors that may cause the user who asked a question to feel they do not belong; discouraging them from participating at all and encouraging lurker behavior [8].

We define *peer parity* as:

When an individual can identify with at least one other peer when interacting in a community.

Peer parity can exist across and within races, genders, experiences, career positions and more. In this work, we study peer parity among identifiable, perceived genders. Based on studies of identity and peer interactions [2], [8], we believe that differences in exposure of peer parity may influence participation in online communities such as Stack Overflow.

We offer the stories of two users with different experiences to give an example of how peer parity may influence participation in the context of Stack Overflow. Olivia is providing an answer to a question she has experienced herself in the past. After looking through the listed answers to a Stack Overflow question she found, Olivia thought the solution she derived worked better. Before she posted an answer she decided to double check other answers that are shown. After she scanned the names and determined that many of the answers are from women, she feels encouraged to see users she can identify with. Olivia then proceeds to post her answer.

Another user Mellie, had a similar scenario, although it did not end as successful as Olivia’s. Mellie also identified a question where she had an alternate answer and expressed interest in posting. Mellie scanned the names of the other user answers and determined many of them may be men. She could not identify any other women on the post, therefore Mellie decided not to post her answer. We want to understand if scenarios such as those Olivia and Mellie encounter effect their interest in posting on Stack Overflow.

To understand if being on the same “race track” as other women can make a difference, we studied the notion of peer parity and how it may effect participation on Stack Overflow. We extracted both post and user data from Stack Overflow. After analyzing the genders of the participants on a question using their display name, just as Olivia and Mellie did, we identified posts where many women are on a thread and compare this to instances where there is only one woman. We then analyzed questions asked by women who were exposed to parity and non-parity and found a significant difference in engagement. To make our findings actionable, we provide implications for researchers and online communities.

A few findings of this work are:

- 1) Of the women we identified, only 32% have ever posted

a question.

- 2) The default user experience for women is low access to peers—most threads only have one woman.
- 3) Women who experience peer parity were more likely to engage sooner.

II. BACKGROUND

This work is guided by research related to posting online as apart of an identity, finding others that share that identity, and how both may encourage posting on programming questions online.

A. Stack Overflow Participation

Stack Overflow has a participation problem and researchers notice. Slag et al. found that the majority of Stack Overflow users make one contribution and stop using the site [9]. The density of poor quality and unguided questions has increased the number of deleted questions over the past two years. Xia et al. noticed the increase in deleted questions and determined an algorithm to predict this occurrence and get ahead of the issue [10]. However, deleting questions does not necessarily solve the participation problem, it actually dismisses it. In our work, we chose to study this participation problem through women and determine whether women asking questions among peers will effect their activity moving forward.

Women participation on Stack Overflow is valuable, yet barely exist according to community builders and researchers. In addition, the site's lack of diverse participation has encouraged much interest in characterizing it with many approaches [11], [12], [3]. All of these studies reach a similar conclusion: women are discouraged to participate on Stack Overflow. In an effort to explain why, researchers have shown that men represent the vast majority of contributors on Stack Overflow [4] and identify barriers as to why Stack Overflow may be so hard for women to contribute [5]. Taking the aforementioned into account, we chose to investigate where the few women in the community are and how and if they are supported by the interactions of each other.

B. Nature of Sharing Identity Online

Sharing an identity online has advantages for different communities. Some people share their real name online as a way to gain social capital in virtual communities [13]. Sarma et al. demonstrated the utility for programmers to share their profiles online in order to increase their visibility and showcase their talents for future careers [14]. In addition, Archdivelli et al. found many cultural differences in sharing identity online across many countries within the same context [15]. Building on this, we study identity through the cultural context of gender and how it may influence programmer contributions online.

Posting questions online can mean users must make themselves vulnerable and admit their knowledge deficit in a public sphere. This can be discouraging to many because they may want to remain anonymous, may be shy about posting, or decide that it may be the wrong group for them to participate

in [6]. However, what can encourage users to post online and use their identity is knowing that they are not the only person being vulnerable [16]. In this work we hypothesize that some women may seek a community of peers that they can identify with.

C. Diverse Peer Influence

Ichinco et al. identified that community members becoming leaders can influence who users identify with in a community [17]. Identifiable leaders can also be helpful as peers now have a personal example of how to approach a similar goal as the leader. One example of this encouragement is through higher education among underrepresented groups. Gershenson et al. found that when Black students who received access to teachers that resembled them, those students were more likely to go to college and seek similar career paths as their teachers [8]. In addition to race, we see similar effects in gender spheres. When people can see women in nontraditional positions, more than just women are likely to aspire for those same roles [18], [19]. In summary, access to a diverse range of examples and role models makes a difference.

Identifying peers can heavily influence how people respond in programming communities, especially among women. Teams consisting of more women are found to be more successful and productive [20]. Intergroup relation theorists have identified that there is a strong cognitive preference for women among women in terms of identity [21]. Ford outlined an application of the Bechdel Test to determine how presence can effect a women on Stack Overflow [22]. In our work, we demonstrate that an in-group preference among women may exist in online programming communities and how recognizing this exist.

III. METHODOLOGY

Our work is guided by the following research question:

How does peer parity influence how women post on Stack Overflow?

Prior work demonstrates individuals who can identify with other members of community will increase interest and engagement while those who do not identify with the community lose interest [2]. On Stack Overflow it is challenging to find identities that may resemble a minority; users have mentioned difficulty finding other women [23]. We hypothesize that peer parity would have an influence on how women posted on Stack Overflow and that women on the parity posts have more activity after participating on a parity thread. To determine if there is a difference among the community, we compare differences in posting activity between women who experience peer parity and those who did not.

A. Data Extracted

We extracted 5,987,284 users and 32,209,817 posts from the Stack Overflow Data Exchange. The oldest post of the data we studied is July 31, 2008 and the latest is September 4, 2016.

B. Identifiable Genders

After we gathered all threads we identified genders of the users based on their display name with Vasilescu et al. genderComputer Tool [12]. The reported precision of genderComputer is about 90%. We modified the tool to review the first name of a user and determine where a variation of that name exist in a list of names with a gender across any country. This results in our tool having higher precision in determining genders. The tool reports the gender of the user as male, female, unisex, or undetermined. In our work, we report females as women and males as men. With our modified tool, we computed the gender of 5,987,284 users and identified 363,133 women; 2,139,305 men; 102,189 unisex; and 3,382,657 undetermined names. More details on the modified gender computing tool are available online [24].

C. Peer Parity Defined

In this work, we determined that peer parity exist if there is more than one distinct woman on a thread. We refer to this as *parity*. Otherwise, we describe threads that only have one distinct woman as *non-parity*. To clarify, we do not specifically isolate posts with only women. From our extracted data we were only able to identify 32% of all identifiable women who have ever posted a question.

IV. FINDINGS ON STACK OVERFLOW PEER PARITY

To investigate our research question, we first randomly selected 1000 women who have ever posted more than one question. Second, we gathered their first question and their second posted activity. Third, we identified the time difference between activity. We selected women who have asked questions to control for a shared first experience on the site. We calculated the gender of all users on a thread and identified whether their first question was on a parity or non-parity thread. We used the time difference between their first question and second activity as a comparison of how soon they re-engage in the community. We also identified the reputation points and number of badges for each of the women. We used the number of reputation points, which are a measure for how much the community trust users, as a measurement for frequency of activity [25]. The number of badges is one way Stack Overflow demonstrates positive user activity [26]. Both the number of badges and reputation points are also visible next to a user's name when posting a question or answer.

Of our 1000 randomly selected women, we identified 452 parity and 548 non-parity threads from their first question. We found a significant difference in type of second activity after participating on a parity or non-parity thread ($p = 2.799e-06$, $\alpha = .05$), which was either posting a question ($N = 833$) or posting an answer ($N = 167$). We found a significant difference in the time between posts for women who asked a question on parity threads in comparison to non-parity threads ($p = 1.83e-05$, $\alpha = .05$). The cumulative time differences by posts are demonstrated in Figure 1. This figure demonstrates that the longest time difference for a parity activity was 1017 days and 1347 days for a non-parity activity. We did not find a

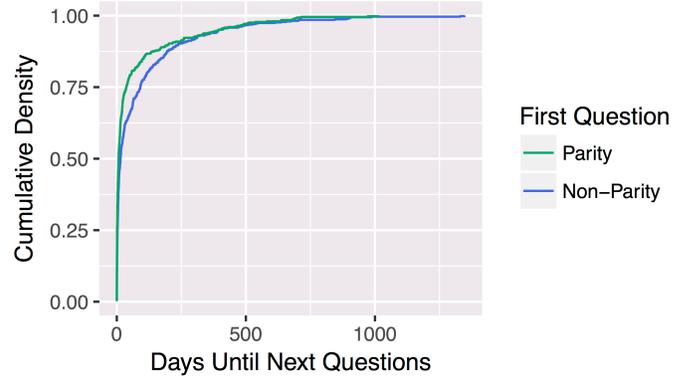


Fig. 1. The time between activity for parity and non-parity threads. The x axis indicates a time difference in days until their next activity after either posting a question on a parity or non-parity thread. The y axis indicates the cumulative frequency of within that time frame. This graph demonstrates that women participating on parity thread have post more immediately after participating on non-parity thread.

significant difference in reputation points or number of badges. We observed a small effect size among the comparison of time differences ($d = .1$), reputation ($d = .1$), and number of badges ($d = .2$).

These statistical differences are supported by all second activity from parity threads being completed before non-parity threads as shown in Figure 1. These results demonstrate that women who are on parity threads engage sooner in Stack Overflow participation activities.

Finding: Women who asked questions on parity posts re-engaged sooner. However, they do not have higher reputation.

V. DISCUSSION

From our findings we determined factors that can be implied from this work through the explicit use of identity and how that can fuel a mentorship program through peer parity. We also provide suggestions on how we can encourage a user's growth and redefine what the best mentors should do based on peer parity [27].

Showcasing success: According to the developer survey, the number of women on Stack Overflow has increased from previous years and we also know that Stack Overflow has been quite vocal in fostering a more inclusive community [3]. The very idea of being transparent about a community problem may have played a factor in the increased interest in the site. Perhaps one way to inspire and increase others to participate is to showcase top-rated questions asked by women. In addition, a series could be launched to showcase these top-rated users where new users can learn how top-rated users have come to be recognized for their contributions to the community. This will not only demonstrate how to post successful questions of Stack Overflow, but also shows the diverse set of users contributing.

Paired guidance: Asking for help can be hard; but it can be easier if a peer is available to help. In our work, we noticed that

women in the presences of peers reengaged sooner. If a user is about to post their first question online, it may be a good opportunity to approach them in the act and let them know there is a group of users willing to offer guidance. As we found that peer parity exist, we recommend mentor-mentee pairing based on peer identification to further expand engagement. Mentorship is a bidirectional relationship—both parties have something to gain. Encouraging users to seek guidance can benefit both the mentor, providing guidance, and the mentee, seeking guidance. Mentors in this community can decide on the varying degrees of their engagement based on their popular use of similar tags that a mentee may add to their question. For example, mentors can be hands on and help users write and edit questions incorporating that tag. In contrast, a mentor can be hands off and offer advice on how to gain the most utility from community users who are answering similarly tagged questions.

Revealing user identity: Our work shows that users may be interested in seeking an identity they can relate to and feel comfortable with online. There has been much objection to Stack Overflow’s stance to not explicitly requesting identifying information such as gender. However, our work builds on the idea that there are features of user identities that users are interested in sharing, but not given the opportunity to. When there is not a platform to support the multiple layers of identity that users have, an organization is saying they are *not* interested in hearing about that identity of a user. This inadvertently demonstrates organizations taking a stance on the topic and saying it does not matter. However, it can be difficult for users to separate from their identity in public spaces, furthermore, they should not have to. The same layer of identity that Q&A resources omit, may be the same feature of identity that effects their lack of diverse community engagement. We should embrace and support users who wish to disclose this information. Similarly to other identity features users can enter in their account profile, such as a birth date and full name, it is should be a common practice to provide the option for users to enter their gender. Allowing users the opportunity to bring their whole self into a community where they seek help may just be the encouragement they need to be active contributors.

VI. THREATS TO VALIDITY

Construct Validity. There are few women on Stack Overflow according to the developer survey and previous research declared it hard to determine the gender of individuals. We cannot say that we captured all the profiles of women on Stack Overflow. Further analysis will have to be done to confirm identities.

Internal Validity. We isolate users in our experiment. This may cause us to focus on threads where women are very active. We may have a more representative experiment had we confirmed the scores of the threads and the number of users on each thread.

External Validity. We may have issues adapting to other programming communities. Stack Overflow has many features

of activity on a thread including visible reputation of users, questions scores, and answer scores that other programming help communities do not offer on the same page.

VII. CONCLUSION

Stack Overflow is a resource that programmers use to fix their programming issues and learn how others work through similar issues. However, there is an difference with who can contribute to those programming solutions, specifically women. To analyze this, we introduce a concept called *peer parity*. We adapt the concept of finding “someone like me” to how women may identify with other women on Stack Overflow. Using first names as identifiers as a gender, we define parity as instances where there are many distinct women on a thread and non-parity as threads that have only one distinct woman. We find that although there are less women participating on parity threads, the women on parity threads reengage sooner in the community. This finding presents a gateway to future mentorship programs. We discuss interpretations of these findings and describe interventions to understand how being among like minded peers can increase engagement in online programming communities. After all, you cannot be what you cannot see.

ACKNOWLEDGMENTS

We thank the alt-code research group and reviewers for their feedback on an earlier version of this work. This material is based upon work supported by the National Science Foundation under Grants No. DGE-1252376 and 1559593.

REFERENCES

- [1] C. Treude, O. Barzilay, and M.-A. Storey, “How do programmers ask and answer questions on the web?: Nier track,” in *Software Engineering (ICSE), 2011 33rd International Conference on*. IEEE, 2011, pp. 804–807.
- [2] M. M. Subramaniam, J. Ahn, K. R. Fleischmann, and A. Druin, “Reimagining the role of school libraries in stem education: Creating hybrid spaces for exploration,” *The Library Quarterly*, vol. 82, no. 2, pp. 161–182, 2012.
- [3] 2017 stack overflow developer survey. [Online]. Available: <https://stackoverflow.com/insights/survey/2017>
- [4] B. Vasilescu, A. Capiluppi, and A. Serebrenik, “Gender, representation and online participation: A quantitative study of stackoverflow,” in *International Conference on Solcial Informatics*. IEEE, 2012, pp. 332–338.
- [5] D. Ford, J. Smith, P. Guo, and C. Parnin, “Paradise unplugged: Identifying barriers for female participation on stack overflow,” *Proceedings of the 24th International Symposium on the Foundations of Software Engineering*, 2016.
- [6] B. Nonnecke, D. Andrews, and J. Preece, “Non-public and public online community participation: Needs, attitudes and behavior,” *Electronic Commerce Research*, vol. 6, no. 1, pp. 7–20, 2006.
- [7] N. Novielli, F. Calefato, and F. Lanubile, “Towards discovering the role of emotions in stack overflow,” in *Proceedings of the 6th International Workshop on Social Software Engineering*, ser. SSE 2014. New York, NY, USA: ACM, 2014, pp. 33–36.
- [8] S. Gershenson, C. Hart, C. Lindsay, and N. W. Papageorge, “The long-run impacts of same-race teachers,” *IZA Institute of Labor Economics*, 2017.
- [9] R. Slag, M. de Waard, and A. Bacchelli, “One-day flies on stackoverflow: Why the vast majority of stackoverflow users only posts once,” in *2015 12th Working Conference on Mining Software Repositories*. IEEE, 2015, pp. 458–461.

- [10] X. Xia, D. Lo, D. Correa, A. Sureka, and E. Shihab, "It takes two to tango: Deleted stack overflow question prediction with text and meta features," in *Computer Software and Applications Conference (COMPSAC), 2016 IEEE 40th Annual*, vol. 1. IEEE, 2016, pp. 73–82.
- [11] B. Lin and A. Serebrenik, "Recognizing gender of stack overflow users," in *Proceedings of the 13th International Conference on Mining Software Repositories*. ACM, 2016, pp. 425–429.
- [12] B. Vasilescu, A. Capiluppi, and A. Serebrenik, "Gender, representation and online participation: A quantitative study," *Interacting with Computers*, p. iwt047, 2013.
- [13] C. A. Steinkuehler and D. Williams, "Where everybody knows your (screen) name: Online games as 'third places,'" *Journal of Computer-Mediated Communication*, vol. 11, no. 4, pp. 885–909, 2006.
- [14] A. Sarma, X. Chen, S. Kuttal, L. Dabbish, and Z. Wang, "Hiring in the global stage: Profiles of online contributions," in *Global Software Engineering (ICGSE), 2016 IEEE 11th International Conference on*. IEEE, 2016, pp. 1–10.
- [15] A. Ardchvili, M. Maurer, W. Li, T. Wentling, and R. Stuedemann, "Cultural influences on knowledge sharing through online communities of practice," *Journal of knowledge management*, vol. 10, no. 1, pp. 94–107, 2006.
- [16] D. Snow, "Collective identity and expressive forms," *Center for the Study of Democracy*, 2001.
- [17] M. Ichinco and C. Kelleher, "Online community members as mentors for novice programmers position statement," in *Blocks and Beyond Workshop (Blocks and Beyond), 2015 IEEE*. IEEE, 2015, pp. 105–107.
- [18] D. E. Campbell and C. Wolbrecht, "See jane run: Women politicians as role models for adolescents," *Journal of Politics*, vol. 68, no. 2, pp. 233–247, 2006.
- [19] M. Fraile and R. Gomez, "Why does alejandro know more about politics than catalina? explaining the latin american gender gap in political knowledge," *British Journal of Political Science*, vol. 47, no. 1, pp. 91–112, 2017.
- [20] B. Vasilescu, D. Posnett, B. Ray, M. G. J. van den Brand, A. Serebrenik, P. Devanbu, and V. Filkov, "Gender and tenure diversity in GitHub teams," in *CHI Conference on Human Factors in Computing Systems*, ser. CHI. ACM, 2015.
- [21] L. A. Rudman and S. A. Goodwin, "Gender differences in automatic in-group bias: Why do women like women more than men like men?" *Journal of personality and social psychology*, vol. 87, no. 4, p. 494, 2004.
- [22] D. Ford, "Recognizing gender differences in stack overflow usage: Applying the bechdel test," in *IEEE Symposium on Visual Languages and Human Centric Computing*. IEEE, 2016, pp. 264–265.
- [23] "Who are the best women on stack overflow?" Feb 2017. [Online]. Available: <https://www.quora.com/Who-are-the-best-women-on-Stack-Overflow>
- [24] "Simple gender computer." [Online]. Available: <https://github.com/alt-code/Research>
- [25] "What is reputation? how do i earn (and lose) it?" [Online]. Available: <https://stackoverflow.com/help/whats-reputation>
- [26] "What are badges?" [Online]. Available: <https://stackoverflow.com/help/what-are-badges>
- [27] A. Tjan, "What the best mentors do," Feb 2017. [Online]. Available: <https://hbr.org/2017/02/what-the-best-mentors-do>