

# Quantifying behavior-based gender discrimination on collaborative platforms

Orsolya Vásárhelyi<sup>1a,b,c,\*</sup> and Balázs Vedres<sup>c,d,e</sup>

<sup>a</sup>Center for Collective Learning, Corvinus Institute for Advanced Studies, Corvinus University of Budapest, Fővám tér 8., Budapest 1093, Hungary

<sup>b</sup>Institute of Data Analytics and Information Systems, Corvinus University of Budapest, Fővám tér 8., Budapest 1093, Hungary

<sup>c</sup>Democracy Institute, Central European University, Nádor utca 13., Budapest 1051, Hungary

<sup>d</sup>Department of Network and Data Science, Central European University, Quellenstraße 51, Vienna 1100, Austria

<sup>e</sup>Oxford Internet Institute, University of Oxford, Q1 St Giles, Oxford OX1 3JS, United Kingdom

\*To whom correspondence should be addressed: Email: [orsolya.vasarhelyi@uni-corvinus.hu](mailto:orsolya.vasarhelyi@uni-corvinus.hu)

Edited By Katherine Ognyanova

## Abstract

Digital collaborative platforms have become crucial venues of career advancement and individual success in many creative fields, from engineering to the arts. Gender discrimination related to behavioral choices of users is a key component to gendered disadvantage on platforms. Such platforms carried the promise of opening avenues of advancement to previously discriminated groups, such as women, as platforms lack managerial gatekeepers with conventional prejudice. We analyzed the extent of behavior-based gender discrimination on two digital platforms, GitHub and Behance, focused on software development and fine arts and design. We found that the main cause of women's disadvantage in attention, success, and survival is largely due to the gender typicality of their behavior that varies between 60 and 90% of the total disadvantage of women. Men and women are penalized if they follow highly female-like behavior, while categorical gender is no longer significant. As platforms employ algorithmic tools and AI systems to manage users' activity and visibility, and recommend new projects to collaborate, stereotypes associated with behavior can have long-lasting consequences.

**Keywords:** gender discrimination, behavior-based discrimination, platforms, GitHub, Behance

## Significance Statement

This study quantifies behavior-based gender discrimination on users' success, visibility, and survival on digital collaborative platforms. Although direct discrimination is not significant, behavior-based discrimination, present in career choices and online activity, significantly disadvantages women and men on both GitHub, a male-dominated platform, and Behance, where women are more prevalent. Behavior-based gender discrimination accounts for 60–90% of the disparity in attention, success, and survival rates between genders. The visibility paradox magnifies these issues: women receive more attention but are not recognized as experts, leading to higher harassment and dropout rates. As algorithmic tools and AI systems manage platforms, they risk consolidating and amplifying biases. Our findings underscore the urgent need for monitoring mechanisms to ensure equitable opportunities for all users.

## Introduction

Platform organizations offer digital affordances to connect producers and consumers, and this organizational form had seen a rapid uptake over the past decade. Today, the world's most valuable businesses (Apple, Microsoft, Amazon, Google, or Facebook) are platforms (1, 2), and a recent report estimated that global digital platforms in 2022 had 371 billion average monthly users (3). Platformization does not appear to slow down, as the annualized growth rate of digital trades (20%) is faster than that of physical products (6%) (4). This acceleration has resulted in an entire new ecosystem (5), which has changed the way we communicate (5), shop (6), travel (7–9), define success (5), work (2, 10), and collaborate (11–13).

Digital collaborative platforms have become crucial tools for independent creative workers, providing opportunities to develop skills, connect with other like-minded people and potential collaborators, and help capture the attention of potential users and buyers (14, 15). GitHub activity, the largest developer platform with more than 100 million users, has been shown to facilitate entry into the labor market while making it more challenging for developers who lack open source competence or status to secure jobs (16). Other platforms, such as Behance, play a similar role for graphic artists and designers, allowing them to build open access portfolios of their visual work, find peers for inspiration, and ultimately capture the attention of buyers and clients (17, 18).

**Competing Interest:** The authors declare no competing interests.

**Received:** June 10, 2024. **Accepted:** January 3, 2025

© The Author(s) 2025. Published by Oxford University Press on behalf of National Academy of Sciences. This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial License (<https://creativecommons.org/licenses/by-nc/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited. For commercial re-use, please contact [reprints@oup.com](mailto:reprints@oup.com) for reprints and translation rights for reprints. All other permissions can be obtained through our RightsLink service via the Permissions link on the article page on our site—for further information please contact [journals.permissions@oup.com](mailto:journals.permissions@oup.com).

The motivations to contribute to open source software (OSS), projects distributed freely with their source code for others to use, have been connected to learning, intellectual stimulus, social interactions, and altruism (19–21). However, the younger generation of OSS developers also uses their activity to promote their careers (22). A verifiable track record with visible skills and earned reputation within the OSS community can increase the probability of being hired, resulting in higher salaries, greater job security, and a competitive advantage to be promoted by demonstrating leadership capabilities (16). Therefore, constant activity in OSS (survival), high attention (large follower base) earned through collaborations, and the number of projects with high reputation (many likes) can be linked to career advancement.

The role of Behance within the design community is primarily self-promotion by allowing users to create their own narrative, displaying their awards, achievements, and client portfolio (18, 23). The platform also promotes a “collective belonging” to the wider design community through social networking features. Users can create “watchlists” of artists that they can follow for inspiration and can comment and like projects. Portfolio building can be integrated into the design process, by making the release of revisions seamless through Adobe integration. Behance users agree that constant activity (survival), popular projects (many appreciations), and a large follower base (high attention) can help to find new customers (24, 25).

The emergence of such portfolio careers (26) carried the promise of extending opportunities to previously disadvantaged groups, among them women (27–29). However, this new reputation-based economy translates social capital and risk taking behavior into value which appears to benefit men more than women (30–33). Platforms often lack features that take gender inequalities into account, and, furthermore, the deepest aspects of their culture, technology, design, and algorithmic management tend to perpetuate gender discrimination (34–39).

Several recent publications set out to chart the gender gap in the digital economy with respect to participation and success (40–44). A key critical point raised about these works is that gender itself is already encoded in the creation of technology (including the algorithms that govern digital platforms), and thus platforms would not be able to mitigate gender inequalities (45, 46). Gender as a social concept assumes that men and women follow their category-specific scripts: norms and behaviors that reinforce societal expectations (47). Role congruity and stereotype fit hypothesis suggest that women and men who specialize in fields which seen as a better “fit” can expect more positive evaluation (48). Consequently, specializing in nongender-typical professions can penalize both genders. Previous research has shown that if women specialize in technical fields where they are stereotypically seen as a better “fit,” and that their presence is more accepted, they are more successful (49).

Particular activities can result in behavior-based discrimination once there is a repeated pattern of a given activity having unequal gender representation and unequal success. Seeing only one activity being associated with a gender—say user interface design being a typical female activity—and at the same time such an activity enjoying lower attention and success is not in itself a sign of discrimination. User interface design might be by chance a less appreciated activity with low potential for success. However, once such activities can be observed as a repeated pattern, we can talk about discrimination: seeing women typically pursue activities that result in lower success. A real-world example of such behavior-based discrimination is a lawsuit against Google, in which women claimed that the company arbitrarily

assigned them to activities (primarily because of their gender and regardless of their education and experience) that had a significantly lower salary bracket.<sup>a</sup> The company paid 118 million dollars to redress past harms resulting from behavior-based discrimination. Our attention to gender typicality intends to capture exactly this repeated pattern across a variety of activities that constitute behavior-based discrimination.

Femininity and masculinity are often in a hierarchical relationship: What is considered feminine tends to be devalued, while masculine behavior is rewarded. Therefore, women can achieve greater success in male-dominated fields when practicing “masculinity” (50). Men who work in stereotypically feminine jobs create strategies to distance themselves from doing “femininity” at work (51). For example, male nurses emphasize the physical parts of their job, such as moving patients between beds, instead of the caring elements.

When masculinity receives higher rewards and leads to greater success in a culture, it eventually leads to inequality. In the field of technology, it is the superiority of men and masculine culture that is considered natural and is therefore beneficial. Direct gender discrimination occurs when decisions and processes are based on an individual’s categorical gender identity, resulting in disadvantages by category. Gender discrimination can also be behavior-based, based on nonsensitive attributes—like activity patterns or specializations—, that are closely linked to gender. There is evidence of direct (37) and behavior-based discrimination against women on platforms (36, 38, 52).

Previous studies found that one’s gender can be predicted fairly accurately based on their collaboration patterns, specialization, and the style of code they produce (38, 52). In a recent study, the gender of users was predicted on Pinterest, where the ratio of women is higher than that of men, based on the content they share (53). In a study exploring the digital music platform The Echo Nest (54), where only 25% of the solo artists are women, authors managed to predict the artist’s gender based on the musical features of their songs with 90% accuracy. These studies did not link gender typicality to success, only aimed to find behavioral characteristics that predict one’s gender.

If users’ activity on digital platforms differ markedly by gender, does this also result in behavior-based gender discrimination? A prior study of GitHub (38) operationalized “femaleness” as the extent of feminine behavior by predicting a user’s inferred gender based on their online behavior. Their results showed that men and women are both penalized if they follow highly female-like behavior, indicating the presence of behavior-based gender discrimination. Wachs et al. (39) could explain gender differences in the popularity and visibility of design projects by the gender typicality of designers’ skills and used visual elements in the Dribbble designer community. Specifically they quantified the “genderedness” of skills by how likely a skill is listed by a woman or a man on their profiles. They also trained a neural network to predict whether a man or a woman generated a shot based on the visual elements used. Both variables had significant relationships with the outcome metrics, indicating that gendered patterns of specialization impact the popularity of users on Dribbble. May et al. (36) found that the gap between men’s and women’s reputation on Stack Overflow, the largest technical Q&A community, is due to the gender typicality of user activity on the site, and differences in how these activities are rewarded by the platform.

This article contributes to quantifying behavior-based gender discrimination in the platform economy (36, 38, 39) in three major ways. First, we provide a multiplatform analysis, as we replicate prior work (38) on Behance, a platform that is on the opposite

end of the creative spectrum in terms of content to GitHub. Since design and technology are similar in having a predominantly masculine culture leading to unequal representation of women within different subspecializations (55), we expect that behavior-based gender discrimination is prevalent on Behance, also. Despite significant differences in the content, all key aspects of programmers' activity on GitHub could be transferred to the context of graphic artists on Behance—likely due to isomorphism in platform design. This enables us to make more general statements about gender inequality on platforms that a single-platform study would not warrant.

Second, we operationalize previously used metric “female-ness” to better capture behavior-based discrimination. Previous studies (38, 39, 54) showed that gendered tie formation and gender homophily are related to success and can improve models to predict one's gender. However, gender differences in the number of men and women with whom one collaborates can be the result of direct gender discrimination. Women might collaborate less with men because men simply do not find them “worthy” to work with (56). To avoid mixing the effects of direct and behavior-based gender discrimination, we have removed the gendered aspects of collaborations from the prediction model that captures “femaleness.”

Finally, third, we add an outcome measure, attention, to augment two measures used before, popularity, and survival. Attention precedes popularity and survival as an initial form of success, in the sense that platform users first need to be noticed before they can succeed along any other dimensions (39). Studies of gender inequalities point to inequalities in being noticed in the workplace as a key dimension of women's disadvantage. Women in highly masculine professions face a paradoxical visibility problem: they attract considerable attention as women, but this does not translate into their acceptance as experts (57, 58). In other words, we expect no direct discrimination in attention (as women often attract attention, especially in fields where they are underrepresented), while it is an open question whether there is discrimination by gender-typical behavior in gaining attention.

In sum, we found that both men and women are penalized if they follow highly female-like behavior: Behavior-based discrimination is significantly associated with gender disparities on digital platforms. In other words, women are penalized more for what they do and not for who they are. This pattern holds for attention, success, and survival, and it is true for both GitHub and Behance. Our findings should be especially alarming, since the public debate is forming around the responsibility of platforms in artificially decreasing the visibility of underrepresented groups (59, 60), resulting in a 30% racial pay gap between influencers (61). Behavior-based discrimination presents a grave risk of deeply rooted, invisible, and stubborn inequality, as it can be baked into the algorithms and culture of online collaborations, greatly magnifying already existing gender inequalities.

## Methods

### Data

GitHub ([www.github.com](http://www.github.com)) is by far the most popular collaborative platform for software projects. It offers online hosting and version control services that allow developers to contribute to software projects from around the world. According to Octoverse, the annual statistical report of GitHub, the platform had more than 100 million user accounts in 2023, regardless of their activity

status (62). Since GitHub provides benefits beyond the recording of contributions and the management of the source code—such as traditional social media functionalities (e.g. following)—, it became the subject of several studies, aiming to understand collaborative activity online (63, 64), team success, diversity (42, 65), and gender inequality in technology (66).

In this article, we use a dataset obtained from [githubarchive.org](https://githubarchive.org), containing individual careers between 2009 February 19 and 2016 October 21. This dataset contains the following information for each individual: the creation of a repository, push to a repository (updating the codebase), opening, closing, and merging pull requests (contributing to others' projects); accompanied by user information using the GitHub API (user names, email addresses, number of followers, number of public repositories, and date they joined GitHub; see Table 2). We collect these datapoints for all users throughout their activity history and generate variables by summing up their total activity by activity type.

Behance ([www.behance.net](http://www.behance.net)) is a digital platform for creative professionals, where they can feature a portfolio of their work, collect, and organize the works of others for inspiration, and become hired as freelancers. Similarly to GitHub, Behance allows users to create relationships via social network features (following, commenting, and appreciating) and share their work in a wide variety of domains, such as photography, graphic design, and user experience (UX) research. Behance is a considerably smaller platform than GitHub, with about 50 million users (according to Behance.net).

The original data source of our study is a randomized sample of the Behance database obtained by Kim (67). Following a commonly used method to sample large graphs randomly (68, 69), they initiated a random walk-based sampling procedure which selected active users as seed sets from public timelines showcasing recent projects in 2016. The process involved picking a user at random from this set and continuing the walk until reaching a target of 50,000 users. The final data contained information on the gender of 37,777 users, specialized topics, number of followers, number of users followed, number of appreciations (likes), number of comments, project views, and stylistic information of the projects, and the total number of projects. We used the official Behance API, to collect more detailed user information about the date of registration, the activity status, and the users' names. This allowed us to evaluate the results of the applied gender inferring method.

In online platforms, users often create accounts without the intention of maintaining an active presence. Users might open an account out of curiosity or to access some features outside of creative work, such as digital storage or viewing designs of others. In order to model users' behavior, we needed subsequent user engagement; thus, in both databases, we filtered users by the level of activity, retaining only those users with at least 10 traces of activity within their careers. Following a common practice in filtering bots on GitHub (70–72), we applied a name-based heuristic approach to remove users who might be bots. We removed users with names containing substrings that classified them as potential artificial agents on GitHub (e.g. “bot,” “test,” “daemon,” “svn2github,” “gitter-badger”). This bot detection method is effective; however, it may lead to an elevated rate of false positives, indicating that nonbot users might have been unintentionally excluded from our analysis. This limitation is preferable to a scenario with a high false negative rate, which would result in a substantial proportion of automated agents within our sample. In the case of Behance, we removed all users whose accounts we could no longer connect to the API and did not have a display

name, which is an indication of being a company. The resulting database contains 1,634,373 GitHub users and 30,186 Behance users.

## Gender inferring

Since none of the data sources lists users' gender, we infer gender from publicly available name information listed by users: first and last names, email addresses, nicknames. Inferring users' perceived gender from public name data based on large-scale gender-name dictionaries has been widely used in computational social science (46, 73). However, it is important to note that these methods can also introduce bias into our results. They perform considerably better with Western names, compared to Asian ones (74, 75) and usually produce binary gender categories (73). We are aware that not everyone has a binary gender identity; however, the name-based gender-inferring methods used are not capable of capturing nonbinary identities. These are important limitations that must be taken into account when discussing results (73); however, we believe that gender-inferring algorithms that attempt to mimic how humans decide about the gender identity of users are valid methods for the purpose of our study. Our study focuses on how *perceived gender* is associated with outcomes on online platforms, and since the public tends to be biased and prefers to categorize people into gender groups (76), our automated method can serve as a suitable proxy for understanding gender inequalities.

In the case of GitHub, we infer first names from display names, usernames, and e-mail addresses using the methods developed by Ref. (38). Behance data were published with inferred gender, using a commercial service called Gender API (<https://gender-api.com/>). Table 1 shows the resulting database by data source and gender. Gender recognition on GitHub yields 11.87% women and 88.13% men out of all users with names, while on Behance the resulting database contained 29.45% women and 70.55% men. After filtering for users with at least 10 traces of activity on both platforms, on GitHub the ratio of women decreases to 5.49%, while on Behance the ratio of active women decreases to 28.39%.

In order to estimate the accuracy of these two gender inferring methods, we took a sample of 200 users for each gender category (female, male, unknown) from each dataset and inferred their gender manually. We compared our classification with the gender inferring methods presented above and also added a third commonly used gender inferring method available as a ready-to-use Python package (Gender Guesser)<sup>b</sup>. We found that among GitHub users our method and the default Python package yielded very similar results, optimized for high male precision. The commercial Gender API used to infer the gender of Behance users

**Table 1.** Data cleaning and gender inferring results.

	GitHub	Behance
N in population	7,798,509	37,777
Women	194,000	11,124
Men	1,441,130	26,653
Unknown	6,163,379	–
N after filtering	1,634,373	30,186
Women	56,731	8,569
Men	977,389	21,617
Unknown	600,253	–
Sample size (by gender)	10,000	6,000

After filtering for users with at least 10 activity points, in GitHub the ratio of women is 5.49%, and on Behance the ratio of active women is 28.39% out of those users whose gender could be inferred.

resulted in higher overall precision, recall, and f-score compared to the default Python package. (See the precision, recall, and F-score of each algorithm by gender in Fig. S1.) To validate the robustness of our results, we run all of our statistical models on datasets with varying levels of gender bias that we introduced by swapping 5, 10, and 25% of the user's gender between male and female.

Finally, to fix the unbalanced nature of our data with regard to gender, we took a biased sample with 10,000 users of each gender (male, female) from the GitHub users and 6,000 each from Behance. We replicated our analysis on five samples; results in the main text are presented based on sample 1 (results for further samples are in our [Supplementary material](#)).

## Identifying specializations

The gender typicality of specializations have been linked to success on digital platforms (38, 39). Therefore, we use data that explain the content of projects on both platforms to identify users' specializations and apply principal component analysis. In both datasets, we created field-specific count variables that measure the frequency at which a user worked with a given programming language on GitHub (e.g. C, Java, Python), or the number of projects where the user indicated a given creative field (e.g. painting, photography, copywriting). For both platforms, we used the 20 most popular programming languages or design fields of those that appeared in at least 1,000 projects. On GitHub, we identified six main specializations; (i) Frontend development, (ii) Developers using Ruby for backend development, (iii) Backend development with high activity in Java, (iv) Data Science, (v) iOS (iPhone Operating System) development, and (vi) PHP projects with frontend focus. In Behance our principal component analysis yielded eight main factors: (i) Photography, (ii) Graphic Design, (iii) Branding, (iv) Art Direction, (v) Digital Art, (vi) Fashion Photography, (vii) Fine Arts, and (viii) Web design- UX. (See Fig. S2 for bar charts showing the explained variance of each factor and the correlation matrices showing the "importance" and the sign of the relationship between the language/design fields in the resulting specialization.)

## Femaleness

We capture the gendered typicality of behavior as the probability of being female, given a pattern of activity. Specifically, we use random forest models to predict whether a user's inferred gender is female. Our features are variables that cover behavioral choices, such as type of engagement (creating and modifying coding repositories, uploading design projects), specialization (programming languages or art categories), and networking (such as the number of people they follow). The resulting prediction score is *femaleness*, which quantifies the female typicality of creative behavior on a scale between 0 (most male-typical behavior) to 1 (most female-typical behavior).

The GitHub Random Forest classification was moderately accurate (AUC = 0.64), on Behance the accuracy was somewhat higher (AUC = 0.69). A key strength of the random forest model is that it can capture nonlinear relationships between variables and enable intuitive ways to quantify the importance of variables (77–79).

Figure 1A and C show how features impact the models' output (Femaleness) for each of our two cases. The dots represent users, and the horizontal axis shows the SHAP (SHapley Additive ExPlanations) value that estimates the contribution of a feature as the difference between expectation without the feature

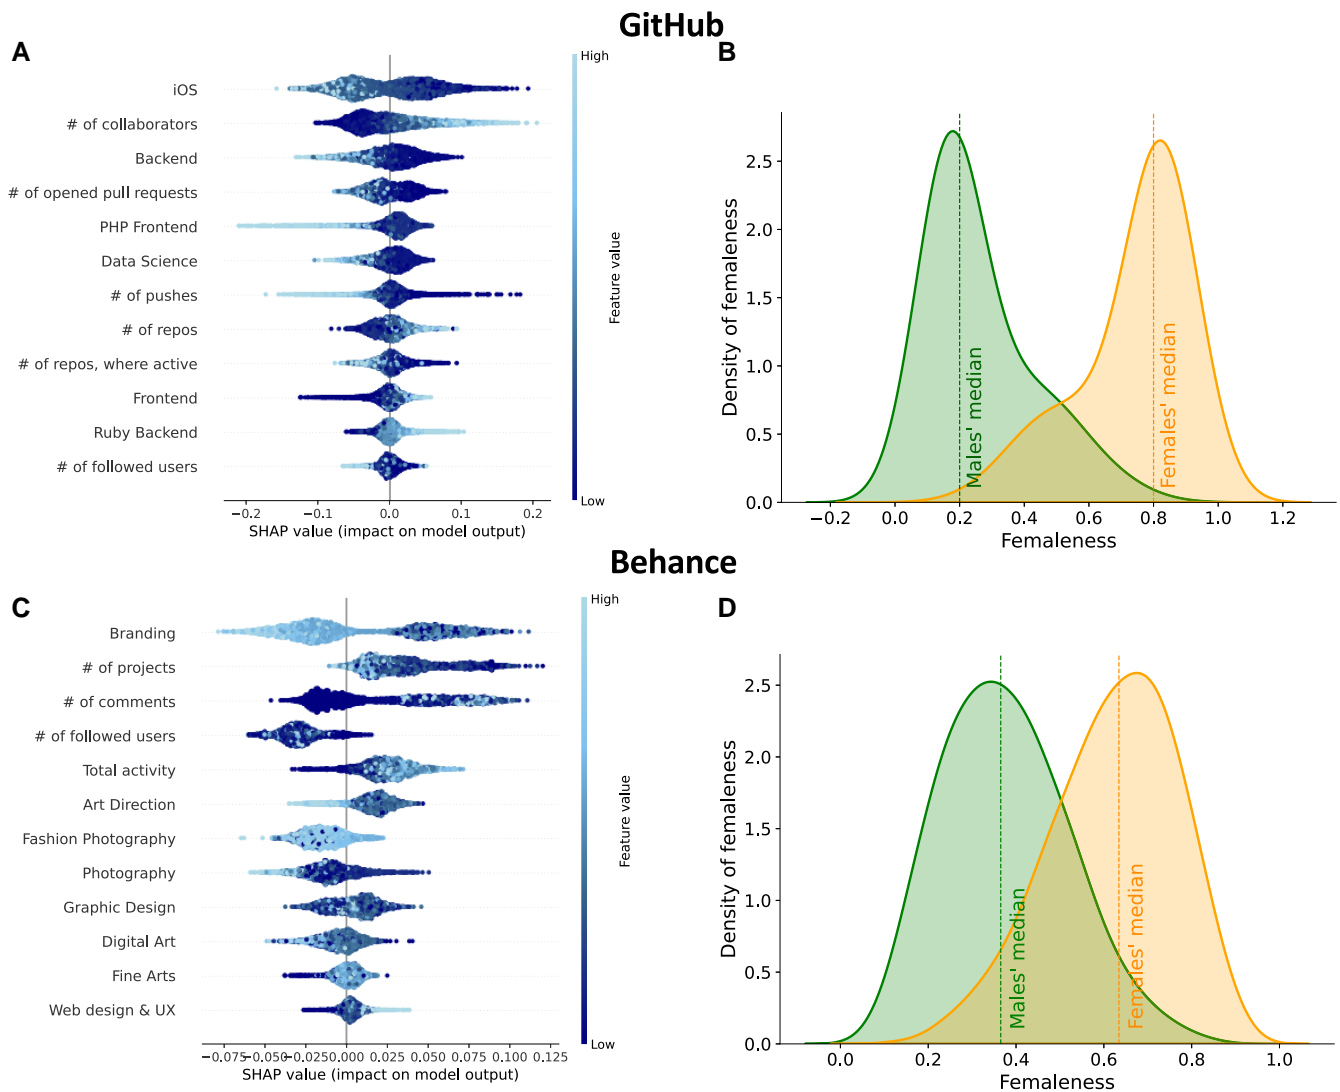


(the mean expected prediction) and the prediction with the feature for each given user. Color is used to display the original value of a feature—light colors indicate high values of the given feature, dark colors low. The features are ordered by their relative importance on the Y axis. For example, in the case of GitHub (Fig. 1A), the most important predictor of being female is developing software for the iOS mobile operating system (labeled “iOS,” first row of and the high values of iOS development (light blue dots in the first row of Fig. 1A) predict a low probability that a user is female (as light blue dots are toward the lower, left-hand side of the x-axis). This indicates that specializing in iOS development is more of a male-typical trait, rather than a female-typical one. Also, according to our model, a high number of collaborators and specialization in “Ruby backend” are more associated with being female, as the light blue dots in these rows are more towards the higher (right-hand) end of the x-axis. On Behance the most important predictor of being female is “Branding,” low values of “Branding” predict higher femaleness, indicating that it is a more masculine specialization. The high “Number of comments,” specializing in “Fashion photography” and “Web design & UX” are the most feminine traits.

Figure 1B and D shows the probability density of femaleness for men (green) and women (orange) on GitHub and Behance. The separation of developers by femaleness is more pronounced on GitHub; the median femaleness for women is 0.2, and for men it is 0.8. On Behance, the probability distributions and medians are closer to each other (0.37 and 0.63 for men and women, respectively). The distributions on Fig. 1B and D indicate that the behavioral pattern does differ by gender, although the distributions of femaleness for men and women do overlap. In other words, while women are often high on femaleness, we do find several women with low femaleness: male typical behavior.

## Models


The three outcomes *attention*, *success*, and *survival* have been previously linked to build a successful portfolio-based career in both design (18) and software (16). A large follower base provides increased visibility which is associated with higher project success (24, 25, 39), in that sense, attention can precede popularity; however, project success can also fuel attention. We measure attention by the number of followers users have, which information



**Fig. 1.** A, C) Beeswarm plots of Femaleness. Each dot represents one data point, where the X axis is determined by the SHAP (SHapley Additive Explanations) value. The features are ordered by their relative importance on the Y axis. Color displays the original value of a feature—light colors indicate high values of the given feature, dark colors low. B, D) Distribution of Femaleness. Graphs represent the probability density of femaleness for males (green), values on, females (orange) on GitHub (B) and Behance (D). Lines indicate median femaleness by gender groups. #, Number of.

**Table 2.** Variables computed for GitHub and Behance users.

	GitHub	Behance
Attention	Number of followers	Number of followers
Success	Number of stars on own repositories	Number of appreciations on own designs
Survival	Activity 1 year after data collection	Activity 1 year after data collection
Tenure	Years since registration	Years since registration
Gender	Inferred from nickname, email, or full name and inferred from a user's name	
Activity	Number of pushes, number of own repositories, number of repositories, where active, number of opened pull requests	Number of projects, number of comments, number of views and appreciations
Networking Fields	Number of collaborators, number of users followed Programming languages used in projects and creative fields designs labeled	Number of users followed

**Fig. 2.** Hypotheses regarding combinations of direct and behavior-based discrimination. Lines show hypothetical marginal prediction of outcomes by gender category. Y axis is the resulting prediction of an outcome, X axis is femaleness, color indicates gender.

was available on both platforms. To become a follower, someone needs to be aware of a user's work and express willingness to keep updated about further works, which requires an action by clicking on a "Follow" button.

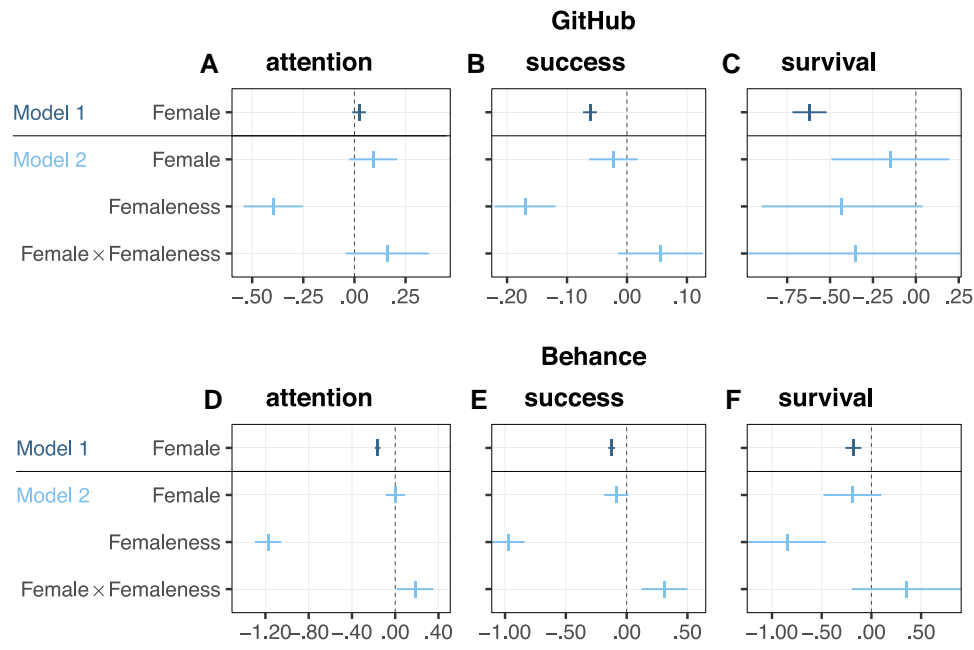
Our success measures sum up projects' popularity, we use on both sites metrics which are community driven, and often used on rankings to list trending developers or designers.<sup>c</sup> On GitHub, we quantify success by the number of stars on users' own repositories, and on Behance, by the number of "appreciations" (likes) on users' own designs. Success is more than merely attention, as it indicates expressed appreciation of quality towards a given piece of work from a user; that is why we do not use project views on Behance. On GitHub, we also use the "Number of accepted pull requests" (the process of merging new code changes into a project, reviewed by other developers) as an alternative success metric. However, because this metric has been shown to exhibit gender bias (37) and cannot be replicated on Behance, we included this part of our analysis only in the [Supplementary material](#) and referenced it in the Results section.

Our third outcome *survival* captures the active participation of users in the platform. Staying active by producing new work and being engaged with the community, is necessary to be noticed by recruiters or potential clients (18, 22). To quantify it, we revisited both platforms 365 days after data collection closed and checked whether the user had additional activity in that 365-day interval. If the user did not leave any trace of activity within this 1-year time window, we marked the user as inactive, otherwise, we marked the user as a survivor.

Because attention and success are considerably skewed to the right, we apply a logarithmic transformation to the number of followers, number of stars, and project appreciations. Thus,  $\log(\text{attention} + 1)$  and  $\log(\text{success} + 1)$  are estimated with linear models. For the estimation of *survival*, since it is a binary variable (users who had activity marked with 1, while dropped outs are marked as 0), we used a logistic regression model.

We enter the same set of control variables in each model, corresponding to relevant alternative explanations for gender differences in outcomes. Due to higher work-family conflict and societal expectations, women generally have less time to maintain their professional presence; therefore, the level of activity might benefit men more than women (80). Men are more likely to join online portfolio sites earlier (81), and users with a longer tenure are more likely to build larger audiences and accumulate more visibility (attention) and success (82, 83). Thus, we control for *tenure* (number of years since registration) and *total activity* (number of repositories or projects, and total activity on sites). Table 2 describes the variables used in modeling the impact of direct and behavior-based discrimination on the three outcomes.

We specified statistical models that examine the relationship between direct and behavior-based discrimination and outcomes. Therefore, our key variables, gender (binary, 1 = Female, 0 = Male), and femaleness, are entered into the models separately and also with their interaction. Figure 2 illustrates our hypotheses, separating the impact of direct discrimination by categorical gender (indicated by color) and behavior-based discrimination by gender typicality (femaleness, on the x-axis) on the outcomes. If individual outcomes were only impacted by direct discrimination (Fig. 2H1), only categorical gender would be a significant predictor in our models, without any slope for femaleness. In the inverse case, with only behavior-based discrimination (Fig. 2H2), femaleness would be a significant predictor in models with a significant slope, without difference between the two gender groups (equally impacting both men and women). The outcomes are likely to be influenced by a combination of direct and behavior-based discrimination. When there is direct discrimination, the prediction lines will have significantly different intercepts by gender, and there will be no significant differences in the slopes due to behavior-based discrimination, which means that it impacts men and women in the same way (Fig. 2H3). Lastly, it is also possible that behavior-based discrimination will have a different



**Fig. 3.** Point estimates of outcomes, with 95% CI, for variables related to gender. *Attention* and *success* models show coefficients from linear models predicting the log. number of stars received and the log. number of project appreciations, while *survival* models show coefficients from logit models predicting survival over a 1 year period following our data collection.

impact by gender, such that, for example, women will be penalized when they follow female-typical behavior, but men will not experience the same behavior-based discrimination (Fig. 2H4). In such a case, the interaction term between direct and behavior-based discrimination will be significant.

## Results

We found that there is a significant baseline gender difference in attention, success on both platforms, and survival on GitHub. The Mann–Whitney  $U$  tests ( $MW$ ) revealed that men have a significantly higher number of followers (GitHub:  $IQR_m = [1, 12]$ ,  $IQR_w = [0, 10]$ ,  $MW_p = 0.000$ —not significant via ordinary least squares (OLS) regression model, Behance:  $IQR_m = [41; 928]$ ,  $IQR_w = [25, 401]$ ,  $MW_p = 0.000$ ) and are more successful (number of stars on GitHub  $IQR_m = [0, 1]$ ,  $IQR_w = [0, 0]$ ,  $MW_p = 0.000$ , number of project appreciations on Behance  $IQR_m = [66, 2264]$ ,  $IQR_w = [47, 980]$ ,  $MW_p = 0.000$ ) on both platforms. Men have a higher average survival rate on GitHub (average survival;  $Avg_m = 0.93$ ,  $Avg_w = 0.88$ ,  $MW_p = 0.000$ ) but their advantage is not significant on Behance ( $Avg_m = 0.45$ ,  $Avg_w = 0.417$ ,  $MW_p = 0.128$ ). (See Table S1 for Mann–Whitney  $U$  test results.)

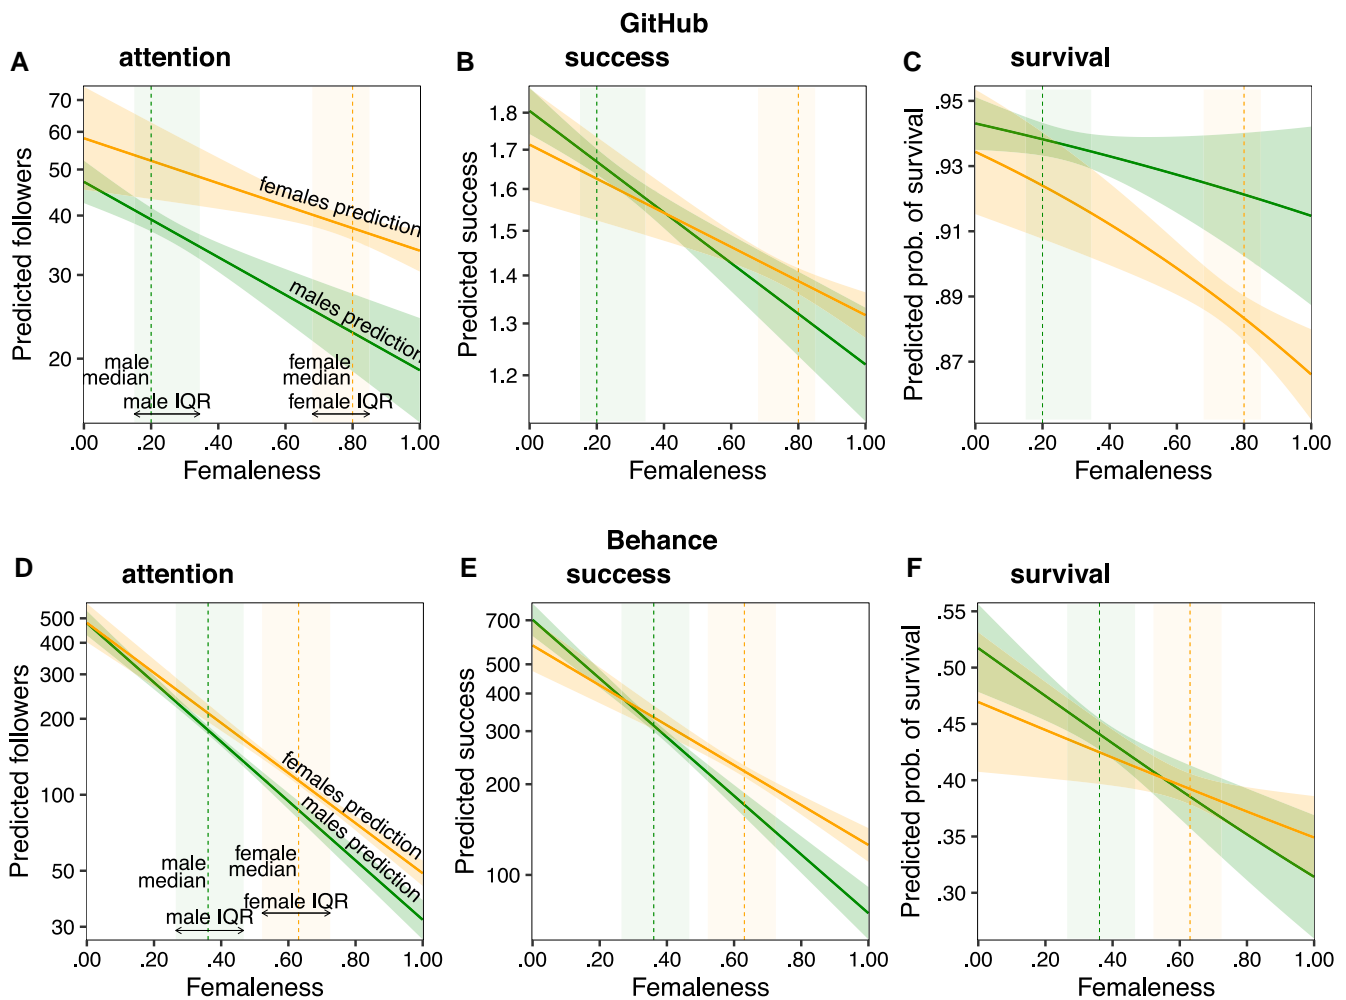
Still considering the gross difference between gender categories (without separating direct and behavior-based discrimination), but also entering controls for activity level, tenure, and fields, we still see a baseline categorical difference for gender in most outcomes. All variables are measured on the scale 0–1, making estimates comparable. Figure 3 model 1 for each of the six panels from A to F shows the relative difference for female developers in all outcomes and platforms, once we take controls into account. With the exception of attention on GitHub, where there is no significant difference ( $P = 0.186$ ), all results show a significant female disadvantage (all  $P = 0.000$ ).

Figure 3 model 2 (across all panels from A to F) shows point estimates after including femaleness in the models. After introducing behavior-based discrimination, the gender category in itself (being female) is no longer significant. (This only means that at the zero

value of femaleness, which is fully male-typical behavior, there is no difference between gender categories in outcomes.) However, femaleness is a significant negative predictor of outcomes in all cases, except survival on GitHub. This indicates that behavior-based discrimination is a significant predictor of differential results between men and women in attention and success on both platforms and in survival on Behance. Our models also indicate that only on Behance in the case of attention and success (Fig. 3D and E) femaleness is associated with men and women significantly differently: Women with high femaleness are predicted to receive more attention and have more project appreciations, suggesting that men are more penalized for exhibiting highly female-like behavior.

Our results are consistent across all five samples. Femaleness remains significant in gender-swapped datasets with error rates 5% in 70–87% of the cases (out of 100 reruns) and with error rates 10%, 55–86% on GitHub and 81–90% on Behance. Simulations with 25% gender swapping are less consistent with significant cases; <50% remains significant. (See Model Tables S2 and S3 for five samples and gender-swapped simulations in [Supplementary material, Table S7](#), and Fig. S6 for a model and marginal prediction with an alternative success metric.)

Figure 4 shows the predicted values of attention (first column), success (second column), and survival (third column) along the range of femaleness by gender categories on GitHub (first row) and Behance (second row). Although the negative impact of femaleness put both men and women in disadvantage in all models (negative slope), in some cases categorical gender predicts outcomes differently. In the case of attention, the difference between women (orange) and men (green) increases along the range of femaleness, predicting a higher level of attention for women users with highly female-like behavior than men. This trend holds for predicting the number of project appreciations on Behance, while there is no significant gender difference between women and men on GitHub by femaleness. Our alternative success metric, available only for GitHub, the number of merged pull requests, yielded similar results: Femaleness is a strong negative predictor, while categorical gender is not a significant variable in the negative



**Fig. 4.** Marginal predictions of outcomes from model 2 from Fig. 3, with fixing all other variables at their means. Vertical dashed lines indicate medians of femaleness, and shaded vertical bars show the interquartile range (IQR).

binomial regression model (see Table S7 and Fig. S6). Categorical gender is not associated with the survival of Behance users, while female GitHub users face a disadvantage compared to men at the higher end of the femaleness spectrum.

Figure 4 allows us to test our hypotheses described in Fig. 2. The presence of behavior-based discrimination is shared across all panels of Fig. 4. There is no clear evidence of direct discrimination; however, the degree of behavior-based discrimination varies by gender and outcome.

To quantify women's disadvantage, we take the prediction of the outcomes of men at their median femaleness and deduct the predicted value of women at their femaleness. On GitHub women have an attention gap of 1.62 followers. Relative to the predicted number of men's followers, at men's median femaleness, it is a 4% gap. The disadvantage is so small, because men suffer more from behavior-based discrimination than women (men lose 17 followers between their and women's femaleness median, while women lose only 15), and women have a direct gender advantage in attention (13 more followers than men at men's median). Women suffer a total attention disadvantage of 26% relative to men on Behance. Although women have a direct gender advantage compared to men, it cannot compensate that they are more affected by behavior-based discrimination.

On GitHub women have a total success disadvantage of 6%, of which 90% is due to behavior-based discrimination and 10% due to

direct. On Behance women have a 37% total disadvantage compared to men's in success, which is entirely caused by behavior-based discrimination and mediated by 20% by a direct advantage of women.

In predicted survival, women suffer a total disadvantage of 6% on GitHub, of which 74% is due to behavior-based discrimination and 26% due to direct discrimination. The trend is similar on Behance with a total of 11% of women's disadvantage, composed of 60% behavior-based and 40% direct discrimination. (See Table SI.4 for prediction results at men's and women's femaleness medians by outcomes, and calculated direct and behavior-based discrimination in exact numbers and percentages.)

## Conclusion

Collaborative platforms show consistent behavior-based gender discrimination, while categorical gender discrimination is only occasional and small. Our findings indicate that behavior-based gender discrimination, present in career choices and online activity, is present both on GitHub (a considerably male-dominated platform) and also on Behance (a platform with higher ratio of women). We found that a significant portion (60–90%) of women's disadvantage in attention, success, and survival can be associated with behavior-based discrimination.

Behavior-based discrimination negatively associated with both genders; furthermore, in the more gender-balanced design



community, men are penalized even more for female-typical behavior, compared to women. Although design careers are often considered more feminine than software development, and have a higher ratio of women professionals (28%), masculine career choices have previously been associated with greater success (39). Although empirical studies suggest that shorter tenure and lower participation rates are the main factor behind lower wages and success in creative fields (31), the case of Behance suggests the opposite: In this case, higher participation of women does not automatically diminish behavior-based gender discrimination.

We found evidence consistent with the visibility paradox of women in technical fields. Women attract considerably more attention, but are often not recognized as experts (57). Although femaleness is negatively related to attention on both fields, and women suffer from even higher behavior-based discrimination on Behance, they have a categorical gender advantage in attention. Female developers on GitHub are extremely visible, they enjoy almost a nine times higher gender advantage compared to men. On Behance, the visibility gap caused by femaleness is only reduced by 45% due to categorical gender.

The higher attention that women attract on online collaborative platforms is a double-edge sword. On the one hand, women could use this increased visibility to build a larger audience and promote their work, which can eventually help them succeed through more visible role models (84, 85). On the other hand, increased attention also has negative consequences for women. Women are more likely to be harassed online and increased visibility could also attract verbal violence (86–88), making women less likely to participate (89) and dropping out at higher rates.

We cannot test the casual hypothesis that attention leads to success and survival (as our platform data do not offer the opportunity for quasiexperimental setups), but we do see indications that survival correlates with both attention and success on both platforms, with a stronger association observed on Behance ( $corr_{GitHub} = 0.51$ ,  $P = 0.00$ ,  $corr_{Behance} = 0.93$ ,  $P = 0.00$ ). (See Fig. S3 for correlations among outcome variables). We applied a quadratic regression model to examine whether there is a U-shaped relationship between attention and survival on Behance. Our findings suggest that users with higher femaleness are more likely to remain active when they have very few followers or a large follower base, regardless of their gender. This pattern may imply that users with highly female-typical behavior could navigate behavior-based gender discrimination by either maintaining a low profile or capitalizing on greater visibility. In contrast, users with lower femaleness have a near-linear relationship between attention and survival, suggesting that high attention does not appear to lead to negative consequences for them (see Fig. S5 and Table S6).

It is important to emphasize that on online collaborative platforms, not only user interests and interactions shape the presence of inequalities, ranking algorithms, and popularity-based recommendation systems can fuel behavior-based segregation (90, 91). On both Behance and GitHub users have two roles: a “creator,” who create art and code, and a “consumer,” who view, like, appreciate, and in the case of GitHub fork (use) projects. This duality needs to serve both personas by exposing users’ work “fairly” and offering inspiring new projects and artists to follow (92). Recommendation systems have been shown to be biased toward more popular items, for which many times the inherit bias of the users is blamed (90). Although both platforms recommend content to expose certain works, the algorithms in place are not transparent. Since most common recommendation algorithms (e.g. content-based and collaborative filtering) employ content similarity and taste homophily to predict projects to watch, they are likely to contribute to behavior-

based segregation of underrepresented groups who have already specialized in subfields (38). Although we cannot eliminate the effect of recommendation algorithms, we are aware that they might contribute to the observed phenomena.

Another aspect of platform design that may influence our results is gamification, which involves integrating game-like features, such as points, rewards, and contests, to encourage user engagement (13). GitHub uses gamification (e.g. visual showing daily activity streak counts in user profiles), which tends to motivate men more than women (36, 93), potentially leading to gender difference in consistency and activity. If this gamified reward system would have a significant disparate influence on men and women’s engagement, we would expect to see categorical gender discrimination as a significant predictor in our survival model in the case of GitHub. Only model 1 which tests for categorical gender discrimination shows a significantly lower survival probability for women compared to men, but when behavior-based discrimination is included in the model, this difference disappears. Furthermore, we see similar trends on Behance, where there is no explicit gamification is implemented. We have two potential explanations, gamification might not impact men and women differently within our sample (active users), or it’s impact is already captured by the gender typicality of behavior and that is why categorical gender is not significant. Our current data does not allow us to further unpack the differential impact of gamification on men and women, since we did not collect information about users’ rewards and achievements.

Since the launch of ChatGPT, large language models have become a key tool for asking programming-related questions and creating digital arts (94). These models were trained on data from the Internet and have been shown to reproduce the biases inherent in their data sources (95). If behavior-based discrimination is highly associated with success in online platforms, products created via such AI systems might prefer solutions and creative outputs generated by users with less female-like behavior. There are already signs that self-learning algorithms would propagate existing gender disparities in the labor market (96), search engines (97), and produce images in a sexist fashion (98, 99). As these models would not have the capacity to recognize and resist gender stereotypes baked into their training data sources, solutions built with them will carry over such stereotypes. As AI solutions become wide spread, we fear that it will become almost impossible to detect and evade behavior-based gender discrimination. The only way to mitigate the impact of AI-amplified behavior-based discrimination would be to put into place a monitoring mechanism that constantly measures direct and behavior-based discrimination to alert the public to intervene against harms to underrepresented groups.

## Notes

<sup>a</sup> See <https://googlegendercase.com/>

<sup>b</sup> <https://pypi.org/project/gender-guesser/>

<sup>c</sup> See <https://github.com/trending> and <https://www.behance.net/search/projects/TRENDING>

## Acknowledgments

The authors thank anonymous reviewers for their valuable suggestions.

## Supplementary Material

Supplementary material is available at PNAS Nexus online.

## Funding

The authors thank the generous support from the MacArthur Foundation who made this research possible. This research was supported by the “Intellectual Themes Initiative” of Central European University, 2016–18. O.V. was funded by the European Union under the European Research Executive Agency project LearnData, 101086712.

## Author Contributions

O.V. and B.V. collected and analyzed the data, wrote and reviewed the manuscript.

## Preprints

An earlier version of this article is published as a preprint at <https://arxiv.org/pdf/2103.01093>.

## Data Availability

The data underlying this article are available in “Behance\_GitHub” repository on GitHub, at [https://github.com/velf/behance\\_github](https://github.com/velf/behance_github).

## References

- De Groen WP, Kilhoffer Z, Lenaerts K, Salez N. 2017. The impact of the platform economy on job creation. *Inter Econ.* 52:345–351.
- Stark D, Pais I. 2020. Algorithmic management in the platform economy. *Sociologica.* 14(3):47–72.
- DinarStandard. Global digital platform powerindex; 2023 [accessed 2025 Feb 4]. <https://www.dinarstandard.com/post/global-digital-platform-powerindex-2023>.
- Stojkoski V, Koch P, Coll E, Hidalgo CA. 2024. Estimating digital product trade through corporate revenue data. *Nat Commun.* 15(1):5262.
- Magaudda P, Marco S. 2020. Platform studies and digital cultural industries. *Sociologica.* 14(3):267–293.
- Xue C, Tian W, Zhao X. 2020. The literature review of platform economy. *Sci Program.* 2020:1–7.
- Cansoy M, Eddy S, Ladegaard I, Schor JB. 2020. Homines diversi: heterogeneous earner behaviors in the platform economy. *Sociologica.* 14(3):143–165.
- Pelzer P, Frenken K, Boon W. 2019. Institutional entrepreneurship in the platform economy: how uber tried (and failed) to change the Dutch taxi law. *Environ Innov Soc Transit.* 33:1–12.
- Schor J. 2016. Debating the sharing economy. *J Self-Gov Manag Econ.* 4(3):7–22.
- Schor JB, Attwood-Charles W, Cansoy M, Ladegaard I, Wengronowitz R. 2020. Dependence and precarity in the platform economy. *Theory Soc.* 49(5-6):833–861.
- Gray ML, Suri S, Ali SS, Kulkarni D. The crowd is a collaborative network. In: *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing.* Association for Computing Machinery, 2016. p. 134–147.
- Mergel I. 2015. Open collaboration in the public sector: the case of social coding on GitHub. *Gov Inf Q.* 32(4):464–472.
- Moldon L, Strohmaier M, Wachs J. How gamification affects software developers: cautionary evidence from a natural experiment on GitHub. In: *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE).* IEEE Press, 2021. p. 549–561.
- Churchill B, Craig L. 2019. Gender in the gig economy: men and women using digital platforms to secure work in Australia. *J Sociol.* 55(4):741–761.
- Grabher G, König J. 2020. Disruption, embedded. A polanyian framing of the platform economy. *Sociologica.* 14(1):95–118.
- Riehle D. 2015. How open source is changing the software developer’s career. *Computer.* 48(5):51–57.
- Halstead S, Serrano HD, Proctor S. 2015. Finding top UI/UX design talent on adobe behance. *Procedia Comput Sci.* 51:2426–2434.
- Scolere L. 2019. Brand yourself, design your future: portfolio-building in the social media age. *New Media Soc.* 21(9):1891–1909.
- Bonaccorsi A, Rossi C. 2003. Why open source software can succeed. *Res Policy.* 32(7):1243–1258.
- Borges H, Valente MT. 2018. What’s in a GitHub star? Understanding repository starring practices in a social coding platform. *J Syst Softw.* 146:112–129.
- Dabbish L, Stuart C, Tsay J, Herbsleb J. Social coding in GitHub: transparency and collaboration in an open software repository. In: *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work.* ACM Press, New York, USA, 2012. p. 1277–1286.
- Gerosa M, et al. The shifting sands of motivation: revisiting what drives contributors in open source. In: *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE).* IEEE, 2021. p. 1046–1058.
- Schembri S, Latimer L. 2016. Online brand communities: constructing and co-constructing brand culture. *J Mark Manag.* 32(7-8):628–651.
- Musienko Y. Why web design companies use Behance? [accessed 2024 Sep 30]. <https://merehead.com/blog/why-we-on-behance/>.
- Susnja L. Behance and Dribbble’s role in my designer success story [accessed 2024 Sep 30]. <https://www.linkedin.com/pulse/behance-dribbles-role-my-designer-success-story-leonardo-susnja/>.
- Neff G, Wissinger E, Zukin S. 2005. Entrepreneurial labor among cultural producers: “cool” jobs in “hot” industries. *Soc Semiot.* 15(3):307–334.
- Ford D, Smith J, Guo PJ, Parnin C. Paradise unplugged: identifying barriers for female participation on stack overflow. In: *Proceedings of the 2016 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering.* Association for Computing Machinery, 2016. p. 846–857.
- Huws U. 2019. The hassle of housework: digitalisation and the commodification of domestic labour. *Fem Rev.* 123(1):8–23.
- van den Berg AC, Giest SN, Groeneveld SM, Kraaij W. 2020. Inclusivity in online platforms: recruitment strategies for improving participation of diverse sociodemographic groups. *Public Adm Rev.* 80(6):989–1000.
- Bogliacino F, Codagnone C, Cirillo V, Guarascio D. Quantity and quality of work in the platform economy. In: *Handbook of labor, human resources and population economics.* Springer International Publishing, 2019.
- Cook C, Diamond R, Hall J, List JA, Oyer P. 2018. The gender earnings gap in the gig economy: evidence from over a million ride-share drivers.
- Poutanen S, Kovalainen A. New economy, platform economy and gender. In: *Gender and Innovation in the New Economy.* Springer, 2017. p. 47–96.
- Silberzahn R, Uhlmann EL, Zhu L. Pay as she goes: for stereotypically male jobs, women tend to be hired by the hour. In: *Academy of Management Proceedings.* vol. 2014, Academy of Management Briarcliff Manor, NY 10510, 2014. p. 16273.

- 34 Brooke SJ. 2021. Trouble in programmer's paradise: gender-biases in sharing and recognising technical knowledge on stack overflow. *Inform Commun Soc.* 24(14):2091–2112.
- 35 Hannák A, et al. Bias in online freelance marketplaces: evidence from taskrabbit and fiverr. In: *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. Association for Computing Machinery, 2017. p. 1914–1933.
- 36 May A, Wachs J, Hannák A. 2019. Gender differences in participation and reward on stack overflow. *Empir Softw Eng.* 24(4):1997–2019.
- 37 Terrell J, et al. 2017. Gender differences and bias in open source: pull request acceptance of women versus men. *PeerJ Comput Sci.* 3:e111.
- 38 Vedres B, Vasárhelyi O. 2019. Gendered behavior as a disadvantage in open source software development. *EPJ Data Sci.* 8(1):25.
- 39 Wachs J, Hannák A, Vörös A, Daróczy B. Why do men get more attention? Exploring factors behind success in an online design community. In: *Proceedings of the International AAAI Conference on Web and Social Media*. Vol. 11. AAAI, 2017. p. 299–308.
- 40 Ford D, Harkins A, Parnin C. Someone like me: how does peer parity influence participation of women on stack overflow? In: *2017 IEEE Symposium on Visual Languages and Human-centric Computing (VL/HCC)*. IEEE, 2017. p. 239–243.
- 41 Horvat E-A, Papamarkou T. Gender differences in equity crowdfunding. In: *Fifth AAAI Conference on Human Computation and Crowdsourcing*. AAAI, 2017.
- 42 Vasilescu B, et al. Gender and tenure diversity in GitHub teams. In: *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. Association for Computing Machinery, 2015. p. 3789–3798.
- 43 Wagner C, Garcia D, Jadidi M, Strohmaier M. It's a man's Wikipedia? Assessing gender inequality in an online encyclopedia. In: *Ninth International AAAI Conference on Web and Social Media*. AAAI, 2015.
- 44 Wagner C, Graells-Garrido E, Garcia D, Menczer F. 2016. Women through the glass ceiling: gender asymmetries in Wikipedia. *EPJ Data Sci.* 5(1):5.
- 45 Grau-Sarabia M, Fuster-Morell M. 2021. Gender approaches in the study of the digital economy: a systematic literature review. *Humanit Soc Sci Commun.* 8(1):1–10.
- 46 Vásárhelyi O, Brooke S. 2022. Computing gender. <https://doi.org/10.31235/osf.io/admc3>, preprint.
- 47 West C, Zimmerman DH. 1987. Doing gender. *GenD Soc.* 1(2):125–151.
- 48 Diekman AB, Brown ER, Johnston AM, Clark EK. 2010. Seeking congruity between goals and roles: a new look at why women opt out of science, technology, engineering, and mathematics careers. *Psychol Sci.* 21(8):1051–1057.
- 49 Kanze D, Conley MA, Okimoto TG, Phillips DJ, Merluzzi J. 2020. Evidence that investors penalize female founders for lack of industry fit. *Sci Adv.* 6(48):eabd7664.
- 50 Katila S, Meriläinen S. 1999. A serious researcher or just another nice girl?: Doing gender in a male-dominated scientific community. *GenD Work Organ.* 6(3):163–173.
- 51 Cross S, Bagilhole B. 2002. Girls' jobs for the boys? Men, masculinity and non-traditional occupations. *GenD Work Organ.* 9(2):204–226.
- 52 Brooke S. 2024. Programmed differently? Testing for gender differences in python programming style and quality on GitHub. *J Comput Mediat Commun.* 29(1):zmad049.
- 53 Mittal S, Gupta N, Dewan P, Kumaraguru P. Pinned it! A large scale study of the pinterest network. In: *Proceedings of the 1st IKDD Conference on Data Sciences*. Association for Computing Machinery, 2014. p. 1–10.
- 54 Wang Y, Horvát E-Á. Gender differences in the global music industry: evidence from musicbrainz and the echo nest. In: *Proceedings of the International AAAI Conference on Web and Social Media*. Vol. 13. AAAI, 2019. p. 517–526.
- 55 Souto VT, Fadel LM, Spinillo CG. 2024. Gender differences in design creativity. *Proc Design Soc.* 4:1125–1134.
- 56 Windsor EJ. Femininities. In: Wright JD, editor. *International Encyclopedia of the social & behavioral sciences*. 2nd ed. Elsevier, 2015. p. 893–897.
- 57 Fernando D, Cohen L, Duberley J. 2018. The problem of visibility of women in engineering and how they manage it. HBR.
- 58 Wang Z, Wang Y, Redmiles D. Competence-confidence gap: a threat to female developers' contribution on GitHub. In: *2018 IEEE/ACM 40th International Conference on Software Engineering: Software Engineering in Society (ICSE-SEIS)*. IEEE, 2018. p. 81–90.
- 59 Albergotti R. 2020. Black creators sue YouTube, alleging racial discrimination. *The Washington Post*.
- 60 Bensinger G, Albergotti R. 2019. YouTube discriminates against LGBT content by unfairly culling it, suit alleges. *The Washington Post*.
- 61 MSL Staff. 2021. MSL study reveals racial pay gap in influencer marketing. MSL [accessed 2022 May 31].
- 62 GitHub. 100 million developers and counting [accessed 2024 May 22]. <https://github.blog/2023-01-25-100-million-developers-and-counting/>.
- 63 Gousios G, Spinellis D. Mining software engineering data from GitHub. In: *2017 IEEE/ACM 39th International Conference on Software Engineering Companion (ICSE-C)*. IEEE, 2017. p. 501–502.
- 64 Zöllner N, Morgan JH, Schröder T. 2020. A topology of groups: what GitHub can tell us about online collaboration. *Technol Forecast Soc Change.* 161:120291.
- 65 Ortu M, et al. 2017. How diverse is your team? Investigating gender and nationality diversity in GitHub teams. *J Softw Eng Res Dev.* 5(1):1–18.
- 66 Imtiaz N, et al. Investigating the effects of gender bias on GitHub. In: *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*. IEEE, 2019. p. 700–711.
- 67 Kim NW. 2017. Creative community demystified: a statistical overview of Behance. *arXiv*, arXiv:1703.00800. <https://doi.org/10.48550/arXiv.1703.00800>, preprint: not peer reviewed.
- 68 Chang S, Kumar V, Gilbert E, Terveen LG. Specialization, homophily, and gender in a social curation site: findings from pinterest. In: *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing*. Association for Computing Machinery, 2014. p. 674–686.
- 69 Leskovec J, Faloutsos C. Sampling from large graphs. In: *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. Association for Computing Machinery, 2006. p. 631–636.
- 70 Lin HY, Thongtanunam P, Treude C, Charoenwet W. Improving automated code reviews: Learning from experience. In: *Proceedings of the 21st International Conference on Mining Software Repositories*. Association for Computing Machinery, 2024. p. 278–283.
- 71 Schueller W, Wachs J, Servedio VDP, Thurner S, Loreto V. 2022. Evolving collaboration, dependencies, and use in the rust open source software ecosystem. *Sci Data.* 9(1):703.
- 72 Wyrich M, Ghit R, Haller T, Müller C. Bots don't mind waiting, do they? Comparing the interaction with automatically and manually created pull requests. In: *2021 IEEE/ACM Third International Workshop on Bots in Software Engineering (BotSE)*. IEEE, 2021. p. 6–10.

- 73 Shugars S, Quintana-Mathé A, Lange R, Lazer D. 2024. Categorizing the non-categorical: the challenges of studying gendered phenomena online. *J Comput Mediat Commun.* 29(1): zmad053.
- 74 Karimi F, Wagner C, Lemmerich F, Jadidi M, Strohmaier M. Inferring gender from names on the web: a comparative evaluation of gender detection methods. In: *Proceedings of the 25th International Conference Companion on World Wide Web - WWW '16 Companion*. International World Wide Web Conferences Steering Committee, 2016. p. 53–54.
- 75 Lockhart JW, King MM, Munsch C. 2023. Name-based demographic inference and the unequal distribution of misrecognition. *Nat Hum Behav.* 7(7):1084–1095.
- 76 Shaw A. 2014. The internet is full of jerks, because the world is full of jerks: what feminist theory teaches us about the internet. *Commun Crit/Cult Stu.* 11(3):273–277.
- 77 Breiman L. 2001. Random forests. *Mach Learn.* 45(1):5–32.
- 78 Parr T, Turgutlu K, Csiszar C, Howard J. 2018. Beware default random forest importances. *March.* 26:2018.
- 79 Strobl C, Boulesteix A-L, Kneib T, Augustin T, Zeileis A. 2008. Conditional variable importance for random forests. *BMC Bioinformatics.* 9(1). <https://doi.org/10.1186/1471-2105-9-307>.
- 80 Ahuja MK. 2002 Mar. Women in the information technology profession: a literature review, synthesis and research agenda. *Eur J Inf Syst.* 11(1):20–34.
- 81 Putzke J, Fischbach K, Schoder D, Gloor PA. 2014. Cross-cultural gender differences in the adoption and usage of social media platforms—an exploratory study of last. FM. *Comput Netw.* 75: 519–530.
- 82 Badashian AS, Stroulia E. Measuring user influence in GitHub: the million follower fallacy. In: *2016 IEEE/ACM 3rd International Workshop on CrowdSourcing in Software Engineering (CSI-SE)*. IEEE, 2016. p. 15–21.
- 83 Jiang J, et al. 2017. Why and how developers fork what from whom in GitHub. *Empir Softw Eng.* 22(1):547–578.
- 84 Eccles JS, Wang M-T. 2016. What motivates females and males to pursue careers in mathematics and science? *Int J Behav Dev.* 40(2): 100–106.
- 85 Kahn S, Ginther D. 2017. Women and stem. Technical report, National Bureau of Economic Research.
- 86 Atir S, Ferguson MJ. 2018. How gender determines the way we speak about professionals. *Proc Natl Acad Sci U S A.* 115(28): 7278–7283.
- 87 Moss-Racusin CA, Rudman LA. 2010. Disruptions in women's self-promotion: the backlash avoidance model. *Psychol Women Q.* 34(2):186–202.
- 88 Yelin H, Clancy L. 2020. Doing impact work while female: hate tweets, 'hot potatoes' and having 'enough of experts'. *Eur J Womens Stud.* 28(1):1350506820910194.
- 89 Coffman KB. 2014. Evidence on self-stereotyping and the contribution of ideas. *Q J Econ.* 129(4):1625–1660.
- 90 Ionescu S, Hannak A, Pagan N. Group fairness for content creators: the role of human and algorithmic biases under popularity-based recommendations. In: *Proceedings of the 17th ACM Conference on Recommender Systems*. Association for Computing Machinery, 2023. p. 863–870.
- 91 Wagner C, et al. 2021. Measuring algorithmically infused societies. *Nature.* 595(7866):197–204.
- 92 Modani N, Jain D, Soni U, Gupta GK, Agarwal P. Fairness aware recommendations on Behance. In: *Advances in Knowledge Discovery and Data Mining: 21st Pacific-Asia Conference, PAKDD 2017, Jeju, South Korea, May 23–26, 2017, Proceedings, Part II 21*. Springer, 2017. p. 144–155.
- 93 Pedro LZ, Lopes AMZ, Prates BG, Vassileva J, Isotani S. Does gamification work for boys and girls? An exploratory study with a virtual learning environment. In: *Proceedings of the 30th Annual ACM Symposium on Applied Computing*. 2015. p. 214–219.
- 94 del Rio-Chanona M, Laurentsyeveva N, Wachs J. 2024. Large language models reduce public knowledge sharing on online Q&A platforms. *PNAS Nexus.* 3(9). <https://doi.org/10.1093/pnasnexus/pgae400>.
- 95 Kotek H, Dockum R, Sun D. Gender bias and stereotypes in large language models. In: *Proceedings of the ACM Collective Intelligence Conference*. Association for Computing Machinery, 2023. p. 12–24.
- 96 Ali M, et al. 2019. Discrimination through optimization: how Facebook's ad delivery can lead to biased outcomes. *Proc ACM Hum Comput Interact.* 3(CSCW):1–30.
- 97 Vlasceanu M, Amodio DM. 2022. Propagation of societal gender inequality by internet search algorithms. *Proc Natl Acad Sci U S A.* 119(29):e2204529119.
- 98 Bolukbasi T, Chang K-W, Zou JY, Saligrama V, Kalai AT. Man is to computer programmer as woman is to homemaker? Debiasing word embeddings. In: Lee D, Sugiyama M, Luxburg U, Guyon I, Garnett R, editors. *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2016. p. 4349–4357.
- 99 Zou J, Schiebinger L. 2018. AI can be sexist and racist—it's time to make it fair.