

Why Do Men Get More Attention? Exploring Factors Behind Success in an Online Design Community

Johannes Wachs,[†] Anikó Hannák,[†] András Vörös,^{*} Bálint Daróczy[‡]

[†]Center for Network Science, Central European University

^{*}Chair of Social Networks, ETH Zurich, Zurich, Switzerland

[‡]Institute for Computer Science and Control (MTA SZTAKI), Hungarian Academy of Sciences

Abstract

Online platforms are an increasingly popular tool for people to produce, promote or sell their work. However recent studies indicate that social disparities and biases present in the real world might transfer to online platforms and could be exacerbated by seemingly harmless design choices on the site (e.g., recommendation systems or publicly visible success measures). In this paper we analyze an exclusive online community of teams of design professionals called Dribbble and investigate apparent differences in outcomes by gender. Overall, we find that men produce more work, and are able to show it to a larger audience thus receiving more likes. Some of this effect can be explained by the fact that women have different skills and design different images. Most importantly however, women and men position themselves differently in the Dribbble community. Our investigation of users' position in the social network shows that women have more clustered and gender homophilous following relations, which leads them to have smaller and more closely knit social networks. Overall, our study demonstrates that looking behind the apparent patterns of gender inequalities in online markets with the help of social networks and product differentiation helps us to better understand gender differences in success and failure.

1 Introduction

Research in the social sciences has shown that both individual and social network attributes impact individual success in education, the workplace, and the job market. However, the specific mechanisms enabling or hindering success are highly dependent on social context: the available channels for social contact, the constraints on social ties, the channels for social influence, group sizes, and other factors clearly influence individual success.

The recent growth in popularity of online platforms for social interactions (e.g. Twitter, Instagram) and job search (e.g. LinkedIn, freelancer.com) changes the social mechanisms that determine individual success. Currently, we know very little about how inequalities emerge in these new types of communities. We do know, however, that the design of the sites matters: some researchers express concerns that the use of algorithms and public feedback might retain or even reinforce inequalities in success based on, especially, race

and gender (Lustig et al. 2016; Lee et al. 2015a). Empirical work in online freelance communities (Teodoro et al. 2014; Thebault-Spieker, Terveen, and Hecht 2015; Hannak et al. 2017) and on collaboration in teams (Vasilescu, Capiluppi, and Serebrenik 2013; De Vaan, Stark, and Vedres 2015) also highlight the presence of gender inequalities.

Some of these communities combine the open nature of online social networks with the professional aspects of real-world labor markets. Users invest in their identities by showing work, exchanging ideas, and collaborating in visible ways. Over time and with great investment users shape permanent identities with reputations and social capital, just like in the real world. Because of the online nature these identities operate in a different social environment. Online ties between people are about sharing access and have much lower costs, and exist at much larger scales. At the same time online platforms shape social interaction are highly structured and governed by algorithms. Our goal is to explore the inequalities that emerge from this combination of a scaled-up social environment with highly structured systems.

In this paper we analyze Dribbble, the most “elite” online community for digital and graphic designers. The site allows designers to showcase their work in web design, illustration, and other creative areas, follow artists whose work they appreciate, discuss design ideas, and work on collaborative projects in teams. Dribbble enjoys high prestige in the worldwide community of digital and graphic designers, as it is invitation-only and provides a good platform for advertising one's work. We crawled the pages of all 994 teams on the site, 6,215 users involved in one of the teams, and finally all 60,406 images created by these teams.

Our questions are *do men and women have different success rates on Dribbble?*, and if yes, *what are the factors contributing to the differences?* We separately analyze the effects of individual user characteristics, activity on the site, production patterns, and social network structure to understand how much each of these factors contributes to success of individuals and where gender differences may be reinforced.

Using the variables extracted from our data set, we define three measures of user success: the average number of views, likes, and responses the works of a user receive. Using regression analysis, we establish that men are more successful according to all measures, even after controlling for basic

individual characteristics extracted from the profile information.

Since skills and social background might determine the kind of work people produce, we next investigate whether this is true on Dribbble and how much of the gender differences can be explained by such factors. We create a measure for skill and image “genderness” using data mining techniques. Interestingly, even though skills are not strongly divisive, our classifier can predict with 0.72 AUC if an image was created by a man or a woman. This suggests that indeed some men and women are creating different art. Once controlling for these variables in our models we find that the relationship between gender and outcome is no longer significant.

Finally, we investigate social behavior and network effects on the social network underlying Dribbble. We run Exponential Random Graph Models (ERGM) and the results suggest that women have fewer ties but more cohesive social networks than men. When we return to our original model of success and control for these network features, we find again that gender no longer has a significant impact on success.

2 Related Work

What determines whether people succeed at school, in the job market, or in the workplace? Social scientists recognized long ago that individual characteristics, especially those related to socio-economic background and behavior are important predictors of performance and achievement (Angrist and Lang 2004; Curcio, Ferrara, and De Gennaro 2006). However, how such individual factors determine success greatly depends on social context and especially on the network of social relations between people (Coleman 1994).

A long line of social research from various disciplines demonstrates that individual outcomes and social networks are dependent on one another and they are shaped by a couple of social mechanisms that are observable across various settings, both online and offline. These include homophily in characteristics such as sex, race, ethnicity or family background (Feld 1982; McPherson, Smith-Lovin, and Cook 2001), triadic closure (Heider 1946; Cartwright and Harary 1956), clustering and hierarchy formation (Davis 1970), and social influence on individual attitudes and behavior (Veenstra and Dijkstra 2011; Marsden and Friedkin 1993; Turner 1991).

Some studies specifically show that people’s position in informal social networks can affect their performance. For example, recent research on peer influence and social networks among academics has found that scientific collaboration impacts academic success (Petersen et al. 2014; Petersen 2015). On the macro level, the centrality of individuals in collaboration networks is positively associated with their success (Sarigl et al. 2014; Servia-Rodríguez et al. 2015).

Thus, in general it seems that having many connections and a cohesive network may promote individual success. This may especially be true in cases when success is closely linked to informal social status: when people can use or mobilize their social relations to “generate success”, such as popularity or expressed appreciation – this is exactly the case in the empirical study described in this paper.

Informal social networks and individual outcomes co-evolve through selection and influence processes (Veenstra and Dijkstra 2011): people select the peers they associate with not independently of individual characteristics (e.g. homophily); in turn, friends, role models, or groups influence how people behave and perform. This links the question of success closely to the emergence of inequalities in social groups. Many studies focus on gender-based inequalities in education systems (Jacobs 1996), hiring (Pager and Shepherd 2008; Clauset, Arbesman, and Larremore 2015), scientific careers (Lee et al. 2013), or work contexts. An important limitation of many empirical studies in the presented research line is the context dependence of their results. While most studies agree that women (and minorities in general) have worse chances of succeeding, the factors highly depend on the community being studied.

The ongoing migration of both professional and social life to online platforms changes the mechanisms related to success and inequalities (Sandvig et al. 2014; Robinson et al. 2015). Demographic characteristics and status signals are visible on users’ online profiles (Tang et al. 2011) while the pathways to success largely depend on website design and invisible algorithms (Lee et al. 2015b). Studies in online social networks find that influence propagates through connections with more trust and more status (Adali et al. 2010; Hannak et al. 2014; Munger 2016; Ajrouch, Blandon, and Antonucci 2005).

A few recent studies explore success and inequalities in online labor markets (Hannak et al. 2017; Thebault-Spieker, Terveen, and Hecht 2015; Ge et al. 2016) and find that the new mechanisms such as public review systems and algorithmic search might amplify inequalities (Fradkin et al. 2015; Pan et al. 2007). In these settings however there is no clearly measurable underlying social network, and thus the relevance of social effects cannot be easily assessed. Finally, works investigating collaborative team-based platforms such as github (Vasilescu, Capiluppi, and Serebrenik 2013) or online video games (De Vaan, Stark, and Vedres 2015) capture the complex interaction of individual and group success.

3 Data and Extracted Features

Dribbble Dribbble, founded in 2009, is an exclusive community for showcasing user-made artwork in graphic design, web design, illustration, photography, and other creative areas. It has an Alexa rank of 1,012 (as of 01-05-2017) making it the second most visited online community for work in digital design after Behance. In contrast with Behance, Dribbble has an invite-only membership system. Users can only upload their work after receiving an invite from a current member. Moreover, users can only post 48 images in a month and five shots in a day. The result is a high standard of work on the platform and the sense of belonging to an “elite” community. Moreover, Dribbble is a true community in the sense that nearly all users share their identities by linking to their social media accounts and personal home pages, and by uploading photo portraits. Dribbble facilitates job matching between companies and users paying for premium accounts. A significant amount of users pay for this service, highlighting Dribbble’s structural importance in this field.

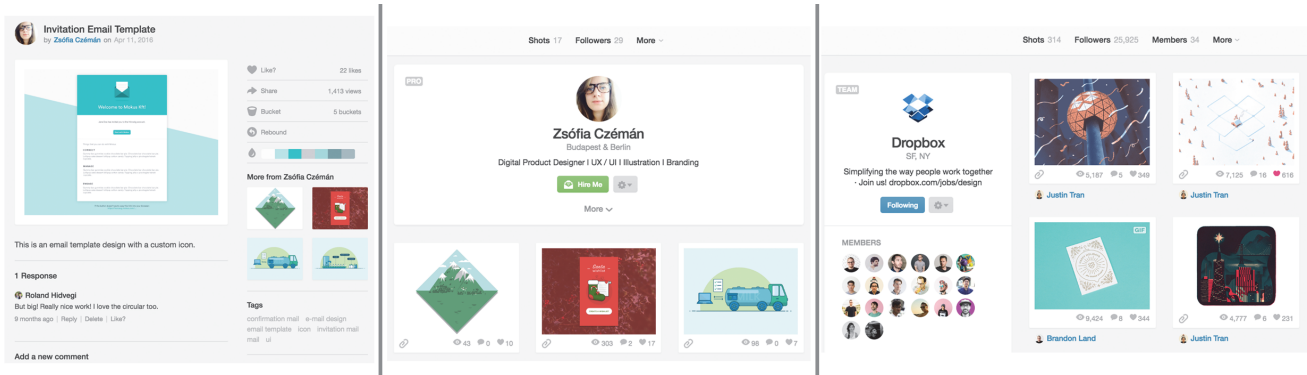


Figure 1: Shot, User and Team Pages on Dribbble.

Shot level	Creation Date, Creating Team/User, Shot, Shot Tags, # of Views, # of Likes, # of Responses
User level	Name, Bio, Skills, Team, Premium Account, # of Shots, # of Followers
Team level	Members, # of Shots, # of Followers

Table 1: Extracted features from Dribbble.

The service has become a platform for some of the most abstract and compelling parts of the creative process. Users can see what their peers are creating and both solicit and give feedback. The site facilitates active use by allowing users to view, like, and comment on work, and to follow others. There is a special subscription that allows organizations to show their work together. In fact many leading IT companies and design boutiques are represented as a “team” on the site. Everything on Dribbble is basketball themed: *players* or *teams* post the art work which they refer to as *shots*. We will follow this convention throughout the paper.

Data Collection Our data collection focuses on three levels of the site: teams, users, and shots. As shown in Figure 1 teams, users, and shots each have an associated profile page. We first crawl the pages of all 994 teams on the site, then we extract the pages all 6,215 users who made a shot while on a team, and finally we collect all 60,406 shots created by teams. About three-fourths of all shots have an individual user authorship tag in addition to the team authorship. All the crawling was done between September and November 2016. During our crawl, we made sure to respect the robots.txt and impose minimal load on Dribbble; we sent a maximum of 1,440 requests per day.

Besides their Dribbble profiles we also collect data about users and teams from their linked Twitter accounts. 86% of teams and 70% of users have twitter accounts. Since it is not possible to obtain the Dribbble follower network directly, we later use Twitter data as a proxy for the social network underlying Dribbble.

Extracted features Table 1 shows the variables we extract from the data set we collected. The three columns show the three types of pages we crawl. For each object (shot, user,

team) we have descriptive variables, such as creation date of a shot, listed skills of a user, or the size of a team. We also extract dynamic features such as number of views, likes, and comments, which we use as measures of success or popularity.

Inferring gender. Since the profiles do not directly list gender, we infer them from the users’ first names using the US baby name (SSA 2016) data set. This is commonly and successfully used method in many studies (Liu and Ruths 2013; Tang et al. 2011; Karimi et al. 2016; CFPB ; Fiscella and Fremont 2006). We are especially confident that users give their real names because of the site’s nature as a platform for designers to advertise their work and gain a following. Our first-pass inference of gender is the probability of being male to each candidate’s name based on its occurrence in the name data set as male. For any user with a name not in the database or an ambiguous gender score (i.e. greater than 10% and less than 90%) we manually check their self-portrait on Dribbble and on linked social media account. We were able to classify all but 77 out of 6,215 users as male or female with high confidence (in the rest of the paper we drop these 77 users). Our final dataset contains 4654 *males* and 1484 *females*.

4 Results

In this section, we investigate the relationship between users’ success and popularity on Dribbble and their gender. Since we find that gender differences in success do exist, we dig deeper in search of user characteristics or social factors which might explain them. Specifically, we focus on two kinds of potential explanations. First, we define various measures that capture how users create products and investigate how these variables affect success and how they relate to gender. Second, we explore the impact of the social aspect of Dribbble by




Figure 2: Cumulative distribution functions for user average views, likes, and responses by gender. (Note: CDF is cut off at 95% of the global average.)

identifying measures of social behavior and network position.

4.1 Gender Differences in Outcomes

Dribbble provides various ways for users to express appreciation for each other’s work, such as liking a shot or leaving a response. These markers of popularity are shown on each product’s page. We can see how many people viewed a shot (this is purely a matter of clicking when the thumbnail shows up in the audience’s feed or following a direct link). We see the number of likes, which is a clear signal of appreciation. Finally, users sometimes comment on an image. The number of responses is a stronger indication of interest, since leaving a public feedback requires more trust or engagement. Although not all comments are positive, we argue they are likely positive in general and that they are a success signal when aggregated on the user profile. We also note that they are non-anonymous. As we want to measure user success, we aggregate these values for all shots created by an individual and use the average view, like, and response counts as three distinct success measures.

At a first glance, women and men have different success rates according to all three measures in this community. Figure 2 shows the gender distributions for the log of the three success measures. (The distribution has a long tail, with only men at very high success rates. For easier interpretation, we cut the CDF at 5% of the global average on either side of the plot.) Men have more views, likes, and responses on average. We find the differences to be significant ($p \leq 0.01$) when applying a Kolmogorov-Smirnov test to compare the distributions. The common language effect sizes between men and women, calculated using the Vargha-Delanay A -measure, are .57, .54, and .54 for views, likes, and responses, respectively. Though the effect sizes are relatively small, we see below that they remain statistically significant even after adding multiple controls.

Considering the user characteristics that could potentially explain success differences, we find that men and women have different levels of participation on Dribbble over time. As shown in Figure 3, male accounts are on average 24% older and post 60% more shots than female accounts. We can

also see that until very recently men posted new work more frequently than women.


Having a premium Dribbble account and identifying as a leader apparently also covary with gender and may be related to success. 38% of men pay for the premium account service or *Pro Badge* compared to 23% of women. We compared user bios based on the self-descriptions of individuals: processed bios containing the strings “founder”, “director”, “manager”, or “partner” were labeled as *self-described leaders*. We find that 19% of men and 10% of women are self-described leaders. There is no significant correlation between having a pro-badge and self-describing as a leader for men or women, while at the same time both variables correlate with higher output.

Arguably, the size of the team a user is part of may also influence his or her success on Dribbble. This is not only because large companies are present with larger teams on the site, but also because of the larger potential for *collegial interactions* and feedback present in larger teams. Although we do not find a statistically significant difference between the average team size of women (mean size = 9.5) and men (mean size = 10), we include it as a control in our models.

We first examine if the described differences in production, incumbency, investment, status and collegial interactions can explain the differences in our success variables using linear regressions. Specifically, we predict the effect of gender on the logarithm of a user’s average views, likes, and responses while controlling for:

- log of number of shots by the user (productivity),
- log of days since first shot of the user (incumbency),
- whether the user has a pro-badge (investment),
- whether the user self-describes as a leader (status),
- log of the size of team (collegial interactions).

Table 2 shows the results for our models setting the *number of views*, *number of likes* and *number of responses* as dependent variable, respectively. From the first column, we see that being experienced, having a Pro Badge, and having a large team all positively contribute to the average number of views



(a) Distribution of the log of shot count of users by gender. (b) Distribution of the age of user accounts in log days by gender. (c) Shots per day per person over time, 90-day moving average by gender.

Figure 3: Summary statistics comparing male and female activity on the site.

User Averages (log):	Views	Likes	Responses
Log(Number of Shots)	0.09***	0.11***	-0.004
Log(Age of Account)	0.03	-0.15***	0.04**
Leadership Word in Bio	0.05	0.08	0.08*
Pro-badge	0.49***	0.44***	0.29***
Log(Team Size)	0.39***	0.21***	0.002***
Is Male	0.19***	0.10*	0.07**
R^2	0.13	0.12	0.06

Table 2: OLS regressions on user-level variables to predict impact of gender on success. (Note: * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$)

Most Male Skills	Most Female Skills
interfaces	calligraphy
productmanagement	copywriting
objectivec	research
iosdev	information
compositing	handlettering
framer	socialmedia
gui	visualcommunication
ruby	gamedesign
apparel	branddev
illustrator	drawing
identity	ecommerce

Table 3: Ten most male and female self-reported skills.

a user receives on their shots. However even after controlling for these variables, males have significantly more views. Next, we look at the second column which shows results for the average number of likes as the dependent variable. The results look fairly similar, except that having joined the site earlier relates to having more likes. Moreover, being male, while still significant, has a smaller positive effect than in case of the views. Finally, we see a slightly smaller effect of being male for the number of responses. We note that while we find that most of our user variables correlate with success, these models do not predict success very well, achieving low R^2 values.

With our three initial regression models, we have established that there are some robust differences in individual success between men and women. Men receive more views, likes, and responses. The main takeaway from this analysis is that while we see differences between men and women in different measures of success, these are not explained by basic user-level characteristics. In the following sections we investigate two social factors that might contribute to the observed gender discrepancies in success: that *men and women are creating different products* and that *men and women have different social network structures*.

4.2 “Genderness” of Skills and Products

In this section we investigate whether women and men have different specializations and create identifiably different products. If so, differences in outcome may come from differences in the taste or size of audiences. For example we observe that

women are significantly more likely to list “copywriting” as a personal skill. If the audience for copywriting-related work is smaller, or even behaves differently, this may explain differences in outcome.

We implement this idea of “genderness” of user production at two levels. At the user level we quantify the extent to which skills are listed by males and females. At the shot level we train a neural network to identify shots made by males and females visually. We augment that model with data from the tags users give their images.

Skills We calculate the maleness of skills using a log likelihood ratio $L(skill, gender)$:

$$L(skill, gender) = \log \left(\frac{P(skill|gender)}{P(skill)} \right)$$

We test for significance by shuffling the gender of users and calculating a 90% confidence interval. We adopt this shuffling approach in order to preserve the co-occurrence of skills. Out of 150 skills¹ listed by at least 10 users, 38 skills are deemed significantly male-dominated, while 15 are considered significantly female-dominated. We share the most significantly male and female skills in Table 3.

We test the impact of this gender difference by adding two binary variables to our original models: whether the user lists a skill categorized as male-dominated, and whether the user

¹We processed the skills using standard text matching and deduplication techniques.

User Averages (log):	Views	Likes	Responses
Log(Number of Shots)	0.08***	0.11***	-0.002
Log(Age of Account)	0.02	- 0.14***	0.04**
Leadership Word in Bio	0.05	0.08	0.08*
Pro-badge	0.45***	0.43***	0.28***
Log(Team Size)	0.38***	0.32***	0.21***
Is Male	0.14***	0.08	0.06
User Lists Male Skill	0.21***	0.04	0.04
User Lists Female Skill	- 0.25***	- 0.16***	- 0.12*
R^2	0.14	0.13	0.06

(a) The effects of user skill “genderness”.

User Averages (log):	Views	Likes	Responses
Log(Number of Shots)	0.11***	0.34***	-0.004
Log(Age of Account)	0.03	- 0.11***	0.04**
Leadership Word in Bio	0.06	0.09	0.08*
Pro-badge	0.47***	0.42***	0.27***
Log(Team Size)	0.40***	0.34***	0.23***
Is Male	0.06	0.02	0.02
Shot Maleness	1.31***	0.66***	0.50***
R^2	0.14	0.12	0.06

(b) The effects of shot “genderness”.

Table 4: OLS regressions predicting success, controlling for skill (a) and image (b) “genderness”. (Note: * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$)

lists a skill categorized as female-dominated. The first set of regressions shown in Table 4 indicate a penalty for users listing female-dominated skills across all measures of success. Users listing male skills receive more views, suggesting that women who list such skills get similar views to their male counterparts.

Images and tags In this subsection, we build a classifier to predict whether a shot has been made by a male or female user. We consider two aspects of the product: the visual content of the image itself and the tags listed by the user.

As an image classification problem, identification of the gender of the author is clearly distinguished from typical object detection classification tasks in which visual content relates to class labels (e.g. is there a cat in the image?). It also differs substantially from the more refined visual concept detection problems (Huiskes, Thomee, and Lew 2010) like detecting calming or frightening moods in images. Feed-forward deep neural networks are widely used in various image related tasks such as image captioning (Vinyals et al. 2016) and are especially popular for object classification (Krizhevsky, Sutskever, and Hinton 2012; He et al. 2015; Szegedy et al. 2015) and object localization (Girshick 2015; Ren et al. 2015). Due the complexity of the latest models (often containing more than 10 million parameters) the training phase demands a lot of data and computational power. Additionally, almost all feedforward-based discriminative networks have low-dimensional flat layers which can be used as a representation. This led us to use an already tuned model. We chose one of the state-of-the-art feed-forward networks, the Inception v3 (Szegedy et al. 2015). Our choice was mainly based on the quality implementation of the specific model in Tensorflow² and the low dimensional representation (2048) of the images before the final, discriminative layer. To sum up, the neural network generates feature vectors from images, encoding visual data in a way amenable to analysis with a classifier.

Authors tag their shots to help users search for and interpret their work. We consider tags as image related annotations and hence as extra information to aid classification. The sparsity of the tag occurrences keep us from using more complex

²<https://github.com/tensorflow/models/tree/master/inception>

Model	Modality	AUC
Logistic Regression (LR)	Tags	0.70
Inception v3 + LR	Visual	0.62
Inception v3 + LR	Mixed	0.72

Table 5: ROC AUC measurements of gender classification via visual and textual content.

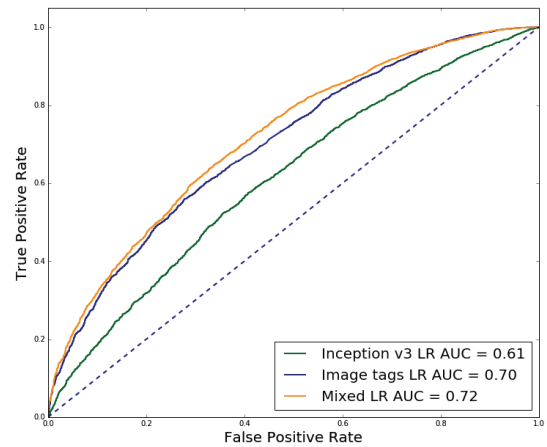


Figure 4: ROC curves of the visual, textual and multimodal image classifiers.

frequency based textual feature generation techniques (e.g. TF-IDF or Okapi BM25 (Robertson and Jones 1976)), thus we extracted the raw image tags and left out both infrequent and author specific tags. The resulting 10,000 dimensional, sparse representation was normalized and combined with the visual representation of the Inception model into a 12,000 dimensional feature space.

We trained three regularized logistic regression models on the image features, the tag features, and all features combined, respectively. We used ten-fold cross-validation (CV) (Kohavi and others 1995) and evaluated the results of the classifiers using the area under the curve (AUC) of the receiver operating characteristic (ROC) averaged across the held-out folds. Our choice of ten-fold CV was driven by the findings

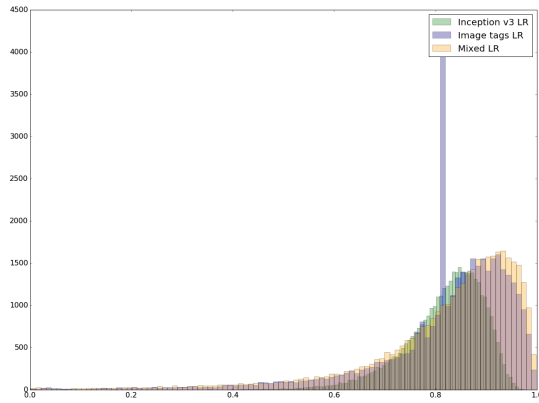


Figure 5: The distributions of the hold-out fold predictions for image maleness for the three classifiers. 1 indicates an image with highly male features. Note that the addition of the image features greatly smooths the distribution of image scores compared to the tag-only model.

in (Ambroise and McLachlan 2002), in which the authors recommended ten-fold cross-validation instead of leave-one-out. The low variance in AUC across the folds for all three modalities (0.01 for the visual-only, 0.006 for image tags and 0.007 for multimodal on average) suggests the results are robust to overfitting. Table 5 shows that the combined multimodal classification model performed the best, while the image tags outperform the visual model. We plot the ROC curves in Figure 4.

We use the score of the multimodal classifier to assign a gender score to each image. We plot the distributions of the classifier outputs in Figure 5, noting that while the increase in AUC of the multimodal classifier over the tag-only classifier is relatively small, it greatly smooths out the predictions of the classifier. The tag-only classifier has difficulty overcoming the sparsity of the tag space.

We incorporate this new shot level measure of gender into our regressions by considering a user’s average shot maleness. The second set of regressions in Table 4 show a significant positive effect of shot maleness on all outcome variables. Moreover, when controlling for this measure the gender of the user is no longer significant for any of the outcomes. Mirroring our findings with the gender of skills, it seems that authors creating more male images, regardless of their gender, are receiving better outcomes.

Summary In this subsection we have presented two ways of quantifying the gender of a user’s output. Both the gender of a user’s skillset and their outputs as defined by our measures have significant relationships with outcomes, indicating that, at least to some extent, differences in gender outcomes on Dribbble are the result of differences in production.

4.3 Social Behavior and Network Position

The most important way for Dribbble users to get informed about new work of others is by following them on the site. Indeed there are two ways a viewer may discover a shot:

User Averages (log):	Views	Likes	Responses
Log(Number of Shots)	-0.55***	-0.50***	-0.41***
Log(Age of Account)	-0.07**	-0.22***	-0.01
Leadership Word in Bio	-0.04	0.001	0.03
Pro-badge	0.03	0.01	-0.01
Log(Team Size)	0.24***	0.19***	0.11***
Is Male	0.04	-0.04	-0.03
Log(Follower Count)	0.72***	0.69***	0.49***
Twitter Reciprocity	-0.04	-0.06	-0.07
Twitter Ego Density	0.26***	0.34***	0.27***
R^2	0.67	0.67	0.50

(a) The effects of network structure on success.

User Averages (log):	Number of Followers
Log(Number of Shots)	0.93***
Log(Age of Account)	0.16***
Leadership Word in Bio	0.13**
Pro-badge	0.66***
Log(Team Size)	0.21***
Is Male	0.21***
R^2	0.58

(b) The effects of user characteristics on follower count.

Table 6: OLS regressions predicting success and number of followers. (Note: * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$)

either the viewer follows the user and the user’s work is then added to the viewer’s chronological feed, or the viewer searches for shots by keyword. The former case is clearly a more dependable source of views, likes, and responses. Thus differences between men and women with regard to success on Dribbble may be partly explained by gender differences in follower networks. Not only the sheer size of the networks (number of followers) may matter for user success, but also network structure: people with more cohesive networks may be more efficient in reaching their audience with new work.

Unfortunately, we do not have access to the full follower network on Dribbble, but we can use the Twitter following network among users as a proxy. This choice is justified by the predictive power of the Dribbble user follower counts and the significant correlation ($\rho = .49$) between Dribbble follower count and Twitter follower count. We therefore consider the effect of two features derived from the Twitter network: the reciprocity of ties of a user and the density of a user’s ego network.³ In the models described in Table 6a we see that the number of followers a user has is a very strong predictor of success. The density of a user’s ego network predicts success as well and the R^2 value of all three models increases drastically over that of previous models. Once controlling for these terms, gender becomes insignificant, yet Table 6b shows that when we use gender to predict the number of followers a user has, we find that males have an

³The reciprocity of a user in a directed network is defined as the ratio of mutual ties to total ties adjacent to a user. A user’s ego network density is defined as number of possible connections between the nodes that the user follows.

advantage despite several controls. This suggests that gender differences in social network position may indeed lie behind success discrepancies between men and women.

Modeling Social Ties To be able to distinguish between the impact of multiple factors shaping the network structure of men and women, we use Exponential Random Graph Models (ERGMs) on the observed Twitter network. ERGMs (Lusher, Koskinen, and Robins 2012) are a multivariate statistical network model family designed exactly for quantifying the contribution of local network configurations (e.g. reciprocated ties, closed triangles) to the global structure of an observed social network, taking into account the inherent dependencies between ties in the network. We estimate ERGMs using the “statnet” package implemented in R (Handcock et al. 2003; Hunter 2007).

Table 7 summarizes our models and findings. The first model includes basic structural effects (such as reciprocity and transitivity) and gender-related sender, receiver, and similarity effects. The second model contains additional effects describing gender differences in following reciprocity and transitivity. The structural effects draw a similar picture of the network in both models: the following network is characterized by positive tendencies for reciprocity and transitivity, while in- and out-degrees, the number of ties received and sent by users, are relatively evenly distributed (as suggested by the negative parameters for degree effects).

With regards to gender differences, our two models suggest slightly different patterns. In the first model, men are significantly less likely to follow other users than women, while same gender ties are more likely to exist than cross-gender ties. However, the second model shows that men are more popular (more likely to receive ties) than women, with a similar tendency for same-sex ties. In this latter model, we also included three effects to further study gender differences, which show that reciprocity is less likely between two men and more likely between two women as compared to mixed gender dyads. Finally, tendencies for transitivity, and thereby clustering, do not seem to differ in same gender triads from mixed gender triads.

The two models together suggest that although men appear to have a larger number of followers than women, female users of Dribbble tend to have a more cohesive network than men. This tendency is expressed in the first model by the negative sender and positive same sex effect: following between users of the same sex is more likely than between those of different sex, but ties between men are less likely to exist than ties between women because of the lower activity of men. This finding is put in a different light in the second model which suggests that while men attract more followers than women, dyads between women are more likely to be reciprocal. Regarding clustering, we do not find evidence that same-gender triads are more likely to exist than mixed-gender triads, pointing out that gender effects may have a more important role on the level of dyads (homophily, reciprocity) than on the level of groups (transitivity).

Strength of Social Ties. As a final test of our theory connecting social network position to Dribbble outcomes, we consider the efficiency of a user, defined as the average num-

Following ties (log odds)		
Edges (Intercept)	-5.263***	-5.682***
Reciprocity (Mutual)	3.110***	3.449***
Transitivity (GWESP, \surd)	0.325***	0.339***
Indegree (GWIn, \surd)	-3.117***	-3.126***
Outdegree (GWOut, \surd)	-2.039***	-2.051***
Sender Sex (Male=1)	-0.169***	-0.037
Receiver Sex (Male=1)	0.012	0.252***
Same Sex	0.191***	0.234***
Reciprocity: Both Female	-	0.367 ⁺
Reciprocity: Both Male	-	-0.641***
Transitivity: Same Sex	-	-0.000
# of Actors	3,765	

Table 7: Exponential Random Graph Models explaining Twitter following networks of users. (Note: ⁺ $p < 0.1$; * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$)

User Average:	Likes per View
Log(Number of Shots)	0.14***
Log(Age of Account)	-0.74***
Leadership Word in Bio	0.13
Pro-badge	-0.08
Log(Team Size)	-0.19***
Is Male	-0.39***
Twitter Reciprocity	-0.14
Twitter Ego Density	0.21
R^2	0.17

Table 8: OLS regression predicting likes per view, controlling for user follower count. (Note: * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$)

ber of likes per view he or she receives, as another dependent variable. We argue this measures how invested the followers of a user are in his or her work, and how effectively the user broadcasts his or her work. Indeed, as seen in Table 8, women have better outcomes according to this measure. We suggest that this is another signal that men and women have structural differences in their social networks and how they navigate them.

5 Discussion

In this paper, we studied the gender differences of individual success in an online community for graphic designers. Our first results showed that men tend to be more successful than women even when controlling for differences including activity, tenure, status markers, and investment, though none of the models created to test these relationships could explain much variation in outcome. Consequently, we turned our attention to how male and female designers may have different audiences for their products or distinct social networks structures, to help better understand the observed results.

When looking at the skills and images users create, we find that there is a significant subset of the top skills that are highly male or female dominated. Using the images and

their descriptive tags we are able to predict the gender of the creator with an AUC of 0.72. Furthermore skill and image “genderness” explain some of the gender differences behind success: women with male skills have similar success rates to males and vice versa. This suggests that part of the gender gap comes from different production patterns.

Finally, analysis on the social network underlying Dribbble shows that the number of followers a user has in the community captures a large part of the variation in individual outcomes. Men produce more work and have more followers, thus they receive more views, likes and responses overall. However, a closer look at the social network structure shows that while men have more followers, women tend to be part of smaller, more densely knit clusters with more reciprocal ties and thus are able to turn their image views into likes and responses more frequently relative to men. As a limitation of our research, we acknowledge that the Twitter follower network might not correspond precisely with the Dribbble social network, and suggest that there is value in extending the data collection effort to the Dribbble network.

Our results demonstrate that there can be multiple potential sources of gender inequalities in online markets. However, this study only puts forth a few simple mechanisms that may be driving these biases. In reality, individual characteristics and social structures are dynamically interrelated: users may adjust their production and self-representation to be able to reach larger audiences, and social relations constantly evolve as a result of users trying to adapt to an ever changing environment. What seems clear, however, is that “gender” remains one of the key categories around which communities produce their understanding of quality and success.

In addition, the described empirical patterns of success also suggest a few lessons for the design of algorithms to present the work of users. Dribbble incorporates common design elements of today’s social and labor market platforms, such as sharing content to social ties, publicly visible feedback/success measures, and content recommendation based on popularity/relevance. For example, if search engine on Dribbble ranks relevant images using views, it indirectly advantages men. If it were to rank images using likes per view, women would be advantaged. This kind of detail is important, given that feedback loops and rich get richer effects can inflate differences in outcomes over time. It certainly merits further investigation by researchers, designers of online platforms, and regulators.

6 Acknowledgments

We thank the anonymous reviewers for their extremely helpful comments. This research was supported in part by XXX. We also thank our colleagues Karl Wachs, Zsófia Czémán, Júlia Koltai, Christo Wilson, and Roberta Sinatra for helpful input and feedback.

References

Adali, S.; Escrava, R.; Goldberg, M. K.; Hayvanovych, M.; Magdon-Ismail, M.; Szymanski, B. K.; Wallace, W. A.; and Williams, G. 2010. Measuring behavioral trust in social networks. In *Intelligence and Security Informatics (ISI), 2010 IEEE International Conference on*, 150–152. IEEE.

Ajrouch, K. J.; Blandon, A. Y.; and Antonucci, T. C. 2005. Social networks among men and women: The effects of age and socioeconomic status. *The Journals of Gerontology Series B: Psychological Sciences and Social Sciences* 60(6):S311–S317.

Ambroise, C., and McLachlan, G. J. 2002. Selection bias in gene extraction on the basis of microarray gene-expression data. *Proceedings of the national academy of sciences* 99(10):6562–6566.

Angrist, J. D., and Lang, K. 2004. Does school integration generate peer effects? evidence from boston’s metco program. *American Economic Review* 1613–1634.

Cartwright, D., and Harary, F. 1956. Structural balance: a generalization of heider’s theory. *Psychological review* 63(5):277.

Using publicly available information to proxy for unidentified race and ethnicity. Consumer Financial Protection Bureau. <https://www2.census.gov/topics/genealogy/2000surnames/surnames.pdf>.

Clauset, A.; Arbesman, S.; and Larremore, D. B. 2015. Systematic inequality and hierarchy in faculty hiring networks. *Science Advances* 1(1):e1400005.

Coleman, J. S. 1994. *Foundations of Social Theory*. Harvard University Press.

Curcio, G.; Ferrara, M.; and De Gennaro, L. 2006. Sleep loss, learning capacity and academic performance. *Sleep medicine reviews* 10(5):323–337.

Davis, J. A. 1970. Clustering and hierarchy in interpersonal relations: Testing two graph theoretical models on 742 sociomatrices. *American Sociological Review* 35(5):843–851.

De Vaan, M.; Stark, D.; and Vedres, B. 2015. Game changer: The topology of creativity1. *American Journal of Sociology* 120(4):1144–1194.

Feld, S. L. 1982. Social structural determinants of similarity among associates. *American Sociological Review* 47(6):797–801.

Fiscella, K., and Fremont, A. M. 2006. Use of geocoding and surname analysis to estimate race and ethnicity. *Health Serv Res*.

Fradkin, A.; Grewal, E.; Holtz, D.; and Pearson, M. 2015. Bias and reciprocity in online reviews: Evidence from field experiments on airbnb. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*.

Ge, Y.; Knittel, C. R.; MacKenzie, D.; and Zoepf, S. 2016. Racial and gender discrimination in transportation network companies. Working Paper 22776, National Bureau of Economic Research.

Girshick, R. 2015. Fast r-cnn. In *Proceedings of the IEEE International Conference on Computer Vision*, 1440–1448.

Handcock, M. S.; Hunter, D. R.; Butts, C. T.; Goodreau, S. M.; and Morris, M. 2003. *statnet: Software tools for the Statistical Modeling of Network Data*. Seattle, WA.

Hannak, A.; Margolin, D.; Keegan, B.; and Weber, I. 2014. Get back! you don’t know me like that: The social mediation of fact checking interventions in twitter conversations. In *ICWSM*.

Hannak, A.; Wagner, C.; Garcia, D.; Mislove, A.; Strohmaier, M.; and Wilson, C. 2017. Bias in Online Freelance Marketplaces: Evidence from TaskRabbit and Fiverr. In *20th ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW 2017)*.

He, K.; Zhang, X.; Ren, S.; and Sun, J. 2015. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*.

- Heider, F. 1946. Attitudes and cognitive organization. *The Journal of Psychology* 21(1):107–112.
- Huiskes, M. J.; Thomee, B.; and Lew, M. S. 2010. New trends and ideas in visual concept detection: the mir flickr retrieval evaluation initiative. In *Proceedings of the international conference on Multimedia information retrieval*, 527–536. ACM.
- Hunter, D. R. 2007. Curved exponential family models for social networks. *Social networks* 29(2):216–230.
- Jacobs, J. A. 1996. Gender inequality and higher education. *Annual Review of Sociology* 153–185.
- Karimi, F.; Wagner, C.; Lemmerich, F.; Jadidi, M.; and Strohmaier, M. 2016. Inferring gender from names on the web: A comparative evaluation of gender detection methods. WWW '16 Companion.
- Kohavi, R., et al. 1995. A study of cross-validation and bootstrap for accuracy estimation and model selection. In *Ijcai*, volume 14, 1137–1145. Stanford, CA.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, 1097–1105.
- Lee, C. J.; Sugimoto, C. R.; Zhang, G.; and Cronin, B. 2013. Bias in peer review. *Journal of the American Society for Information Science and Technology* 64(1):2–17.
- Lee, M. K.; Kusbit, D.; Metsky, E.; and Dabbish, L. 2015a. Working with machines: The impact of algorithmic and data-driven management on human workers. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*.
- Lee, M. K.; Kusbit, D.; Metsky, E.; and Dabbish, L. 2015b. Working with machines: The impact of algorithmic and data-driven management on human workers. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*.
- Liu, W., and Ruths, D. 2013. Whats in a name? using first names as features for gender inference in twitter.
- Lusher, D.; Koskinen, J.; and Robins, G. 2012. *Exponential Random Graph Models for Social Networks: Theory, Methods, and Applications*. Cambridge University Press.
- Lustig, C.; Pine, K.; Nardi, B.; Irani, L.; Lee, M. K.; Nafus, D.; and Sandvig, C. 2016. Algorithmic authority: the ethics, politics, and economics of algorithms that interpret, decide, and manage. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, 1057–1062. ACM.
- Marsden, P. V., and Friedkin, N. E. 1993. Network studies of social influence. *Sociological Methods & Research* 22(1):127–151.
- McPherson, M.; Smith-Lovin, L.; and Cook, J. M. 2001. Birds of a feather: Homophily in social networks. *Annual Review of Sociology* 27:415–444.
- Munger, K. 2016. Tweetment effects on the tweeted: Experimentally reducing racist harassment.
- Pager, D., and Shepherd, H. 2008. The sociology of discrimination: Racial discrimination in employment, housing, credit, and consumer markets. *Annual review of sociology* 34:181.
- Pan, B.; Hembrooke, H.; Joachims, T.; Lorigo, L.; Gay, G.; and Granka, L. 2007. In Google We Trust: Users' Decisions on Rank, Position, and Relevance. *J. Comp. Med. Comm.* 12(3).
- Petersen, A. M.; Fortunato, S.; Pan, R. K.; Kaski, K.; Penner, O.; Rungi, A.; Riccaboni, M.; Stanley, H. E.; and Pammolli, F. 2014. Reputation and impact in academic careers. *Proceedings of the National Academy of Sciences* 111(43):15316–15321.
- Petersen, A. M. 2015. Quantifying the impact of weak, strong, and super ties in scientific careers. *Proceedings of the National Academy of Sciences* 112(34):E4671–E4680.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Neural Information Processing Systems (NIPS), 2015*.
- Robertson, S. E., and Jones, K. S. 1976. Relevance weighting of search terms. *Journal of the American Society for Information science* 27(3):129–146.
- Robinson, L.; Cotten, S. R.; Ono, H.; Quan-Haase, A.; Mesch, G.; Chen, W.; Schulz, J.; Hale, T. M.; and Stern, M. J. 2015. Digital inequalities and why they matter. *Information, Communication & Society* 18(5):569–582.
- Sandvig, C.; Hamilton, K.; Karahalios, K.; and Langbort, C. 2014. Auditing algorithms: Research methods for detecting discrimination on internet platforms. In *Proceedings of "Data and Discrimination: Converting Critical Concerns into Productive Inquiry", a preconference at the 64th Annual Meeting of the International Communication Association*.
- Sarigl, E.; Pfitzner, R.; Scholtes, I.; Garas, A.; and Schweitzer, F. 2014. Predicting scientific success based on coauthorship networks. *EPJ Data Science*.
- Servia-Rodríguez, S.; Noulas, A.; Mascolo, C.; Fernández-Vilas, A.; and Díaz-Redondo, R. P. 2015. The evolution of your success lies at the centre of your co-authorship network. *PLoS one* 10(3):e0114302.
2016. Baby names from social security card applications-national level data. data.gov. <https://catalog.data.gov/dataset/baby-names-from-social-security-card-applications-national-level-data>.
- Szegedy, C.; Vanhoucke, V.; Ioffe, S.; Shlens, J.; and Wojna, Z. 2015. Rethinking the inception architecture for computer vision. *arXiv preprint arXiv:1512.00567*.
- Tang, C.; Ross, K. W.; Saxena, N.; and Chen, R. 2011. What's in a name: A study of names, gender inference, and gender behavior in facebook. In *DASFAA Workshops*.
- Teodoro, R.; Ozturk, P.; Naaman, M.; Mason, W.; and Lindqvist, J. 2014. The motivations and experiences of the on-demand mobile workforce. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work and Social Computing*.
- Thebault-Spieker, J.; Terveen, L. G.; and Hecht, B. 2015. Avoiding the south side and the suburbs: The geography of mobile crowdsourcing markets. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work and Social Computing*.
- Turner, J. C. 1991. *Social influence*. Milton Keynes, UK: Open University Press.
- Vasilescu, B.; Capiluppi, A.; and Serebrenik, A. 2013. Gender, representation and online participation: A quantitative study. *Interacting with Computers* iwt047.
- Veenstra, R., and Dijkstra, J. K. 2011. Transformations in adolescent peer networks. *Relationship pathways: From adolescence to young adulthood* 135–154.
- Vinyals, O.; Toshev, A.; Bengio, S.; and Erhan, D. 2016. Show and tell: Lessons learned from the 2015 mscoco image captioning challenge. *IEEE transactions on pattern analysis and machine intelligence*.