Effects of Attention and Recognition on Engagement, Content Creation and Sharing: Experimental Evidence from an Image Sharing Social Network

Justin T. Huang and Sridhar Narayanan*

December 2020

Abstract

In this study, we examine the impacts of attention and recognition received by a user's content on a social network on that user's subsequent engagement on the network, content creation and content sharing. The study of the impact of attention and recognition is typically challenging because they are not randomly assigned. Systematic differences within and across users in the degree of attention and recognition received by content shared by them makes the identification of effects difficult. To solve this identification problem, we implemented a field experiment in collaboration with an art-sharing social network, where we experimentally manipulated attention and recognition by selectively featuring users' content. A unique aspect of our experimental context is that we are able to observe both on-network and off-network activity of the individuals concerned. The main results of our experiment are that our manipulation shifting attention and recognition on the network increases engagement, tie-formation, posting of creative output and the usage of underlying software tools used to create content. We explore the temporal variation, heterogeneity, and mediation in these effects.

Keywords: User Generated Content, Attention effects, Recognition Effects, Social Networks, Causal Effects, Field Experiments

*Justin is an Assistant Professor of Marketing at the Ross School of Business at the University of Michigan and Sridhar is an Associate Professor of Marketing at the Graduate School of Business at Stanford University. They can be reached at huangjt@umich.edu and sridhar.narayanan@stanford.edu, respectively. The authors would like to thank participants at the UT Dallas FORMS Conference 2016, Marketing Science Conference 2016, McGill Marketing Camp 2016, University of Iowa Marketing Camp 2017, CMU Digital Marketing Conference 2018, London Business School Marketing Seminar, Marketing Science Doctoral Consortium 2019, Marketing Science Conference 2020 and the Stanford Marketing Work-in-Progress Seminar Series for their comments. They are also grateful to Yesim Orhun for her insightful comments. All errors are their own.

1 Introduction

Social media usage has exploded in the past decade. There were estimated to be about 3.6 billion regular users of social networks worldwide in the second quarter of 2020.[1] The most prominent social network, *Facebook*, has expanded from 6 million monthly active users in 2005 to nearly 2.7 billion users by the second quarter of 2020[2]. Unlike traditional media outlets, social media networks rely on their user base to generate the content that other users consume on the network. This content creation and sharing is key to user engagement on the network, and consequently to advertising and other revenue streams for the networks. Yet, the motivations and drivers of social media content creation and sharing has been underexplored, largely due to the endogeneity of content sharing behavior, the consequent challenges in obtaining causal estimates using observational data, and the difficulty in implementing experiments on large social networks. In this paper, we examine two factors that potentially motivate users to share content on social network - the attention and recognition they receive from other users on the network when their content is displayed by the network. We present the results of an experiment that we conducted in cooperation with an art-sharing social network called *Behance*, which is owned and operated by *Adobe Systems, Inc.* In the experiment, we utilized a key lever available to social networks to direct user attention and recognition to content - the featuring of content on the website. Thus, we generate exogenous variation in the attention and recognition received by different user generated content, and this allows us to examine the effects of attention and recognition on subsequent content creation and engagement of the users. This study is unique in that the experimental design allows for direct causal estimation of the effect of receiving attention and recognition from a social network on a user's subsequent behavior. A further unique feature of our experimental context is that we are able to observe not just the activity of the user on the network itself, but also on the software tools used to create the content that is posted on the network. We are thus able to obtain a very detailed view into the effects of attention and recognition on the network on users' subsequent content creation, sharing and engagement on a social network.

Attention and recognition are important motivators for behavior, and in particular creating and sharing content with others. Lerner and Tirole [2002] identified peer attention as one of the important benefits that contributors derive when sharing content that does not generate immediate monetary rewards. This can be a purely intrinsic payoff, in that it increases the satisfaction and self worth that the contributor experiences from sharing content, and can also have an extrinsic component through the potential to signal the contributor's abilities to potential future employers. The positive affect that attention by peers may generate can increase creativity and creative output (Isen et al. 1987, Baas et al. 2008). Additionally, when the content, and thereby the contributor have the attention of others, the marginal value of future contributions can be higher because of

[1] Source: https://www-statista-com.stanford.idm.oclc.org/statistics/278414/number-of-worldwide-social-network-users/ last accessed 30 Sep 2020

[2] https://www-statista-com.stanford.idm.oclc.org/statistics/264810/number-of-monthly-active-facebook-users-worldwide/ last accessed 30 Sep 2020.

the potential for greater attention on the network for these future contributions. These mechanisms suggest increasing attention may increase future content generation and sharing. Likewise, recognition has been studied as a major motivator for effort in a variety of contexts, including contribution to open source software (Wu, Gerlach, Young 2013), online question-and-answer communities (Jin, Li, Zhong, Zhai 2015), and salesforce motivation (Larkin 2011).

On the other hand, both attention and recognition are extrinsic motivators which could have a crowding out effect on intrinsic desire to perform creative or effortful tasks (Frey and Jegen 2001, Titmuss 1970). Further, if there are diminishing marginal returns to attention and recognition (Toubia and Stephen 2013), greater attention and recognition in the form of increased followers can decrease the marginal benefits from future contributions. Finally, increasing attention may increase the stakes for the user in terms of the quality of subsequent content they produce and share, and this might decrease the volume of content contributed in the future in the presence of quantity vs. quality tradeoffs. Thus, the overall effect of attention and recognition on future content sharing of the user is ambiguous and warrants further empirical examination.

While the literature has tended to study attention and recognition separately, the distinction is less relevant in social media contexts where the two cannot be easily manipulated independently. Common levers available to the platform to highlight a user's contribution, such as front page featuring or increasing visibility on search rankings, direct both peer attention in terms of views and peer recognition in the form of likes and comments. Indeed, it is hard to argue in our context - or even more generally - that public recognition can be given without generating attention. Likewise, directing attention to content carries implicit recognition from the platform and in many cases may lead some members of the augmented audience to endorse the content where they would be otherwise unaware of its existence. Lastly, many online content creators perceive audience size (attention) to be a form of recognition in itself - for instance, being shown organically on the front page or reaching 1 million followers are considered major milestones and commonly celebrated by social media influencers. This linkage is reinforced by platform notifications (Huang et al. 2019) or even physical trophies[3] that recognize content creators for achieving audience size thresholds.

The effects of attention and recognition are hard to measure causally because they are typically endogenously assigned. A cross-sectional analysis of these effects across users is confounded by the systematic differences across users. For instance, users who are more prolific in creation of content are likely to also receive more attention and recognition. Thus, correlations between the two and subsequent content creation and sharing may merely reflect the baseline differences in rates of content creation and sharing of users who receive more vs. less attention and recognition rather than a causal effect of attention and recognition on subsequent user behavior. The availability of panel data would not in general solve this problem because of the potential for time varying correlations between attention and recognition and content sharing. For instance, if users have sprees of content sharing, they are likely to generate and share more content during those sprees, and also receive more attention and recognition during those sprees. As an example, users on Facebook

[3]https://www.youtube.com/creators/awards/ Last accessed 01 December 2020.

may share more content when they are on vacation, and might also generate more attention and recognition for their content during that period. Thus, positive correlations between attention and recognition and subsequent content sharing may reflect those content sharing sprees rather than a causal relationship between these variables. Conversely, it is also possible that content creation is an inherently effort-inducing process, with creators sharing content and taking breaks between sharing episodes. Their content may then generate attention and recognition during the period between their sharing activities, and thus, one may see a negative correlation between sharing and attention and recognition even in the absence of any causal relationships between these variables. Thus, the causal effects of attention and recognition on subsequent content creation are hard to estimate using observational data. We solve this problem by conducting a field experiment on a social network, where we exogenously manipulate the degree to which specific pieces of content are promoted on the platform and thereby directed attention and recognition. By randomly assigning different pieces of content to treatment and control groups, we are able to obtain causal effects of attention and recognition on users' subsequent behavior.

Next we describe the field setting for our experiment, *Behance*, and discuss how it is a good context to study the effects of attention and recognition on content creation and sharing by users. *Behance* is an online social network for artists to share their work with peers in the form of digital albums (termed 'Projects') that contain one or more images. These projects can contain a variety of artistic media, including but not limited to photography, digital illustration, or animation. Users on the network form unidirectional following relationships with other users, and can view, like (termed 'Appreciate' on *Behance*) and comment on projects shared by other users. We exploit an important aspect of the network to facilitate discovery of, direct attention to, and recognize creative content by users on the network at large - the featuring of curated content on the main homepage of the network and in other places throughout the website.

Our experiment takes a subset of the users who own a project in the feature queue (an internal tool listing projects identified by curators to be of sufficient quality to be featured) and randomly assigns them into the treatment or control group. For each user, we randomly select one of their projects in the queue and assign it to be featured or not featured based on whether the user was in the treatment or control group, respectively. Our manipulation controls the timing of featuring to guarantee that the projects in the treatment group are featured at some time during the experimental period of approximately two months, whereas the projects in the control group are guaranteed to not be featured during this period. When a project is featured, it is prominently displayed on the website, and followers of the featured user are notified that the project was featured. This in turn leads to increased peer attention and recognition to the featured content from the network in the form of views and appreciations left by users. Thus, we manipulate attention and recognition that are otherwise hard to study observationally in real-world social networks, and our manipulation utilizes a lever common to a variety of social networks (including *Flickr*, *YouTube*, and *Twitter*) that use front pages and curated lists to promote content.

A unique aspect of our experimental context is that we are able to link the social network accounts for a subset of the users to the accounts used for the cloud-connected software utilities used to create content that is shared on the network. When users use these software utilities to create content, their activity on the software is logged due to the cloud-connected nature of the *Adobe Creative Cloud* software utilities. This gives us a unique insight into the off-network content-creation activities of these individuals, going beyond their activities on the network itself. Sharing of content on the network when consumers have the attention of others may simply result from the user's incentive to push out previously produced content at a time when others are likely to see it. Secondly, if a user has heightened activity on the network, it might lead to lower activity on content creation itself. Since the Creative Cloud suite of products are the primary software tools used for content creation and editing, we are able to get a more complete picture than a look at on-network activity alone would provide. This is enabled by our unique dataset linking on-network and off-network activity for users.

From our experimental analysis, we find that increased attention and recognition as a result of featuring on the network have a positive causal effect on subsequent content creation. Treated users were more likely to engage with other users and post content on the network after treatment. What is interesting is that there is a positive effect not just on the posting of content, but also on content creation, as measured by activity in the *Adobe Creative Cloud* software that are linked with *Behance* network accounts. These lifts in engagement and content creation are large and significant, yet attenuate more rapidly than the lift in peer attention and recognition from being featured on the network. This suggests that users may become accustomed to a new, higher level of attention and recognition while their activity reverts to a previous baseline. We also examine the effects of treatment on tie creation and find that it leads to creation of both inbound and outbound ties. While it is intuitive that attention and recognition lead to more inbound ties (other individuals following the treated individual), the fact that it leads to more outbound ties is somewhat surprising, and could result from both greater engagement on the network by the featured individual, as well as the increased level of inbound tie creation leading to reciprocal outbound tie creation. Overall, we find that users receiving attention and recognition for one of their projects engage more with the network subsequently, share more content and spend more time on content creation. We also find that these effects are heterogeneous, namely that users who have received higher past recognition (in the form of appreciations) were less affected by the treatment than users with lower past recognition. Further, through our mediation analysis, we rule out that featuring or potential satisfaction derived from the website's endosement of the user's content was the driving factor behind these changes in user behavior. The changes in user behavior, including engagement, content creation, and tie formation, were completely mediated by the increased attention and recognition received from peers as a result of our manipulation.

The rest of the paper proceeds as follows: Section 2 provides a background on the related literature, and provides details on the institutional setup of the *Behance* social network and *Adobe Creative Cloud* suite of software utilities. Section 3 describes the experimental design, procedure, and describes randomization and manipulation checks we conducted. Section 4 presents our main results - the effect of featuring on both the project owner's network engagement, content creation, and usage of the underlying creative tools and mediation analysis to support the role of attention and recognition in driving these changes in user behavior. Section 5 concludes.

2 Background

2.1 Relationship to Prior Literature

Social networks such as *Facebook*, *Twitter* and *Behance*, the last one being the network where we conduct our study, depend on user generated content. The ability of a network to develop user generated content is crucial to attracting an audience (and vice versa), with profound implications for the network's evolution and success. A number of studies have examined this both theoretically (see for instance Zhang and Sarvary 2014, Iyer and Katona 2016) and empirically (Dae-Yong Ahn and Mela 2016). Relatedly, the development of open source software depends on contributions from users. In these contexts, users contribute content without any expectation of immediate financial reward, and hence the question of why users contribute content has been the focus of study across multiple streams of literature. A seminal paper in the economics literature on open source is Lerner and Tirole [2002], which describes the immediate and delayed costs and benefits to contributors, and the consequent economic incentives for them to work on contributions that do not provide them immediate monetary benefits. The contributors incurs an immediate opportunity cost of time that could be spent on other activities that provide monetary compensation. Furthermore, time spent on the open source contributions could be spent on activities that advance the person's opportunities to earn in the future. On the benefits side, the paper talks about immediate benefits to the contributor, including the opportunity to learn from this unpaid contribution and thereby increase monetary compensation on other paid work by the same user, and the psychological benefits that accrue from working on a 'cool' project. The long term benefits include the ability to advance one's career prospects and future monetary gains through future job offers, shares of future commercial exploitation of the open source content, and access to capital for the contributors' ventures in the future. A second type of long-term benefit comes from the individual's desire for attention and recognition from peers.

A number of subsequent studies have empirically examined contributors' motivation to share content in contexts that do not involve immediate monetary reward, such as open source software, social networks and knowledge repositories such as Wikipedia, amongst others. One stream of literature has used survey data to explore the reasons why users contribute to open source software projects and in discussion communities related to them (see Hars and Ou 2002, Kankanhalli et al. 2005, Wasko and Faraj 2005, Lakhani and Wolf 2005, Nov 2007, Oreg and Nov 2008, Zhang et al. 2013). The literature documents the presence of both intrinsic and extrinsic motivators (Deci and Ryan 1985) for creating and sharing content. However, the studies in this literature are subject to the critique that there is a difference between the stated motivations of the contributors, and the motivations revealed by their actual behaviors. Further, the extent of the relationship between attention or recognition and future behavior of contributing users is not studied in this literature. We address these issues by implementing a field experiment with revealed behaviors and quantify the relationship between attention or recognition and subsequent creation.

Other studies such as Thorsten Hennig-Thurau and Gremler [2004], Ghose and Han [2011], Moqri et al. [2018] and Zeng and Wei [2012] have taken an observational approach to study content creation. They utilize longitudinal observational data, and econometric panel data analysis methods to study motivations for content creation, but are subject to critiques on the causal inferences that can be drawn when using observational data.

More closely related to this paper are a few streams of literature that attempt to causally estimate the effects of underlying motivators for content creation. For instance, Shriver et al. [2013] study the effect of network ties on content sharing and vice versa via natural variation in wind conditions on a social network for wind surfers. Similarly, in a study of Wikipedia in China, Zhang and Zhu [2011] utilize the blocking of the Chinese language version of Wikipedia in mainland China as a natural experiment to study the effects of audience size on contributions to the platform by contributors located outside mainland China, who continued to have access to it. These studies document the positive relationship between network ties or audience size and contributions, supporting the idea that contributors are motivated to share content because of the social benefits they derive from the usage of that content by others. However, in both of these studies, the exogenous variation is a source outside the control of the firm, so actionable implications for managers of these platforms are limited. Further, the shock in the study on Chinese wikipedia was an unusually large one, in that a large majority of the audience disappeared overnight. Our study, by contrast, looks at a lever to direct peers' attention and recognition that is within the control of the network and typical to the daily consumer experience.

Kummer [2015] studies attention effects in the German Wikipedia platform, again using a natural experiment. This study uses exogenous shocks to attention that result from attention generated by neighboring articles. These shocks include large newsworthy events that increase attention to a particular page on Wikipedia and the page being chosen to be featured on the front page of the platform. The exogeneity comes from the fact that there are spillovers from the pages that experience attention shocks to other pages. These other pages thus experience an exogenous increase in attention. The study finds small effects of such increased attention on further content creation on these pages. These findings are reinforced by Aaltonen and Seiler [2015], who find that within articles on Wikipedia, past attention leads to greater current attention and thereby content growth. While these studies, like ours, examine the relationship between attention and content creation, they do so in a context where the motivations to share content in response to increased attention are presumably weaker (the articles' authorship is shared by multiple contributors, and the identities of contributors are not easily observable to viewers). Also, a critique of Kummer [2015] is that the exogenous news spike or featuring on the front page of Wikipedia may directly affect attention and content creation for the pages being studied. By contrast, we experimentally manipulate attention and recognition directly and in a meaningful fashion for the content being analyzed, and hence have clean causal estimates of the effects of attention and recognition on future behavior of the featured users.

Also closely related to our study is Toubia and Stephen [2013], which examines the non-financial motivators for sharing of content by users on *Twitter*, specifically focusing on the intrinsic vs. image-related utilities for sharing on the platform. The main idea in this paper is that intrinsic utility, which is the utility derived from the viewership that the content receives from the user's followers, requires that users post content in order to be seen and heard; by contrast, image-related utility, which is the utility derived from being seen as having many followers, is incurred regardless of whether a user posts or not. This distinction in the effect of number of followers is utilized to set up an experiment where treatment group users are followed by fake accounts created by the experimenters, and are compared to a control group of users that are not followed by the fake accounts. The main finding of this paper is that while both sets of motivators for content creation are present, the image-related motivators dominate for most users of the platform. Our study differs from this study in some important respects. Notably, we manipulate real attention and recognition received from peers on the social network rather than the number of followers as the main manipulation. Our manipulation of featuring content and thereby directing peer attention and recognition is a common policy lever utilized by social networks, while to our knowledge major social networks do not create fake accounts to motivate user activity. Second, the addition of fake followers directly changes the network structure, and hereby may not be consistent with users' existing beliefs about the relationship between their follower count and the level of attention and recognition they should be receiving (fake followers contribute to follower counts, but never retweet or like experimental users' posts). By contrast, we manipulate attention in a belief-consistent manner, in that both control group and treatment group content in our experiment come from the same list of content to be featured. Users do not know that their content is on the list, and do not have expectations of when it would be featured. We merely change the timing of featuring to achieve our experimental manipulation. Further, our study directly measures the effects of motivators for content creation, rather than in that study's case, depending on predictions about the relationships between the number of followers and motivation to create content that rely on underlying assumptions.

Finally, a few field experiments have examined related questions in the context of online content generation and consumption. Salganik et al. [2006] and Muchnik et al. [2013] study the role of signals of peer attention in driving content consumption in music market and news aggregator contexts, respectively. In both of these randomized field studies, the authors document the self-reinforcing nature of attention, finding that information on peers' interaction with content spur other consumers to also interact with it. Huang et al. [2019] conducted a series of studies, including a large scale field experiment, to examine how feedback message framings influence content creation, and highlight competition amongst creators as a novel motivator. Other studies have looked at financial incentives (Cabral and Li 2015), social comparisons (Chen et al. 2010) as motivators for content creation by users.

This, to the best of our knowledge, is the first study content creation by directly manipulating the levels of attention and recognition received through a field experiment intervention on a large social network. Our experimental manipulation involving featuring directly changes the amount of attention and recognition a given piece of user generated content receives on the social network, and does this in a belief-consistent manner. Unlike prior studies that either used natural experiments involving belief-changing large shocks to the number of users on the network as a whole, or added fake followers to the treatment group users, our experiment involves a manipulation that is belief-consistent. We do not feature content and increase attention and recognition for content that would otherwise not be featured, or remove attention or recognition from content that would otherwise be featured. We merely change the order in which content is featured, where consumers do not know the order, or even the fact that the content was in the list to be featured. A second contribution of this study is that it is the first one to examine not just users' activity on the network, but also off-network on the software tools used to create the content shared on the network. This lets us more comprehensively examine the effects of attention and recognition not just on sharing, but also content generation efforts outside the network. In this study, we also examine temporal variation in the effects of attention and recognition - specifically whether and at what rate they attenuate over time, and heterogeneity in these effects across users with a varying degrees of past recognition on the network. Thus, we comprehensively and causally examine an important lever available to social networks to motivate the creation and sharing of user generated content.

2.2 The *Behance* image sharing social network

The setting for the experiment is the social network called *Behance*. *Behance* is an online portfolio sharing network for creative amateurs and professionals in the broad area of digital art. This US-based social network was founded in 2006 and has since acquired over 5 million users who collectively posted over 9 million projects. The social network was acquired by *Adobe Systems, Inc.* in 2012, roughly coinciding with the transition of *Adobe*'s suite of software utilities targeted at creative professionals to the *Adobe Creative Cloud* (moving customers from single-purchase software licenses to a cloud-based subscription model for creative tools such as *Photoshop*, *Lightroom* and *Illustrator*). This integration allowed *Behance* and *Adobe Creative Cloud* (henceforth referred to as *Creative Cloud*) accounts to share a common identifier and easier uploading from within an *Creative Cloud* application straight onto a *Behance* profile. As a result of this integration, we have the unique ability to directly link *Creative Cloud* product usage and purchases to a *Behance* social network account for a subset of users in our sample[4].

The core of the *Behance* user experience is the creation of an individual's portfolio profile page. Profiles may include a photograph of the creative professional as well as details on the professional's focus areas, education, employment, and other social media links (*Facebook*, *Twitter*, *Instagram, LinkedIn*, etc). Users may also upload projects, which are collections of digital images that typically share a common theme. The array of projects uploaded to the site is quite diverse,

[4]Users were not forced to link the accounts for *Creative Cloud* and *Behance* but had incentives to do so, and a subset of users chose to link their accounts.

spanning photography, illustrations, architecture, sculpture, animations, 3-D modeling, advertising, and branding. These projects have often been made using *Creative Cloud* products, and users have the ability to tag their works with the creative tools used to create them.

Users on the website connect with other users by forming outbound ties or 'following' relationships. A following relationship is a unidirectional tie initiated by the following user in order to receive updates on the activities of the followed user on *Behance*. Thus, in a similar way to *Twitter* and other social networks, the presence of many inbound ties or 'followers' is seen as a measure of popularity and influence on the network. Users interact with works and each other in a variety of ways, including viewing, appreciating and commenting on projects and sending direct messages to each other. The activities of uploading a new project and appreciating or commenting on a project are pushed out as news stories to all of a user's followers and appears on a chronological *Activity* feed.

In addition to social network promotion of projects on the site, *Behance* has another mechanism of filtering and featuring quality works. *Behance* has a team of curators, who are employed by it to review all new works on the site and create a list of exemplary projects to be featured. The network features approximately 50 projects per day across a variety of focus areas (Branding, Illustration, Photography, etc). These featured works push a news story to all of their owners' followers that the work has been featured, are displayed more prominently in the *Activity* and *Discover* sections of the website, and are distinguished with a star-shaped badge marking their excellence and the date the project was featured. An example of a featured work can be seen in figure 1, and display of featured works in the *Activity feed* can be seen in figure 2. Our experiment utilizes this curated featuring mechanism to create exogeneous variation in attention and recogntion for estimation of effects of attention and recognition on user behavior and creative output.

2.3 *Creative Cloud* suite of software products

A particularly unique feature of our experiment is that due to the ownership of *Behance* by *Adobe Systems, Inc.*, we are able to obtain data not just for activities on the network itself, but also in the suite of cloud-based software products sold by *Adobe* referred to as *Creative Cloud*. *Creative Cloud* is a subscription-based software service that encompasses 11 different creative products, including *Photoshop, Lightroom, InDesign, Illustrator, Acrobat, After Effects, Premiere Pro, Audition, Flash Pro, Dreamweaver, Adobe Muse,* and *InCopy*. These products include creative tools that allow users to design and edit documents, images, animations, videos, and websites. Users can purchase subscription licenses online at a rate of $9.99 per month for *Photoshop* and *Lightroom*, $19.99 per month for any other single application license, or $49.99 per month for a package that includes all 11 products. Each subscription to *Creative Cloud* allow users to take advantage of cloud computing performance enhancements, text reading and rendering tools, a marketplace of stock images, mobile and tablet-based supplementary applications, and cloud storage that allows content to be synchronized across devices. According to 2018 *Adobe* investor releases, *Creative Cloud* services accounts for over $5 billion in annualized recurring revenue.

Following the acquisition of *Behance* by *Adobe* in late 2012, it became possible for users on *Behance* to link their accounts with *Creative Cloud*. This allows the linkage of *Behance* social network activity with *Creative Cloud* usage. We observe login counts, application launches, application session lengths, purchases and cancellations within *Creative Cloud*.

3 Data & Experiment Details

We aim to measure the causal effect of attention and recognition received by a user's content on subsequent behavior by that user both on the *Behance* social network and within *Creative Cloud*. To achieve this goal, we implement a randomized controlled experiment which adds exogeneous variation to the featuring process and thereby shifting attention and recognition, and then compare treatment and control groups across a variety of metrics that capture user creative production, level of engagement and attention on the network, and tie formation. The fully randomized design avoids several potential identification issues that would be present in an observational study, including correlation between level of attention and quality of creative output, endogeneous tie formation, and endogeneity in network actions.

3.1 Selection and Randomization

We collaborated with the *Behance* curation team in order to implement this experiment featuring projects on the *Behance* social network. At the start of the experiment, a total of 8,921 projects collectively owned by 3,625 users were identified by the curation team to meet a high standard of quality, creativity, and uniqueness. These works were selected from 5 creative focus areas (Illustration, Photography, Branding, Typography, and Web Design) and were chosen in part because of their alignment with the *Creative Cloud* suite of products[5]. The randomization proceeded as follows: from the 3,625 eligible users, we randomly selected 329 users for each of the treatment and control groups. For each user in the treatment and control group, we then randomly selected one of their projects that was in the 8,921 projects that met the curation team's quality standard. By this process, we collected treatment and control groups of 329 projects each, for a final count of 658 user-projects in the experiment.

3.2 Manipulation

Over the experimental period of August 4, 2015 through September 30, 2015, projects sorted into the treatment group via the aforementioned process were regularly featured on the *Behance* social network at a rate of roughly 6 per day. Thus, the treatment group projects were pushed towards the front of the featuring queue at the rate of roughly 6 randomly selected projects per day. Control group projects were pushed backwards sufficiently in the queue such that they would be guaranteed not to be featured during the experimental period until after October 21, 2015. Thus, nothing is

[5]An implicit motivation behind *Adobe*'s acquisition of the *Behance* social network was to promote the use of the *Creative Cloud* products

altered as a result of the experiment other than the timing of the featuring for the experimental projects. Treatment group projects were featured during the period of the experiment, and control group projects were featured after the experiment based on their position in the queue. A user's inclusion in either experimental or control group did not affect the probability that other works may be featured during the experimental period, and thus the manipulation can be thought of as adding one marginal project feature to subjects in the treatment group. Users were unaware that there was an experiment running or that their work had been selected as part of an experiment by the curation team, and thus would not be able to predict if or when their work would be featured. The total number of projects featured on *Behance* was unchanged by the experiment. Treatment takes place on different days for different users as the rate of featuring for experimental projects was about 6 a day, and we had a set of 329 projects to be featured. The experiment thus ran for a period of about 60 days to allow all the treated projects to be featured. The timeline of the experiment is depicted graphically in figure 3. As can be seen in the figure, different projects were featured on different days, with the pre-treatment and post-treatment periods varying across projects. The analysis of the experimental data leverages this variation in timing and the panel structure of the data in a difference-in-difference style measure rather than a simple mean comparison between treatment and control groups.

Featuring involves making the cover images of these projects, and links to the project pages visible on the landing webpage of the network. Additionally, all users who follow the user whose project is featured receive notifications about the featuring of that content. Featured projects receive a permanent badge indicating that they were featured on the network. Projects are featured chronologically, with the latest project to be featured appearing at the top of the list of featured projects, and getting pushed down as subsequent projects get featured. As we show in the analysis, this featuring directs significant attention and recognition from peers.

We would like to note that the experiment did not alter the process of featuring on the network. The process by which the projects that were at the front of the queue were scheduled for featuring, their getting the featured status, being displayed on the front pages of *Behance*, and the sending of notifications to followers of the featured user remained unchanged. The only thing that the experiment manipulated was the timing of featuring of the experimental projects.

One potential concern with an experimental manipulation on a social network is whether the treatment directly affects the control group because of the interconnected nature of the treatment and control group users. In our case, this is unlikely to be a significant concern. There were over 5 million users on the *Behance* network at the time of the experiment, with over 9 million projects posted. Since only a handful of projects are featured on a given day, there is a very small probability that any given project will be featured. Thus, there is no significant expectation amongst users that their work would be featured, and no expectation of when it would be featured even if they believe a given project is likely to be featured. The fact that the control group users do not have their work featured on a given day is therefore unlikely to directly affect the control group.

3.3 Data

We collected data on users in the experiment (n=658) and followers of those users (n=482,097) for the period from May 1, 2015 through January 31, 2016 - thus we have data for several weeks before the start of the experiment, and several weeks after the end of the experiment. The extent of pre-treatment and post-treatment data for the users in the treatment group varies because projects are featured on different days during the experimental period. However, we have a long enough data period to allow us to conduct analysis of pre-experimental data (to conduct randomization checks) and post-treatment data to allow us to examine temporal effects of treatment over a period of time.

The dataset provides a variety of user-day metrics from the *Behance* network, including new creative works published, inbound and outbound tie formation, appreciations, comments, and views. Further, for a fraction of users in the experiment (n=138) we also observe *Creative Cloud* timestamped session data, including session start, session duration, and products used. Both of these datasets were merged via a common identifier and aggregated on a user-day level, allowing us to implement a panel data analysis with both user and time fixed effects.

3.4 Pre-Treatment Activity

Table 1 describes pre-treatment activity levels for the 658 users in the experiment. From the table, we note that project views are the most common *Behance* network action taken by users in the experiment, followed by appreciations given, comments given, inbound ties gained, and then project postings. Users in the experiment receive more inbound actions (such as inbound ties or comments received) than outbound (outbound ties or comments given), which is expected given that these users were selected from a group of content creators for their quality works.

We collect three dependent measures related to *Creative Cloud* usage: total session length, a daily usage indicator, and total session length conditional on being active on the given day. The typical user in the dataset spends an average 4 hours and 25 minutes of session time within *Creative Cloud* products daily and logs in an average of 34% of days.

Columns 2 and 3 provide within-user and across-user standard deviations for these activity variables. Across user variation tends to be greater than within user variation, except for projects published and outbound ties gained.

3.5 Randomization Checks

Table 2 presents treatment and control group means for pre-treatment variables and results from two randomization checks. Column 3 in the table displays corresponding p-values for a t-test for difference of means between the test and control groups. Since the p-values for the t-tests show no significant difference between treatment and control groups, we can be reasonably confident that our randomization procedure was valid and yielded the desired results. Similar randomization checks comparing followers of the treatment and control groups similarly show no differences between those users.

4 Results

4.1 Empirical Methodology

Our primary analysis utilizes the panel structure of our dataset to estimate average treatment effects via least squares linear regression using a difference-in-difference style identification strategy. We choose to utilize a panel data model rather than a pure user-level analysis to lend additional statistical power to our results. Users in the treatment group are featured at varying times throughout the 276-day observation period, and, prior to being featured, these treatment group users are expected to behave similarly to the control group due to random assignment. The panel data model can be thought of as buying additional control group observations over a user-level analysis, in addition to allowing for controls for day-level network-wide shocks. Further, the treatment group has repeated day-level observations of user activity metrics before and after treatment, as our observation period runs from May 1, 2015 through January 31, 2016, while the experiment occurs between Aug 4, 2015 and September 30, 2015. Using the difference-in-difference strategy, the measured treatment effect is the interaction between being in the treatment group and being post-treatment, with the modification that treatment occurs at different timings during the experimental period for each user. We include both user and day fixed effects respectively to account for systematic differences between users and network-wide daily activity shocks or trends. Specifically, we estimate the following model:

$$\begin{aligned} activity_{it} = \alpha_1 \cdot treatment_i \times week1 of feature_{it} + \alpha_2 \cdot treatment_i \times week2 of feature_{it} + \\ \alpha_3 \cdot treatment_i \times week3 of feature_{it} + \alpha_4 \cdot treatment_i \times week4 of feature_{it} + \eta_i + \xi_t + \epsilon_{it} \end{aligned} \quad (1)$$

where $activity_{it}$ is the dependent variable metric of interest, users are indexed from $i = 1, ...N$, and days from $t = 1, ...T$ representing our observation period, $treatment_i$ is an indicator variable capturing whether user i is in the treatment group, and $weekX of feature_{it}$ is an indicator for whether t is in the Xth week of i's project feature date. Depending on the analysis, the activity variables are measures of attention and recognition (in the case of manipulation checks), variables indexing the level of engagement of the user with the network, variables indicating project posting activity on the network, or even user activity in the *Creative Cloud* software products. The coefficients $\{\alpha_1, \alpha_2, \alpha_3, \alpha_4\}$ capture the treatment effect in the first through fourth weeks, respectively, following the feature date. η_i are user-level fixed effects that capture heterogeneity in user activity and ξ_t are day-level fixed effects. We allow for the potential that the errors ϵ_{it} may be serially correlated within user and day and compute cluster-robust standard errors accordingly. Treatment group observations greater than 4-weeks post-treatment are dropped, as these are not comparable to the pre-treatment or control observations and would otherwise bias the treatment group's user-level fixed effects. In our day-level analysis, we utilize a modified version of this model to examine day-level effects in the first week, substituting the interactions of treatment and week dummies with the interactions of treatment and day dummies for the 7 days after being featured. We also subset users by their quartile of previous appreciations and repeat this panel analysis on the subsets in order to explore heterogeneity in treatment effects.

Figure 4 depicts the data structure for the panel regressions. As previously described, each user in the experiment has a project selected by *Behance* curators which meets the quality worthy of being featured. For users in the control group, that project is ensured not to be featured during the experiment and subsequent observation period. For the treatment group, the project is featured at a random within the experimental period. Thus, observations of the treatment group can be divided into a pre-treatment and post-treatment period.

We further implement mediation checks to examine the degree to which attention and recognition mediate the changes in user behavior surrounding content creation, engagement on the network, tie formation, and usage of *Creative Cloud*.[6] As previously discussed, attention and recognition are not entirely separable in social network contexts, not least due to the mechanical relationships between the two and tendency among content creators and platforms for attention to be interpreted as recognition and vice-versa. However, we still attempt to create a common sense delineation of attention versus recognition measures on the network. We looked for an attention measure that best captured the visibility of the user's works or profile. Views received best matched this goal. We similarly looked for a measure of recognition which would capture the aggregate praise or endorsement that a creator received from their peers, settling on appreciations received. These findings are detailed in Section 4.6.

4.2 Manipulation Checks

In this subsection, we conduct a series of checks to establish that our experimental manipulation of featuring led to increased attention and recognition for the featured users and their content. In Table 3, we compare the treated and control group users on a set of variables that indicate the degree of attention and recognition they receive on the network. These variables include the number of views received for their projects, the number of appreciations and comments their projects received from other users, as well as the number of outbound ties or inbound ties from peers that they gained[7]. We find that featuring increases all of these variables, with views being significantly different for treated and control users at the 90% significance level, and the other variables being significantly different at the 95% level or higher.

This analysis shows that there was a significant effect of our featuring manipulation on the amount of attention and recognition that featured users received for their content on the network, and in the number of inbound ties they had. Furthermore, these differences were large in relative terms. There was a 34% increase in views received over the experimental period, an approximately 37% increase in appreciations received, 31% increase in comments received and an almost 72% increase in number of inbound ties gained. The manipulations are even bigger in magnitude than

[6]We thank the editors at *Marketing Science* for this suggestion.

[7]Users almost never lost existing inbound ties in our dataset so we did not specifically study the loss of inbound ties

these numbers would indicate because they are mean comparisons across the entire experimental period, even though treatment for the treated users was spread across an approximately two month period. Thus, included in the sample are users for whom featuring happened later in the experimental period, even on the last day, and their observations include periods when they had not yet been featured, and thus were similar to the control group.

We next estimated panel regressions described in subsection 4.1 to more thoroughly examine the effects of the manipulation on attention and recognition measures. As described in that subsection, the panel regression includes day level data for the period before the experiment and also after the experiment concluded, fixed effects for users to account for user-level differences in attention and recognition measured, and day fixed effects to account for time varying activity on the network. We break down the effects of the manipulation into 4 week-level effects for the 4 week period after being featured[8] to examine the temporal differences in these effects. The results of these regressions for views received, appreciations received, comments received, and inbound ties gained are reported in Table 4. The finding in this table are consistent with those discussed earlier in this section and reported in Table 3.

Specifically, we observe significant effects of treatment on all measures of peer engagement, including views, appreciations and comments received as well as inbound ties. These effects are highest in the first week after treatment, and attenuate as expected over time. The effects on views received persists for the entire 4 week period, with a significant increase in views even four weeks after being featured. The effect on inbound ties also remains significant at the 95% level four weeks after being featured. Appreciations received are significantly higher for weeks 1, 3 and 4 after treatment at the 95% level, and at the 90% level in the second week after treatment. Comments received are higher for the treatment group for the first three weeks after treatment, all at the 95% significance level except for week 2 which is at the 90% level. In terms of magnitude, these effects are large. For instance, the lift in views received for the treated users over the control group users is 43% in the first week, 31% in the second week, 31% in the third week, and 25% in the fourth week after treatment. For inbound ties, the lifts are 43%, 25%, 33% and 27% respectively for the first four weeks after treatment. Appreciations received go up by 66%, 22%, 31% and 28% respectively in the first four weeks after treatment, and comments received go up by 87%, 26% and 33% in the first three weeks respectively.

In Table 5, we examine these same effects at the day level for the first week. We find that the effects of being featured on appreciations received, views received and inbound ties are all positive and significant. The effect on comments received is significant for the first three days after being treated. The magnitudes of these effects are again large. For instance, in the first day after being treated, appreciations received go up by about 169%, comments received go up by 288%, views received go up by 71% and inbound ties go up by 94% over baseline,

To summarize, we find that our manipulation of featuring increases attention and recognition

[8] Note that the featuring happened on different days for different users in the treatment group, and therefore, the four-week window starting on the day of treatment differs by user

received by the featured users' content in a significant way, both in terms of statistical significance, and in terms of the magnitude of the changes in attention and recognition.

4.3 Treated Users' Engagement

In this subsection, we discuss how featuring, and the consequent attention and recognition from peers affects treated users' engagement with the *Behance* network. We define engagement as variables that indicate the user's activity on the network other than the posting of projects (which we will examine in the next subsection along with other metrics of content creation). In our data, we have information on users' views of project pages, appreciations given, comments given, and number of users they follow, i.e. their outbound ties. Thus, we examine how featuring affects these metrics of engagement on the network. For this purpose, we employ the panel regressions described in subsection 4.1, with the dependent measures being these engagement metrics described earlier.

Table 6 details effects of treatment on levels of engagement for the treated user. We see that featured users engage more with the network after being treated, but the treated user response is much less persistent than the network response. All four engagement metrics - appreciations given, comments given, views given and outbound ties - are significantly higher for the treated group than the control group in the first week after treatment. However, these differences are not signifiant in the subsequent weeks. In terms of magnitude, these effects are substantial. Appreciations given by the treated users go up by 39% relative to baseline, comments given go up by 56%, views of pages on the network go up by 46% and outbound ties go up by 96% on average. Thus, there is a significant increase in engagement on the network by the treated users, although this increase is not long-lasting. Whole not reported in the table, we find that all of these metrics are also significantly greater over the four week period following treatment, indicating that the increase in engagement is not neutralized by a drop in engagement in the following weeks. We can also see that in the table from the fact that while there is an increase in engagement metrics in the first week, there is no significantly negative effect in subsequent weeks for any of the engagement variables.

Next, we examine the effects of featuring on engagement for the treated user over the first seven days after treatment. These effects are reported in Table 7. Our main finding is that engagement metrics for the treated users are significantly higher mainly for the first three days after treatment. They view more pages on the network, give more comments and appreciations, and create more outbound ties for the first three days after treatment. The effects for the subsequent days are insignificant, indicating that there is no drop in engagement in subsequent days to counterbalance the significant increase in engagement for the first few days after treatment. The magnitudes of the effects in the first few days are expected to be larger than those for the week-level effects. For instance, appreciations given increase between 69% and 92% relative to control, comments given increase between 90% and 169%, views given increase between 72% and 96%, and outbound ties increase between 184% and 231% relative to control in the first three days. Thus, there is a net increase in engagement metrics in the first week, which is large in magnitude and statistically significant.

4.4 Content Creation/Sharing

We next examine the response to featuring in terms of users' content creation activity outside the network and content sharing activity on the network. As discussed previously, the provision of user generated content is critical to the success of social networks. The publication of projects could be differentiated from other activities on the *Behance* social network, such as the engagement metrics studied in the previous sub-section, due to the level of commitment and planning required to produce and publish a polished work. We use two sets of metrics to study the effects of attention and recognition derived from our featuring manipulation on content creation and sharing. First, we examine if there are causal effects on project publication by the treated user on the network. Uniquely, we have information for a subset of users about their usage of *Creative Cloud* products. Specifically, we observe whether a user used *Creative* Cloud products on a given day and their total session length in seconds for every day (this includes zero seconds for users who did not use the products on that day).

Since only a subset of users in our experiment also had linked *Creative Cloud* accounts, we first present a comparison of the full sample with those for whom the *Behance* and *Creative Cloud* accounts were linked to examine any differences between them. This is reported in Table 8. Users which could be matched to a *Creative Cloud* account tended to publish more projects on the network, were more active, and received more attention and recognition from their peers.

Table 9 provides results for the effects of featuring and the consequent attention and recognition it generates, on projects published and *Creative Cloud* activity metrics. We conduct this analysis using panel regressions similar to the ones used for examining the engagement effects. This involves regressions with interactions between treatment and week-level dummies for the four weeks after treatment for each user, and user and day level fixed effects. We find that on average, users who are featured have more project postings in the week after having a project featured on the network. This is statistically significant at the 99% level, and is large in magnitude, with the daily project posting level increasing by about 100% relative to baseline. While the effects are not significant for the subsequent three weeks, we also do not see a significant dip in the number of projects in the subsequent three weeks.

Not just does their content sharing activity see an effect, as measured by their project posting activity, there is a signficant increase in their activity on *Creative Cloud* products. We examine this for two reasons. First, effects on sharing on the network might indicate the users posting content they might have already created in order to take advantage of the attention they have due to the treatment. An examination of *Creative Cloud* activity allows us to examine effects on content creation activity as opposed to merely sharing. Second, to the extent that time spent on the network and time spent creating content outside the network (for instance in *Creative Cloud* software products) come at the cost of each other, the effect of treatment on activity on the network could be substitutes. This substitution could plausibly lead to negative effects of treatment on creative effort if it increases the time spent by the user on the network. By examining if the effects of treatment on emphCreative Cloud usage is positive, we can take a more comprehensive look at this issue.

We find that treated users' total session length increases by about 28% in the first week after treatment compared to the baseline and this increase is statistically significant at the 90% level[9]. The daily usage indicator for *Creative Cloud* products increases in a statistically significant manner only for the third week after being featured. This increase is 22% relative to baseline. Next, we examine the effect of the treatment on the intensive margin of *Creative Cloud* product usage. The usage of *Creative Cloud* products conditional on login into the suite of products is significant (at the 90% level) in the first week after treatment, but not in subsequent weeks. The main conclusion we draw from the analysis of *Creative Cloud* usage is that attention and recognition seem to increase usage of the products, and that in the immediate aftermath of treatment, this is at the intensive margin. Taken together, we find that effect of treatment on *Creative Cloud* usage is positive - thus the effect of treatment on network activity does not seem to decrease content creation activity - in fact it increases it. The fact that this comes at the intensive margin in the first week suggests that it is not merely a reminder effect of the treatment on *Adobe* products in general.

In Table 10, we report the effects on content creation and sharing for the first seven days after being treated. Users have an increase in projects published for three of the seven days in the first week, and all of these effects are significant at the 95% level. Furthermore, in terms of magnitude, these effects range between 141% on the first day to 202% on the second day after treatment. The usage of *Creative Cloud* products increases on four of the seven days in the first week after treatment, with two of these being significant at the 95% level and two others at the 90% level. The magnitudes of these effects range between 34% on the fourth day after treatment to 46% on the second day after treatment.

In sum, we have found that being featured increases the content sharing and content creation activities of the treated users in statistically significant and meaningful ways.

4.5 Heterogeneity in effects

In this sub-section, we examine potential heterogeneity in effects based on the past stock of appreciations users have received. We use past appreciations as a measure of recognition received by the users from their peers on the network. The rationale for examining heterogeneity based on past recognition is that users potentially post their projects on the *Behance* network to generate recognition for their creations. For users who do not have a large stock of past recognition, additional recognition obtained as a result of having one of their projects being featured could potentially be meaningful. By contrast, users who have a large stock of past recognition may not gain very much in recognition by being featured. We considered examining the stock of past featuring for users to examine heterogeneity, but based on conversations with the curators this measure would be zero for the large majority of users in the study and additionally this variable is not available in the dataset

[9]Note from Table 9 that the sample sizes for the regressions for the *Creative Cloud* variables is an order of magnitude lower than that for project posting activity on the *Behance* network, due to the partial nature of the matches between *Behance* and *Creative Cloud* user profiles. The lower levels of statistical significance are therefore not surprising.

we obtained. Thus, we use past appreciations as a close substitute. While we present heterogeneity results based on past appreciations received (recognition), users' historical appreciation stocks are highly correlated with comments and views, and therefore, analysis based on past attention or recognition variables would lead to very similar conclusions as those using appreciations. In the interest of brevity, we present only the results for heterogeneity based on the stock of past appreciations.

To conduct this analysis, we first divided the set of users into four quartiles based on their stock of past appreciations. We then conducted separate panel regressions, as presented in the last two sub-sections, for the same set of dependent variables in the case of engagement metrics, and project postings in the case of content creation/sharing variables. We were unable to conduct this quartile-level analysis in the case of *Creative Cloud* variables because of the smaller sample sizes in the case of these variables. Recall that the sample size for the matched *Creative Cloud* dataset is an order of magnitude lower than the full dataset, due to the partial nature of the matching between *Behance* and *Creative Cloud* user profiles. We first present summary statistics for these four quartiles in Table 11. As can be seen from the table, there are large differences between these quartiles on all the variables. Unsurprisingly, the higher quartiles receive more views, appreciations and comments than the lower quarter. have more inbound ties. In terms of their own activities, the table presents a more complex picture. For instance, users in quartile 4 give more appreciations on average than those in quartile 3, but the difference in percentage terms is much lower than in the case of appreciations received. The users in quartile 4 in fact have fewer outbound ties on average than those in quartile 3, and not much more so than those in quartile 2.

We present the results of the quartile-level analysis in Tables 12 through 15. As expected, we find more statistically significant effects for users in the lower two quartiles than for those in the upper two quartiles. Users in quartiles 3 and 4, do not have statistically significant effects of featuring on any of the variables at the 95% significance level. In the case of quartile 3, there are no significant effects at the 90% level either and for quartile 4, there are only two significant effects at the 90% level - they view more projects and publish more projects subsequent to being featured. For quartile 1, there are signficant increases in comments given, views, outbound ties and project published in the first week, with effects for the first three of these variable being significant at the 99% level, and projects published at the 90% level. The magnitudes of these effects are large, with lifts of 871% in comments given, 118% in views given, 175% for outbound ties and 140% for projects published. Users in quartile 2 have an increase in appreciations given, views given and projects published in the first week after being featured at the 99% level, with comments given and outbound ties increasing at the 90% level for that week. Once again, the magnitudes of these effects are large, with 278% in appreications given, 234% for comments given, 71% for views given, 115% for outbound ties, and 194% for projects published. The only other significant effects, albeit only at the 90% level are for quartile 4, and the magnitudes of the effects are smaller - 34% increase in views given, and 52%.

In summary, we find significant heterogeneity in the effects of being featured on users' engagement and content sharing on the network. Users in lower quartiles, with a lower stock of past appreciations have a larger set of significant effects, and have large magnitudes of these effects.

Users in higher quartiles, who have a greater stock of past appreciations have fewer significant effects, and the effect sizes are smaller in relative terms that for lower quartiles. This suggests that there are diminishing returns to recognition - users who have received more of it in the past are less likely to be affected by new recognition they receive for the projects in the experiment. But users who have received less of it in the past are more likely to be influenced by the recognition they receive for their projects that are part of the experiment.

4.6 Mediation

We conduct a mediation analysis to examine the roles of attention and recognition in driving engagement on the network, and determine whether *Behance*'s endorsement of the content has any direct effect beyond its role in shifting peer attention and recognition. We identified 'Views Received', the number of users who viewed any of an artist's works in a given day, as the best measure of the level of peer attention a user has received. Similarly, we identified 'Appreciations Received', the number of peers who clicked 'Appreciate' and thereby endorsed and gave their approval to a user's work on a given day, as the best measure of the recognition that a user has received.

To adapt our data to the classic mediation model (Baron and Kenny 1986) utilizing a single observation per individual rather than a panel structure, we first average daily activities on a user level in the one week immediately prior to and one week after the user's date of treatment.[10] We then difference these weekly averages to create the change in activity levels pre- and post-experiment. As we demonstrated in our mediation checks, the differences in views and appreciations received as a result of our manipulation is both large in magnitude and highly significant. To reduce the role of outliers, as some users received thousands of extra daily views as a result of the manipulation while others received only tens, we winsorize our attention and recognition received measures using 5% and 95% thresholds.

We run a single mediation analysis twice per dependent measure rather than a mediation model with both the attention and recognition mediators (MacKinnon 2008) due to the high correlation between views and appreciations received, which would otherwise introduce collinearity. Our mediation equations are given as follows:

$$\Delta DependentMeasure = i_1 + \alpha_1 * treatment + \epsilon \tag{2}$$

$$\Delta DependentMeasure = i_2 + \alpha_1' * treatment + \alpha_2 * \Delta ViewsReceived + \epsilon \tag{3}$$

$$\Delta DependentMeasure = i_3 + \alpha_1'' * treatment + \alpha_3 * \Delta AppreciationsReceived + \epsilon \tag{4}$$

Table 16 shows the results of our mediation analysis on user engagement, tie formation, and content creation. The total effects of Treatment (α_1) in columns (1), (4), (7), and (10) on Views

[10]For users in the control group, we assign each user a hypothetical treatment date randomly drawn over the experimental period.

Given, Appreciations Given, Outbound Ties, and Projects Published are all large and significant (p-values: <0.001, 0.059, 0.014, 0.028, respectively), confirming the results from our previous analyses utilizing the panel dataset in the difference-in-difference framework. When decomposed into direct and indirect effects, we see no evidence of a direct effect of Treatment on user engagement, tie formation, or content creation. In each case, the coefficients on Treatment decrease dramatically in magnitude and become non-significant when controlling for peer attention or recognition (all p-values > 0.1). Instead, nearly all of the total effect of the manipulation comes through the indirect effects of attention (Views Received) and recognition (Appreciations Received). These α_2 and α_3 are large and highly significant for each dependent measure, and, when taken with the results on Treatment coefficients, indicate the mediation of user behavior effects by attention and recognition is complete.

4.7 Discussion

Our results indicate that the use of featuring to direct attention and recognition towards content is effective in motivating users' content creation and engagement on the social network, as they have direct causal effects on content creation and creative product usage for treated users. These effects are limited to the short- and mid-term, and treated users return to baseline levels of activity despite still receiving increased levels of attention and recognition. The ability to tie social network activity to direct measurements of creative software usage is unique to our knowledge in the literature, and allows us to directly measure behaviors related to effort. One explanation for the increase in projects published could be that users in the context of *Behance* accelerate posting of a backlog of already completed work, leading to no new work being created. We can rule out this explanation by utilizing *Creative Cloud* data to demonstrate effects on creative software usage and noting the lack of post-treatment slump in project published. Thus, we can be relatively certain that users are committing additional effort to content creation in response to being featured.

We see evidence of diminishing returns to attention and recognition received on *Behance*, both in our main results and through examining heterogeneity of treatment effects. Treated users in the experiment see a much larger percentage lift in attention and recognition measures than engagement, and the attenuation of engagement occurs even while attention and recognition are significantly elevated above baseline. This provides suggestive evidence that users are becoming accustomed or satiated to the increased level of attention and recognition from their peers. We also observe that users in the highest quartile of appreciations, who are presumably accustomed to a high level of recognition, see a smaller lift in engagement as a result of the intervention than users in the lower two quartiles. Interestingly, the second quartile of users saw the highest lifts in absolute terms. This segment of users may have struck a balance between seeing enough attention to be engaged with *Behance* yet not so popular as to be accustomed to the spotlight of recognition represented by having their work featured. Still, it should be noted that even users in the lowest quartile of previous appreciations within our experiment are not representative of the "average" user on the network. Thus, these results should be interpreted as a treatment effect on the selection of users whose work meets the standard of potentially being featured.

Finally, we examine the mediating role of attention and recognition, and answer the crucial question of whether our effects on user engagement, tie formation, and content creation are a result of the platform endorsement through the receipt of a distinguishing badge rather than peer effects. We find this is not the case. Our mediation analysis shows the strong mediating role of peer attention and recognition in driving treated user behavior changes. The process can be described as complete mediation, with no significant direct effect of featuring on subsequent user behavior.

These mediation analysis results are subject to a couple of caveats, which we acknowledge and highlight. First, we discussed earlier how attention and recognition are intrinsically linked and often difficult to separate on social networks. This is underlined in our dataset by the fact that $\Delta_{ViewsReceived}$ and $\Delta_{AppreciationsReceived}$ have a correlation coefficient of r = 0.829 and should arguably be considered a single construct, despite the tendency of the literature to treat these as separable constructs in non-social media contexts. Secondly, we cannot rule out the possibility of omitted variable bias, for example if the effort an artist dedicates to the *Behance* platform is correlated with both their reception by the platform upon featuring and their subsequent increase in engagement post-featuring.

5 Conclusion

This paper presents the results of a collaboration with a large social network to examine the effects of online attention and recognition on users' subsequent engagement and content creation. Attention and recognition have been identified as key motivators of content creation, and are of particular interest to social networks that both rely on user generated content to attract users and have the ability to direct peer attention and recognition through various design levers. Due to the endogeneity of attention and recognition by peers, their effects have traditionally been difficult to study. We solve this challenge by implementing an experiment that exogeneously manipulates the attention and recognition a user's works receives through front page featuring.

The manipulation is highly effective at directing attention and recognition, seamlessly utilizing an existing mechanism on the site. Treated users published nearly twice as much content on the network in the following week relative to the pre-experiment baseline. This increase in content publishing on *Behance* is mirrored by increased usage of the underlying creative tools on *Creative Cloud*, suggesting that new content creation was occurring rather than intertemporal substitution. Treated users became more engaged with the site - viewing, appreciating, and commenting at increased rates. However, these engagement and creation effects attenuate quickly, returning to baseline levels by the second week despite still receiving increased attention and recognition from the network. Lastly, mediation analysis confirms the role that increased attention and recognition play in driving user behavioral changes, with no significant direct effect of featuring by the platform itself.

These findings are highly relevant to managers of social networks and other online platforms, which rely on user generated content. In particular, they demonstrate the value of attention and recognition-directing levers such as front pages and featuring in motivating content creation. That attention and recognition show diminishing returns and users become satiated suggests that limited attention and recognition directed by front page space could best be allocated to users who have received less of it in the past, and reinforces the view that display mechanisms should play a role in facilitating discovery of underappreciated content. In this paper, our focus has been on the effect of featuring and consequent attention and recognition from peers on the quantity of content created by users of a social network. In subsequent work, we hope to build on these results by applying image recognition techniques to project images from this study to understand how the subject matter and quality of work changes in response to featuring, attention, and recognition.

References

A. Aaltonen and S. Seiler. Cumulative growth in user-generated content production: Evidence from wikipedia. *Management Science*, 62(7):2054–2069, 2015.

M. Baas, C. K. W. D. Dreu, and B. A. Nijstad. A meta-analysis of 25 years of mood-creativity research: Hedonic tone, activation and regulatory focus. *Psychological Bulletin*, 134(6):779–806, 2008.

R. M. Baron and D. A. Kenny. The moderator–mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *Journal of Personality and Social Psychology*, 51(6):1173–1182, 1986.

L. Cabral and L. Li. A dollar for your thoughts: Feedback-conditional rebates on ebay. *Management Science*, 61(9):2052–2063, 2015.

Y. Chen, F. M. Harper, J. Konstan, and S. X. Li. Social comparisons and contributions to online communities: A field experiment on movielens. *American Economic Review*, 100(4):1358–1398, 2010.

J. A. D. Dae-Yong Ahn and C. F. Mela. Managing user-generated content: A dynamic rational expectations equilibrium approach. *Marketing Science*, 35(2):284–303, 2016.

E. Deci and R. M. Ryan. *Intrinsic Motivation and Self Determination in Human Behavior*. Springer, 1985.

A. Ghose and S. P. Han. An empirical analysis of user content generation and usage behavior on the mobile internet. *Management Science*, 57(9):1671–1691, 2011.

A. Hars and S. Ou. Working for free? motivations for participating in open-source projects. *International Journal of Electronic Commerce*, 6(3):25–39, 2002.

N. Huang, G. Burtch, B. Gu, Y. Hong, C. Liang, K. Wang, D. Fu, and B. Yang. Motivating user-generated content with performance feedback: Evidence from randomized field experiments. *Management Science*, 65(1):327–345, 2019.

A. M. Isen, K. A. Daubman, and G. P. Nowicki. Positive affect facilitates creative problem solving. *Journal of Personality and Social Psychology*, 52(6):1122–1131, 1987.

G. Iyer and Z. Katona. Competing for attention in social communication markets. *Management Science*, 62(8):2304–2320, 2016.

A. Kankanhalli, B. C. Y. Tan, and K.-K. Wei. Contributing knowledge to electronic knowledge repositories: An empirical investigation. *MIS Quarterly*, 29(1):113–143, 2005.

M. E. Kummer. Spillovers in networks of user generated content. Working paper, Georgia Institute of Technology, 2015.

K. R. Lakhani and R. G. Wolf. Why hackers do what they do: Understanding motivation and effort in free/open source software projects. In J. Feller, B. Fitzgerald, S. A. Hassam, and K. R. Lakhani, editors, *Perspectives on Free and Open Source Software*, pages 3–21. MIT Press, Cambridge, 2005.

J. Lerner and J. Tirole. Some simple economics of open source. *Journal of Industrial Economics*, L (2):197–234, 2002.

D. P. MacKinnon. *Introduction to Statistical Mediation Analysis.* Taylor and Francis Group/Lawrence Erlbaum Associates, 2008.

M. Moqri, X. Mei, L. Qiu, and S. Bandyopadhyay. Effect of "following" on contribuitions to open source communities. *Journal of Management Information Systems*, 35(4):1188–1217, 2018.

L. Muchnik, S. Aral, and S. J. Taylor. Social influence bias: A randomized experiement. *Science*, 341(6146):647–651, 2013.

O. Nov. What motivates wikipedians. *Communications of the ACM*, 50(11):60–64, 2007.

S. Oreg and O. Nov. Exploring motivations for contributing to open source initiatives: The roles of contribution context and personal values. *Computers in Human Behavior*, 24(5):2055–2073, 2008.

M. J. Salganik, P. S. Dodds, and D. J. Watts. Experimental study of inequality and unpredictability in an artificial cultural market. *Science*, 311:854–856, 2006.

S. K. Shriver, H. S. Nair, and R. Hofstetter. Social ties and user-generated content: Evidence from an online social network. *Management Science*, 59(6):1425–1443, 2013.

G. W. Thorsten Hennig-Thurau, Kevin P. Gwinner and D. D. Gremler. Electronic word-of-mouth via consumer-opinion platforms: What motivates consumers to articulate themselves on the internet? *Journal of Interactive Marketing*, 18(1):38–52, 2004.

O. Toubia and A. T. Stephen. Intrinsic vs. image-related utility in social media: Why do people contribute content to twitter? *Marketing Science*, 32(3):368–392, 2013.

M. Wasko and S. Faraj. Why should i share? examining social capital and knowledge contribution in electronic networks of practice. *MIS Quarterly*, 29(1):35–57, 2005.

X. Zeng and L. Wei. Social ties and user content generation: Evidence from flickr. *Information Systems Research*, 24(1):71–87, 2012.

C. Zhang, J. Hahn, and P. De. Continued participation in online innovation communities: Does community response matter equally for everyone? *Information Systems Research*, 24(4):1112–1130, 2013.

K. Zhang and M. Sarvary. Differentiation with user-generated content. *Management Science*, 61 (4):898–914, 2014.

X. Zhang and F. Zhu. Group size and incentives to contribute: A natural experiment at chinese wikipedia. *American Economic Review*, 101(4):1601–1615, 2011.

Figures

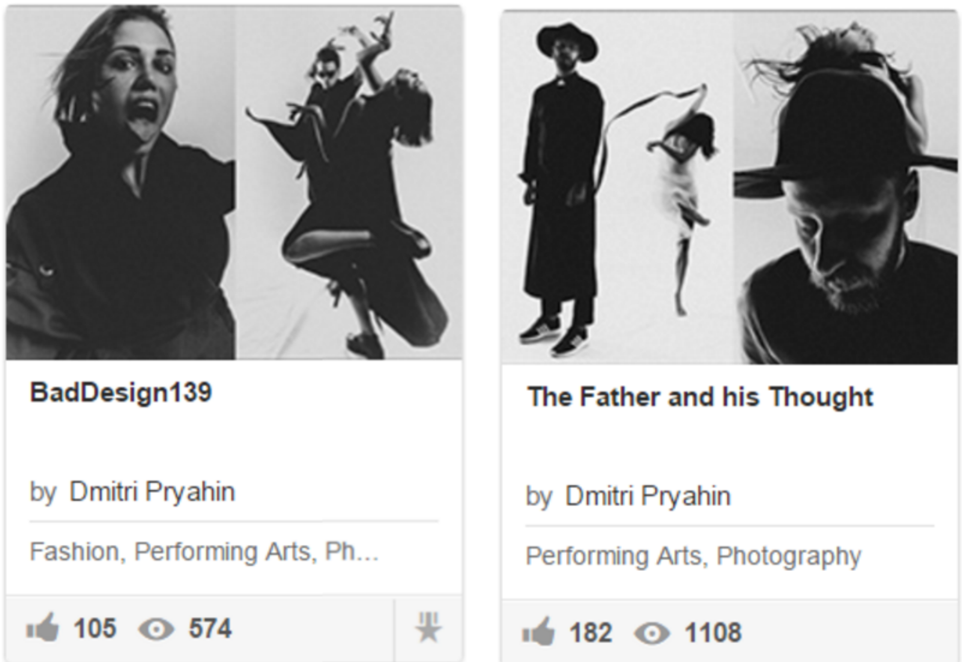

work

Figure 1: Featured vs Not-Featured Projects. Two projects by the same user. Featured projects are distinguished with the star badge indicator in the bottom right next to the Appreciations and Views the project has received. Featured projects push a notification to followers of the owner of the featured project and appear more prominently in feed rankings and results

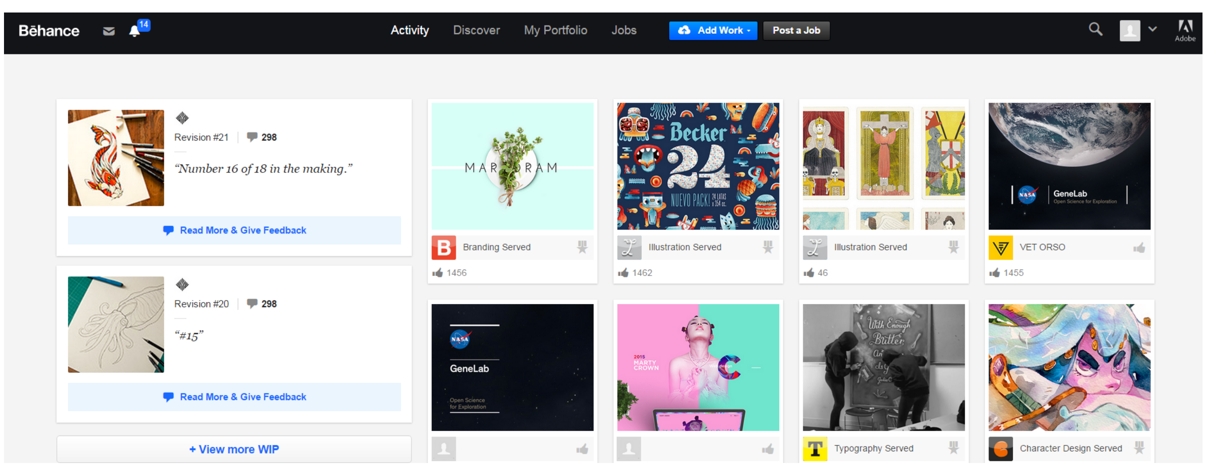

Figure 2: *Activity Feed*. Featured works are denoted with a star badge in the *Activity feed* of the *Behance* social network.

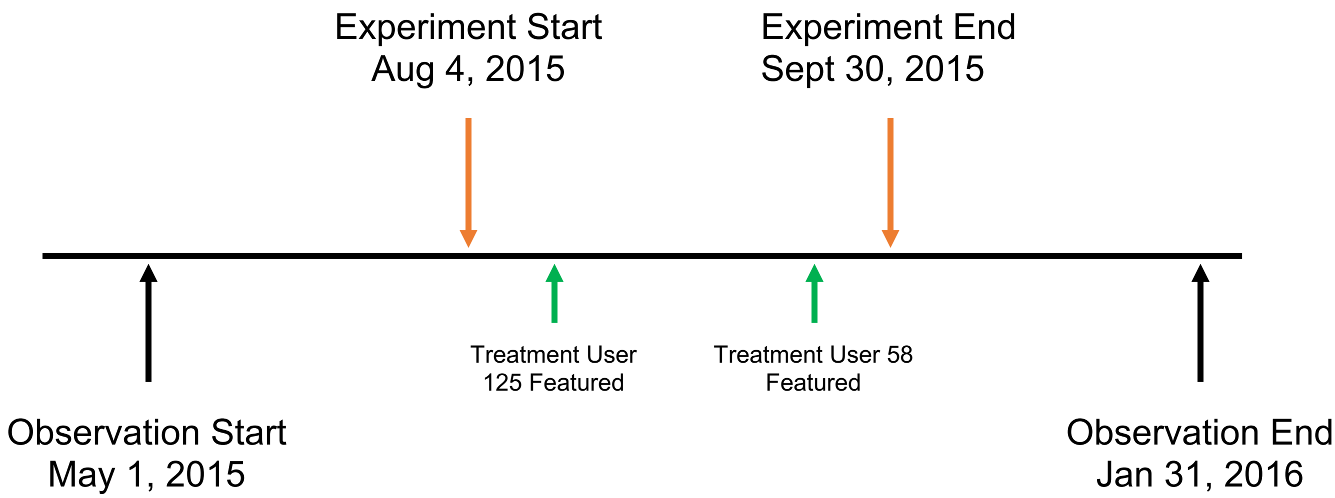

Figure 3: Schematic Diagram Depicting the Timeline of the Experiment

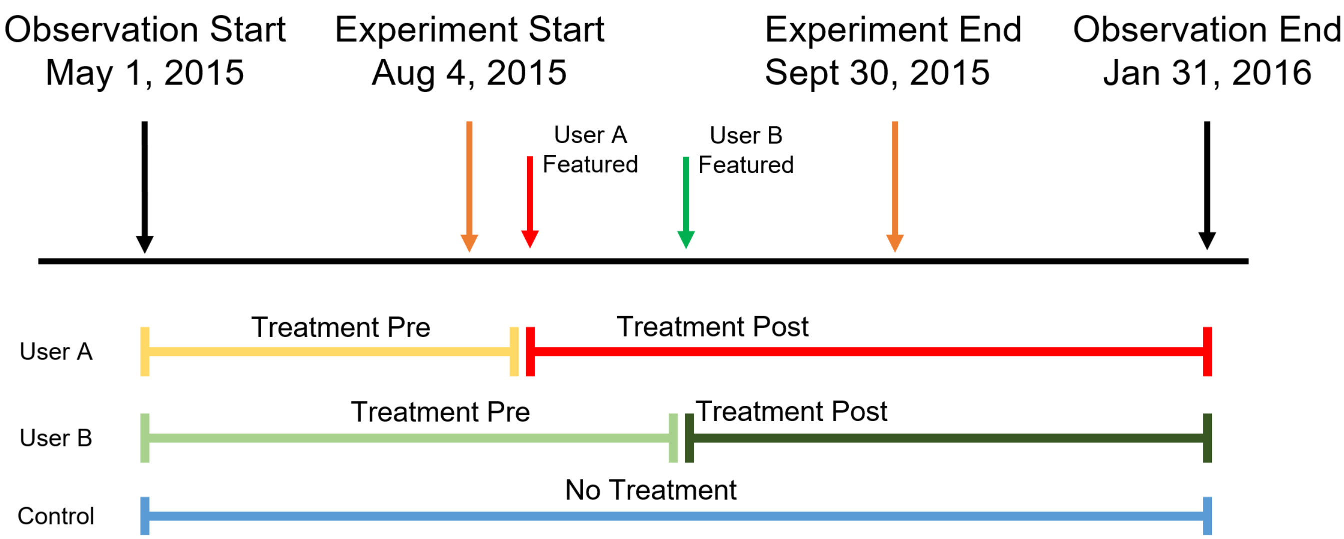

Figure 4: Schematic Diagram Depicting the Data Structure for the Panel Regressions

Tables

	Mean	Within-user Std Dev	Across-user Std Dev
Appreciations Received	10.719	16.701	18.737
Comments Received	0.697	1.446	1.573
Views Received	115.562	145.015	222.030
Inbound Ties	4.527	5.708	7.837
Appreciations Given	1.018	1.877	4.251
Comments Given	0.270	0.576	1.292
Views Given	5.842	9.698	10.978
Outbound Ties	0.217	0.829	0.558
Projects Published	0.0213	0.122	0.050
CC Session Time (s)	15890	18930	17360
CC Daily Login	0.342	0.341	0.282

Table 1: Summary statistics on pre-treatment variables. Quantities represent day-level measures.

	Mean of Control	Mean of Treatment	p: Two-sided t-test
Appreciations received	10.233	11.203	0.513
Comments received	0.600	0.794	0.119
Views received	111.626	119.486	0.655
Inbound ties	4.163	4.890	0.241
Appreciations given	0.972	1.064	0.783
Comments given	0.237	0.303	0.516
Views given	5.707	5.976	0.757
Outbound ties	0.194	0.240	0.299
Projects published	0.018	0.024	0.123
CC session time (s)	16860	14930	0.516
CC daily login	0.337	0.347	0.845
Note:		*p<0.1; **p<0.05; ***p<0.01	

Table 2: Randomization Checks on Pre-experiment Variables. Quantities represent day-level measures averaged across pre-treatment days by user and then by experimental group.

	Mean of Control	Mean of Treatment	p: two-sided t-test
Views received	96.450	129.150	0.065*
Appreciations received	8.0776	11.069	0.046**
Comments received	0.566	0.740	0.016**
Inbound ties gained	3.903	6.712	0.018**
Note:		*p<0.1; **p<0.05; ***p<0.01	

Table 3: Manipulation Checks - Means Comparison. Compares attention and engagement measures by user for treatment vs. control groups over the experimental period. Daily quantities are averaged by user across the experimental period August 4, 2015 through September 30, 2015 and then averaged by experimental group.

Treatment interacted with	Dependent variable:			
	Appreciations Received	Comments Received	Views Received	Inbound Ties
Week1 post-feature	6.687***	0.604***	47.000***	1.784***
	(0.977)	(0.105)	(10.652)	(0.345)
Week2 post-feature	2.243*	0.183*	34.313**	1.053**
	(1.204)	(0.108)	(13.687)	(0.480)
Week3 post-feature	3.147**	0.231**	34.479**	1.378**
	(1.601)	(0.114)	(15.947)	(0.698)
Week4 post-feature	2.903**	0.125	27.012**	1.127**
	(1.378)	(0.094)	(13.283)	(0.558)
User fixed effects	Yes	Yes	Yes	Yes
Day fixed effects	Yes	Yes	Yes	Yes
Observations	132,510	132,510	132,510	132,510
R^2	0.316	0.310	0.406	0.382
F-test (p-value)	<0.001***	<0.001***	<0.001***	<0.001***
Note:			*p<0.1; **p<0.05; ***p<0.01	

Table 4: Manipulation Checks 2 - Panel Regression. Covariates capture the effect of being in the treatment group within the n-th week following the date on which a user's work was featured. Standard errors shown in parenthesis are clustered at the user and day level.

Treatment interacted with	Dependent variable:			
	Appreciations Received	Comments Received	Views Received	Inbound Ties
Day1 post-feature	17.218***	1.711***	78.201***	3.887***
	(1.389)	(0.145)	(9.412)	(0.517)
Day2 post-feature	12.872***	1.167***	78.390***	2.990***
	(1.594)	(0.136)	(11.749)	(0.472)
Day3 post-feature	6.481***	0.459***	48.346***	1.687***
	(1.830)	(0.144)	(13.598)	(0.534)
Day4 post-feature	2.494**	0.117	27.180**	0.917**
	(1.157)	(0.106)	(12.180)	(0.366)
Day5 post-feature	2.350*	0.301	27.715*	0.804*
	(1.235)	(0.187)	(14.334)	(0.413)
Day6 post-feature	3.249***	0.272	35.169**	1.335***
	(1.261)	(0.173)	(16.239)	(0.453)
Day7 post-feature	2.480**	0.167	40.336**	1.009**
	(1.116)	(0.116)	(17.780)	(0.430)
User fixed effects	Yes	Yes	Yes	Yes
Day fixed effects	Yes	Yes	Yes	Yes
Observations	125,790	125,790	125,790	125,790
R^2	0.313	0.289	0.401	0.384
F-test (p-value)	<0.001***	<0.001***	<0.001***	<0.001***
Note:			*p<0.1; **p<0.05; ***p<0.01	

Table 5: Manipulation Checks 3 - Panel Regression. Covariates capture the effect of being in the treatment group within the n-th day following the date on which a user's work was featured. Standard errors shown in parenthesis are clustered at the user and day level.

Treatment interacted with	Dependent variable:			
	Appreciations Given	Comments Given	Views Given	Outbound Ties
Week1 post-feature	0.372**	0.135**	2.536***	0.179**
	(0.183)	(0.060)	(0.560)	(0.072)
Week2 post-feature	−0.078	0.064	0.667	−0.010
	(0.147)	(0.063)	(0.517)	(0.051)
Week3 post-feature	−0.144	0.040	0.343	−0.013
	(0.168)	(0.056)	(0.554)	(0.048)
Week4 post-feature	−0.017	0.059	0.693	0.079
	(0.126)	(0.055)	(0.503)	(0.077)
User fixed effects	Yes	Yes	Yes	Yes
Day fixed effects	Yes	Yes	Yes	Yes
Observations	132,510	132,510	132,510	132,510
R^2	0.368	0.329	0.341	0.058
F-test (p-value)	0.087*	0.046**	<0.001***	0.019**
Note:			*p<0.1; **p<0.05; ***p<0.01	

Table 6: Treated User Response - Engagement Metrics. Covariates capture the effect of being in the treatment group within the n-th week following the date on which a user's work was featured. Standard errors shown in parenthesis are clustered at the user and day level.

Treatment interacted with	Dependent variable:			
	Appreciations Given	Comments Given	Views Given	Outbound Ties
Day1 post-feature	0.663**	0.248*	5.338***	0.342***
	(0.326)	(0.136)	(1.040)	(0.119)
Day2 post-feature	0.886***	0.405***	4.012***	0.430***
	(0.292)	(0.126)	(1.058)	(0.165)
Day3 post-feature	0.749***	0.215***	4.127***	0.348**
	(0.275)	(0.078)	(1.308)	(0.138)
Day4 post-feature	0.295	0.066	1.443*	0.058
	(0.281)	(0.085)	(0.758)	(0.071)
Day5 post-feature	−0.063	−0.014	0.338	0.057
	(0.291)	(0.067)	(0.612)	(0.087)
Day6 post-feature	0.112	0.026	0.653	−0.034
	(0.342)	(0.070)	(0.835)	(0.063)
Day7 post-feature	−0.139	−0.016	1.449	0.049
	(0.267)	(0.067)	(1.048)	(0.071)
User fixed effects	Yes	Yes	Yes	Yes
Day fixed effects	Yes	Yes	Yes	Yes
Observations	125,790	125,790	125,790	125,790
R^2	0.364	0.312	0.342	0.059
F-test (p-value)	<0.001***	<0.001***	<0.001***	<0.001***
Note:			*p<0.1; **p<0.05; ***p<0.01	

Table 7: Treated User Response - Engagement Metrics. Covariates capture the effect of being in the treatment group within the n-th day following the date on which a user's work was featured. Standard errors shown in parenthesis are clustered at the user and day level.

	Matched with *Creative Cloud*			Unmatched with *Creative Cloud*		
	Mean	Within-user Std Dev	Across-user Std Dev	Mean	Within-user Std Dev	Across-user Std Dev
Appreciations received	13.344	17.744	26.435	9.235	12.780	16.816
Comments received	1.015	1.687	3.087	0.552	1.148	1.063
Views received	137.953	144.828	302.959	102.889	106.986	203.965
Inbound ties	5.255	5.715	9.302	4.133	4.310	8.369
Appreciations given	1.320	1.836	6.792	0.747	1.482	2.253
Comments given	0.437	0.627	2.221	0.180	0.387	0.766
Views given	5.935	8.949	11.003	4.910	8.199	8.225
Outbound ties	0.236	0.813	0.986	0.167	0.592	0.401
Projects published	0.033	0.142	0.101	0.014	0.094	0.026

Table 8: Comparison of matched and unmatched users. This table presents summary statistics for users whose *Behance* and *Creative Cloud* accounts were matched, with those for whom it was not.

Treatment interacted with	Dependent variable:			
	Projects Published	CC Total Session Length (seconds per day)	CC Daily Usage Indicator	CC Session Length (conditional on login)
Week1 post-feature	0.020***	4525*	0.006853	6377*
	(0.006)	(2542)	(0.03702)	(3866)
Week2 post-feature	0.001	3145	0.06544	2156
	(0.005)	(2859)	(0.04317)	(3570)
Week3 post-feature	0.005	2982	0.07684*	2118
	(0.005)	(3257)	(0.04529)	(4824)
Week4 post-feature	−0.001	357.7	0.02380	2767
	(0.004)	(3328)	(0.05864)	(3689)
User fixed effects	Yes	Yes	Yes	Yes
Day fixed effects	Yes	Yes	Yes	Yes
Observations	132,510	10,251	10,251	3701
R^2	0.046	0.411	0.407	0.455
F-test (p-value)	0.018**	0.064*	0.130	0.072*
Note:	*p<0.1; **p<0.05; ***p<0.01			

Table 9: Treated User Response - Engagement Metrics. Covariates capture the effect of being in the treatment group within the n-th week following the date on which a user's work was featured. Standard errors shown in parenthesis are clustered at the user and day level.

Treatment interacted with	Dependent variable:			
	Projects Published	CC Total Session Length (seconds per day)	CC Daily Usage Indicator	CC Session Length (conditional on login)
Day1 post-feature	0.030**	6491*	0.07893	3837
	(0.014)	(3458)	(0.05330)	(6449)
Day2 post-feature	0.043**	7314**	0.06978	7813
	(0.018)	(3301)	(0.05333)	(5754)
Day3 post-feature	0.006	3977	-0.007486	3854
	(0.010)	(3115)	(0.04806)	(7630)
Day4 post-feature	0.034**	5468*	0.005393	9253
	(0.016)	(3223)	(0.04988)	(7212)
Day5 post-feature	0.002	3980	0.008296	8033
	(0.009)	(3427)	(0.05714)	(6172)
Day6 post-feature	0.011	3169	-0.007944	6885
	(0.008)	(3468)	(0.06400)	(7377)
Day7 post-feature	0.010	6484**	0.02906	14,860***
	(0.010)	(3271)	(0.05392)	(5286)
User fixed effects	Yes	Yes	Yes	Yes
Day fixed effects	Yes	Yes	Yes	Yes
Observations	125,790	9,041	9,041	3224
R^2	0.042	0.408	0.409	0.440
F-test (p-value)	<0.001***	<0.001***	0.331	0.008***
Note:	*p<0.1; **p<0.05; ***p<0.01			

Table 10: Treated User Response - Content Creation/Sharing. Covariates capture the effect of being in the treatment group within the n-th day following the date on which a user's work was featured. Standard errors shown in parenthesis are clustered at the user and day level.

	mean (std. deviation)			
	Quartile 1	Quartile 2	Quartile 3	Quartile 4
Appreciations received	0.379	2.362	7.794	36.792
	(0.297)	(0.941)	(2.488)	(35.082)
Comments received	0.029	0.162	0.476	2.544
	(0.038)	(0.145)	(0.350)	(3.248)
Views received	6.360	24.186	78.895	374.894
	(9.975)	(15.637)	(46.598)	(419.279)
Inbound ties	0.188	0.983	3.684	14.590
	(0.167)	(0.555)	(2.565)	(13.772)
Appreciations given	0.291	0.667	1.219	1.860
	(0.710)	(1.333)	(3.779)	(6.433)
Comments given	0.014	0.109	0.246	0.802
	(0.036)	(0.290)	(0.833)	(2.446)
Views given	2.561	4.105	7.105	10.318
	(3.729)	(5.377)	(11.596)	(15.667)
Outbound ties	0.093	0.246	0.339	0.259
	(0.161)	(0.642)	(1.403)	(0.393)
Projects published	0.010	0.017	0.024	0.065
	(0.019)	(0.027)	(0.035)	(0.179)

Table 11: Heterogeneity - Summary Statistics. Users were divided into four quartiles based on their pre-experiment number of appreciations received.

Treatment interacted with	Dependent variable:				
	Apps. Given	Comments Given	Views Given	Outbound Ties	Projects Published
Week1 post-feature	0.322	0.122***	3.030***	0.163***	0.014*
	(0.228)	(0.039)	(0.853)	(0.056)	(0.007)
Week2 post-feature	−0.078	0.034	1.550*	0.120	0.016
	(0.200)	(0.024)	(0.910)	(0.109)	(0.013)
Week3 post-feature	−0.140	0.003	0.840	0.025	0.015
	(0.176)	(0.013)	(0.733)	(0.055)	(0.017)
Week4 post-feature	−0.002	0.018	1.067*	0.016	0.002
	(0.130)	(0.012)	(0.612)	(0.052)	(0.007)
User fixed effects	Yes	Yes	Yes	Yes	Yes
Day fixed effects	Yes	Yes	Yes	Yes	Yes
Observations	35,100	35,100	35,100	35,100	35,100
R^2	0.100	0.034	0.179	0.045	0.023
F-test (p-value)	0.280	0.053*	<0.001***	0.079*	0.106
Note:					*p<0.1; **p<0.05; ***p<0.01

Table 12: Heterogeneity in Treated User Response - Quartile 1. Users were classified into four quartiles based on past appreciations received. Covariates capture the effect of being in the treatment group within the n-th week following the date on which a user's work was featured. Standard errors shown in parenthesis are clustered at the user and day level.

Treatment interacted with	Dependent variable:				
	Apps. Given	Comments Given	Views Given	Outbound Ties	Projects Published
week1 post-feature	0.810**	0.255*	2.927***	0.282*	0.033***
	(0.363)	(0.131)	(0.928)	(0.147)	(0.011)
week2 post-feature	−0.259	0.002	−0.757	−0.108**	−0.008
	(0.238)	(0.050)	(0.678)	(0.044)	(0.006)
week3 post-feature	−0.340	−0.071**	−0.814	−0.065	−0.001
	(0.284)	(0.032)	(0.720)	(0.045)	(0.006)
week4 post-feature	−0.492**	−0.062	−1.002	−0.094**	−0.002
	(0.232)	(0.039)	(0.642)	(0.046)	(0.007)
User fixed effects	Yes	Yes	Yes	Yes	Yes
Day fixed effects	Yes	Yes	Yes	Yes	Yes
Observations	32,160	32,160	32,160	32,160	32,160
R^2	0.213	0.156	0.162	0.090	0.020
F-test (p-value)	0.020**	0.061*	0.002***	0.093*	0.031**
Note:			*p<0.1; **p<0.05; ***p<0.01		

Table 13: Heterogeneity in Treated User Response - Quartile 2. Users were classified into four quartiles based on past appreciations received. Covariates capture the effect of being in the treatment group within the n-th week following the date on which a user's work was featured. Standard errors shown in parenthesis are clustered at the user and day level.

Treatment interacted with	Dependent variable:				
	Apps. Given	Comments Given	Views Given	Outbound Ties	Projects Published
Week1 post-feature	0.169	0.127	1.940	0.073	−0.001
	(0.271)	(0.106)	(1.266)	(0.170)	(0.007)
Week2 post-feature	0.183	−0.027	1.424	−0.077	−0.003
	(0.344)	(0.114)	(1.324)	(0.146)	(0.007)
Week3 post-feature	−0.071	0.001	0.190	−0.054	−0.008
	(0.313)	(0.138)	(1.410)	(0.155)	(0.007)
Week4 post-feature	0.421	0.013	1.849	0.352	−0.004
	(0.271)	(0.130)	(1.365)	(0.277)	(0.009)
User fixed effects	Yes	Yes	Yes	Yes	Yes
Day fixed effects	Yes	Yes	Yes	Yes	Yes
Observations	32,653	32,653	32,653	32,653	32,653
R^2	0.247	0.143	0.337	0.064	0.020
F-test (p-value)	0.374	0.286	0.153	0.455	0.191
Note:					*p<0.1; **p<0.05; ***p<0.01

Table 14: Heterogeneity in Treated User Response - Quartile 3. Users were classified into four quartiles based on past appreciations received. Covariates capture the effect of being in the treatment group within the n-th week following the date on which a user's work was featured. Standard errors shown in parenthesis are clustered at the user and day level.

Treatment interacted with	Dependent variable:				
	Apps. Given	Comments Given	Views Given	Outbound Ties	Projects Published
Week1 post-feature	0.139	0.013	2.400*	0.205	0.034*
	(0.528)	(0.151)	(1.378)	(0.145)	(0.020)
Week2 post-feature	−0.231	0.235	0.451	0.018	0.004
	(0.354)	(0.199)	(0.967)	(0.082)	(0.010)
Week3 post-feature	−0.001	0.241	1.475	0.071	0.016
	(0.483)	(0.152)	(1.199)	(0.086)	(0.012)
Week4 post-feature	0.111	0.287	1.358	0.067	0.003
	(0.331)	(0.178)	(1.080)	(0.125)	(0.012)
User fixed effects	Yes	Yes	Yes	Yes	Yes
Day fixed effects	Yes	Yes	Yes	Yes	Yes
Observations	32,597	32,597	32,597	32,597	32,597
R^2	0.423	0.446	0.393	0.057	0.085
F-test (p-value)	0.399	0.208	0.104	0.391	0.076*
Note:					*p<0.1; **p<0.05; ***p<0.01

Table 15: Heterogeneity in Treated User Response - Quartile 4. Users were classified into four quartiles based on past appreciations received. Covariates capture the effect of being in the treatment group within the n-th week following the date on which a user's work was featured. Standard errors shown in parenthesis are clustered at the user and day level.

	Dependent variable:											
	$\Delta_{ViewsGiven}$			$\Delta_{AppsGiven}$			$\Delta_{OutboundTies}$			$\Delta_{ProjectsPublished}$		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Treatment	3.139***	1.231	0.574	0.602*	0.220	0.006	0.113**	0.051	0.047	0.017**	0.008	−0.007
	(0.849)	(0.927)	(0.903)	(0.318)	(0.352)	(0.345)	(0.046)	(0.051)	(0.050)	(0.008)	(0.009)	(0.008)
$\Delta_{ViewsReceived}$		0.040***			0.008**			0.001***			0.0002**	
		(0.008)			(0.003)			(0.0005)			(0.0001)	
$\Delta_{AppsReceived}$			0.369***			0.086***			0.009***			0.003***
			(0.054)			(0.021)			(0.003)			(0.0005)
Constant	−1.282**	−1.256**	−1.139*	−0.507**	−0.502**	−0.474**	−0.043	−0.042	−0.040	−0.004	−0.004	−0.003
	(0.600)	(0.590)	(0.580)	(0.225)	(0.224)	(0.222)	(0.032)	(0.032)	(0.032)	(0.005)	(0.005)	(0.005)
Observations	658	658	658	658	658	658	658	658	658	658	658	658
R^2	0.021	0.054	0.087	0.006	0.015	0.031	0.009	0.021	0.024	0.008	0.016	0.081
Note:										*p<0.1; **p<0.05; ***p<0.01		

Table 16: Mediation analysis results. Standard errors shown in parenthesis.