



Management Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

The Impact of Social Nudges on User-Generated Content for Social Network Platforms

Zhiyu Zeng, Hengchen Dai, Dennis J. Zhang, Heng Zhang, Renyu Zhang, Zhiwei Xu, Zuo-Jun Max Shen

To cite this article:

Zhiyu Zeng, Hengchen Dai, Dennis J. Zhang, Heng Zhang, Renyu Zhang, Zhiwei Xu, Zuo-Jun Max Shen (2023) The Impact of Social Nudges on User-Generated Content for Social Network Platforms. Management Science 69(9):5189-5208. <https://doi.org/10.1287/mnsc.2022.4622>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2022, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

The Impact of Social Nudges on User-Generated Content for Social Network Platforms

Zhiyu Zeng,^a Hengchen Dai,^b Dennis J. Zhang,^c Heng Zhang,^d Renyu Zhang,^{e,*} Zhiwei Xu,^f Zuo-Jun Max Shen^{g,h}

^aDepartment of Industrial Engineering, Tsinghua University, Beijing 100000, China; ^bAnderson School of Management, University of California, Los Angeles, California 90095; ^cOlin Business School, Washington University in St. Louis, St. Louis, Missouri 63130; ^dW. P. Carey School of Business, Arizona State University, Tempe, Arizona 85287; ^eDepartment of Decision Sciences and Managerial Economics, The Chinese University of Hong Kong, Hong Kong, China; ^fIndependent Contributor, Beijing 100000, China; ^gDepartment of Industrial Engineering and Operations Research, University of California, Berkeley, Berkeley, California 94720; ^hDepartment of Civil and Environmental Engineering, University of California, Berkeley, Berkeley, California 94720

*Corresponding author

Contact: zengzhiy18@mails.tsinghua.edu.cn, <https://orcid.org/0000-0003-0002-876X> (ZZ); hengchen.dai@anderson.ucla.edu, <https://orcid.org/0000-0001-7640-6558> (HD); denniszhang@wustl.edu, <https://orcid.org/0000-0002-4544-775X> (DJZ); hengzhang24@asu.edu, <https://orcid.org/0000-0002-6105-6994> (HZ); philipzhang@cuhk.edu.hk, <https://orcid.org/0000-0003-0284-164X> (RZ); rickyzhiwei@gmail.com (ZX); maxshen@berkeley.edu, <https://orcid.org/0000-0003-4538-8312> (Z-JMS)

Received: July 8, 2021

Revised: March 3, 2022

Accepted: April 19, 2022

Published Online in Articles in Advance:
December 9, 2022

<https://doi.org/10.1287/mnsc.2022.4622>

Copyright: © 2022 INFORMS

Abstract. Content-sharing social network platforms rely heavily on user-generated content to attract users and advertisers, but they have limited authority over content provision. We develop an intervention that leverages social interactions between users to stimulate content production. We study *social nudges*, whereby users connected with a content provider on a platform encourage that provider to supply more content. We conducted a randomized field experiment ($N = 993,676$) on a video-sharing social network platform where treatment providers could receive messages from other users encouraging them to produce more, but control providers could not. We find that social nudges not only immediately boosted video supply by 13.21% without changing video quality but also, increased the number of nudges providers sent to others by 15.57%. Such production-boosting and diffusion effects, although declining over time, lasted beyond the day of receiving nudges and were amplified when nudge senders and recipients had stronger ties. We replicate these results in a second experiment. To estimate the overall production boost over the entire network and guide platforms to utilize social nudges, we combine the experimental data with a social network model that captures the diffusion and over-time effects of social nudges. We showcase the importance of considering the network effects when estimating the impact of social nudges and optimizing platform operations regarding social nudges. Our research highlights the value of leveraging co-user influence for platforms and provides guidance for future research to incorporate the diffusion of an intervention into the estimation of its impacts within a social network.

History: Accepted by Victor Martínez-de-Albéniz, operations management.

Funding: H. Dai thanks the University of California, Los Angeles (UCLA) [Hellman Fellowship and Faculty Development Award] for funding support. R. Zhang is grateful for financial support from the Hong Kong Research Grants Council [Grant 16505418].

Supplemental Material: The data files and online appendix are available at <https://doi.org/10.1287/mnsc.2022.4622>.

Keywords: content production • platform operations • social network • field experiment • information-based intervention

1. Introduction

Online content-sharing social network platforms such as Facebook and TikTok, where users create and consume content, are playing an increasingly important role in society. As of January 2021, an estimated 4.2 billion people, 53.6% of the world's population, were using these platforms.¹ They have evolved into powerful marketing tools, reshaping the global economy. For example, advertising spending on these types of platforms is expected to reach U.S. \$230.30 billion in 2022.² User-generated

content (UGC) on these platforms can exert considerable influence on consumer decision making, affecting sales of products and services (see, e.g., Chen et al. 2011).

These platforms, by nature, rely heavily on UGC to engage and retain users and advertisers alike. However, because users who generate organic content (“content providers”) are not paid workers and UGC is essentially a public good, platforms have limited control over how often users produce content, how much, and at what quality level (Yang et al. 2010, Gallus 2017). Hence, the

underprovision of UGC has been a challenge that interests both practitioners (Pew Research Center 2010) and academics (Burtch et al. 2018, Huang et al. 2019, Kuang et al. 2019). Understanding drivers of content production and devising effective operational levers to motivate content production are vital for content-sharing social network platforms—this is the focus of our research.

A prominent feature of these platforms is that users have intensive social interactions with each other. The platforms can leverage the connections between users to stimulate UGC supply, as well as to help solve other operational problems. We study a novel kind of intervention that utilizes existing connections between users, capitalizes on psychological principles about when people are motivated to exert effort, and contains no financial incentives. Specifically, we study *social nudges* implemented by a user's neighbors on a platform (i.e., platform users who are connected to this user) to explicitly encourage her to supply more content on the platform.³ We propose that by taking the time to explicitly encourage the user to produce more, neighbors convey that they value the user and her existing work and at the same time, communicate their interest in viewing more of the user's future content. This may make the user feel more competent and valued (Ryan and Deci 2000) and increase her confidence in her future work receiving continued appreciation, which further motivates content provision (Grant and Gino 2010, Bradler et al. 2016).

Prior psychological and management research suggests that recognition from managers, companies, or platforms (Ashraf et al. 2014a, b; Bradler et al. 2016; Banya 2017; Gallus 2017) can boost recipients' production and retention. However, scant research has causally examined the motivating power of pure *peer recognition* that is not accompanied by financial incentives; moreover, this limited work has presented mixed evidence for the effectiveness of peer recognition in boosting production (Restivo and van de Rijt 2014, Gallus et al. 2020). Also, prior research has been silent about how interactions on a platform and its underlying social network could reinforce the effects of an intervention on production. Taking a more holistic perspective, we implemented large-scale field experiments to not only estimate the direct effects of our intervention (social nudges) on recipients' content production but also, assess how being exposed to the intervention facilitates the spread of the intervention, which further stimulates additional recipients' content production. We then incorporated empirical findings from these field experiments into a social network model to estimate the impact of our intervention on content production over the entire social network.

Specifically, we conducted two randomized field experiments on a large-scale video-sharing social network platform (hereafter "Platform O" to protect its identity). As on Facebook, each user on Platform O can play two roles: content provider and content viewer.

Users can follow other users and be followed. In this setting, we refer to a user's followers and to the users whom the user herself follows as *neighbors*.

We study social nudges sent by one type of neighbor: a user's followers. For users involved in our experiments, their followers could send them a message to convey the interest in seeing their videos and nudge them to upload more videos. Users in our experiments were randomly assigned to either the treatment or the control condition. The only difference introduced by our experimental manipulation between the two conditions was whether users could actually receive social nudges; treatment users could receive social nudges sent by their neighbors, but control users could not. Because the difference between the two groups of users is in their roles as providers and our primary focus was content production, we hereafter refer to users involved in our experiments as *providers*. We conducted our main experiment—the focus of this paper—from September 12 to 14, 2018 and our second replication experiment from September 14 to 20, 2018.

Analyses about 993,676 providers in our main experiment yield several important insights. To begin with, we present four main findings about the effects of social nudges on recipients' content production (direct effects of social nudges on production). First, receiving social nudges boosted the number of videos that treatment providers uploaded on the day they received the first nudges by 13.21%, without causing providers to alter their video quality. This in turn increased consumption of treatment providers' content by 10.42%. Second, receiving a social nudge yielded a larger immediate boost in production when a provider and the follower who sent the nudge had a two-way tie (i.e., the provider was also following the follower; 17.39%) than when they had a one-way tie (i.e., the provider was not following the follower; 9.37%), suggesting that stronger ties between users strengthen the effect of social nudges on production. Third, the effect of receiving social nudges on production declined over time but remained significant within three days of receiving social nudges (a relative increase of 13.21% on the day of receiving social nudges versus 5.29% and 2.54% on the first and second days afterward, respectively). Fourth, leveraging data from another experiment on Platform O that studied nudges sent to providers *by the platform*, we find suggestive evidence that social nudges from peer users can more effectively boost production than platform-initiated nudges.

Next, we examine whether providers receiving social nudges became more likely to send nudges to users they follow, which if holding true, could further boost production on the platform (indirect effects of social nudges on production). We present three key findings about nudge diffusion. First, treatment providers sent 15.57% more social nudges on the day of receiving social nudges

relative to control providers. Second, receiving a social nudge had a stronger effect on providers' willingness to send social nudges when they got a nudge from a two-way tie (29.97%) versus from a one-way tie (2.87%). Third, the diffusion effect of social nudges declined over time and was significant within two days of receiving social nudges (a 15.57% increase on the day of receiving social nudges versus a 7.87% increase on the following day).

The diffusion of social nudges by nudge recipients as well as the over-time effects of social nudges impose challenges for estimating the impact of social nudges on production and in turn, optimizing platform operational strategies regarding social nudges in different scenarios. We refer to the stationary effect of social nudges on content production on the entire social network—where every user could receive and send social nudges—as the *global effect* of social nudges. To precisely estimate this effect, we propose an infinite-horizon stochastic social network model. We model the social network embedded on Platform O as a directed graph, in which each user is a *node* and each following relationship is an *edge*. Based on our empirical evidence, the actual number of nudges sent on an edge in a period depends on both (1) the baseline number of nudges that would be sent without the influence of nudge diffusion and (2) the number of nudges its origin has received (i.e., the diffusion of nudge). Each user's production boost in a period is determined by all the social nudges she has received.

We also incorporate the time-decaying effect of both direct and indirect effects of social nudges with estimated decaying factors. Leveraging such a social network model, we provide a framework to estimate the global effect of social nudges on production boost, and we show that simply comparing the number of videos uploaded by treatment versus control providers right after they were sent social nudges during the field experiment severely underestimates the global effect of social nudges. Moreover, based on this model, we devise a variant of the Bonacich centrality for edges (BCE), and we further develop the social nudge index (SNI) of each edge that quantifies the total production boost attributed to this edge. Via simulation, we showcase that platforms can use the SNI to optimize operational decisions, such as optimal seeding and provider recommendation for new users, highlighting this model's potential to improve platform performance in various settings.

In summary, we study a low-cost, behaviorally informed intervention that is initiated by neighbors on online platforms and can be widely applied to content providers on a platform. Empirically, we document both its direct production-boosting effect and its diffusion by intervention recipients. Theoretically, we develop a model to incorporate its diffusion into a social network model, thus allowing for a precise estimate of its global effect on production over the entire platform, as well as optimization

of its overall effectiveness. Methodologically, our work provides guidance to future researchers for more comprehensively estimating an intervention's causal effects on a social network. Practically, our proposed low-cost, psychology-based intervention is valuable to online content-sharing social network platforms for managing their UGC, and our model can be a useful tool for platforms to evaluate and optimize the strategy for increasing the global effect of an intervention on a social network.

The rest of the paper proceeds as follows. Section 2 reviews the relevant literature. Section 3 introduces our field setting, experimental design, and data. Sections 4 and 5 present the direct effects of social nudges on content production and the diffusion of nudges, respectively. Section 6 describes the social network model, counterfactual analyses, and two practical applications illustrating the operational implications of our model. In Section 7, we discuss practical implications of our research and directions for future research.

2. Literature Review

Our research builds primarily on four streams of literature: production, peer effects and social networks, information-based interventions, and platform operations.

2.1. Production

Our work is most closely connected to research that seeks to motivate content generation on online content-sharing platforms. The interventions examined in prior work include financial incentives (e.g., rewarding content providers with money) (Cabral and Li 2015, Burtch et al. 2018, Kuang et al. 2019), social norms (e.g., informing content providers about what most of their peers do) (Chen et al. 2010, Burtch et al. 2018), performance feedback (e.g., informing content providers about their performance) (Huang et al. 2019), hierarchies (e.g., ranking content providers based on their contributions to a website) (Goes et al. 2016), symbolic awards (e.g., giving content providers badges based on their recent activities) (Ashraf et al. 2014a, Restivo and van de Rijt 2014, Gallus 2017), and a combination of these tools (Burtch et al. 2018, 2022).

Our contribution to this literature is threefold. First, we study a novel intervention (social nudges) that leverages individual to individual peer recognition, contains no material incentives, and is applicable to all content providers on a platform. Apparently, social nudges differ fundamentally from financial incentives, social norms, performance feedback, and hierarchies. Additionally, although social nudges are related to symbolic awards in the sense that both convey recognition without monetary incentives, awards must be given to a select body of users who deserve them (e.g., users who recently contributed UGC, top-performing users) in order to maintain their prestige and meaning,

and thus, their scope is more limited than that of social nudges.

Second, the nascent literature that examines recognition-based interventions (Frey and Gallus 2017) has mostly studied recognition communicated by authoritative figures such as managers and organizations (Ashraf et al. 2014a, Gallus 2017). The scant work examining the causal effect of peer recognition without financial incentives (Restivo and van de Rijt 2014, Gallus et al. 2020) presents inconclusive evidence for whether peer recognition can increase users' contributions. Specifically, Restivo and van de Rijt (2014) conducted a field experiment among the top 10% of providers to Wikipedia. They found that peer recognition increased production only among the most productive 1% providers but did not affect other providers who were relatively less productive (those at the 91st to 99th percentiles). If anything, the treatment *reduced* retention of providers at the 91st to 95th percentiles. Such negative effect of peer recognition might occur because providers who were not the most prolific (e.g., those at the 91st to 99th percentiles) did not see themselves as sufficiently qualified to receive the recognition given that they had not received any recognition before and the recognition in the experiment came from experimenters who pretended to be peer users. In a field experiment among the workforce at the National Aeronautics and Space Administration (NASA), Gallus et al. (2020) found a null effect of peer recognition on individuals' contributions to a NASA crowdsourcing platform. Peer recognition may fail to motivate in this context because NASA employees did not perceive the recognized activity as part of their core work and thus, did not view peer recognition as legitimate or meaningful. Thus, it remains an open question whether an intervention that conveys peer recognition can boost recipients' effort provision on a UGC social network platform. We speak to this open question by implementing large-scale field experiments to test the effectiveness of an intervention that conveys peer recognition.

Third, prior studies have focused on testing the effects of an intervention on targets' content production, but they have rarely focused on whether and how the intervention diffuses (i.e., how a user, upon receiving the intervention, spreads and applies it to influence other users). We take a critical first step in this direction by not only empirically examining the diffusion of social nudges but also, incorporating the diffusion process into our social network model to more accurately estimate the impact of our intervention on content production over the entire social network.

Within the production literature, our research is also related to prior studies on how to lift productivity in service and manufacturing settings. These studies have focused on four types of interventions for increasing productivity: those that (1) are based on workers' economic

considerations (Lazear 2000, Celhay et al. 2019), (2) offer workers training (De Grip and Sauermann 2012, Konings and Vanormelingen 2015) or introduce information technology (Tan and Netessine 2020), (3) assign workers to various staffing or workload settings (Tan and Netessine 2014, Moon et al. 2022), and (4) capitalize on workers' psychological needs and tendencies (Kosfeld and Neckermann 2011, Roels and Su 2014, Song et al. 2018). These interventions are usually implemented by firms or managers. Extending this line of work, we develop and test a novel psychology-based intervention that does not originate from firms or managers but instead, leverages peer recognition to motivate effort provision and production.

2.2. Peer Effects and Social Networks

Research about peer effects (Zhang et al. 2017, Bramoullé et al. 2020) often investigates how schoolmates (Sacerdote 2001, Whitmore 2005), coworkers (Mas and Moretti 2009, Tan and Netessine 2019), family members (Nicoletti et al. 2018), residential neighbors, and friends (Kuhn et al. 2011, Bapna and Umyarov 2015) affect someone's own behaviors, ranging from mundane consumption and product adoption to consequential outcomes about education, health, and career.

We extend this literature about peer effects in two ways. First, prior research usually estimates peer effects without distinguishing whether peers exert influence passively (e.g., peers' choices are observed by others who then feel pressure to choose accordingly) or actively (e.g., peers persuade others to make certain choices). We clearly assess the active impact of peers by examining a novel kind of interaction initiated by peers because of their intention to influence others (i.e., peers send nudges to others in the hope of boosting others' production). Second, whereas prior research has normally focused on the effects of peers' outcomes (or behaviors) on another person's outcomes (or behaviors) in the same domain, our work simultaneously examines how peers actively influence another person's production via sending a social nudge as well as how the nudged person subsequently "learns," adopts the same tactic, and spreads this form of active influence via sending nudges to more peers.

Besides peer effects, we also speak to the literature that optimizes operational objectives based on social network models, such as identifying key users (Ballester et al. 2006), seeding (Zhou and Chen 2016, Candogan and Drakopoulos 2020, Gelper et al. 2021), pricing (Candogan et al. 2012, Papanastasiou and Savva 2017, Cohen and Harsha 2020), and advertising (Bimpikis et al. 2016). Drawing insights from this literature, we propose an infinite-horizon stochastic social network model to characterize user interactions in a social network that allows for the precise calculation and optimization of an intervention's global effect. Our work takes this literature one step further by leveraging

causal estimates from field experiments to calibrate model parameters, leading to an end to end implementation of such an optimization strategy.

2.3. Information-Based Interventions

Our work adds to the emergent operations management literature that empirically tests the effectiveness of information-based interventions in solving operational problems. This literature has examined such interventions as offering customers more information about firms and the market (Buell and Norton 2011, Parker et al. 2016, Cui et al. 2019, Li et al. 2020, Mohan et al. 2020, Xu et al. 2021) and offering service providers more information about customers (Buell et al. 2017, Cui et al. 2020a, Zeng et al. 2022). These interventions have been shown to increase customers' engagement with firms and perceived service value as well as to improve service speed and capacity. We contribute to this literature by designing a novel information-based intervention that originates from neighbors within a social network and then, causally demonstrating its production-boosting effect and diffusion.

2.4. Platform Operations

Finally, our research extends the growing literature that addresses operations problems on online platforms. This literature has examined how to build effective systems for pricing (Cachon et al. 2017, Bai et al. 2019, Bimpikis et al. 2019, Zhang et al. 2020), recommendations (Banerjee et al. 2016, Mookerjee et al. 2017), staffing rules (Gurvich et al. 2019), and optimization of content production (Caro and Martínez-de Albéniz 2020); it has also studied how to estimate and leverage the spillover effects across platform users (Zhang et al. 2019, 2020) and how to ensure service quality (Cui et al. 2020b, Kabra et al. 2020). We contribute to this literature by empirically demonstrating that allowing platform users to send social nudges—a low-cost, easy to implement strategy—could lift content production and in turn, total capacity and consumption on content-sharing platforms.

3. Field Setting, Experiment Design, and Data

3.1. Field Setting and Experimental Design

To empirically examine the impact of social nudges, we collaborated with Platform O, where each user can play two roles simultaneously—content provider and content viewer. Content providers (1) can upload videos for distribution on Platform O, (2) can decide when and what to post, and (3) do not get paid by Platform O for uploading videos. Content viewers can watch videos for free. Platform O, like most online content-sharing platforms, generates revenue primarily through online advertising (i.e., disseminating advertising videos to users).

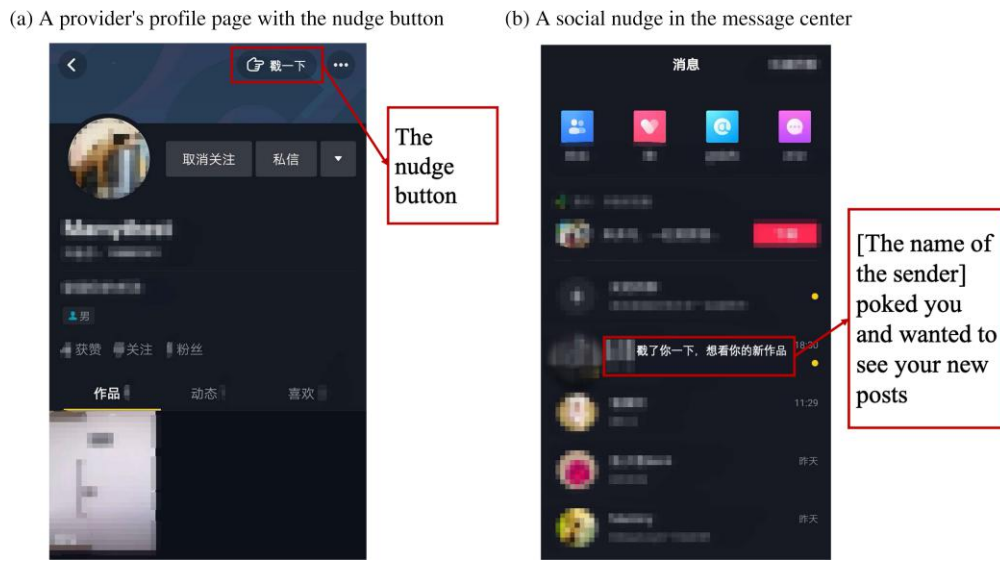
Videos on Platform O are usually short, typically just a few seconds to a few minutes. Popular subjects include daily lives (e.g., views of a nearby park, work scenes, kids, pets), jokes or funny plots, performance (e.g., dancing, singing, making art), and know-how (e.g., cooking or makeup tips). Video content is usually displayed to users on one of three pages: (1) the page of videos uploaded by providers they follow, (2) that of popular videos recommended by Platform O, and (3) that of videos from providers who are geographically close to a given user.

When watching a video, users can leave comments beneath the video and upvote it by clicking the *like* button. The only way for users to privately and directly communicate with each other on Platform O is through the private message function. To establish closer relationships, users can follow others by clicking the “follow” button (available at the top of a video or on other users' profile page).

We conducted two randomized field experiments to causally test how social nudges from neighbors affected users' video production. Our first experiment lasted from 2 p.m. on September 12, 2018 to 5 p.m. on September 14, 2018. This is our main study. Our second field experiment, which replicates the first experiment, lasted from 5 p.m. on September 14, 2018 to the end of September 20, 2018. This experiment (see Online Appendix B for the data and results) targeted a smaller, nonoverlapping group of providers but lasted longer.

For providers involved in our experiments, their followers could send them a standard message to nudge them to upload new videos if they had not published videos for one or more days.⁴ To do so, followers simply clicked a button on the provider's profile page that read “Poke this provider” (ChuoYiXia in Chinese) (see Figure 1(a)).⁵ We refer to this behavior as “sending a social nudge.”

Providers in our experiments were randomly assigned to either the treatment or the control condition. The *only* factor that we manipulated between the two conditions was whether providers were able to view social nudges sent to them. Specifically, treatment providers could see social nudges sent to them in their message center along with other kinds of messages, whereas control providers could *not* see the social nudges in their message center. The standard social nudge message to all providers said “[name of the sender] poked you and wanted to see your new posts” (see Figure 1(b)).⁶ If treatment providers clicked on a social nudge message, they would be directed to a list of all nudges that had ever been sent to them. On that page, newer nudges were displayed closer to the top. There, each social nudge message read “[name of the sender] poked you [time when the nudge was sent] and wanted to see your new posts.” We designed these social nudges to be bare bones, simple, and standardized so as to examine as cleanly as possible the basic effect of being nudged by a neighbor.

Figure 1. (Color online) How Social Nudges Are Sent by Neighbors and Displayed to Treatment Providers

3.2. Data and Randomization Check

For the main analyses, our sample of providers ($N = 993,676$) included all treatment providers and control providers who satisfied two criteria; (1) at least one of their followers sent them a social nudge during our experiment, and (2) they had never received any social nudges before the experiment.⁷ Treatment and control providers in our sample preserved the benefits of random assignment because our random assignment of providers into the treatment condition versus the control condition had no way of affecting whether and when their neighbors sent them the first social nudge during the experiment. To confirm the success of randomization among our sample of providers, we compared treatment providers ($n = 496,976$) and control providers ($n = 496,700$) in their gender, basic network characteristics, and preexperiment production statistics. As shown in Table 1, treatment and control providers in our sample had similar proportions of female providers, number of users who were following them (“number of followers”) on the day prior to the experiment, and number of users they were following (“number of following”) on the day prior to the experiment, as well as the number of videos they uploaded and the number of days when they uploaded any video during the week prior to the experiment. These results confirm that the treatment and control providers in our sample were comparable, suggesting that any difference between conditions after the experiment started should be attributed to our experimental manipulation—that is, whether providers could actually receive social nudges.

To protect Platform O’s sensitive information,⁸ we standardized all continuous variables used in our

analyses to have a unit standard deviation. To help readers better understand our empirical context, we report the scaled or standardized distributional information of relevant variables and network features in Online Appendix G. We also provide the code for our empirical and simulation analyses in a GitHub repository.⁹

4. Direct Effects of Social Nudges on Content Production

Our investigation began by examining the effects of receiving social nudges on the recipient’s content production (i.e., the direct effects of social nudges on content production). The time unit we focused on was one day, which matches the granularity of our data offered by Platform O. Platform O cares about aggregate daily metrics (e.g., daily active providers, daily new videos), which break down to daily metrics at the individual level (e.g., on a given day, whether a user uploaded any video, how many videos she uploaded). In addition, 79% of providers in our sample had median intervals of video postings¹⁰ at least one day, further confirming the appropriateness of using one day (rather than a smaller time window, such as one hour) as the time unit.

4.1. Direct Effects of Social Nudges on Content Production on the First Reception Day

We first tested whether social nudges had a positive effect on content production on the first day when a provider could be affected—that is, the day a provider was sent the first social nudge during the experiment; we refer to it as the providers’ *first reception day*. Most (97%) providers in our sample were sent only *one* social nudge on the first reception day, suggesting that the

Table 1. Randomization Check

	Treatment providers (1)	Control providers (2)	<i>p</i> -Value of two-sample proportion test or <i>t</i> test (3)
Statistics on the day prior to the experiment			
<i>Proportion of Females</i>	51.34%	51.38%	0.82
<i>Number of Followers</i>	0.0622	0.0605	0.38
<i>Number of Following</i>	0.8485	0.8480	0.81
Statistics during one week prior to the experiment			
<i>Number of Uploaded Videos</i>	0.3674	0.3693	0.33
<i>Number of Days with Videos Uploaded</i>	0.5057	0.5078	0.30

Notes. All variables, other than whether a provider is a female, were standardized to have a unit standard deviation. To calculate the proportion of females, we excluded the 8,895 providers (~0.9%) with missing gender information.

effects of our intervention on the first reception day were driven mostly by receiving one social nudge. Our unit of analysis was a provider on her first reception day; we analyzed 993,676 observations, with each provider contributing one observation.

We used the following ordinary least squares regression specification with robust standard errors to causally estimate the effects of social nudges on the first reception day:

$$\text{Outcome Variable}_i = \beta_0 + \beta_1 \text{Treatment}_i + \epsilon_i, \quad (1)$$

where *Outcome Variable_i* is detailed later and *Treatment_i* is a binary variable indicating whether provider *i* was in the treatment (versus control) condition.

For each provider *i*, we first examined the number of videos she uploaded on the first reception day (*Number of Videos Uploaded_i*). Column (1) of Table 2 reports the result of a regression that follows specification (1) to predict *Number of Videos Uploaded_i*. The positive and significant coefficient on treatment indicates that receiving social nudges immediately had a positive effect on the nudge recipient’s production. Specifically, receiving social nudges increased the number of videos uploaded on the first reception day by 0.0262 standard deviations (*p* < 0.0001), a 13.21% increase relative to the average in the control condition.

Two underlying forces may drive this production-boosting effect: (1) providers became more willing to upload at least one video on the first reception day, and (2) providers who decided to upload at least one video on the first reception day uploaded more videos that day. To test the presence of the first force, for each provider *i*, we examined whether she uploaded at least one video on the first reception day (*Upload Incidence_i*). To test the presence of the second force, we examined the number of videos uploaded on the first reception day among providers who uploaded at least one video that day (*Number of Videos Uploaded Conditional on Uploading Anything_i*).

We used regression specification (1) to predict *Upload Incidence_i* and *Number of Videos Uploaded Conditional on Uploading Anything_i*. Column (2) of Table 2 shows that receiving social nudges lifted the average probability of

providers uploading any videos on the first reception day by 0.94 percentage points (*p* < 0.0001), a 13.86% increase relative to the average probability in the control condition. However, as shown in column (3) of Table 2, *Number of Videos Uploaded Conditional on Uploading Anything_i* did not statistically significantly differ between conditions (*p* = 0.3533). Altogether, these results suggest that the boost in video supply on the first reception day was mainly driven by the first force—that is, providers became more willing to upload something after receiving social nudges.

Inspired by the social network literature (e.g., Jackson 2005), we next examined whether social nudges from closer peers could be more motivating. To answer this question, we tested whether the direct effects of social nudges on content production became stronger if a provider was also following the follower who sent her a nudge (in which case we refer to the relationship between the provider and the nudge sender as a two-way tie) than if the provider was not following that follower (in which case we refer to their relationship as a one-way tie). For each provider *i* on her first reception day, we identified the follower who sent the first social nudge to provider *i* (i.e., the *first social nudge sender*). We constructed a binary variable, *Two-Way Tie_i*, which equals one if provider *i* was also following her first social nudge sender and zero otherwise. We used the following regression specification with robust standard errors to predict *Number of Videos Uploaded_i*, where each observation was a provider on her first reception day:

$$\begin{aligned} \text{Outcome Variable}_i = & \beta_0 + \beta_1 \text{Treatment}_i + \beta_2 \text{Two-Way Tie}_i \\ & + \beta_3 \text{Treatment}_i \times \text{Two-Way Tie}_i + \epsilon_i. \end{aligned} \quad (2)$$

Column (4) of Table 2 shows that the coefficient on the interaction between *Treatment_i* and *Two-Way Tie_i* is significant and positive (*p* < 0.001). This suggests that, consistent with the social network literature (Jackson 2005), receiving social nudges increased a provider’s content production to a greater extent when the provider and the follower who sent the nudge had a two-way tie than

Table 2. Direct Effects of Social Nudges on Content Production on the First Reception Day

Outcome variable	Main treatment effects			Heterogeneous treatment effect
	Number of Videos Uploaded (1)	Upload Incidence (2)	Number of Videos Uploaded Conditional on Uploading Anything (3)	Number of Videos Uploaded (4)
Treatment	0.0262**** (0.0020)	0.0094**** (0.0005)	-0.0168 (0.0181)	0.0186**** (0.0025)
Two-Way Tie				0.0700**** (0.0027)
Treatment × Two-Way Tie				0.0159**** (0.0041)
Relative effect size, %	13.21	13.86		
Observations	993,676	993,676	71,883	993,676

Notes. Continuous variables (*Number of Videos Uploaded* and *Number of Video Uploaded Conditional on Uploading Anything*) were standardized to have a unit standard deviation before entering the regressions. The unit of analysis for all columns was a provider on her first reception day. Columns (1), (2), and (4) include all providers in our sample. Column (3) includes the providers who uploaded at least one video on their first reception day. Robust standard errors are reported in parentheses.

*** $p < 0.001$; **** $p < 0.0001$.

when they had a one-way tie. Specifically, receiving a social nudge from a follower with a one-way tie boosted the number of videos uploaded on the first reception day by 0.0186 standard deviations ($p < 0.0001$), whereas receiving a social nudge from a follower with a two-way tie boosted the number of videos uploaded by 0.0345 (i.e., $0.0186 + 0.0159$) standard deviations ($p < 0.0001$). The relative effect sizes, compared with the average number of videos uploaded in the control condition, are 9.37% (one-way tie) and 17.39% (two-way tie), respectively.

4.2. Direct Effects of Social Nudges on Content Consumption and Content Quality

Beyond video production, how do social nudges affect overall video consumption and video quality? To evaluate the direct effects of social nudges on video consumption, we focused on the total number of views each provider engendered that could be attributed to videos they uploaded on the first reception day. Following Platform O's common practice, for each video uploaded on a provider's first reception day, we tracked the total number of views it received over the first week since its creation. Platform O normally uses the views each video accumulates during the first week after its creation to capture the short-term consumption it brings because videos on Platform O are usually watched much more frequently during the first week and attract fewer views as time goes by. Then, for each provider i , $Total\ Views_i$ equals the total number of views within one week across all videos that provider i uploaded on the first reception day. If provider i did not upload videos on the first reception day, $Total\ Views_i$ equals zero, which reflects the fact that no views were engendered by provider i as a result of her production effort on the first reception day. To address outliers, we winsorized $Total\ Views_i$ at the 95th percentile of nonzero values.¹¹

We used regression specification (1) to predict $Total\ Views_i$. As shown in column (1) of Tables 3 and 4, receiving social nudges increased the total views providers contributed to the platform as a result of their production effort on the first reception day by 0.0171 standard deviations, a 10.42% increase relative to the average in the control condition.¹²

To assess video quality, for every video uploaded by provider i on her first reception day, we collected four quality measures based on viewer engagement during the following week. Then, for provider i , we calculated the average of each quality measurement across these videos: (1) the average percentage of times viewers watched a video until the end (*Complete View Rate_i*), (2) the average percentage of viewers who gave likes to a video (*Like Rate_i*), (3) the average percentage of viewers who commented on a video in the comments section beneath it (*Comment Rate_i*), and (4) the average percentage of viewers who chose to follow provider i while watching a video (*Following Rate_i*).

We used regression specification (1) to predict *Complete View Rate_i*, *Like Rate_i*, *Comment Rate_i*, and *Following Rate_i*. Columns (2), (4), and (5) of Table 3 indicate that social nudges did not significantly alter the complete view rate, comment rate, and following rate of videos uploaded on the first reception day (all p -values are > 0.4). Column (3) suggests that videos uploaded by treatment providers on the first reception day were less likely to receive likes by 0.0174 standard deviations (1.48%) relative to videos uploaded by control providers ($p < 0.05$). To explore this difference in like rates, we further compared historical like rates between treatment and control providers who uploaded any videos on their first reception day. *Historical Like Rate_i* equals the total number of likes provider i received from January 1, 2018 to the day prior to the experiment divided by the total number of views provider i received during that same period.

Table 3. Effects of Social Nudges on Video Consumption and Quality: Main Treatment Effects

Outcome variable	Total Views (1)	Complete View Rate (2)	Like Rate (3)	Comment Rate (4)	Following Rate (5)
Treatment	0.0171**** (0.0020)	0.0007 (0.0075)	-0.0174* (0.0075)	-0.0068 (0.0075)	0.0041 (0.0075)
Observations	993,676	71,634	71,634	71,634	71,634
Relative effect size, %	10.42		-1.48		

Notes. All continuous variables were standardized to have a unit standard deviation before entering the regressions. The unit of analysis for all columns was the provider level. Column (1) includes all providers in our sample. Columns (2)–(5) include providers whose videos uploaded on their first reception day were watched at least once in the following week. Robust standard errors are reported in parentheses.

* $p < 0.05$; **** $p < 0.0001$.

Column (1) of Table 4 shows that among these providers who uploaded videos on the first reception day, treatment providers’ historical like rates were significantly lower than control providers’ historical like rates by 0.0522 standard deviations (3.48%). This difference in historical like rates between treatment and control providers who uploaded videos on the first reception day could lead the like rates for videos uploaded on the first reception day to be lower in the treatment condition than in the control condition. In fact, when we predicted *Like Rate_t* while controlling for *Historical Like Rate_t*, the coefficient on treatment was no longer significant (column (2) in Table 4). Altogether, we find that social nudges did not directly cause providers to increase or decrease video quality.

4.3. Direct Effects of Social Nudges on Content Production over Time

So far, we have shown that social nudges significantly lifted providers’ willingness to upload videos on the first reception day, which in turn, led them to contribute more views to the platform but did not change video quality. Next, we explored how the effect of

Table 4. Effects of Social Nudges on Video Consumption and Quality: Investigating Why Treatment Providers Had Lower Like Rates than Control Providers

Outcome variable	Historical Like Rate (1)	Like Rate (2)
Treatment	-0.0522**** (0.0085)	0.0081 (0.0062)
Historical Like Rate		0.5185**** (0.0070)
Observations	69,825	69,594
Relative effect size, %	-3.48	

Notes. All continuous variables were standardized to have a unit standard deviation before entering the regressions. The unit of analysis for all columns was the provider level. Columns (1) and (2) include providers whose videos uploaded on their first reception day were watched at least once in the following week and whose earlier videos were watched at least once between January 1, 2018 and the day prior to the experiment (September 11, 2018). Robust standard errors are reported in parentheses.

**** $p < 0.0001$.

receiving social nudges on content production changed over time. We compared the number of videos uploaded each day between treatment and control providers from the first reception day until the first day when the difference between conditions was not statistically significant. Specifically, for each day t starting from the first reception day (where t equals $1, 2, \dots$ and $t = 1$ refers to the first reception day itself), we predicted the number of videos uploaded that day using regression specification (1).

Table 5 shows that the effect of receiving social nudges on content production was largest on the first reception day and decreased as time elapsed, but it was positive and significant for a couple of days. Specifically, the number of videos uploaded was higher in the treatment condition than in the control condition by 13.21% on the first reception day (0.0262 standard deviations; $p < 0.0001$) (column (1) of Table 5), by 5.29% on the day after the first reception day (0.0129 standard deviations; $p < 0.0001$) (column (2) of Table 5), and by 2.54% on the second day after the first reception day (0.0065 standard deviations; $p < 0.0001$) (column (3) of Table 5). The effect of receiving social nudges on the nudge recipient’s production was not significant on the third day after the first reception day ($p = 0.7644$) (column (4) of Table 5).

4.4. Additional Analyses About the Direct Effects of Social Nudges

This subsection is devoted to further discussions and analyses to supplement our main results.

4.4.1. Control Providers’ Resentment. One potential alternative explanation for our observed difference in video production between treatment and control providers is that control providers somehow realized that they could not receive the social nudges sent by their followers, which made them resent the platform and thus, reduce their production. Given that the private message function is the only way for connected users to directly and privately communicate with each other on Platform O, this function is likely the only channel via which followers told control providers about social

Table 5. Over-Time Direct Effects of Social Nudges on Content Production

Outcome variable	Number of Videos Uploaded			
	On day 1 (first reception day) (1)	On day 2 (2)	On day 3 (3)	On day 4 (4)
<i>Treatment</i>	0.0262**** (0.0020)	0.0129**** (0.0020)	0.0065** (0.0020)	0.0006 (0.0020)
Relative effect size, %	13.21	5.29	2.54	
Observations	993,676	993,676	993,676	993,676

Notes. *Number of Videos Uploaded* was standardized to have a unit standard deviation before entering the regressions. The unit of analysis for all columns was a provider on day t relative to the first reception day, where $t = 1$ means the first reception day. Columns (1)–(4) include all providers in our sample. Robust standard errors are reported in parentheses.

*** $p < 0.01$; **** $p < 0.0001$.

nudges they sent. Thus, we conducted two sets of additional analyses about the private message function to address this alternative explanation (see Online Appendix C.1). First, we used the difference-in-differences method to examine whether receiving private messages from followers who sent them social nudges during the experiment negatively affected control providers' content production. Second, we tested whether the treatment effect of social nudges on production differed between providers who received any private message from their first social nudge sender during the experiment versus providers who did not. For both analyses, we find no evidence supporting the alternative explanation based on control providers' resentment.

4.4.2. Role of Likes and Comments. Because receiving social nudges could boost video production, nudge recipients might also receive more likes and comments because of the increased number of videos uploaded, which could in turn motivate nudge recipients to produce more. We tested how much the immediate increase in likes and comments because of the receipt of social nudges contributed to the effect of receiving social nudges on content production after the first reception day (see Online Appendix C.2). We find that the increased numbers of likes and comments are neither the only reason nor the primary reason why the effect of receiving social nudges on content production lasted for days. Indeed, the magnitude of the production-boosting effect of social nudges after the first reception day was decreased only by a slight to moderate amount when we controlled for the quantity of likes and comments providers obtained earlier in the experiment. This observation suggests that receiving social nudges per se is sufficient to boost video production beyond the first reception day, even without additional positive feedback from likes and comments.

4.4.3. Effects of Social Nudges Across Providers with Different Baseline Productivity. Restivo and van de Rijdt (2014) found that a peer recognition intervention motivated only the most productive 1% of content

providers but not providers ranked at the 91st to 99th percentiles. We actually observe that receiving social nudges boosted production among the most productive 1% of providers, the providers ranked at the 91st to 99th percentiles, and the providers ranked below the 91st percentile (see Online Appendix C.4). These results suggest that receiving social nudges is generally effective in motivating content provision across users with different levels of productivity.

4.4.4. Comparison with Platform-Initiated Nudges. To motivate content provision, a platform may also directly nudge its users. To explore whether social nudges from peers are more effective than nudges sent by the platform, we leveraged another randomized field experiment where content providers were randomly assigned to either receive or not receive nudges from Platform O (see Online Appendix C.5). Adopting similar empirical analyses as described in Sections 4.1 and 4.3, we find that social nudges boosted providers' production to a larger extent than platform-initiated nudges.

5. Indirect Effects of Social Nudges on Production via Nudge Diffusion

Going beyond social nudges' direct impact on content production, we next turn to the diffusion of social nudges. Inspired by the diffusion phenomenon in the social network literature (e.g., Zhou and Chen 2016), we focus on how receiving social nudges could affect the number of social nudges sent by the recipient to other providers they were following.

5.1. The Effects of Social Nudges on Nudge Diffusion on the First Reception Day

We began our investigation by testing how receiving social nudges facilitated nudge diffusion on the first reception day—the first day when a provider could be affected by social nudges during our experiment. Our unit of analysis was a provider on her first reception day, and we analyzed 993,676 observations, with each provider contributing one observation. We examined the number of social nudges sent by each provider i to

other providers on the first reception day (*Number of Social Nudges Sent_i*). Similar to how we addressed outliers earlier, we winsorized *Number of Social Nudges Sent_i* at the 95th percentile of nonzero values. We used regression specification (1) to predict *Number of Social Nudges Sent_i*. Column (1) of Table 6 shows that, on average, receiving social nudges increased the number of social nudges providers sent to others on the first reception day by 0.0325 standard deviations (15.57%; $p < 0.0001$).

Next, we tested whether social nudges from closer peers could more effectively facilitate nudge diffusion. Similar to how we analyzed the heterogeneous treatment effect for the direct production-boosting effect of social nudges (Section 4.1), here we examined the heterogeneous treatment effects for nudge diffusion based on whether a provider and the follower sending her a nudge had a two-way tie or a one-way tie. Specifically, we used regression specification (2) to predict *Number of Social Nudges Sent_i*.

Column (2) of Table 6 shows that the coefficient on the interaction between *Treatment_i* and *Two-Way Tie_i* is significant and positive ($p < 0.0001$), suggesting that receiving a social nudge motivated a provider to diffuse social nudges to a greater extent when the provider and the follower who sent the nudge had a two-way tie than when they had a one-way tie. Specifically, receiving a social nudge from a follower with a one-way tie boosted the number of social nudges a provider sent on the first reception day by 0.0060 standard deviations ($p < 0.05$), whereas receiving a social nudge from a follower with a two-way tie boosted the number of social nudges sent by 0.0625 (i.e., $0.0060 + 0.0565$) standard deviations ($p < 0.0001$). The relative effect sizes, as compared with the average number of social nudges sent in the control condition, are 2.87% (one-way tie) and 29.97% (two-way tie). Combining these results with the findings in Section

4.1, we find that receiving social nudges both increased a provider’s own content production to a greater extent and yielded a larger diffusion effect when the provider and the nudge sender were following each other than when only the nudge sender was following the provider, suggesting that social nudges from closer peers were more influential.

5.2. Effects of Social Nudges on Nudge Diffusion over Time

Going beyond the first reception day, we next examined how receiving social nudges affected nudge diffusion over time. Similar to how we analyzed the direct effect of social nudges on content production over time, we compared the number of social nudges providers sent each day between treatment and control conditions from the first reception day on until the first day when the difference between conditions was not statistically significant. Specifically, for each day t starting from the first reception day (where t equals $1, 2, \dots$ and $t = 1$ refers to the first reception day itself), we predicted the number of social nudges sent that day using regression specification (1).

Table 7 shows that the effect of receiving social nudges on the number of social nudges sent was largest on the first reception day and decreased as time elapsed. Specifically, the number of social nudges sent to others was higher in the treatment condition than in the control condition by 15.57% on the first reception day (0.0325 standard deviations; $p < 0.0001$) (column (1) of Table 7) and by 7.87% on the day after the first reception day (0.0139 standard deviations; $p < 0.0001$) (column (2) of Table 7). This effect of receiving social nudges on nudge diffusion was not significant on the second day after the first reception day ($p = 0.1686$) (column (3) of Table 7).

Table 6. Effect of Social Nudges on Nudge Diffusion on the First Reception Day

Outcome variable	<i>Number of Social Nudges Sent</i>	
	(1)	(2)
<i>Treatment</i>	0.0325**** (0.0020)	0.0060* (0.0023)
<i>Two-Way Tie</i>		0.1304**** (0.0028)
<i>Treatment × Two-Way Tie</i>		0.0565**** (0.0041)
Relative effect size, %	15.57	
Observations	993,676	993,676

Notes. *Number of Social Nudges Sent* was standardized to have a unit standard deviation before entering the regressions. The unit of analysis for all columns was a provider on her first reception day. Columns (1) and (2) include all providers in our sample. Robust standard errors are reported in parentheses.

* $p < 0.05$; **** $p < 0.0001$.

6. A Social Network Model

The reduced-form results reported in Sections 4 and 5 describe the transient and local impacts of social nudges. Platforms may be interested in evaluating the global effect of social nudges: the total impact of social nudges on production in the counterfactual scenario where every user on the platform can send and receive nudges. They may also be interested in optimizing various operational decisions regarding social nudges, such as seeding and recommending providers to new users. However, the over-time effects and diffusion of social nudges, which we document in Sections 4 and 5, impose challenges for these tasks. To tackle these challenges, we propose a novel social network model to capture both the over-time effects and diffusion of social nudges. Applying this model allows us to quantify both the direct and indirect effects of social nudges

Table 7. Effects of Social Nudges on Nudge Diffusion over Time

Outcome variable	Number of Social Nudges Sent		
	On day 1 (first reception day) (1)	On day 2 (2)	On day 3 (3)
Treatment	0.0325**** (0.0020)	0.0139**** (0.0020)	0.0028 (0.0020)
Relative effect size, %	15.57	7.87	
Observations	993,676	993,676	993,676

Notes. Number of Social Nudges Sent was standardized to have a unit standard deviation before entering the regressions. The unit of analysis for all columns was a provider on day t relative to the first reception day, where $t = 1$ means the first reception day. Columns (1)–(3) include all providers in our sample. Robust standard errors are reported in parentheses.

**** $p < 0.0001$.

on content production over time and thus, more accurately estimate the global effect of social nudges.

6.1. The Model and the Global Effect

We model Platform O as a social network, denoted as $G = (V, E)$, in which $V := \{1, 2, 3, \dots, |V|\}$ is the set of nodes (i.e., users on Platform O who can be viewers and providers) and $E := \{1, 2, 3, \dots, |E|\}$ is the set of directed edges (i.e., the “following” relationship on Platform O). We use i, j and e, ℓ to denote nodes and edges, respectively. Let e_o and e_d be the origin and destination, respectively, of edge $e \in E$, so viewer i following provider j is represented as $e = (i, j)$, $e_o = i$, and $e_d = j$. The dynamics of social nudges and their effects on providers’ production are captured using a discrete-time stochastic model with an infinite time horizon. We use t to index the discrete time period (a single day in our empirical context, which is consistent with the business practice of Platform O), where $t = 1$ refers to the period when the social nudge function first becomes available to all users on the platform. In Figure 2, we illustrate the structure of the social network model. If e_o sends e_d a nudge, the recipient, e_d , will not only (1) increase her production but also, (2) send more nudges to other providers she is following, which could further boost other providers’ production. We summarize the notations involved in the social network model in Table 8.

We first model the over-time direct effect of social nudges on production. Let $x_i(t)$ denote the boost of

provider i ’s production in period t because of the social nudges she has received before and during period t . We use $y_e(t)$ to denote the number of nudges sent on edge e (from e_o to e_d) in period t . Let p_e denote the expected additional number of videos provider e_d would upload as a result of receiving one social nudge from viewer e_o on the day the nudge is received. Section 4 shows that, in our field experiment on Platform O, the direct effect of receiving social nudges on production gradually wears off over time. Thus, we capture the dynamic of production increment by the following dynamic equation:

$$x_i(t) = \sum_{1 \leq s \leq t} \alpha_p^{t-s} \sum_{e \in E: e_d = i} p_e y_e(s) + \epsilon_i^x(t), \quad \forall i \in V, \quad (3)$$

where $\alpha_p \in (0, 1)$ denotes the time-discounting factor of social nudges’ direct production-boosting effect. We denote the random noise of production boost for provider $i \in V$ in period t as $\epsilon_i^x(t)$, independent across different providers and periods with zero means.

We next model the diffusion of social nudges. Motivated by the empirical results in Section 5, we assume that the number of social nudges sent on an edge e in period t is driven by two additive factors. First, we let μ_e denote the expected number of nudges sent on edge e that are not affected by the number of nudges e_o herself has received. We refer to μ_e as the expected number of organic nudges and denote $\boldsymbol{\mu} := (\mu_e : e \in E)$. Second,

Figure 2. (Color online) How Social Nudges Influence Users on a Network

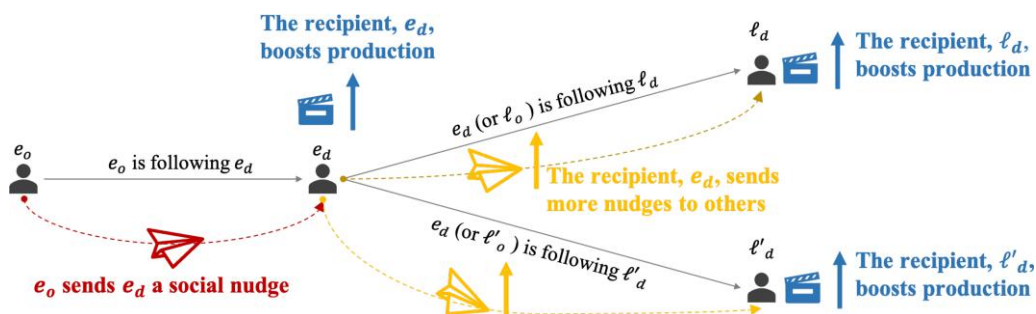


Table 8. Notations Involved in the Social Network Model

Notations	Interpretations
$G = (V, E)$	The network in which V is the set of nodes and E is the set of directed edges
$x_i(t)$	The boost of node i 's production in period t because of nudges node i has received before (including) period t
$y_e(t)$	The number of nudges sent from e_o to e_d in period t
p_e	The additional number of videos provider e_d would be expected to upload in period t as a result of receiving one social nudge from viewer e_o in period t
μ_e	The number of nudges that e_o sends to e_d without being affected by the nudges that e_o has received
$d_{\ell e}$	The expected increase in the number of nudges sent on edge e in period t because of one additional nudge e_o receives in period t from edge ℓ (i.e., $\ell_d = e_o$)
$\epsilon_i^x(t), \epsilon_i^y(t)$	The independent and identically distributed random noises with a zero mean and a bounded support
α_p, α_d	The time-discounting factors corresponding to p_e and $d_{\ell e}$, respectively

the diffusion effect described in Section 5 suggests that when a provider receives a nudge, she tends to send more nudges to other providers she follows. We refer to these social nudges engendered through the diffusion process as *diffused nudges*. Combined, the dynamic of social nudges on the network G is captured by

$$y_e(t) = \mu_e + \sum_{1 \leq s \leq t} \alpha_d^{t-s} \sum_{\ell \in E: \ell_d = e_o} d_{\ell e} y_\ell(s) + \epsilon_e^y(t), \quad \forall e \in E. \tag{4}$$

Here, the second term in Equation (4) embodies the diffusion effect. In particular, $d_{\ell e}$ captures the intensity of social nudge diffusion (i.e., the expected increase in the number of nudges sent on edge e in a given period because of one additional nudge e_o receives in the same period on edge ℓ directing to e_o (that is, $\ell_d = e_o$)). Similar to $\alpha_p, \alpha_d \in (0, 1)$ denotes the time-discounting factor of nudge diffusion, which captures the extent to which the diffusion effect that resulted from a single nudge decays over time, as discussed in Section 5. We denote the random noise of social nudges sent on edge e in period t as $\epsilon_e^y(t)$, independently distributed across different edges and periods with zero means.

Equations (3) and (4), built on the well-established models to study social interactions in the literature (e.g., Ballester et al. 2006, Candogan et al. 2012, Zhou and Chen 2016) and the key empirical observations from our experimental data, are the backbones of our social network model and together, capture the over-time effects and diffusion of social nudges. As we will show in Section 6.2 and Online Appendix E.4, both the estimation of the model parameters (Table 9) and that of different terms (the direct and indirect effects) in the global effect of social nudges (Table 10) are fairly consistent with respect to data from different experiments on Platform O. Such consistency provides further evidence that our model could reasonably capture the interactions observed in our network data.

To quantify the global effect of social nudges, we characterize the long-run steady state of the system defined by Equations (3) and (4). Theorem 1, whose proof is in

Online Appendix D.2, shows that the expected production and nudge quantities converge to a well-defined limit. We define $d_{\ell e} = 0$ if $\ell_d \neq e_o$, and the matrix $\mathbf{D} := (d_{\ell e} : (\ell, e) \in E^2)$. The matrix \mathbf{D} with nonnegative entries therefore captures the first-order diffusion on all edge pairs of the social network. We further define $\eta_e := p_e / (1 - \alpha_p)$ and $\boldsymbol{\eta} := (\eta_e : e \in V)$. We use \mathbf{I} to denote the identity matrix of appropriate dimension. The total production increment in period t is $x(t) := \sum_{i \in V} x_i(t)$. Define a matrix series:

$$\mathbf{M}(k) := \mathbf{I} + \sum_{i=1}^k \frac{1}{(1 - \alpha_d)^i} \mathbf{D}^i, \quad \text{for } k \in \mathbb{Z}_+.$$

A key condition we need here is the convergence of $\mathbf{M}(k)$ to a finite-valued matrix, as $k \rightarrow \infty$. In this case, we say that (α_d, \mathbf{D}) satisfies Condition \mathcal{C} . Note that, because \mathbf{D} is nonnegative, $\mathbf{M}(k)$ is component-wise increasing in k , so $\lim_{k \rightarrow +\infty} \mathbf{M}(k)$ is well defined if and only if $\mathbf{M}(k)$ is component-wise bounded from above. Also, note that Condition \mathcal{C} holds if the ℓ_∞ matrix norm of $1/(1 - \alpha_d)\mathbf{D}$ is strictly below one (Horn and Johnson 2012). Indeed, for the real social network of Platform O, we verify that $\|1/(1 - \alpha_d)\mathbf{D}\|_\infty < 1$, which implies that Condition \mathcal{C} holds (see Online Appendix D.1 for details). Inspired by the classical Bonacich centrality measure

Table 9. Estimation of Parameters in the Social Network Model

Parameter	Estimation results using data from the experiments	
	Main Experiment (1)	Replication Experiment (2)
p_e	0.05492	0.05156
α_p	0.6345	0.6945
d_e	0.0008436	0.0009200
α_d	0.3750	0.3378

Notes. To protect Platform O's sensitive information, we are not permitted to disclose the raw estimates of p_e and d_e . The values of p_e and d_e reported here equal the raw estimates of p_e and d_e multiplied by a fixed constant. We report α_p and α_d using the raw estimates.

Table 10. Estimation of the Global Effect of Social Nudges

	Naïve approach using data from the experiment	Network-modeling approach using data from the experiments	
	Main Experiment (1)	Main Experiment (2)	Replication Experiment (3)
Direct effect	48.65	130.08	One day: 47.55; 146.06 beyond one day: 82.53
Indirect effect		10.59	12.24
Global effect		140.67	166.30
Ratio of indirect effect to direct effect, %		8.14	8.38

Note. When reporting the direct effect estimated by the network-modeling approach, we present the estimated overall direct effect over time (e.g., 130.08 for the first experiment), and we separately show the estimated direct effect on the day of receiving nudges (e.g., 47.55) and the estimated direct effect beyond that day (e.g., 82.53).

defined for nodes in the network economics literature (e.g., Ballester et al. 2006), we define the following *Bona-fide centrality for edges*.

Definition 1. Given the social network G and the associated diffusion matrix \mathbf{D} , we define the BCE measure on E with respect to vector v as

$$\mathcal{BE}(\mathbf{D}, v) := \left(\mathbf{I} - \frac{1}{1 - \alpha_d} \mathbf{D} \right)^{-1} v, \quad (5)$$

where v is real valued with compatible dimension, provided that (α_d, \mathbf{D}) satisfies Condition \mathcal{C} .

We remark that Condition \mathcal{C} guarantees that $\mathbf{I} - (1/(1 - \alpha_d))\mathbf{D}$ is invertible,¹³ so $\mathcal{BE}(\mathbf{D}, v)$ is well defined for any v . The following theorem shows that the global effect of social nudges in the long-run steady state can be characterized by the BCE measure.

Theorem 1. If (α_d, \mathbf{D}) satisfies Condition \mathcal{C} , it then follows that $\lim_{t \rightarrow \infty} \mathbb{E}[x(t)] = x^*$ and $\lim_{t \rightarrow \infty} \mathbb{E}[y(t)] = y^*$, where x^* and y^* satisfy $x^* = \boldsymbol{\eta}^\top y^*$ and

$$y^* = \mathcal{BE}(\mathbf{D}, \boldsymbol{\mu}). \quad (6)$$

In brief, Theorem 1 takes into account the over-time effects and the diffusion of social nudges. Importantly, for any $e \in E$, the BCE measure $\mathcal{BE}_e(\mathbf{D}, \boldsymbol{\mu})$ quantifies the total expected number of nudges user e_o sends to e_d , including both the organic nudges and the diffused nudges. The factors $1/(1 - \alpha_d)$ in Equation (5) and $1/(1 - \alpha_p)$ in the definition of $\boldsymbol{\eta}$ materialize the diffusion and production-boosting effects, respectively, that accumulate over time. As we will show in Section 6.2, under Condition \mathcal{C} , the BCE measure bears a natural expansion with a clear economic interpretation that $\mathcal{BE}(\mathbf{D}, \boldsymbol{\mu})$ can be decomposed according to the radius of nudge diffusion.

6.2. Approximation and Estimation of the Global Effect

By Equation (5), an exact evaluation of the global effect of social nudges on providers' production involves

inverting the $|E|^2 - \text{dimensional}$ matrix $\mathbf{I} - (1/(1 - \alpha_d))\mathbf{D}$. For Platform O, the dimension of $\mathbf{I} - (1/(1 - \alpha_d))\mathbf{D}$ is roughly at the magnitude of 10^{32} , so its inverse is computationally infeasible to obtain. Therefore, we resort to an approximation scheme to quantify the steady-state (daily) number of social nudges between viewers and providers (i.e., y^*) and the (daily) production boost from these nudges (i.e., x^*).

Toward this goal, we note, by Lemma 2 in Online Appendix D.1, that if (α_d, \mathbf{D}) satisfies Condition \mathcal{C} , the inverse of $\mathbf{I} - (1/(1 - \alpha_d))\mathbf{D}$ is given by $\mathbf{I} + \sum_{i=1}^{\infty} (1/(1 - \alpha_d))^i \cdot \mathbf{D}^i$ (Equation (14) in Online Appendix D.1). Motivated by this formula, we define a sequence of (approximate) BCE measures, indexed by $k \in \mathbb{Z}_{+,+}$ as

$$\widetilde{\mathcal{BE}}(\mathbf{D}, v, k) := \mathbf{M}(k) \cdot v = \left(\mathbf{I} + \sum_{i=1}^k \frac{1}{(1 - \alpha_d)^i} \mathbf{D}^i \right) v. \quad (7)$$

Thus, we can develop approximates of the steady-state social nudge vectors, $\tilde{y}(k)$, and total production boost from nudges, $\tilde{x}(k)$:

$$\tilde{y}(k) := \widetilde{\mathcal{BE}}(\mathbf{D}, \boldsymbol{\mu}, k) \text{ and } \tilde{x}(k) := \boldsymbol{\eta}^\top \tilde{y}(k). \quad (8)$$

The following result, which is a corollary of Theorem 1 and Lemma 2 in Online Appendix D.1, validates using $\tilde{y}(k)$ and $\tilde{x}(k)$ to approximate y^* and x^* , respectively.

Corollary 1. Assume that (α_d, \mathbf{D}) satisfies Condition \mathcal{C} . We have (a) $\lim_{k \uparrow +\infty} \tilde{y}(k) = y^*$ and $\lim_{k \uparrow +\infty} \tilde{x}(k) = x^*$; (b) $\tilde{y}_e(k)$ is increasing in k for any $e \in E$, and so is $\tilde{x}(k)$ increasing in k . Therefore, for each $k \in \mathbb{Z}_{+,+}$, $\tilde{y}_e(k) \leq y_e^*$ for all $e \in E$, and $\tilde{x}(k) \leq x^*$.

Economically, the approximate BCE, $\widetilde{\mathcal{BE}}(\mathbf{D}, \boldsymbol{\mu}, k)$, is the expected total number of nudges sent on each edge in E if the diffusion radius is at most k . Because the diffusion matrix \mathbf{D} has an extremely high dimension, we introduce two important approximations to make the estimation of the global effect of social nudges computationally tractable. First, we adopt the approximation scheme (8) with $k = 1$, thus ignoring the effect of nudge diffusion beyond radius 1. As we will show, such

approximation will only incur a relative error of less than 1% for the global effect of social nudges on Platform O. Second, we adopt another layer of approximation by down sampling a subset of providers from V (denoted as \tilde{V}). We estimate the total production boost of the providers in \tilde{V} brought by the social nudges they receive, denoted as \hat{w}_0 , as well as the total production boost caused by the social nudges the providers in \tilde{V} send out as a result of the social nudges they receive (i.e., the diffusion of nudges), denoted as \hat{w}_1 . Hence, \hat{w}_0 captures the direct effect of social nudges, and \hat{w}_1 captures the indirect effect in the steady state per period. Both \hat{w}_0 and \hat{w}_1 take into account the over-time effects of social nudges. Scaling these estimates by a factor of $\frac{|V|}{|\tilde{V}|}$ would, therefore, yield unbiased estimates of the true direct and indirect global effects. Therefore, we devise $\frac{|V|}{|\tilde{V}|}(\hat{w}_0 + \hat{w}_1)$ as an unbiased estimate for $\tilde{x}(1)$.¹⁴ We summarize the detailed estimation procedure as Algorithm 1 in Online Appendix D.3.

Based on Algorithm 1 in Online Appendix D.3, quantifying the global effect for Platform O involves estimating the following four sets of parameters: (1) the expected number of organic social nudges for each edge (i.e., μ_e for $e \in E$); (2) the effect of receiving one social nudge on boosting the nudge recipient's production (i.e., p_e for $e \in E$); (3) the intensity of social nudge diffusion (i.e., $d_{e\ell}$ for $e, \ell \in E$ and $e_d = \ell_o$); and (4) the time-discounting factors (i.e., α_p and α_d). Our estimation of μ_e is based on observational data, whereas that of p_e , $d_{e\ell}$, α_p , and α_d relies on experimental data. The estimation results of the model parameters based on data from different experiments are provided in Table 9. We relegate the estimation details to Online Appendix E.

Before presenting the estimate for the global effect of social nudges on production using Algorithm 1 in Online Appendix D.3, we first describe a naïve benchmark that directly uses data from our experiment to calculate the difference in the number of videos uploaded by treatment versus control providers on the first day when they are sent a social nudge. Then, we scale this difference to the entire population on the platform by the average number of providers who are sent social nudges on the platform per day, which can be estimated by (1) the number of providers in the analysis sample of our experiment who received social nudges on a day divided by (2) the ratio of the number of providers targeted by the experiment to the total number of providers on the platform.

Following the naïve approach and using data from our main experiment, we first estimate that the total boost of video uploads caused by social nudges among 1,000,000 providers is 48.65 per day. Then, following Algorithm 1 in Online Appendix D.3, we approximate the total production boost of social nudges on the entire network on a given day in the steady state by

down sampling a subset of providers \tilde{V} , where $|\tilde{V}| = 1,000,000$. To protect sensitive data, we only report the boost on \tilde{V} without rescaling it back to the entire platform (i.e., $\hat{w}_0 + \hat{w}_1$). The estimation results using data from the main experiment are presented in Table 10, column (1). For those 1,000,000 randomly sampled providers in \tilde{V} , the accumulated direct production boost is $\hat{w}_0 = 130.08$ videos per day, and the accumulated indirect production boost from social nudge diffusion is $\hat{w}_1 = 10.59$ videos per day, yielding a total production boost of $\hat{w}_0 + \hat{w}_1 = 140.67$ videos per day. Therefore, our results suggest that the indirect production boost from nudge diffusion accounts for at least 8.14% of the direct effect (i.e., $10.59/130.08$).

In addition, we remark that the estimation results discussed suggest that using $\tilde{x}(1)$ is a reasonable approximation of x^* . Specifically, because the (first-order) indirect effect from nudge diffusion is about 8.14% of the direct effect, the production boost from second- and higher-order diffusion accounts for only about 0.72% (i.e., $\frac{0.0814^2}{1-0.0814}$) of the direct effect. Thus, ignoring the diffusion with radius 2 or beyond will introduce only fairly small additional errors.

It is clear that our social network model could help address the substantial underestimate of the naïve approach to predict the social nudges' total production boost. The more precise estimation of social nudges' global effect over the entire user population using our social network model (140.67 per day for 1,000,000 providers) is 2.89 times as large as the naïve estimate (48.65 per day for 1,000,000 providers). Such a huge gap comes from two factors. (1) The social network model incorporates the over-time accumulation of the direct boosting effect of social nudges on recipients' production, which yields a 167% (i.e., $(130.08 - 48.65)/48.65$) increase compared with the naïve estimation. (2) The model also captures the diffusion of nudges, which accounts for another 22% (i.e., $10.59/48.65$) increase. We obtain similar results based on data from the replication experiment, as shown in Table 10, column (2). This robustness check, along with another one based on a different random sample of \tilde{V} (see Online Appendix E.4), confirms the robustness of our estimation and validates the accuracy of our model in quantifying the global effect of social nudges on production boost on Platform O. Above all, our social network model provides a framework to causally quantify the global effect of our intervention (including its direct and indirect effects), which will be underestimated by the naïve estimation method.

6.3. Operational Implications

In this section, we demonstrate the operational implications of our social network model with two important practical applications: (1) seeding and targeting for the social nudge function and (2) recommendation of content

providers to new users. To this end, we first leverage the BCE measure to construct the SNI that assigns a metric to each (existing or new) edge that quantifies its value in production boost through social nudges.

For each edge $e \in E$, we define its SNI as the expected per-period total production boost on the entire network that can be attributed, either directly or indirectly through diffusion, to the *organic nudges* sent by e_o to e_d . Denote $\boldsymbol{\mu}_e \in \mathbb{R}^{|E|}$ as a vector with all entries equal to zero, except for that of edge $e \in E$ being μ_e . Define the SNI of edge $e \in E$ as

$$v_e := \boldsymbol{\eta}^T \cdot \mathcal{BE}(\mathbf{D}, \boldsymbol{\mu}_e), \text{ provided that } (\alpha_d, \mathbf{D}) \text{ satisfies Condition } \mathcal{C}, \quad (9)$$

where $\mathcal{BE}(\mathbf{D}, \boldsymbol{\mu}_e)$ is given in Definition 1. As discussed, exactly computing $\mathcal{BE}(\mathbf{D}, \boldsymbol{\mu}_e)$ is computationally infeasible for a large-scale social network such as Platform O. Instead, we can bound v_e from below, leveraging the approximate BCE as follows:

$$\tilde{v}_e(k) := \boldsymbol{\eta}^T \cdot \widetilde{\mathcal{BE}}(\mathbf{D}, \boldsymbol{\mu}_e, k), \text{ provided that } (\alpha_d, \mathbf{D}) \text{ satisfies Condition } \mathcal{C}, \quad (10)$$

where $\widetilde{\mathcal{BE}}(\mathbf{D}, \boldsymbol{\mu}_e, k)$ is given by Equation (7). Similar to evaluating the global effect of social nudges, we focus on the case $k = 1$ in the computational simulation to balance accuracy and tractability. Therefore, of particular importance is the approximate SNI with diffusion radius $k = 1$ (so diffusion of order 2 or higher is ignored):

$$\begin{aligned} \tilde{v}_e(1) &= \boldsymbol{\eta}^T \cdot \widetilde{\mathcal{BE}}(\mathbf{D}, \boldsymbol{\mu}_e, 1) \\ &= \frac{\mu_e p_e}{1 - \alpha_p} + \sum_{\ell: \ell_o = e_d} \frac{\mu_e d_{e\ell} p_\ell}{(1 - \alpha_p)(1 - \alpha_d)}, \text{ for } e \in E. \end{aligned} \quad (11)$$

The approximate SNI (i.e., Equation (11)) offers insights on the property of a high-value edge; it either generates a high volume of organic nudges (the first term) or promotes a high volume of diffusion (the second term). For a wide range of practical applications, the key is to target the edges on a social network whose organic nudges boost provider production over the entire platform the most. With our social network model, this problem is equivalent to selecting the edges in E with the highest social nudge indices. In the case in which computing the (exact) SNIs is intractable, we can further reduce this problem to a simpler one of finding the edges $e \in E$ with the largest $\tilde{v}_e(1)$'s as a reasonable approximation. Next, we briefly illustrate how (approximate) SNIs can be used to address the seeding problem and the content provider recommendation problem for content-sharing social network platforms. The details are relegated to Online Appendix F.

6.3.1. Optimal Seeding. To boost content production, a content-sharing platform may use operational levers to prompt users to send social nudges. For example,

the platform can use push notifications or private messages that encourage viewers to send out social nudges to specific providers. Sensibly, any type of operational lever would require user attention, whereas users only have limited attention and patience (Dukas 2004). Therefore, the platform must carefully control the intensity of such interventions to avoid disturbing or upsetting its users.

Considering the limited number of levers that the platform could use at once without causing annoyance, the usage of one lever means forgoing the opportunity of implementing another lever. In this sense, when seeking to get more viewers to send out social nudges, the platform is faced with a capacity constraint, has to decide on which edge to exert influence via a given lever, and has to select a set of n edges $K \subset E$ to target. We denote that for each $e \in K$, the average number of social nudges sent on this edge per day will increase by a relative effect of δ_μ after e_o receives the motivation from the platform (i.e., from μ_e to $\mu_e(1 + \delta_\mu)$). The platform could control the strength of its encouragement for users to send more social nudges by adopting the appropriate lever. In our model, this is captured by the platform being able to change the parameter δ_μ according to its need. For example, besides targeting push notifications or private messages to selected viewers, the platform can modify the app user interface of some viewers to highlight the social nudge function for certain providers they are following. Based on our conversation with Platform O, the latter approach is likely to have a greater impact on users' behavior but requires much greater resources to set up compared with the former one.

Next, we explore how the platform should optimize the global effect of social nudges and estimate the extent to which the optimal strategy outperforms a random dissemination strategy in increasing the global effect of social nudges.

The global effect of social nudges with respect to the selected edges, K , is $\boldsymbol{\eta}^T \mathcal{BE}(\mathbf{D}, \boldsymbol{\mu}_K) \delta_\mu$, where $\boldsymbol{\mu}_K \in \mathbb{R}^{|E|}$ represents a vector with an entry of edge $e \in K$ ($e \notin K$) equal to μ_e (zero). Such production boost can be reasonably approximated by $\delta_\mu \cdot \sum_{e \in K} \tilde{v}_e(1)$. Thus, it is (approximately) "optimal" to select n edges in E with the highest (approximate) SNIs (i.e., the n edges with the largest $\tilde{v}_e(1)$). As a benchmark, the platform may adopt the simple, straightforward strategy of randomly targeting a subset of edges $K \subset E$ ($|K| = n$) and encouraging the users to nudge more on these edges (i.e., the random strategy). By simulation, we calculate the relative improvement of the "optimal" strategy over the random strategy in the total production boost of social nudges. We find that the "optimal" strategy substantially outperforms the random strategy regardless of the effectiveness of the platform's encouragement for users to send additional nudges δ_μ , especially

when the size of selected target providers n is small. See Online Appendix F.1 for details.

6.3.2. Content Provider Recommendation for New Users. An important strategy for a platform to engage and retain newly registered users is to recommend to them some providers who they can follow and potentially nudge afterward. Considering users' limited attention, the platform needs to decide the ranking of the provider list, after which it sequentially recommends the listed content providers to new users. After receiving the recommended list of providers, a new user may follow some or all of them. These new following links will in turn enable the new user to send social nudges to these providers and boost their content production. The platform seeks to maximize the total production boost from the nudges sent by new users.

We denote the set of newly registered users as N . For each new user $i \in N$, let us assume that the set of existing providers this user chooses to follow is U_i and the associated set of new following relationships is $E_i := \{(i, u) : u \in U_i\}$. Define $E' := \cup_{i \in N} E_i$ as the set of new edges. Then, the additional production boost attributed to the social nudges sent by the new users is given by $\sum_{i \in N} (\sum_{e \in E_i} v_e)$ (Proposition 2 in Online Appendix D.4); it can be reasonably approximated by $\sum_{i \in N} (\sum_{e \in E_i} \tilde{v}_e(1))$. Hence, the content provider recommendation of each new user can be optimized separately.

For a new user $i \in N$, given the potential content provider list M_i to recommend, the platform selects $V_i \subset M_i$ with $|V_i| = m$ and recommends the providers in V_i to the new user in a sequential manner. To avoid overly interrupting users, m is generally not too large (i.e., at the magnitude of a few dozen). Denote the probability that a new user will follow the j th provider recommended to her as c_j , where $c_1 \geq c_2 \geq \dots \geq c_m$. Let $\pi(j)$ refer to the provider ranked in the j th position. Then, we get the (approximate) additional production boost from the social nudges sent by new user i as $\sum_{j=1}^m c_j \tilde{v}_{(i, \pi(j))}(1)$. Therefore, the (approximate) "optimal" strategy is to select m providers in M_i with the highest induced (approximate) SNIs and rank them in descending order of induced (approximate) SNI. Similar to optimal seeding, we compare the SNI-based provider recommendation with the benchmark random recommendation, which recommends the content providers based on a random permutation of M_i . By simulation, we also find that the "optimal" strategy significantly outperforms the random strategy in production boost, especially when the recommended provider list length m is small. See Online Appendix F.2 for details.

7. Conclusions and Discussion

In two field experiments on a large online content-sharing social network platform, we consistently find

that social nudges not only directly boosted nudge recipients' production but also, stimulated overall content provision by motivating nudge recipients to send more nudges to others. These effects were amplified when nudge recipients and nudge senders had stronger ties, and they persisted beyond the day nudges were sent.

Inspired by these results, we developed a novel social network model that incorporates the diffusion and over-time effects of social nudges into the estimation of their global effect. We find that the naive approach simply based on experiments underestimates social nudges' total production boost, but our model helps address this issue. Moreover, via simulation examples, we demonstrate that another advantage of adopting our social network model is to find strategies to optimize platform operations regarding social nudges.

Our research offers important practical implications for content-sharing social network platforms. First, social nudges can be a cost-effective intervention for these platforms to lift production on the supply side and consequently, increase consumption on the demand side. Platforms are naturally eager to control costs. Compared with financial incentives, social nudges require minimal costs on the platform's end. In fact, because of the success of social nudges observed in our experiments, after the second experiment, Platform O scaled up this function, enabling all users to receive and send social nudges as long as they (or the target they want to nudge) have not uploaded any video for a day or more.

As we noted in Section 2, prior research suggests that peer recognition may not enhance production and could even harm motivation if people do not view the recognized activity as core work in a given setting, doubt the credibility of peer recognition, or see themselves as not qualified for the recognition (Restivo and van de Rijt 2014, Gallus et al. 2020). Those are not concerns in our empirical context. For one thing, providing content is providers' core activity on the platform, and viewers naturally hold the authority to judge providers' content. Thus, recognition from viewers is meaningful to providers. For another thing, because all social nudges on Platform O are spontaneously initiated by followers (rather than being imposed by researchers on providers who might not believe their own qualifications, as in Restivo and van de Rijt 2014), providers who receive social nudges may naturally feel qualified for this form of recognition. In fact, we find that receiving social nudges boosted production among providers with different levels of productivity, including providers who were not very prolific (Section 4.4 and Online Appendix C.4). We are hopeful that on content-sharing platforms, nudges from social neighbors could avoid the pitfalls of peer recognition observed in previous research and instead, boost production across a broad set of providers.

Second, this work highlights the value of leveraging co-users' influence. Content-sharing social network

platforms connect users and facilitate transactions or relationships between users; thus, they have the advantage of influencing users through interactions between social neighbors, although they have limited power to directly control its providers to produce more content. Thus, platforms can guide co-users to influence each other as a way to improve overall user engagement on platforms.

Likes and positive comments are a prevalent form of co-user influence that may also boost production on a content-sharing platform, but they differ from social nudges in two aspects. One is that whereas viewers send social nudges because they intentionally want to encourage providers to produce more content, viewers who leave likes or positive comments do not necessarily intend to *actively* influence providers to produce more, and even if they do, their intentions are not clearly conveyed by likes and comments. The other difference is that social nudges are sent by social neighbors, which is not necessarily the case for likes and comments on many content-sharing social network platforms. Prior research has shown that social neighbors are powerful in changing people's behaviors (Bapna and Umyarov 2015, Wang et al. 2018). In our experiments, we also find that stronger ties between neighbors strengthen the effect of social nudges on production, which suggests that the power of social relationships may contribute to the success of social nudges.

Considering these distinctions between likes/comments and social nudges, we speculate that viewers use social nudges differently than likes and positive comments and that social nudges may work on top of likes/comments. As suggestive evidence for our speculation, an additional analysis reveals that sending nudges to providers did not decrease viewers' use of likes and comments (see Online Appendix C.3); as shown in Section 4.4, social nudges boosted production beyond the first reception day, even when we controlled for the increased likes and comments received by providers, which suggests that providers are motivated by social nudges beyond the influence of likes and comments.

Third, by showcasing that the diffusion of social nudges is crucial for measuring and optimizing the effects of social nudges on production, our work reveals how important it is for platforms to consider the diffusion of an intervention when they decide whether to scale up the intervention and how to maximize its effectiveness. Furthermore, by exploring strategies to maximize the global effect of social nudges—including the optimal seeding strategy and the optimal provider recommendation strategy for new users—our method may inspire platform managers to leverage a model such as ours to enhance the power of an intervention.

The limitations of our research open up interesting avenues for future research. For one, the type of social nudge we examined is simple, private, and subtle. It

was standardized across users, contained simple content, and leveraged no additional psychological principles. It was visible only to recipients in the message center. Also, as more messages arrived in the message center, earlier social nudge messages were pushed down, often off the front page of the message center, and they become less visible. Using such a light-touch, bare-bones social nudge allows us to provide a clean test of the effect of being nudged, but future research could examine how to design social nudges to produce stronger, longer-lasting effects—for example, by incorporating persuasion techniques and additional psychological insights into nudge messages, allowing senders to write personalized messages, or displaying social nudges publicly in a dedicated area. Another limitation of our research is that we could not causally study the effects of repeatedly receiving social nudges because the number of social nudges sent to each provider was not exogenous. Future research could randomly assign people to receive varying numbers of nudges and causally estimate their various effects based on the number of nudges received.

Acknowledgments

The authors thank the Department Editor Prof. Victor Martínez-de-Albéniz, the anonymous associate editor, and four referees for their very helpful and constructive comments, which have led to significant improvements in both the content and exposition of this study. They also thank the industry partner for their support of conducting the experiments and sharing the data.

Endnotes

¹ See <https://datareportal.com/reports/digital-2021-global-overview-report>.

² See <https://www.statista.com/outlook/dmo/digital-advertising/social-media-advertising/worldwide>.

³ The word *nudge* is a behavioral science concept for describing interventions that intend to change individuals' behaviors without altering financial incentives or imposing restrictions (Thaler and Sunstein 2009). Nudges are usually implemented by managers, marketers, and policy makers. We coin the term *social nudges* to refer to nonfinancial, nonrestrictive interventions that are intentionally implemented by neighbors within a social network to influence peers.

⁴ Most providers could satisfy this requirement. For example, on the first day of the experiment among all providers on Platform O who uploaded any videos in the past 30 days, 88% had not posted a video for 1 or more days.

⁵ To protect Platform O's identity, we digitally altered the app interface of a widely used video-sharing platform in China to obscure some nonessential details and reflect where the nudge button and social nudges are and what they look like on Platform O. Platform O has a similar app interface to Figure 1.

⁶ In the message center, the most recent message appears at the top. Messages about social nudges were not given a higher priority over other types of messages. In general, messages disappear only when providers delete them.

⁷ In the few months before our first experiment, social nudges were being tested and developed; as a result, some providers in our experiment received social nudges before the experiment. We removed those providers, per our second selection criterion, in order to estimate how social nudges change behavior when a platform starts to implement the social nudge function. Our results are qualitatively unchanged if we remove the second criterion and include all providers whose followers sent them at least one social nudge during our experiment (see Online Appendix A.1).

⁸ The authors have a nondisclosure agreement with Platform O.

⁹ See https://github.com/ZhiyuZeng-Public/the_impact_of_social_nudges_on_UGC.

¹⁰ For each provider, we calculated the interval (in days) between any two videos she successively uploaded (which equaled zero if two videos were uploaded on the same day) from January 1, 2018 to the day before the main experiment; then, we calculated her median interval of video postings across all pairs of successively uploaded videos.

¹¹ Because the majority of providers produced no videos on the first reception day and consequently, had a value of zero for $Total Views_i$, the 95th percentile of the raw values of $Total Views_i$ was small. Because we wanted to address extreme outliers caused by a small number of videos that went viral, we winsorized at the 95th percentile of nonzero values. That is, we replaced values of $Total Views_i$ that were greater than the 95th percentile of nonzero values with the 95th percentile of nonzero values. The result is robust if we winsorize at the 99th percentile of nonzero values.

¹² The positive effect of social nudges on content consumption is robust if we use the total views a provider obtained on her first reception day (as opposed to within the first week of her first reception day) as the outcome variable.

¹³ See Lemma 2 in Online Appendix D.1 for a formal proof.

¹⁴ See Proposition 1 in Online Appendix D.3 for a formal proof.

References

- Ashraf N, Bandiera O, Jack BK (2014a) No margin, no mission? A field experiment on incentives for public service delivery. *J. Public Econom.* 120:1–17.
- Ashraf N, Bandiera O, Lee SS (2014b) Awards unbundled: Evidence from a natural field experiment. *J. Econom. Behav. Organ.* 100:44–63.
- Bai J, So KC, Tang CS, Chen X, Wang H (2019) Coordinating supply and demand on an on-demand service platform with impatient customers. *Manufacturing Service Oper. Management* 21(3):556–570.
- Ballester C, Calvó-Armengol A, Zenou Y (2006) Who's who in networks. wanted: The key player. *Econometrica* 74(5):1403–1417.
- Banerjee S, Sanghavi S, Shakkottai S (2016) Online collaborative filtering on graphs. *Oper. Res.* 64(3):756–769.
- Banya BS (2017) *The Relationship Between Simple Employee Recognition and Employee Productivity in Business Organizations. A Case Study* (Anchor Academic Publishing, Hamburg, Germany).
- Bapna R, Umyarov A (2015) Do your online friends make you pay? A randomized field experiment on peer influence in online social networks. *Management Sci.* 61(8):1902–1920.
- Bimpikis K, Candogan O, Saban D (2019) Spatial pricing in ride-sharing networks. *Oper. Res.* 67(3):744–769.
- Bimpikis K, Ozdaglar A, Yildiz E (2016) Competitive targeted advertising over networks. *Oper. Res.* 64(3):705–720.
- Bradler C, Dur R, Neckermann S, Non A (2016) Employee recognition and performance: A field experiment. *Management Sci.* 62(11):3085–3099.
- Bramoullé Y, Djebbari H, Fortin B (2020) Peer effects in networks: A survey. *Annual Rev. Econom.* 12:603–629.
- Buell RW, Norton MI (2011) The labor illusion: How operational transparency increases perceived value. *Management Sci.* 57(9):1564–1579.
- Buell RW, Kim T, Tsay C-J (2017) Creating reciprocal value through operational transparency. *Management Sci.* 63(6):1673–1695.
- Burch G, He Q, Hong Y, Lee D (2022) How do peer awards motivate creative content? Experimental evidence from Reddit. *Management Sci.* 68(5):3488–3506.
- Burch G, Hong Y, Bapna R, Griskevicius V (2018) Stimulating online reviews by combining financial incentives and social norms. *Management Sci.* 64(5):2065–2082.
- Cabral L, Li L (2015) A dollar for your thoughts: Feedback-conditional rebates on eBay. *Management Sci.* 61(9):2052–2063.
- Cachon GP, Daniels KM, Lobel R (2017) The role of surge pricing on a service platform with self-scheduling capacity. *Manufacturing Service Oper. Management* 19(3):368–384.
- Candogan O, Drakopoulos K (2020) Optimal signaling of content accuracy: Engagement vs. misinformation. *Oper. Res.* 68(2):497–515.
- Candogan O, Bimpikis K, Ozdaglar A (2012) Optimal pricing in networks with externalities. *Oper. Res.* 60(4):883–905.
- Caro F, Martínez-de Albéniz V (2020) Managing online content to build a follower base: Model and applications. *INFORMS J. Optim.* 2(1):57–77.
- Celhay PA, Gertler PJ, Giovagnoli P, Vermeersch C (2019) Long-run effects of temporary incentives on medical care productivity. *Amer. Econom. J. Appl. Econom.* 11(3):92–127.
- Chen Y, Wang Q, Xie J (2011) Online social interactions: A natural experiment on word of mouth vs. observational learning. *J. Marketing Res.* 48(2):238–254.
- Chen Y, Harper FM, Konstan J, Li SX (2010) Social comparisons and contributions to online communities: A field experiment on MovieLens. *Amer. Econom. Rev.* 100(4):1358–1398.
- Cohen MC, Harsha P (2020) Designing price incentives in a network with social interactions. *Manufacturing Service Oper. Management* 22(2):292–309.
- Cui R, Li M, Li Q (2020b) Value of high-quality logistics: Evidence from a clash between SF Express and Alibaba. *Management Sci.* 66(9):3879–3902.
- Cui R, Li J, Zhang DJ (2020a) Reducing discrimination with reviews in the sharing economy: Evidence from field experiments on Airbnb. *Management Sci.* 66(3):1071–1094.
- Cui R, Zhang DJ, Bassamboo A (2019) Learning from inventory availability information: Evidence from field experiments on Amazon. *Management Sci.* 65(3):1216–1235.
- De Grip A, Sauermann J (2012) The effects of training on own and co-worker productivity: Evidence from a field experiment. *Econom. J.* 122(560):376–399.
- Dukas R (2004) Causes and consequences of limited attention. *Brain Behav. Evolution* 63(4):197–210.
- Frey B, Gallus J (2017) *Honours Vs. Money: The Economics of Awards* (Oxford University Press, Oxford, UK).
- Gallus J (2017) Fostering public good contributions with symbolic awards: A large-scale natural field experiment at Wikipedia. *Management Sci.* 63(12):3999–4015.
- Gallus J, Jung O, Lakhani KR (2020) Recognition incentives for internal crowdsourcing: A field experiment at NASA. Harvard Business School Technology & Operations Management Unit Working Paper No. 20-059, Harvard Business School, Boston.
- Gelper S, van der Lans R, van Bruggen G (2021) Competition for attention in online social networks: Implications for seeding strategies. *Management Sci.* 67(2):1026–1047.
- Goes PB, Guo C, Lin M (2016) Do incentive hierarchies induce user effort? Evidence from an online knowledge exchange. *Inform. Systems Res.* 27(3):497–516.
- Grant AM, Gino F (2010) A little thanks goes a long way: Explaining why gratitude expressions motivate prosocial behavior. *J. Personality Soc. Psych.* 98(6):946–955.

- Gurvich I, Lariviere M, Moreno A (2019) Operations in the on-demand economy: Staffing services with self-scheduling capacity. Hu M, ed. *Sharing Economy: Making Supply Meet Demand* (Springer International, Berlin), 249–278.
- Horn RA, Johnson CR (2012) *Matrix Analysis* (Cambridge University Press, Cambridge, UK).
- Huang N, Burtch G, Gu B, Hong Y, Liang C, Wang K, Fu D, Yang B (2019) Motivating user-generated content with performance feedback: Evidence from randomized field experiments. *Management Sci.* 65(1):327–345.
- Jackson MO (2005) The economics of social networks. Working paper, California Institute of Technology, Pasadena, CA.
- Kabra A, Belavina E, Girotra K (2020) Bike-share systems: Accessibility and availability. *Management Sci.* 66(9):3803–3824.
- Konings J, Vanormelingen S (2015) The impact of training on productivity and wages: Firm-level evidence. *Rev. Econom. Statist.* 97(2):485–497.
- Kosfeld M, Neckermann S (2011) Getting more work for nothing? Symbolic awards and worker performance. *Amer. Econom. J. Microeconomics* 3(3):86–99.
- Kuang L, Huang N, Hong Y, Yan Z (2019) Spillover effects of financial incentives on non-incentivized user engagement: Evidence from an online knowledge exchange platform. *J. Management Inform. Systems* 36(1):289–320.
- Kuhn P, Kooreman P, Soetevent A, Kapteyn A (2011) The effects of lottery prizes on winners and their neighbors: Evidence from the Dutch postcode lottery. *Amer. Econom. Rev.* 101(5):2226–2247.
- Lazear EP (2000) Performance pay and productivity. *Amer. Econom. Rev.* 90(5):1346–1361.
- Li Y, Lu LX, Lu SF (2020) Do social media trump government report cards in influencing consumer choice? Evidence from U.S. nursing homes. Preprint, submitted January 24, https://www.researchgate.net/publication/338853735_Do_Social_Media_Trump_Government_Report_Cards_in_Influencing_Consumer_Choice_Evidence_from_US_Nursing_Homes.
- Mas A, Moretti E (2009) Peers at work. *Amer. Econom. Rev.* 99(1):112–145.
- Mohan B, Buell RW, John LK (2020) Lifting the veil: The benefits of cost transparency. *Marketing Sci.* 39(6):1105–1121.
- Mookerjee R, Kumar S, Mookerjee VS (2017) Optimizing performance-based Internet advertisement campaigns. *Oper. Res.* 65(1):38–54.
- Moon K, Bergemann P, Brown D, Chen A, Chu J, Eisen E, Fischer G, Loyalka PK, Rho S, Cohen J (2022) Manufacturing productivity with worker turnover. *Management Sci.*, ePub ahead of print August 9, <https://doi.org/10.1287/mnsc.2022.4476>.
- Nicoletti C, Salvanes KG, Tominey E (2018) The family peer effect on mothers' labor supply. *Amer. Econom. J. Appl. Econom.* 10(3):206–234.
- Papanastasiou Y, Savva N (2017) Dynamic pricing in the presence of social learning and strategic consumers. *Management Sci.* 63(4):919–939.
- Parker C, Ramdas K, Savva N (2016) Is it enough? Evidence from a natural experiment in India's agriculture markets. *Management Sci.* 62(9):2481–2503.
- Pew Research Center (2010) Online product research (September 29), <https://www.pewresearch.org/internet/2010/09/29/online-product-research-2/>.
- Restivo M, van de Rijt A (2014) No praise without effort: Experimental evidence on how rewards affect Wikipedia's contributor community. *Inform. Comm. Soc.* 17(4):451–462.
- Roels G, Su X (2014) Optimal design of social comparison effects: Setting reference groups and reference points. *Management Sci.* 60(3):606–627.
- Ryan RM, Deci EL (2000) Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. *Amer. Psych.* 55(1):68–78.
- Sacerdote B (2001) Peer effects with random assignment: Results for Dartmouth roommates. *Quart. J. Econom.* 116(2):681–704.
- Song H, Tucker AL, Murrell KL, Vinson DR (2018) Closing the productivity gap: Improving worker productivity through public relative performance feedback and validation of best practices. *Management Sci.* 64(6):2628–2649.
- Tan TF, Netessine S (2014) When does the devil make work? An empirical study of the impact of workload on worker productivity. *Management Sci.* 60(6):1574–1593.
- Tan TF, Netessine S (2019) When you work with a superman, will you also fly? An empirical study of the impact of coworkers on performance. *Management Sci.* 65(8):3495–3517.
- Tan TF, Netessine S (2020) At your service on the table: Impact of tabletop technology on restaurant performance. *Management Sci.* 66(10):4496–4515.
- Thaler RH, Sunstein CR (2009) *Nudge: Improving Decisions About Health, Wealth, and Happiness* (Penguin, New York).
- Wang C, Zhang X, Hann I-H (2018) Socially nudged: A quasi-experimental study of friends' social influence in online product ratings. *Inform. Systems Res.* 29(3):641–655.
- Whitmore D (2005) Resource and peer impacts on girls' academic achievement: Evidence from a randomized experiment. *Amer. Econom. Rev.* 95(2):199–203.
- Xu Y, Armony M, Ghose A (2021) The interplay between online reviews and physician demand: An empirical investigation. *Management Sci.* 67(12):7344–7361.
- Yang J, Wei X, Ackerman M, Adamic L (2010) Activity lifespan: An analysis of user survival patterns in online knowledge sharing communities. *Proc. Internat. AAAI Conf. Web Social Media*, vol. 4 (AAAI Press, Palo Alto, CA), 186–193. <https://ojs.aaai.org/index.php/ICWSM/article/view/14010>.
- Zeng Z, Clyde N, Dai H, Zhang D, Xu Z, Shen Z-JM (2022) The value of customer-related information on service platforms: Evidence from a large field experiment. Preprint, submitted October 25, <http://dx.doi.org/10.2139/ssrn.3528619>.
- Zhang DJ, Allon G, Van Mieghem JA (2017) Does social interaction improve learning outcomes? Evidence from field experiments on massive open online courses. *Manufacturing Service Oper. Management* 19(3):347–367.
- Zhang DJ, Dai H, Dong L, Wu Q, Guo L, Liu X (2019) The value of pop-up stores on retailing platforms: Evidence from a field experiment with Alibaba. *Management Sci.* 65(11):5142–5151.
- Zhang DJ, Dai H, Dong L, Qi F, Zhang N, Liu X, Liu Z, Yang J (2020) The long-term and spillover effects of price promotions on retailing platforms: Evidence from a large randomized experiment on Alibaba. *Management Sci.* 66(6):2589–2609.
- Zhou J, Chen Y-J (2016) Targeted information release in social networks. *Oper. Res.* 64(3):721–735.