

# DO ONLINE COMMUNITIES BENEFIT FROM APPOINTING VOLUNTEER MODERATORS? EVIDENCE FROM A REGRESSION DISCONTINUITY DESIGN

Jens Foerderer and Armin Heinzl<sup>1</sup>

<sup>1</sup> Business School, University of Mannheim, Germany.  
{foerderer,heinzl}@uni-mannheim.de

**Abstract.** We study whether online communities can benefit from appointing users to “volunteer moderators”, in terms of community members or leaders who are supposed to encourage participation and plant the seeds of community. We exploit a quasi-experiment on Stack Exchange, a network of more than 160 online communities, over the period from 2010 to 2017. These communities regularly hold democratic elections of moderators. By focusing on elections decided by a narrow vote margin, we can exploit a regression discontinuity that yields a quasi-random assignment of moderatorship. We find that online communities can significantly benefit from volunteer moderatorship: closely elected moderators contribute not only more content, they also provide more extensive content to the community. Furthermore, the direct outcomes of moderatorship appear to have positive spillover effects on community discourse: moderators spark more extensive discussions, which are rated higher and bookmarked more often.

**Keywords:** Online communities, community governance, moderators, quasi-experiment, regression discontinuity

## 1 Introduction

The “1% rule” describes a fundamental problem of user participation in online communities [1]: only 1% of an online community’s users account for most contributions (“key users”), 9% contribute from time to time (“commenters”), and the majority of 90% of users free-ride (“lurkers”). For example, of the more than 20 million registered users on Wikipedia, less than .05% of them are estimated to create and edit the majority of articles. Similar estimations have been documented for participation on the question-and-answer site Stack Exchange, YouTube, and SAP’s developer forums. The “1% rule” illustrates the fundamental challenge of online community governance [e.g., 2–5]: what governance instruments are effective in stimulating user contributions? While scholars have identified various economic, psychological, and sociological instruments that may engender contributions of the “lurkers” and “commentators” [e.g., 6–8], few studies systematically assess

Multikonferenz Wirtschaftsinformatik 2018,  
March 06-09, 2018, Lüneburg, Germany

instruments that allow community owners to keep the 1%—in terms of the “core that contributes content and protects the boundary of the community” [9]—engaged. This issue is particularly critical given that scholars found that many online communities are sustained by these “1%” of members [10, 11] and how easily communities can tip toward failure if key users turn over.

One governance instrument that has found widespread but controversial resonance concerns promoting key users to volunteer “moderators”, “community leaders”, or “advocates”. Moderators shall “facilitate discussion on a voluntary basis [...] to engender trust, encourage participation, and plant the seeds of community” [8]. By appointing moderators, community owners hope to boost key users to participate even more actively in a community, while having positive spillover effects on other users and reducing a community owner’s costs of managing a community [8, 12]. To achieve this goal, community owners grant moderators extensive control to create and enforce policies, promote moderatorship as the highest social status within a community, and manage relations to them. However, the effectiveness of moderatorship as a governance instrument is controversial. Instead of promoting contributive behavior, much anecdotal evidence suggests moderatorship to discourage contributions [e.g., 13, 14]. For example, Other anecdotes described that moderators shifted their efforts on policing the contributions of other users rather than contributing. Other examples describe how moderatorship triggered even deceptive and collusive behaviors of key users such as Wikipedia moderators who revised articles in the favor of third-parties or the various communities in the Reddit network that have become notorious for their moderators’ misbehavior [12, 14]. Of course, these examples may reflect isolated instances of the “wrong people” being appointed moderators. Nevertheless, an evaluation of the consequences of moderatorship is so far missing.

In this paper, we address two important yet unresolved questions about the implications of appointing volunteer moderators. The first question we study is whether moderatorship indeed discourages key users to contribute to a community. The debate on this question is characterized by two opposing theoretical perspectives that suggest governance instruments like moderatorship as a catalyst or inhibitor of contributive user behavior, respectively. The inhibitor argument—put forward by proponents of negative effects of devolving control to users and awarding community members with social status—is that moderatorship has no or even a detrimental effect on a user’s contributions [15]. Quite the opposite, once moderatorship is granted to users, they lose the incentives to contribute [15]. Following this logic, moderatorship is merely an effective instrument and may withdraw key users from the discourse, which is a reason why community owners should avoid its widespread adoption. By contrast, the catalyst argument says that moderatorship might motivate greater investments in the community by instrumental, sociologically-based, and psychologically-motivated mechanisms. Instrumentally, moderatorship provides a user with greater control from which a user may draw to contribute [15, 16]. Sociologically, moderatorship is a status that may motivate greater contributions to the community [e.g., 6]. Psychologically, moderatorship may trigger greater attachment to and identification with a community [e.g., 17]. If moderatorship indeed

catalyzes key users to contribute more by one of these mechanisms, then community owners may substantially benefit from implementing volunteer moderators. By establishing the causal effects of moderatorship and untangling its underlying mechanisms we seek to shed light on this debate. The second question we study is whether volunteer moderator behavior has spillover effects on community discourse, in terms of leading to more extensive interactions and contributions that are also perceived as more valuable by community members. Despite the prevalence of this question, previous studies provide little insight on this question.

The key empirical problem here is that moderatorship is usually not randomly assigned but results from a two-stage selection process. Users first strive for moderatorship based on unobserved characteristics such as motivation or information, and then community owners screen and appoint candidates based on another set of unobserved characteristics such as skills or personal impression [12]. Because of this two-stage selection, naïve comparisons of regular users and moderators are strongly biased. Unless differences in omitted variables can plausibly be removed, attributing user behavior to moderatorship is not possible. To overcome the selection problem, we exploit a regression discontinuity design [18, 19] in the context of Stack Exchange, a network of online communities that comprises more than 30 million posts and 7 million users. Many of these communities decide on their moderators by holding secret ballot elections, which are remarkably similar to real-world democratic election processes. Comparing moderator candidates who closely won an election with candidates who closely lost provides us with a quasi-random assignment of moderatorship [19, 20]. Consistent with our identification assumption, we find that candidates close to the cutoff behave similarly in the period prior to the election. Our dataset combines ballot data of all elections in Stack Exchange communities with user-level data on 1,012 moderatorship candidates across more than 70 communities over the period from 2010 to 2017.

## **2 Theoretical Background**

Online communities have become an important aspect of digital environments [21–23]. We refer to online communities as “virtual space[s] where people come together with others to converse, exchange information or other resources, learn, play, or just be with each other” [3]. Communities emerged as an important tool of firms to benefit from the ideas of outsiders [2, 24], foster brand building [25], or provide user-to-user support [26, 27]. Communities are also relevant outside the business world. Various interest groups and associations run online communities for fostering exchanges among like-minded people [5]. Communities are also the backbone of the open source software movement as well as of the various user-generated content sites like Wikipedia or Stack Exchange [28]. Regardless of their purpose, all communities have in common that they draw on the participation and contribution of their members.

The governance of online communities has proven to be a complex undertaking due to the inherent public goods problem: users can consume the contributions of others without having to contribute themselves [3, 7, 29, 30]. For example, Wikipedia

users can consume the articles contributed by others without having to contribute back. Similarly, users of Stack Overflow can receive programming help without having to give advice themselves. If the public goods problem is not addressed by proper governance mechanisms, communities will suffer from an undersupply of content due to insufficient incentives for users to contribute, and eventually fail [29]. Traditional economic models are less informing about the design of such mechanisms as they mostly turn toward pricing schemes, which cannot directly be applied in online communities as online communities rely on voluntary participation and contribution of time and effort rather than the exchange of money [31, 32]. Thus, the key challenge to community owners is the design of mechanisms that create incentives for users to contribute and that motivate core participants to sustain and improve their contribution [5, 22, 23, 32, 33].

Volunteer moderators represent a widespread instrument of online community governance [12, 34]. There exists no coherent definition of the role of a moderator but for the purpose of this study, we refer to moderators as “members who facilitate discussion on a voluntary basis [...] to engender trust, encourage participation, and plant the seeds of community” [8]. While the tasks of a moderator differ from community to community, they usually encompass the posting of new expert content, supporting the flow of discussions, and keeping a high quality of content contributions going, if necessary, by revising and deleting of content, the warning, suspending, and banning of users [12].

Our study of volunteer moderators relates to a broader discussion on three theoretical mechanisms that are suggested to be key to influence user participation, namely devolving control to users [8, 33], sociologically-driven status incentives [7, 13, 35], and psychologically-based community identification [17, 36]. First, moderatorship, in terms of the position or “office” of a moderator, grants a user substantial control, in terms of decision-making rights, to create and enforce community policies. As such, moderatorship relates to a broader array of literature that studies how online collectives govern, organize, and coordinate the actions of individuals to achieve collective outcomes [11, 33, 37]. Second, moderatorship represents a social status. Social status refers to an actor’s standing in a group when standing is based on prestige, honor, or deference [38]. Work in this stream is focused on understanding whether community owners can foster user contributions by awarding user contributions with trivial symbols of status and graduating contributors in the status hierarchy [e.g., 6, 7, 35, 39]. For example, in the Stack Exchange communities, contributors can earn symbolic bronze, silver, or gold badges along with fictitious titles like “Copy Editor”, “Explainer”, or “Teacher”. Finally, moderatorship may stimulate contributive behavior by triggering greater identification, affiliation, and attachment to a community in terms of a psychological motive. As such, moderatorship relates to work that outlined psychological identification with a community as a driver of contributive behavior [17, 36, 40].

### 3 Empirical Strategy and Data

Our empirical context is Stack Exchange, a network of online communities in a question-and-answer format. Communities cover various interests including computer programming, personal finance, or graphic design. According to Stack Exchange statements, moderators are expected to be “patient and fair, led by example, show respect for their fellow community members in their actions and words, are open to some light but firm moderation to keep the community on track and resolve (hopefully) uncommon disputes and exceptions”. Moderators are also expected to contribute their expert knowledge to the community. On Stack Exchange, moderatorship grants users both status and control. Moderators can exert control by revising or deleting content as well as the warning or banning of users. Moderator status is the highest status granted to users, which is publicly displayed next to a user’s name in the form of a diamond symbol (◆). With their status and control, moderators can directly influence the success of a community.

Stack Exchange moderatorship is decided by a community election, which closely mimics real-world democratic elections. Elections are held within individual communities and usually take place once a year. A typical election consists of three phases, nomination, primary, and final election. During the seven-day nomination phase, users can propose candidates, also themselves. If a user gets nominated by other users, the nominated user must accept the nomination before proceeding to the subsequent phase. Nominees are required to write a short, freely editable summary of why they might make a good community moderator. Most communities hold public Q&As with nominees, during which candidates must publicly stand questions of other users. Typical questions demand more details on the candidate’s availability per day, reaction time, and prior experience, but also seek to challenge the user by showing potential misbehaviors in the past. Examples of such questions include, “[...] you have posted 5 questions on Meta (that still survive today), plus another 20 on Meta.SE. Do you plan to become more active?”, “You recently hammer-closed a question as a duplicate [...]. However, a long comment thread and disputatious edits from OP followed your action, and the question is now reopened. How else might you have handled the situation, in retrospect?”, “Based on my short encounter with you, I doubt you'd be a good moderator”.

During primary, the top 30 nominees by reputation score advance to preliminary community voting. Primary candidates are displayed in random order and votes are public. Voters have one up or down vote per candidate. After 4 days, the top 10 candidates based on primary vote score proceed to the final election. The election usually takes between two to three weeks. The election is based on a single transferable vote system. Under this voting system, a voter has a single vote that is initially allocated to their most preferred candidate. Once a candidate has more votes than the quota, the candidate is elected. Throughout the counting procedure and as candidates are elected, the vote is transferred according to the voter's preferences, proportional to any surplus or discarded votes. During the election, candidates are displayed in random order on the ballot to avoid “donkey voting” or strategic gaming.

Intermediate voting results are not displayed. After the election, moderatorship is granted for life.

The election setup has two features that allow us to overcome the outlined two-stage selection bias and other inferential challenges. The first feature of our research design is that it allows for causal inference by exploiting very close election outcomes in a regression discontinuity design [18, 41]. In such a design, we compare the behavior of users that closely won such an election with users that closely lost the election. Such “close call” elections are akin to a random assignment of moderatorship to users, the smaller the margin of votes. The identifying assumption of a regression discontinuity (RD) design is that, in addition to candidate, community, or time factors, there is some randomness that determines the outcome of a close election. Intuitively, the difference between a user that is granted status with 50.1% of the votes and a user that loses the election with 49.9% of the votes seems negligible [18, 41]. Close call elections have been used as an identification strategy in various studies in political science and economics [20, 42].

Stack Exchange held the first elections in December 2010. We identified all elections held on Stack Exchange since then until January 2017 and collected the published ballot data. The ballot data contained identifiers of the candidates and the vote preferences they received, as well as descriptive information on the time, seats, and quota of each election. In total, our sample includes 154 elections across 70 communities. Voter turnout in these elections ranged from 3% to 42%, with the median election having a turnout of 13%. On average, 10,984 voters were eligible, with 972 votes casted on average. On average, 6.54 candidates competed for 2.44 seats. Aside from one election in 2010, elections are almost uniformly distributed over the years we consider. When inspecting the density function of vote shares, we see that many observations are close to the cutoff (1), supporting the use of a RD design in our context.

We matched ballot data with candidate-specific information provided directly by Stack Exchange. This data included static demographics of candidates as well as their posting, reputation, and editing behavior. To assess changes in candidate behavior, we consider a post-election window of six months after the election. We stop observing candidate behavior six months after the election to provide enough variance for our estimations but also to avoid confounding events like preceding or succeeding elections. Our final dataset contains 1,012 candidates. To assess pre-election differences between candidates, we collect candidate-level data also until six months before the election. Our sample covers elections in 70 communities. The communities cover a wide range of topics, including various computer-related topics like operating systems, databases, and programming but also mundane topics like cooking, home improvement, photography, or music-making. On average, communities in our sample listed more than 267,000 questions and 3.5 Mio answers contributed by 177,000 users. The average community in our sample has existed for 4.7 years and held 2.1 elections over the observation period.

We measure candidate contributions in terms of the number of posts they contribute (POSTS) and the word count of the posts (POSTLENGTH), the number of upvotes casted (UPVOTES), and the number of downvotes casted (DOWNVOTES).

We measure community members’ perceptions of moderators by their rating behavior, in terms of the rating they give to a candidate’s questions (QUESTION RATING), the rating of a candidate’s answers (ANSWER RATING), the number of views for a candidate’s questions (QUESTION VIEWS), and the number of times a candidate’s question is bookmarked (i.e., marked as favorite) by others (BOOKMARKS). Table 1 describes the sample. Overall, the sample includes 1,012 candidates over the period from 2010 to 2017.

**Table 1.** Summary Statistics

		<i>Mean</i>	<i>S.D.</i>	<i>Min.</i>	<i>Max.</i>	<i>P25</i>	<i>Median</i>	<i>P75</i>
1	Posts	8.88	19.46	0	164	0.00	2.00	8.00
2	Post length	895.89	897.25	0	12186	0.00	787.38	1140.15
3	Log(Answers)	0.10	0.37	0	8	0.00	0.01	0.06
4	Log(Comments)	265.45	799.93	0	10354	0.00	0.00	192.25
5	Log(Bookmarks)	0.03	0.15	0	4	0.00	0.00	0.02
6	Post rating	0.14	0.52	0	10	0.00	0.01	0.08
7	Badges	46.54	19.34	0	87	17.00	50.00	62.00
8	Months since joined	18.91	30.03	0	429	0.00	10.00	23.00

For each candidate  $i$ ,  $Y_{i,t+1}$  denotes the outcome of interest in  $t+1$ , while  $share_{i,t}$  is the “running variable” that determines treatment assignment. Candidates are assigned to the treatment group (1) if their vote share in  $t$  exceeds the defined quota and otherwise are assigned to the control group (0):

$$W_i = f(share_{i,t}) = \begin{cases} 1 & \text{if } share_i > quota \\ 0 & \text{if } share_i < quota \end{cases} \quad (1)$$

Gelman and Imbens [43] outline higher order polynomials as a potential source of bias to RD estimates. Therefore, we fit second order polynomials. Following the suggestions of Imbens and Lemieux [41] and Lee and Lemieux [19], we use a triangular kernel. The robustness section shows that the choice of these parameters matters little for our dataset, as the estimates remain comparable across various parameter choices. Accordingly, we rely on the standard RD in the following form:

$$Y_{i,t+1} = \alpha + \beta W_i + f(share_{i,t}) + community_i + \epsilon_i \quad (2)$$

In this equation,  $f(\cdot)$  is the function of the vote share outlined above,  $share_{i,t}$  is the continuous vote share,  $W_i$  is an indicator of whether the candidate was elected,  $community_i$  controls for community-specific effects, and  $\epsilon_i$  denotes the error term. The coefficient of interest is  $\beta$ , which gives the effect of being treated with moderator status. In the model, we assume that  $W_i$  is randomly distributed over candidates near the cutoff  $c$ .

## 4 Results

**Direct effects.** To assess the effects of moderatorship on the focal user, we compare candidates who were closely elected as moderators (winners) to candidates who closely lost at the ballots (losers). We therefore estimate equation (2) for increasingly small bandwidths around the vote share cutoff. Table 2 reports estimates of POSTS for close winners and losers for small bandwidths around the cutoff. Columns (1) - (4) indicate that, for increasingly small bandwidths around the cutoff, a consistently significant discontinuity in POSTS. The regressions estimate the increase between 10.5 and 15.3 posts. Columns (5) – (8) report the discontinuity in POST LENGTH. The models estimate an increase of POST LENGTH between 8.42 and 8.87 words per post. Taken together, both findings confirm what the above graphical analysis indicated: moderatorship triggers a strong increase in content contributed to the community, which stands in stark contrast to the suggestion that moderatorship discourages contributions.

**Spillover effects.** The findings so far indicate that moderators contribute more content. The natural follow-up question is whether these effects translate into greater discourse and exchange among other community members. To assess such spillover effects, we compare how community members react to a moderator’s posts and how they participate in discussions initiated by a moderator. We would conclude in favor of positive effects on community discourse, if we were to observe that other community members contribute more posts as a response to a candidate’s content (RESPONSES), if other members comment more actively the postings of moderators (COMMENTS), and if other members perceive a candidate’s postings as more valuable, in terms of bookmarking it more often (BOOKMARKS) or rating it higher (RATING).

**Table 2.** Regression Discontinuity Estimates of Direct Moderatorship Effects

	<i>Posts</i>				<i>Post length</i>			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Bandwidth	15%	12.5%	10%	7.5%	15%	12.5%	10%	7.5%
Elected	10.46** (3.76)	12.37*** (3.38)	12.63** (4.21)	15.32** (4.84)	8.42* (3.61)	8.54* (3.80)	7.81 (4.15)	8.87* (4.42)
N	340	313	283	254	340	313	283	254

Note: \*, \*\*, \*\*\* indicate significance at the 5%, 1%, and .1% levels, respectively.

Data in Table 3 provides strong evidence in favor of such a spillover effect. Columns (1) - (4) estimate that community members respond more actively to a moderator’s postings. Columns (5) to (8) estimate that community members comment a moderator’s posts more intensively, which appears less marked than the discontinuity in RESPONSES. Columns (9) to (12) and columns (13) to (16) show strong discontinuities in ratings by other community members. Users bookmark a moderator’s postings more often and rate it more positively. In sum, the regressions provide strong evidence for a spillover effect: moderatorship triggers a more desirable



behavior of moderators, which has ripple effects on the contributions of other community members.

**Table 3.** Regression Discontinuity Estimates of Spillover Moderatorship Effects

	<i>Responses</i>				<i>Comments</i>			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Bandw.	15%	12.5%	10%	7.5%	15%	12.5%	10%	7.5%
Elected	9.05** (3.46)	11.86*** (3.55)	13.23*** (3.57)	11.96*** (3.30)	5.50* (2.19)	6.88** (2.29)	7.42*** (2.25)	6.61** (2.04)
N	340	313	283	254	340	313	283	254

  

	<i>Bookmarks (weighted)</i>				<i>Post rating (weighted)</i>			
	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Bandw.	15%	12.5%	10%	7.5%	15%	12.5%	10%	7.5%
Elected	13.53** (4.12)	13.94*** (4.16)	12.79** (4.04)	10.06** (3.73)	29.68** (9.27)	31.99*** (9.34)	32.25*** (9.12)	28.36*** (8.60)
N	340	313	283	254	340	313	283	254

Note: \*, \*\*, \*\*\* indicate significance at the 5%, 1%, and .1% levels, respectively.

**Robustness and rival explanations.** We run four substantial robustness checks, which we can detail only shortly due to the space limitations. The RD approach in this paper makes two identifying assumptions [18, 44]. First, candidates are balanced on all covariates prior the election. Second, we do not find evidence that candidates are able to manipulate election outcomes. Third, we assure that the effects are not driven by losing candidates becoming alienated from a community as a result of the election. Finally, we run placebo regressions in which we randomly distribute moderator status—we do not observe significant effects, which supports the robustness of our findings.

## 5 Discussion and Conclusion

This paper contributes to the ongoing debate on the design and effectiveness of online community governance by providing new causal evidence about the effect of granting community members decision-making rights over a community on their behavior. To isolate the effects, we exploit the quasi-random assignment of moderatorship that occurs in close-call elections of moderators on Stack Exchange, one of the largest networks of online communities. We find that moderatorship catalyzes behavior desirable for online communities. They contribute more content, and their contributions are perceived as more helpful by community members. These lasting effects are likely to result from the psychological effects of holding moderator status within a community rather than resulting from newly granted instruments or resources. Whom should a community appoint as moderators? The evidence here identifies users with a below-median standing in a community to react strongest to moderatorship. In this sense, moderatorship should not necessarily be seen as a

“crown” that is bestowed as a finishing touch of a long-standing and well-established community membership but rather as an “office” that should also be open to less-established community members.

These findings have important implications for theory and policy design. First, the finding runs counter to the paradigm that moderatorship inhibits or even demolishes desirable user behavior. Instead, our findings support those who argue that moderatorship acts as an important catalyst for even further, desirable contributions to a community. Second, the findings demonstrate that moderatorship catalyzes user behavior through psychological mechanisms rather than due to instruments and resources that accompany such an appointment. This suggests that fostering the psychological mechanisms of moderatorship—for example by praising the position among community members, signaling moderatorship by badges or other means, and creating high-awareness processes for moderatorship applications—might be beneficial to realize the full gains from this form of self-governance. Third, the fact that the effects of moderatorship are rather heterogeneous suggests that more work is needed to better understand how the effects of moderatorship vary across users. We hope that our study stimulates further research examining governance forms and decision-makers in digital environments.

## References

1. Arthur, C.: What Is the 1% Rule? *The Guardian*. (2006).
2. Jeppesen, L.B., Frederiksen, L.: Why Do Users Contribute to Firm-Hosted User Communities? The Case of Computer-Controlled Music Instruments. *Organization Science*. 17, 45–63 (2006).
3. Kraut, R.E., Resnick, P.: *Building Successful Online Communities: Evidence-Based Social Design*. MIT Press, Cambridge, MA (2012).
4. Lakhani, K.R., Von Hippel, E.: How Open Source Software Works: “Free” User-to-User Assistance. *Research Policy*. 32, 923–943 (2003).
5. Wasko, M.M., Faraj, S.: Why Should I Share? Examining Social Capital and Knowledge Contribution in Electronic Networks of Practice. *MIS Quarterly*. 29, 35–57 (2005).
6. Chen, Y., Harper, F.M., Konstan, J., Xin Li, S.: Social Comparisons and Contributions to Online Communities: A Field Experiment on MovieLens. *The American Economic Review*. 100, 1358–1398 (2010).
7. Goes, P.B., Guo, C., Lin, M.: Do Incentive Hierarchies Induce User Effort? Evidence from an Online Knowledge Exchange. *Information Systems Research*. 27, 497–516 (2016).
8. Williams, R.L., Cothrel, J.: Four Smart Ways to Run Online Communities. *MIT Sloan Management Review*. 41, 81–91 (2000).
9. Faraj, S., Lakani, K., Monteiro, E., von Krogh, G.: Online Community as Space for Knowledge Flows. *Information Systems Research*. 27, 668–684 (2016).
10. Butler, B., Sproull, L.K., Kiesler, S.: Community Effort in Online Groups: Who Does the Work and Why. In: *Leadership at a Distance: Research in Technologically-supported Work*. pp. 171–194. Lawrence Erlbaum Associates, Mahwah, NJ (2007).

11. Von Krogh, G., Spaeth, S., Lakhani, K.R.: Community, Joining, and Specialization in Open Source Software Innovation: A Case Study. *Research Policy*. 32, 1217–1241 (2003).
12. Matias, N.J.: Civic Labor of Online Moderators. In: *Oxford Internet, Policy, and Politics Conference*. , Oxford, UK (2015).
13. Chen, A.: When the Internet's 'Moderators' Are Anything But, <https://www.nytimes.com/2015/07/26/magazine/when-the-internets-moderators-are-anything-but.html>.
14. Peckham, M.: Reddit Demotes Technology Section to Punish Lazy Moderators, <http://time.com/71644/reddit-demotes-technology-section-to-punish-lazy-moderators/>.
15. Sauder, M., Lynn, F., Podolny, J.M.: Status: Insights from Organizational Sociology. *Annual Review of Sociology*. 38, 267–283 (2012).
16. Azoulay, P., Stuart, T., Wang, Y.: Matthew: Effect or Fable? *Management Science*. 60, 92–109 (2014).
17. Ma, M., Agarwal, R.: Through a Glass Darkly: Information Technology Design, Identity Verification, and Knowledge Contribution in Online Communities. *Information Systems Research*. 18, 42–67 (2007).
18. Angrist, J.D., Pischke, J.-S.: *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton Press, Princeton, NJ (2009).
19. Lee, D.S., Lemieux, T.: Regression Discontinuity Designs in Economics. *Journal of Economic Literature*. 48, 281–355 (2010).
20. Lee, D.S., Moretti, E., Butler, M.J.: Do Voters Affect or Elect Policies? Evidence from the U. S. House. *The Quarterly Journal of Economics*. 119, 807–859 (2004).
21. Armstrong, A., Hagel, J.: The Real Value of Online Communities. *Knowledge and Communities*. 74, 85–95 (2000).
22. Butler, B.S.: Membership Size, Communication Activity, and Sustainability: A Resource-Based Model of Online Social Structures. *Information Systems Research*. 12, 346–362 (2001).
23. Faraj, S., Jarvenpaa, S.L., Majchrzak, A.: Knowledge Collaboration in Online Communities. *Organization Science*. 22, 1224–1239 (2011).
24. Bayus, B.L.: Crowdsourcing New Product Ideas over Time: An Analysis of the Dell IdeaStorm Community. *Management Science*. 59, 226–244 (2013).
25. Porter, C.E., Donthu, N.: Cultivating Trust and Harvesting Value in Virtual Communities. *Management Science*. 54, 113–128 (2008).
26. Huang, P., Mithas, S., Tafti, A.: Platform Sponsor's Investments and User Contributions in Knowledge Communities: The Role of Knowledge Seeding," accepted for publication. *MIS Quarterly*. (2017).
27. Kane, G.C., Fichman, R.G., Gallagher, J., Glaser, J.: Community Relations 2.0. *Harvard Business Review*. 87, 45–50 (2009).
28. Aaltonen, A., Seiler, S.: Cumulative Growth in User-Generated Content Production: Evidence from Wikipedia. *Management Science*. 62, 2054–2069 (2016).
29. Varian, H.R.: *Microeconomic Analysis*. Norton & Company, New York, NY (1992).
30. Wasko, M.M., Teigland, R., Faraj, S.: The Provision of Online Public Goods: Examining Social Structure in an Electronic Network of Practice. *Decision Support Systems*. 47, 254–265 (2009).

31. Lerner, J., Tirole, J.: Some Simple Economics of Open Source. *The Journal of Industrial Economics*. 50, 197–234 (2002).
32. Von Krogh, G., Von Hippel, E.: The Promise of Research on Open Source Software. *Management Science*. 52, 975–983 (2006).
33. O'Mahony, S., Ferraro, F.: The Emergence of Governance in an Open Source Community. *Academy of Management Journal*. 50, 1079–1106 (2007).
34. Ren, Y., Kraut, R.E.: A Simulation for Designing Online Community: Member Motivation, Contribution, and Discussion Moderation. Working Paper. (2011).
35. Hann, I.-H., Roberts, J.A., Slaughter, S.A.: All Are Not Equal: An Examination of the Economic Returns to Different Forms of Participation in Open Source Software Communities. *Information Systems Research*. 24, 520–538 (2013).
36. Bateman, P.J., Gray, P.H., Butler, B.S.: Research Note—The Impact of Community Commitment on Participation in Online Communities. *Information Systems Research*. 22, 841–854 (2011).
37. McWilliam, G.: Building Stronger Brands Through Online Communities. *Sloan Management Review*. 41, 43 (2000).
38. Berger, J., Cohen, B.P., Zelditch Jr, M.: Status Characteristics and Social Interaction. *American Sociological Review*. 37, 241–255 (1972).
39. Roberts, J.A., Hann, I.-H., Slaughter, S.A.: Understanding the Motivations, Participation, and Performance of Open Source Software Developers: A Longitudinal Study of the Apache Projects. *Management Science*. 52, 984–999 (2006).
40. Spaeth, S., von Krogh, G., He, F.: Research Note—Perceived Firm Attributes and Intrinsic Motivation in Sponsored Open Source Software Projects. *Information Systems Research*. 26, 224–237 (2014).
41. Imbens, G.W., Lemieux, T.: Regression Discontinuity Designs: A Guide to Practice. *Journal of Econometrics*. 142, 615–635 (2008).
42. Eggers, A.C., Fowler, A., Hainmueller, J., Hall, A.B., Snyder, J.M.: On the Validity of the Regression Discontinuity Design for Estimating Electoral Effects: New Evidence from Over 40,000 Close Races. *American Journal of Political Science*. 59, 259–274 (2015).
43. Gelman, A., Imbens, G.: Why High-order Polynomials Should not be Used in Regression Discontinuity Designs. National Bureau of Economic Research (2014).
44. McCrary, J.: Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test. *Journal of Econometrics*. 142, 698–714 (2008).