

Effects of Peer Feedback on Contribution: A Field Experiment in Wikipedia

Haiyi Zhu, Amy Zhang, Jiping He, Robert E. Kraut, Aniket Kittur
Carnegie Mellon University

5000 Forbes Avenue, Pittsburgh, PA 15213

haiyiz@cs.cmu.edu, {amyz, jiping}@andrew.cmu.edu, robert.kraut@cmu.edu, nkittur@cs.cmu.edu

ABSTRACT

One of the most significant challenges for many online communities is increasing members' contributions over time. Prior studies on peer feedback in online communities have suggested its impact on contribution, but have been limited by their correlational nature. In this paper, we conducted a field experiment on Wikipedia to test the effects of different feedback types (positive feedback, negative feedback, directive feedback, and social feedback) on members' contribution. Our results characterize the effects of different feedback types, and suggest trade-offs in the effects of feedback between the focal task and general motivation, as well as differences in how newcomers and experienced editors respond to peer feedback. This research provides insights into the mechanisms underlying peer feedback in online communities and practical guidance to design more effective peer feedback systems.

Author Keywords

Peer Feedback; Field Experiment; Online Communities; Wikipedia.

ACM Classification Keywords

H.5.3 [Information Interfaces and Presentation]: Group and Organization Interfaces – Collaborative computing, Computer-supported cooperative work, Web-based interaction; K.4.3 [Computers and Society]: Organizational Impacts – Computer supported collaborative work.

General Terms

Experimentation; Management.

INTRODUCTION

Over the past decade, thousands of volunteers working in online communities have created complex products important to society. Examples include Wikipedia, the world's largest and most popular general reference work on the Internet [17]; Linux, the leading operating system on servers [15]; and Apache, the open source web server that hosts half of the world's web pages [13]. One significant and prevalent problem of such online communities is

increasing members' contributions so as to keep active and healthy.

Empirical research has demonstrated that peer feedback predicts the amount and quality of recipients' subsequent contributions. On the one hand, feedback can increase contribution; for example, Choi et al. showed that initial communication (e.g., welcoming messages, assistance and constructive criticism) with new project members was associated with increased contribution and member retention [5]. Feedback can also decrease motivation: Halfaker et al. found that members whose edits were reverted reduced their contribution, with this effect greatest for newcomers [8]. A more nuanced view may be required which takes into account the type of feedback given. For example, Zhu et al. [19, 20] characterized four different types of peer feedback (i.e., positive, negative, directive and social) and found they were associated with either increases or decreases in contribution [4].

A key problem with these and other similar studies investigating the effects of peer feedback in online settings [5, 8, 4, 12, 14, 19] are their correlational nature, which limits the causal inferences one can draw. For example, Wikipedians who receive negative feedback from peers likely behave differently from those who receive positive feedback and these pre-existing differences may account for outcomes researcher want to attribute to the feedback. In this paper, we address this problem of endogeneity through a true experiment in which Wikipedians who write new articles randomly receive different feedback messages. We investigate whether feedback messages have different impact depending on the receivers' experience, and whether they affect performance on a specific task that the messages explicitly target as well as on more general contribution.

This research seeks to understand the mechanisms underlying peer feedback in online communities and provides practical guidance to design more effective peer feedback systems.

THEORY AND HYPOTHESES

Since Zhu et al.'s feedback categorization is nuanced and general, we follow their feedback classification [19, 20] in this paper. In this section, we first briefly summarize the four feedback types. Then we use feedback intervention theory (FIT) to predict the effects of feedback on people's performance on the specific task the messages were designed to influence and on general contribution.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2013, April 27–May 2, 2013, Paris, France.

Copyright © 2013 ACM 978-1-4503-1899-0/13/04...\$15.00.

Feedback Types

Feedback is defined as “actions taken by (an) external agent(s) to provide information regarding some aspect(s) of one’s task performance” [11]. Feedback can come from school teachers, company managers, peers, boarding school counselors, or even computer programs and in online setting can come from formal leaders, peers and bots.

Zhu et al. identified four types of feedback from peer leaders¹ in Wikipedia [20]. Table 1 shows the feedback types, description, and example messages from Wikipedia. Specifically, they identified positive comments to a contributor (positive feedback), criticisms or reprimands to a contributor for not complying to Wikipedia guidelines (negative feedback), giving directions to correct an error (directive feedback), and feedback with a sociable, person-focused, friendly, and supportive tone (social feedback).

Zhu et al. [19] investigated the effectiveness of these four types of feedback in Wikipedia by comparing changes in editing behavior of those editors who received feedback messages to those who did not receive messages [19]. Their findings suggest that receiving feedback messages changed recipients’ subsequent editing in Wikipedia. Positive feedback and social feedback can motivate members to contribute more, whereas negative feedback decreases members’ contributions. Furthermore, feedback messages from core members (i.e., Wikipedia administrators) are more influential on recipients’ subsequent behaviors than the feedback sent from peripheral members.

However, the previous research has three limitations. First and most importantly, to the best of our knowledge, research examining peer-feedback and peer-influence in online communities [5, 8, 4, 12, 14, 19] has been correlational. As previously indicated, one cannot conclude from correlational research that leadership feedback actually changes the behavior of those who receive it. Zhu et al.’s research [19] used propensity score matching to try to equate pre-existing characteristics, such as the amount of articles Wikipedia editors have previously edited and the types of communication they have received in the past, of those who received and did not receive messages. However, omitting variables, such as politeness or extraversion that potentially predict both the type of messages people receive and their subsequent behavior can still undermine causal inferences. In contrast, randomly assigning someone to receive a particular type of feedback message or not ensures that within the limits of chance

¹ Zhu et al. theorized feedback types as leadership behaviors [19, 20]. Although leadership behavior, defined as the behavior of persuading and influencing other people to pursue a common goal [18], is not identical to feedback intervention, the concept overlaps substantially. Many leadership influence attempts are interpreted by the recipient as feedback from others about their prior behavior.

those two groups were equivalent on both measured and unmeasured variables before the intervention.

Moreover the use of a true experiment allows us to investigate limits of feedback only hinted at in prior research. While previous research suggests that feedback messages have stronger effects when delivered by formal leaders [19], it failed to examine how effectiveness varies with differences among people who receive them. There is correlational evidence showing that newcomers and established members react differently when receiving feedback. For example, Halfaker et al. suggests that having one’s edits reverted has stronger negative effect for newcomers [8]. We tested this distinction more definitively in the experiment reported here.

Finally, most prior research only investigated how feedback affected receivers’ general motivation to work (e.g., total number of edits). It remains unknown about how feedback affects people’s performance on the specific tasks which the feedback explicitly targets. Research in education demonstrates that negative feedback motivates students to reduce the gap between current and desired understandings and thus enhances learning on the domain which feedback targets [9]. Other research on peer influence in Wikipedia suggests that while receiving negative feedback depresses motivation, it can improve task performance among editors who continue to participate [8]. We used Feedback Intervention Theory (FIT), described in more detail below, to predict how different types of feedback influence both general work motivation and specific task performance.

Feedback Types	Description and Message Examples
Positive Feedback	Feedback intended to energize people through acknowledging work and provides rewards <i>“I’m so impressed. This is a very fine article!”</i>
Negative Feedback	Feedback intended to regulate people through negative messages, warnings and reprimands. <i>“If you continue in this manner you will be blocked from editing without further warning.”</i> <i>“...there is a concern that the rationale you have provided for using this image under ‘fair use’ may be invalid”</i>
Directive Feedback	Feedback intended to direct people through issuing instructions, commands, assigning tasks, setting goals. <i>“Please read the instructions at... Using one of the template...”</i> <i>“... do you think you could take some pictures at Mission Mill? I’d like to spruce up the article but it really needs some photos...”</i>
Social Feedback	Feedback intended to maintain close social relationships, support group cohesion, and develop subordinates’ self-confidence and skills. <i>“Drop me a line on my talk page sometime, we’ll get a coffee...”</i> <i>“Hi XX. Welcome to WikiProject XXX! We are glad to have you. If I can help at all let me know :) ...”</i>

Table 1. Feedback types, description, and example messages.

Effects of feedback

Feedback Intervention Theory (FIT)

The first key assumption of FIT is that feedback is processed hierarchically [11]. To simplify the presentation, the hierarchy can be divided into two levels: *meta-task processes* involving the self (e.g., self-goals and self-beliefs) and *task processes* involving the focal task and the detail of the task². Processes at the higher level (i.e., the meta-task processes) can supervise the performance in the lower level (task level processes). The processes in the lower level may also divert attention up the hierarchy and influence higher level process.

The second key assumption of FIT [11] is that people use feedback to evaluate their performance relative to their standards, often referred to as feedback-standard comparisons. When they note a discrepancy between performance and standard, people are motivated to reduce it. Typically people choose to eliminate the discrepancy by attempting to attain the standard.

Based on these two assumptions, we can predict people's reaction towards four types of feedback – providing positive feedback, negative feedback, directive feedback and social feedback.

First, positive feedback, negative feedback and directive feedback are all task-oriented feedback and focus on details and progress towards a focal task. Negative feedback, which signals that performance falls short of a standard, leads people to increase effort towards the focal task. Directive feedback, which provides instructions to either achieve standards or raise standards, will also lead people to invest more effort in the focal task and improve performance. In contrast, positive feedback signals that performance exceeds the standard. Therefore, when people receive positive feedback, they typically maintain their effort or even reduce it [11]. In contrast, social feedback focus on the person level rather than the task level, and therefore should have little effect on people's performance on specific task.

H1. Providing negative feedback and directive feedback can increase people's effort on focal task and improve task performance; while positive feedback and social feedback should have less effect on focal task performance.

² The hierarchy can be more complex and contain more sublevels. Actually in the Kluger and DeNisi's original work [11], they divided the hierarchy into three levels: meta-task, task-motivation (involving focal task) and task-learning (involving task details). Hattie and Timperley added one more level—self-regulation level [9]. However, the two-level abstraction is already enough to explain the mechanisms of how four types of feedback work.

Although positive feedback tends to have little effect on performance of specific tasks, it has its effects at the meta-task level, influencing people's view of themselves. Positive feedback might increase people's self-efficacy and self-esteem and thus increase their general motivation to work. This increased motivation might spill over to non-focal tasks [11], lead to persistence in an activity and increase self-report interest in the activity [7]. Similarly, although social feedback does not affect specific task performance, it can help to develop people's self-confidence, build commitment toward the community and thus increase general motivation. In contrast, negative feedback might be perceived as a threat to self-esteem and decrease motivation. Directive feedback does not draw attention to the self-level and should not influence motivation.

H2. Providing positive feedback and social feedback can increase people's general motivation to work; directive feedback has limited effects on general motivation, while providing negative feedback might decrease people's motivation.

Moderating effects of prior experience.

The effect of feedback is substantially influenced by the willingness of the recipients to respond to the feedback [10]. Specifically, prior experience is an important variable moderating the reaction to feedback interventions [10]. People with little experience in a task are less certain about standards and their abilities. In conventional organizations, newcomers, in contrast to more established members, have greater uncertainty regarding role requirements. As a result, they are especially eager to try to learn the beliefs, values, orientations, behaviors, skills, and so forth necessary to fulfill their new roles and function effectively within an organization [1]. Therefore, we expect that newcomers will be particularly susceptible to influence, compared to experienced members [1]. Therefore, we propose our third hypothesis.

H3. Positive, negative, directive and social feedback is more influential on newcomers than experienced users.

METHOD

To test the hypotheses, we conducted a field experiment in Wikipedia. In the experiment, we sent different types of feedback messages or no message at all to editors who had recently created new articles.

Participants

Research participants were the original authors of newly created Wikipedia articles that were 2 to 10 days old. They were randomly selected without replacement via a computer script from Wikipedia's new article list. The lower bound of 2 days old ensured compliance with Wikipedia's policy to give users, especially new ones, some time to revise their articles before critiquing them. The upper bound of 10 days old ensured that the user we chose was still involved in the article's development.

We excluded participants whose new articles were deleted or tagged for a “speedy deletion” by other Wikipedians. We also excluded editors whose new articles were under 1500 characters, because these articles contained too little content on which to give feedback. We also excluded authors of pages that were simply lists of other pages. For example, the Wikipedia page “Science (disambiguation)” is a list of articles associated with science, and would have been excluded. The authors of other pages that were clearly tagged as “under major restructuring” or indicated that Wikipedia administrators were changing a page for copyright reasons, were excluded because the authors already understood the page’s shortcomings. Finally, each new article was evaluated on several dimensions to insure that potential feedback messages were relevant to it. If the article was not relevant to at least one feedback template, the author was excluded. For example, authors of new articles with nothing explicitly incorrect were excluded, because that editor could not be randomly assigned to receive negative feedback or not. Similarly, editors of an article that contains nothing praiseworthy were dropped because it could not randomly receive positive feedback.

Seven-hundred and three editors were included in the experiment. The experiment period lasted from August 2011 to November 2011.

Design

Eighty percent of selected Wikipedia editors were randomly assigned to receive a message, and the remaining twenty percent who did not receive a message served as a control group. All messages contained some common content (the base). The additional components - positive feedback, negative feedback, directive feedback, and a sociable wrapper (including a social greeting and a social closing) - each had a 50% chance of inclusion. We used a 2 (social wrapper vs not) x 2 (positive feedback vs not) x 2 (negative feedback vs not) x 2 (direction feedback vs not) between subjects factorial design for the 80% who received a message. To understand the effects of different types of messages, we measured the users’ contribution to the particular article we gave feedback to (efforts on focal task) as well as their contributions to any Wikipedia articles (general motivation) over the following month.

Materials

All messages contained some or all of the following components.

“*[Social Greeting] + [Base Message] + [Positive Feedback] + [Negative feedback] + [Directive Feedback] + [Social Closing] + [Signature]*”.

Figure 1 is an example which contains all the components. All messages contained a base and signature. In order to provide experimental control, a computer script randomly decided whether to include the additional components - positive feedback, negative feedback, directive feedback, or a sociable wrapper (social greeting plus social closing).

The base message provides an addition control. Comparing the behavior of editors who received only the base message with those receiving no-message at all shows the effects of getting a non-feedback message that did not contain any explicit positive, negative, directive or sociable components. The base message simply acknowledged the editor’s new article (e.g., “I just thought I'd let you know that I saw your article [[title]] in the New Articles list”). All feedback messages included the base message component, drawn randomly from one of four base messages, and the researcher’s signature, automatically generated by using four tildes “~~~~”.

We created twelve templates for positive feedback, ten templates for negative feedback, nine templates for directive feedback, four templates for social greeting and eight templates for social closing. Table 2 shows two examples of each message component, and Figure 1 shows an example of a message assembled from the components.

To generate both positive feedback and negative feedback components, a script was used to run through the various templates in a random order, asking the researcher if a specific positive or negative template applied to the article. This ensured that the aspect was both appropriate and randomly chosen. Note that the negative feedback only politely critiqued the editor’s work by pointing out an error, but was not directive, such as requesting that the editor make a particular change.

Hello <u>[[participant's username]]</u> , I just thought I'd let you know that I saw your article <u>[[title]]</u> in the New Articles list--	The information is
Social greeting	Base message
<u>presented clearly and is easy to understand.</u>	<u>However, I noticed the article contains an error: this article currently does not contain any</u>
Positive feedback	Negative feedback
<u>references. As a new article, the most important thing is to find reliable references for all existing information..</u>	<u>It would be great if you</u>
Negative feedback	
<u>could also upload a picture for the related article <u>[[title]]</u>.</u>	<u>Kind regards and happy editing!</u> Jipinghe (talk) 19:20, 30 November 2011
Directive feedback	Social closing Signature

Figure 1. An example message containing all the elements.

Messages assigned to the social condition contained both the social greeting and social closing component. The social greeting component contained a personal greeting, such as “hello”, followed by the editor’s Wikipedia username. The social closing component contained phrases of encouragement, or wished the editor a good day, etc. A script was used to randomly select a social greeting and social closing component from our selection of Sociable message templates.

Directive feedback asked for the editor’s help with improving a related article without being positive or negative about the new article that the user created. We used Suggestbot [6] to help find related articles that needed work. SuggestBot matched keywords within the new article, and generated the top five choices that were most closely related to the article. For those participant who were randomly assigned to directive conditions by our computer script, the researcher chose the most relevant, appropriate, and valid article suggested by SuggestBot to be incorporated into the message.

Research Ethics

We designed this experiment with the twin goals of observing how different types of messages naturally affect Wikipedia editors while at the same time minimizing potential risk to Wikipedia editor-participants and the Wikipedia community as a whole.

First, we made sure that the feedback messages sent to Wikipedia editors who have created a new page were

natural and appropriate. The researchers posting the messages are members of the New Page Patrol, a collection of Wikipedia editors who evaluate and comment on new articles. They both had experience editing in Wikipedia. Furthermore, all the component templates sent to editors were based on observations of messages on Wikipedia, suggestions by senior Wikipedia editors, and the guidelines of civility in Wikipedia. Thus, these messages are very similar to those that Wikipedia users might encounter in their everyday interactions on the website, although perhaps more polite.

In particular, negative feedback components in the experiment are milder than negative feedback messages often sent between editors. In the wild, some editors use intimidation, threat and harsh language to decrease undesired behaviors from targets. Here are two examples: “If you continue in this manner you will be blocked from editing without further warning” and “Blech. This really needs [[WP:TNT]],” which is Wikipedia’s jargon for “Blow it up and start over.” In our experiment design, negative feedback consisted only of constructive criticism.

To insure that recipients had natural reactions to the feedback messages, messages did not indicate that they were sent as part of an experiment, and recipients were not asked to sign consent forms. If participants receiving messages knew they were in an experiment, we would not be able to generalize the results to understand how editors react to messages from other Wikipedians. We believe these messages will have their effects in part because they imply a relationship between the sender and recipient. This relational meaning would be undercut if the recipients believed they received feedback simply because they were part of an experiment. However, we did not try to hide the experimenters’ research affiliation. The researchers who sent the messages had a brief description of the experiment and links to a fuller description on their user pages. Their user pages also contained instructions on how to withdraw from participation, for users who wanted to opt out. However, none of our participants have opted-out.

If message recipients initiated further interaction with the experimenters (e.g., asking follow-up questions), experimenters responded as briefly as possible, and pointed the recipients to appropriate pre-existing Wikipedia help pages. This prevented participants from feeling confused because the researcher initiated contact, but failed to maintain it. The brevity and links to other relevant Wikipedia was intended to protect against confounds from additional feedback from the researchers.

The experiment was approved by the Carnegie Mellon Institutional Review Board, as well as the Wikipedia research committee. Information about the experiment was posted on public Wikipedia pages and received unanimous agreement of active discussants from the Wikipedia community [16].

Component Type	Template 1	Template 2
Social Opening	<i>Hi XX,</i>	<i>Hey XX,</i>
Base Message	<i>I'm posting this message on your talk page because you've recently created the new article XX --</i>	<i>I saw your article XX in the new articles list --</i>
Positive Feedback	<i>The content seems well-organized.</i>	<i>There is a good number of citations and references.</i>
Negative Feedback	<i>However, I noticed the article contains an error: this article currently does not contain any references. As a new article, the most important thing is to find reliable references for all existing information.</i>	<i>However, I noticed the article contains an error: the article does not contain any Wikilinks, and so doesn't follow Wikipedia style guidelines.</i>
Directive Feedback	<i>It would be great if you could also improve the related article XX.</i>	<i>It would be great if you could also clean-up the related article XX.</i>
Social Closing	<i>Happy editing! Hope your day is going well and you are having fun.</i>	<i>It's always nice to see users contributing to make Wikipedia better!</i>

Table 2. Example templates for message components.

ANALYSIS

The goal of the analysis was to measure the effects of messages on participants' efforts on focal task and general motivation.

Dependent variables

- *Efforts on focal task.* To measure participants' performance on their focal task, we calculated the number of edits they made in the month after receiving a message on the article that was the target of the message. Note that for participants who received a directive message asking them to improve a related article, efforts on focal task also included edits on that related article³.
- *General motivation.* To measure the effects of messages on participants' general motivation to work, we calculated the number of edits the participants made in the month after receiving our feedback messages on any Wikipedia articles excluding the focal article(s) which the feedback messages target.

Independent variables

- *Positive feedback.* This dummy variable indicates whether the participant received a message with the positive feedback component (1) or without this component (0).
- *Negative feedback.* This dummy variable indicates whether the participant received a message with the negative feedback component (1) or without this component (0).
- *Directive feedback.* This dummy variable indicates whether the participant received a message with the directive feedback (1) or without this component (0).
- *Social feedback.* This dummy variable indicates whether the participant received a message with the social components (1) or without the components (0).
- *Receiver is a newcomer.* This dummy variable indicates whether the receiver is a newcomer (1) or not (0). We define newcomers as editors with less than six months experience in Wikipedia and had received fewer than four messages before receiving our message.

³ We also tested the effects of directive feedback on the two types of articles separately and found that directive messages have positive effects on both the related article and the original one. Directive messages led participants (both experienced members and newcomers) to make on average 0.4 edits on the related article, which is significantly more than zero ($p < .02$). Furthermore, newcomers receiving directive messages made 2.8 more edits than those who did not receive directive messages on the original articles ($p < .01$). Directive messages have no statistically significant effects for the experienced members on the original articles.

	Newcomers	Experienced editors
Number of people	132	473
Efforts on focal task Unit: # of edits	$M = 2.1; SD = 7.6$	$M = 1.3; SD = 3.7$
General motivation Unit: # of edits	$M = 128; SD = 25$	$M = 403; SD = 959$
# of people receiving messages	106	362
# of people receiving positive feedback	45	183
# of people receiving negative feedback	48	164
# of people receiving directive feedback	47	126
# of people receiving social feedback	61	194

Table 3. Descriptive statistics of the participants

Analysis strategy

We excluded participants who were blocked by Wikipedia or those whose new articles were deleted within one month after receiving our messages. Finally, we included six hundred and five participants in the analysis. The statistics of these participants is shown in Table 3.

Because the dependent variables (the number of edits editors made on particular target articles and other WPO articles) are count data and because editors might not log in to Wikipedia and have a chance to see the messages during the time window (one month after receiving the message), we analyzed the data using a zero-inflated negative binomial regression.

Zero-inflated negative binomial regression is often used when the dependent variable is a count value and is over dispersed, with more zeros than predicted by a regular binomial distribution (i.e., variance is much larger than the mean). The basic idea is that the excess zeros can be generated by a separate process that can be modeled independently. In our case, the goal is to predict whether reading the messages changes participants' behavior. Some recipients might not have been influenced by the message, because they were not persuaded by its content. However, others might have failed to log in recently and hadn't actually seen the message meant for them. To model these two separate processes, the zero-inflated negative binomial analysis has two stages. In the first stage, we used a logit regression to predict the excess zero (i.e., the likelihood of not seeing the message). In the second stage, given the likelihood of being exposed to the message, we predicted the effects of messages on the number of edits. Specifically, we used the following two estimates of editors' recent activity to predict the likelihood of their seeing the message.

- *Number of edits one day before receiving our message.* The more edits the participant did in the 24 hours before we sent them messages, the more active they were and the more likely they were to have seen our message.

- *Number of days between last edit and receiving our message.* Similarly, we included the number of days between the last edit the participant made and the time we sent our message.

Analysis results

The results of zero-inflated negative binomial regression are shown numerically in Table 4 and graphically in Figures 2 to 5. The error bars in Figures 2 to 5 indicate 95% confidence interval. The analysis focuses on the participants who received our messages. We treat the behavior for participants who didn't receive messages at all as the baseline and indicate the baseline with dotted lines in Figures 2 to 5. We report the main effects of receiving a particular type of feedback component. For example, in the figures, the condition of "with social feedback components" includes "social" and "social + positive" and "social + negative" etc; the condition of "without social feedback components" includes "base" and "positive" and "negative" etc. We did not find significant interaction effects between different types of feedback components.

The bottom panel of Table 4 indicates that the likelihood ratio test of $\alpha = 0$ is significantly different from zero. This suggests that our data is overdispersed and that a zero-inflated negative binomial model is more appropriate than a

zero-inflated Poisson model. The Vuong test suggests that the zero-inflated negative binomial model is a significant improvement over a standard negative binomial model. Furthermore, the predictor "number of days between last edit and receiving our message" in the part of the logit model predicting excessive zeros is statistically significant. The longer the time between the last edit and the time receiving our message, the less likely the participant was to edit, suggesting that he or she missed our message. These results suggested that we used the right statistics model.

The top panel of Table 4 shows analyses testing hypotheses 1-3. Model 1 tests whether receiving feedback message led editors to edit more on the article the feedback message targets (focal task). Model 2 tests whether receiving feedback message increased editors' activities in general. Each coefficient represents the change in the log of the expected number of edits the editor will produce when increasing the independent variable by one unit, when other variables in the model are held constant at zero. For ease of interpretation, we also included the change in edit counts in the original units. Thus, the intercept indicates that old-timers who received base messages (i.e., baseline) can be expected to make 1.7 ($e^{.51}$) edits to the focal article (see Intercept in Model 1). Newcomers made edits 0.3 ($(e^{-1.2})$)

Dependent variable	Focal task			General motivation		
	Model 1			Model 2		
Predictors	Coef	S.E.	Change in # of edits	Coef	S.E.	Change in # of edits
Intercept	.51	(.24)		6.2**	(.20)	
Positive feedback	.09	(.23)	1.09	-.051	(.19)	0.95
Negative feedback	.04	(.24)	1.04	-.15	(.19)	0.86
Directive feedback	-.09	(.25)	0.91	-.036	(.19)	0.96
Social feedback	.06	(.24)	1.06	-.13	(.19)	0.88
Receiver is newcomer	-1.2*	(.66)	0.30	-4.5**	(.50)	0.01
Newcomer * Positive feedback	-.48	(.70)	0.62	1.3**	(.54)	3.67
Newcomer * Negative feedback	1.4**	(.65)	4.06	-.26	(.54)	0.77
Newcomer * Directive feedback	2.2**	(.66)	9.03	.60	(.51)	1.82
Newcomer * Social feedback	.25	(.68)	1.28	2.2**	(.50)	9.03
Inflate						
Number of edits during one day before receiving our message						
Number of days between last edit before receiving our message and the time they receive the message						
Alpha			3.3			2.6
Likelihood-ratio test of $\alpha=0$	chibar2(01) = 434; Pr>=chibar2 = 0.0000			chibar2(01) = 2.4e+5; Pr>=chibar2 = 0.0000		
Vuong test of zinb vs. standard negative binomial	z = 3.41 Pr>z = 0.0003			z = 1.67 Pr>z = 0.05		

Table 4. Model 1 measures the effects on efforts to the focal task: participants' editing activities on the particular articles which our feedback message targets; Model 2 measures the effects on general motivation: the effects on participants' general contributions excluding the particular articles which the feedback message explicitly targets. ** p<0.05; * p<0.1 .

times compared to experienced editors. Therefore, newcomers who received base messages make 0.51 edits (1.7*0.3) to the focal article. Newcomers who received negative feedback made 2.1 edits (1.7*0.3*4.06) to the focal article.

For experienced editors, receiving any type of feedback message has no significant impact on their subsequent editing behavior, either for the specific articles on which we gave feedback (focal task) or any other articles (general motivation).

Model 1 shows that messages had significant effects on newcomers' subsequent editing of the target as our hypotheses predict. Receiving negative feedback and directive feedback increased their editing in the target article. The coefficient of newcomer*negative feedback is 1.4, indicating that newcomers who received negative feedback are estimated to make edits on focal articles approximately four times compared to newcomers who did not receive negative feedback. The coefficient of newcomer*directive feedback is 2.2, indicating that

newcomers who received directive feedback messages are estimated to make edits on focal articles approximately nine times compared to newcomers who did not receive directive feedback messages. Positive feedback and social feedback do not have effects on local tasks.

Results of Model 2 confirm our hypothesis about the effects of messages on editors' general motivation. In contrast to Model 1, negative feedback and directive messages do not have effects on general motivation. Instead, messages with positive feedback and social feedback components substantially increase newcomers' general work motivation. The coefficient of newcomer*positive feedback is 1.3, indicating that positive feedback causes 3.67 times change in number of edits for newcomers. The coefficient of newcomer*social feedback is 2.2, indicating that messages with social feedback cause 9.03 times change in number of edits for newcomers.

DISCUSSION

The results basically confirm our hypotheses: 1) negative feedback and direction increase people's efforts on focal




Figure 2. The effects of receiving messages on newcomers' efforts on focal task. The dotted line indicates the baseline: newcomers receiving no messages (1.4).




Figure 3. The effects of receiving messages on newcomers' general motivation. The dotted line indicates the baseline: newcomers receiving no messages (7.3).




Figure 4. The effects of receiving messages on experienced members' efforts on focal task. The dotted line indicates the baseline: experienced members receiving no messages (1.0).




Figure 5. The effects of receiving messages on experienced members' general motivation. The dotted line indicates the baseline: experienced members receiving no messages (500).

tasks; 2) positive feedback and social messages increase people's general motivation to work; 3) the effects are stronger for newcomers (in Table 4, the effects are only significant for newcomers).

Experienced Members' Reaction

Although we predict that the effects should be stronger for newcomers because they are particularly susceptible to influence, we are still surprised to see that the messages had no significant effects at all on experienced members. From the messages the participants sent back to the researchers' user pages, we found evidence that experienced members might have psychological reactance to our messages. Psychological reactance was originally proposed by Brehm, in which a person has a negative emotional response in reaction to being persuaded, and thus chooses the option which is being advocated against [3]. For example, some participants wrote to us and said that:

"Well, er, yes, I am not new here and the stub tag was intended as a cheerful acknowledgement of the effort's insufficiency." – P1.

"There are plenty of external references on that page for John Hess (journalist) for the information given. I can show you plenty of pages that do not have any external references - worry about those first..." – P2.

"You're still wet behind the ears and have too little experience to have perspective." – P3.

Participants' comments above suggest that experienced members might perceive negative feedback and directive message as a challenge to their knowledge and expertise (P1 and P2), especially when noticing that the message senders⁴ have less experience than themselves (P1 and P3). Previous research shows that when people perceived feedback as self-threatening, they might avoid exposure to the feedback [2] or even abandon the entire task [11]. It is possible that experienced editors chose not to follow what their newbie colleagues suggested, so as to preserve positive self-belief about their expertise. The results suggest that although any member can try to send feedback to others in online communities, the relative status of the sender might still matter. Therefore, to ensure the effectiveness of feedback, it is probably better to have senior community members to deliver feedback messages.

Focal and general

Our results demonstrate the tradeoff of different types of feedback on recipients' focal task performance and general

work motivation. Negative feedback and directive feedback benefit focal task performance but do not have effects on general work motivation, while positive feedback and social feedback can positively influence general work motivation but do not have effects on focal tasks. Practitioners can consider their primary goal (e.g., accomplishing current task or encouraging long-term motivation) and then design feedback messages accordingly.

One possible way to optimize both the focal and the general effects is to combine multiple feedback types. Unfortunately, although the study was originally designed as a full factorial experiment, the sample size is small because the effects seem to be only significant for the newcomers. Therefore we have limited statistical power to draw any reliable conclusions about the interaction effects between different feedback types. In future work, we will increase the sample size to explore the optimal feedback intervention strategy.

Effects of negative feedback

Our results suggest that negative feedback caused newcomers to work harder on the target article and no effects on general motivation, while Zhu et al. [19] and Halfaker et al. [8] found negative feedback reduced motivation. Remember that the negative feedback messages in our experiment were intentionally designed to be milder than negative feedback messages actually sent between Wikipedia editors (as examined in Zhu et al.'s study [19]). Furthermore, the negative feedback Halfaker and his colleagues investigate is the action of reverting, which is stronger than textual messages since "Actions speak louder than words". The different ways of conveying the same meaning might influence how the recipients interpret and perceive the negative feedback, and thus lead to different reactions. One direction of future work is to conduct qualitative study to investigate how people interpret different types of feedback messages (especially the negative feedback) to further understand how feedback messages take effect.

CONCLUSION

We conducted a field experiment in Wikipedia to examine how different types of feedback affect receivers' focal task performance and general work motivation, moderated by receivers' prior experience. This research furthers our understanding of the effects of feedback in online communities and provides practical guidance to design effective feedback systems.

ACKNOWLEDGE

We thank the members of the Social Computing Group at Carnegie Mellon University for helpful feedback. This research was supported by NSF grants IIS-11-11124, IIS-09-68484, OCI-09-43148, IIS-11-49797, and the Center for the Future of Work.

REFERENCES

1. Ashforth, B.E., and Saks, A. M. (1996) Socialization Tactics: Longitudinal Effects on Newcomer Adjustment.

⁴ Both researchers who sent out messages are relatively new in Wikipedia. One has 2 months experiences and 133 edits before conducting the experiment; and the other has 6 months experiences and 65 edits before conducting the experiment. It is easy for experienced Wikipedia editors to find out such information in Wikipedia.

- The Academy of Management Journal*, Vol. 39, No. 1 (Feb., 1996), pp. 149-178.
2. Baumeister, R.F.; Cairns, K.J. (1992) Repression and self-presentation: When audiences interfere with self-deceptive strategies. *Journal of Personality and Social Psychology*, Vol 62(5), May 1992, 851-862.
 3. Brehm, J.W. (1966) A Theory of Psychological Reactance. Oxford, England: Academic Press. (1966). x 135 pp.
 4. Brzozowski, M. J., Sandholm, T., & Hogg, T. (2009, May). Effects of feedback and peer pressure on contributions to enterprise social media. In *Proceedings of the ACM 2009 international conference on Supporting group work* (pp. 61-70). ACM.
 5. Choi, B. R., Alexander, K., Kraut, R. E., & Levine, J. M. (2010). Socialization Tactics in Wikipedia and Their Effects *CSCW'10: Proceedings of the ACM Conference on Computer-Supported Cooperative Work* (pp. 107-116). New York: ACM Press
 6. Cosley D., Frankowski D., Terveen L., and Riedl J. (2007) SuggestBot: using intelligent task routing to help people find work in wikipedia. In *Proceedings of the 12th international conference on Intelligent user interfaces (IUI '07)*. ACM, New York, NY, USA, 32-41.
 7. Deci, E.L.; Koestner, R., Ryan, R.M. (1999) A meta-analytic review of experiments examining the effects of extrinsic rewards on intrinsic motivation. *Psychological Bulletin*, 125, 627-668.
 8. Halfaker, A., Kittur A., and Riedl J. (2011) Don't bite the newbies: how reverts affect the quantity and quality of Wikipedia work. In *Proceedings of the 7th International Symposium on Wikis and Open Collaboration (WikiSym '11)*.
 9. Hattie, J., & Timperley, H. (2007). The power of feedback. *Review of educational research*, 77(1), 81-112.
 10. Ilgen D.R., Fisher, C.D., & Taylor, M.S.(1979) Consequences of individual feedback on behavior in organization. *Journal of Applied Psychology*, 64, 349-371.
 11. Kluger, A. N., DeNisi, A. (1996) The effects of feedback interventions on performance: A historical review, a meta-analysis, and a preliminary feedback intervention theory. *Psychological Bulletin*, Vol 119(2), Mar 1996, 254-284.
 12. Lampe, C., & Johnston, E. (2005, November). Follow the (slash) dot: effects of feedback on new members in an online community. In *Proceedings of the 2005 international ACM SIGGROUP conference on Supporting group work* (pp. 11-20). ACM.
 13. March 2012 Web Server Survey <http://news.netcraft.com/archives/2012/03/05/march-2012-web-server-survey.html>
 14. Moon, J. Y., & Sproull, L. S. (2008). The role of feedback in managing the Internet-based volunteer work force. *Information Systems Research*, 19(4), 494-515.
 15. Wikipedia editors. Linux. <http://en.wikipedia.org/wiki/Linux>
 16. Wikipedia editors. Research talk:Effects of Feedback on Participation in Wikipedia http://meta.wikimedia.org/wiki/Research_talk:Effects_of_Feedback_on_Participation_in_Wikipedia
 17. Wikipedia editors. Wikipedia. <http://en.wikipedia.org/wiki/Wikipedia>
 18. Yukl, G. A. (1998). *Leadership in organizations* (4th ed.). Upper Saddle River, N.J.: Prentice Hall.
 19. Zhu, H., Kraut, R., & Kittur, A. (2012). Effectiveness of shared leadership in online communities. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work* (pp. 407-416). ACM.
 20. Zhu, H., Kraut, R. E., Wang, Y. C., & Kittur, A. (2011). Identifying shared leadership in Wikipedia. In *Proceedings of the 2011 annual conference on Human factors in computing systems* (pp. 3431-3434). ACM.