Encouraging Online Knowledge Contributions: Evidence from a Field Experiment

Yikun Jiang *

First Version: June 1, 2024; This Version: June 28, 2025

Abstract

With the prominence of user-generated content platforms, online knowledge platforms have experienced substantial growth. However, it is unclear why individuals voluntarily contribute knowledge, and how platform strategies can facilitate individuals' contributions through non-monetary incentives. I develop a stylized model of individual knowledge contributions where individuals have social motivation to gain reputation and instrumental motivation to obtain functionality privileges on the platform. Based on the model, I design and implement a large-scale field experiment, involving 12,182 individuals on one of the largest online question-and-answer platforms. I sample and manually treat participating platform individuals daily over the course of four and a half months. The treatment gives one anonymous upvote to eligible answers, exogenously shifting individuals' social and instrumental motivations. I then track comprehensive data on individuals' subsequent behavior on the platform daily for four months. I find that the treatment significantly increases an individual's probability of contributing additional answers by around 15% of the baseline, and the difference between the control and treatment groups persists over time. The treatment effect is the most pronounced for individuals with low-to-moderate answering experience or reputation and is slightly stronger for those who are close to obtaining additional privileges after the treatment. The overall quality and effort of future answers remain stable. Using data from the field experiment, I structurally estimate the model of contribution decisions to quantify the relative importance of social and instrumental motivation. Simulation results suggest that social motivation is more important, and platform strategies that boost social motivation are more effective in encouraging contributions.

Keywords: Consumer Dynamics; Innovation; Field Experiment; User-Generated Contents; Online Knowledge Platforms; Online Reputation; Platform Design; Social Preferences

*University of California, Berkeley, Department of Economics. Email: jiangyk@berkeley.edu.

Previously titled, Dynamics in Online Knowledge Contributions - Evidence from a Field Experiment. I thank advisors, Stefano DellaVigna, Matthew Osborne, Nano Barahona, Benjamin Handel, and Olivia Natan. I appreciate helpful discussions with Ned Augenblick, Luisa Cefala, Nan Chen, Yan Chen, Bryan Chu, John Conlon, Avi Goldfarb, Ali Goli, David Holtz, Kei Kawai, Jura Liaukonyte, Emaad Manzoor, Nitin Mehta, Sarah Moshary, Ricardo Perez-Truglia, Carolyn Stein, Avner Strulov-Shlain, Steven Tadelis, Dmitry Taubinsky, J. Miguel Villas-Boas, Nathan Yang, participants at UC Berkeley seminars, Spring School in Behavioral Economics at UC San Diego, Annual Meeting of the Society for the Science of Motivation, ISMS Marketing Science Conference, ISPOC (Information Systems Student Presentations Over the Cloud) Webinar Series, and Behavioral Industrial Organization & Marketing Symposium. The experiment received IRB approval from the University of California, Berkeley, before implementation. The IRB approval number is 2023-04-16245. The pre-registration ID at AEA RCT Registry is AEARCTR-0011615. All errors are my own.

1 Introduction

In recent years, user-generated content platforms have experienced rapid growth, with online knowledge platforms such as Quora and Stack Overflow attracting substantial traffic. These platforms rely on individuals to contribute knowledge as a public good, while monetizing that content through advertising and related services. Sustaining user contributions is critical to their long-term viability. For example, Yahoo Answers—launched in 2005 as the first online Q&A platform—saw a steady decline in content generation starting in 2011 and ultimately shut down in 2021 (Guo et al. 2023). Understanding how to sustain users' content contributions is therefore essential for the design and survival of such platforms.

A common feature of successful online knowledge platforms is the use of non-monetary rewards to incentivize user contributions. These rewards typically fall into two categories: social recognition from other users, such as reputation points and virtual badges, and platform-granted instrumental privileges, including editing rights or moderation access, which users can obtain once they accumulate sufficient recognition from others.[1] By offering these rewards, platforms aim to sustain engagement and content generation without direct monetary compensation. Besides knowledge platforms, non-monetary incentives are also widely employed across user-generated content ecosystems, from review sites to educational platforms and social media.[2]

A key open question is whether user contributions on content platforms are driven primarily by social motivation—the desire for recognition, social status, or the intrinsic satisfaction of sharing with others—or by instrumental motivation, in which users contribute

[1] Such platform-granted privileges are common. For example, on Twitter, Community Notes contributors must first rate a sufficient number of notes before gaining the ability to write their own. On Wikipedia, experienced users with a strong edit history can become administrators, granting access to tools such as page deletion, content protection, and user blocking.

[2] On TripAdvisor, reviewers receive badges and earn higher levels by sharing their travel experiences, leading to increased influence and recognition. On Goodreads, book reviewers can earn likes, comments, and followers by contributing reviews, lists, and recommendations. Khan Academy also uses non-monetary rewards to encourage contributions in their course discussion forum - individuals who answer questions, contribute to discussions, or practice skills receive badges and energy points that build their online reputation. Social media platforms like Instagram, Twitter, TikTok, and Facebook let individuals click "like" on others' posts as a way to encourage future posts.

to gain access to platform-specific privileges or functionalities. This distinction has important implications for platform design. If social motivation dominates, enhancing visibility or promoting recognition may increase engagement. If instrumental motives prevail, making recognitions and privileges too easily attainable may reduce effort by weakening the incentive to contribute. Understanding the relative strength of these motivations is crucial to designing effective incentives on content platforms.

In this paper, I present evidence from a large-scale field experiment on Stack Overflow, a leading online Q&A platform for programmers, to quantify the relative importance of social versus instrumental motivations in user contributions. On Stack Overflow, users earn reputation when their posts are upvoted, and specific privileges are unlocked at defined reputation thresholds. I experimentally increase the reputation of a randomly selected group of contributors, shifting both their social and instrumental motivations. This design allows me to quantify the effect of each motivation on contribution behavior.

To motivate the experimental design, I begin with a model of online knowledge contribution in which individuals derive utility from both reputation gains and access to platform privileges. The model allows individuals to form expectations about the reputation points associated with each contribution and incorporates heterogeneity in the valuation of reputation and privileges. I then describe the design of the field experiment, which was conducted over several months and included a substantial number of new and less active users—an understudied segment with significant potential to become frequent contributors.

Between August 31, 2023, and January 10, 2024, I conducted a large-scale field experiment on Stack Overflow, involving 12,182 individuals through daily sampling and manual treatment over four and a half months. The treatment consisted of adding a single anonymous upvote—via one of my own accounts—to a recent eligible answer, thereby increasing the recipient's reputation by 10 points. I then tracked detailed daily behavior for each individual through July 1, 2024, capturing post-treatment activity over an extended period. The experiment was carefully designed and implemented to ensure full compliance with platform

policies and to avoid any form of deception.[3]

Do exogenously assigned upvotes affect the decision to contribute knowledge online? I find that a single additional anonymous upvote substantially increases individuals' propensity to contribute further answers. At the extensive margin, the probability of contributing again rises by approximately 15% relative to baseline. At the intensive margin, the number of answers contributed within three weeks increases by about 6%. These effects persist for at least four months, suggesting they are not driven by short-term intertemporal substitution. Treatment effects are heterogeneous: the impact at the extensive margin is largest among individuals with low to moderate prior answering experience and reputation. Importantly, there is no evidence that treatment reduces the quality or effort of subsequent contributions. In fact, suggestive evidence shows that individuals with lower past answer quality and effort—who are more responsive to the treatment—tend to increase both when they do contribute.

Experimental upvotes may influence future contributions at the extensive margin through two main channels: by shifting beliefs related to social motivation, and by increasing reputation, thereby altering an individual's proximity to the next privilege threshold—capturing the instrumental component. To assess the importance of instrumental motivation, I estimate treatment effects separately for two groups. The first group crosses a reputation threshold due to the upvote, and thus becomes further from the next threshold; if instrumental incentives are dominant, their contribution rates may decline. The second group moves closer to the next threshold, potentially increasing contributions if instrumental motivation matters. I find that the treatment effect is positive and significant for both groups, with a slightly larger effect for those who move closer to the next threshold. This suggests that instrumental incentives contribute to behavior, but are not the sole driver—social motivation also plays a meaningful role.

Do exogenously assigned upvotes affect users' overall platform activity or other forms of

[3]The experiment received IRB approval prior to implementation.

participation? I find no significant effect on whether individuals log in within 21 days of treatment, but the treatment increases the number of login days by 0.55—approximately 3.9% relative to baseline. The upvotes do not significantly affect the likelihood of posting new questions or the number of questions posted. However, the treatment increases the probability of posting comments by 3.3 percentage points and raises the number of comments within 21 days by approximately 5.8% relative to the control group.

How do exogenously assigned upvotes influence individuals' evaluations of others' content? I find that the treatment significantly affects voting behavior within 21 days. The probability of casting at least one upvote increases by 7.3% relative to baseline, and the total number of upvotes rises by approximately 5%. For downvoting, the treatment increases the likelihood of making at least one downvote by about 20% of baseline, and the total number of downvotes increases by 2.1%.

The experimental results suggest that both recognition and access to privileges motivate knowledge contribution, measured at the extensive margin as the decision to contribute additional online knowledge—the primary outcome of interest. However, the relative importance of social motivation (e.g., reputation gains driven by status, altruism, or warm glow) versus instrumental motivation (i.e., earning privileges tied to reputation thresholds) remains unclear. If social motivation dominates, platforms might increase engagement by enhancing opportunities for recognition. In contrast, if instrumental motivation is key, making recognition easier could reduce effort by lowering the perceived difficulty of reaching privilege thresholds. Research on observational data supports this concern: contributions significantly decline after users gain platform incentives (Goes et al. 2016). In such cases, platforms might instead reinforce instrumental motivation—for example, by adding new privileges or adjusting how difficult they are to obtain.

To quantify the influence of both motivations on contribution decisions, I structurally estimate a model of individual behavior with heterogeneous utility parameters. The estimates reveal substantial variation in how individuals value social motivation versus instrumental

platform privileges.

Using estimates from the structural model, I conduct a series of counterfactual simulations to assess the relative importance of social and instrumental motivations. When social motivation is removed, the share of individuals willing to contribute falls to 26% of the baseline. In contrast, removing instrumental motivation results in a smaller decline, with contribution rates remaining at 84% of the baseline. This stark difference suggests that social motivation plays a more central role in driving contribution behavior.

I then evaluate two platform strategies aimed at strengthening social motivation: (1) amplifying visibility by tripling the number of upvotes for answers that already have at least one, and (2) awarding an additional upvote to high-quality answers. The simulations show that the first strategy increases contributions by 15% relative to baseline, while the second boosts contributions by 25%, underscoring the potential of targeted recognition to sustain user engagement.

This paper contributes to the broader literature on social preferences and public goods provision (e.g., Lerner and Tirole (2002); Andreoni (2007)). In the context of online knowledge-sharing platforms, several studies document the influence of social motivations. On Chinese Wikipedia, Zhang and Zhu (2011) finds that reducing group size exogenously lowers contributions from non-blocked users, due to diminished social benefits. Similarly, Wang et al. (2019) shows that expanding the user base of a major review platform increases both the volume and quality of reviews. Gallus (2017) finds that symbolic awards improve volunteer retention on Wikipedia, and Chen et al. (2023) shows that better alignment between recommended articles and contributor expertise increases the length and quality of comments.

This paper also contributes to the literature on incentivizing open-source innovation and public goods creation (e.g., Athey and Ellison (2014)). Conti et al. (2023), using a difference-in-differences design, finds that monetary incentives can crowd out social motivation among open-source developers, diverting their efforts away from community-oriented tasks. This issue is especially salient in the age of AI: Burtch et al. (2023) documents that the introduc-

tion of ChatGPT led to large declines in Stack Overflow activity, particularly in areas where the model had strong training data coverage.[4]

This study builds on and extends Xu et al. (2020), which examines a group of highly active Stack Overflow users invited to apply for jobs through Stack Overflow Jobs. Using a difference-in-differences design, they show that career incentives partially motivate contribution behavior. In contrast, this paper focuses on a broader and less active population—newly registered users who represent the majority of contributors but remain underexamined in the literature. It investigates how online reputation and access to platform-granted privileges influence knowledge contributions in this understudied group.[5] Moreover, the experiment was conducted while Stack Overflow Jobs was discontinued, minimizing potential career-related confounds.[6]

This paper also contributes to the broader literature on the motivations behind user-generated content (UGC) online.[7] Prior work has shown that monetary incentives can bias reviews (Cabral and Li 2015), though combining financial and social incentives can improve outcomes (Sun et al. 2017, Burtch et al. 2018). Several observational studies explore the role of non-monetary incentives in UGC settings. For example, Jin et al. (2015) finds that peer recognition is positively correlated with contributions in a Chinese Q&A community. Ahn et al. (2016) develops and estimates a dynamic rational expectations model using forum data. Goes et al. (2016) shows that glory-based incentives can lead to short-term contribution increases but long-term declines after users reach symbolic reward thresholds.[8] Deolankar et al. (2023) finds that negative peer feedback can increase subsequent commenting on Reddit, and

[4] Burtch et al. (2023) finds that Stack Overflow question volume declined significantly after ChatGPT's release, with larger drops in topics where the model had access to extensive public training data.

[5] Among the 1,301 users analyzed in Xu et al. (2020), the average number of answers per user is 255.59. In comparison, the 12,182 individuals in this study average 18.15 answers, with a median of 4. Notably, 99.02% of sampled users had contributed fewer than 255.59 answers at the time of sampling.

[6] Stack Overflow Jobs was sunset on March 31, 2022. See: https://meta.stackoverflow.com/questions/415293/sunsetting-jobs-developer-story. A redesigned version was released on May 8, 2024. See: https://meta.stackexchange.com/questions/399440/testing-a-new-version-of-stack-overflow-jobs.

[7] For a comprehensive review, see Chen (forthcoming).

[8] See also Lacetera and Macis (2010) for similar evidence in blood donation contexts, where nonlinear symbolic rewards affect donation timing.

Paridar et al. (2024) shows that peer versus platform rewards can have opposing effects on posting behavior in a gaming forum.

Moving beyond observational evidence, a growing body of field experiments examines non-monetary incentives for UGC. Chen et al. (2010) shows that social comparisons boost contributions from below-median users on a movie-rating site. Burtch et al. (2022) finds that anonymous awards increase Reddit participation. Toubia and Stephen (2013) documents that follower counts have heterogeneous effects on posting behavior on Twitter. Eckles et al. (2016) finds that receiving feedback increases Facebook activity. Jiménez Durán (2021) shows that reporting hateful content does not reduce subsequent posts. Huang et al. (2022) finds that exogenous variation in attention on an image-sharing site alters content creation behavior. More recent work—e.g., Mummalaneni et al. (2023), Srinivasan (2023), Zeng et al. (2023), and Zhang and Luo (2024)—confirms that social recognition and user engagement can significantly increase content production across platforms.

This study contributes to the literature on user-generated content, non-monetary incentives, and public goods provision by implementing a large-scale field experiment on Stack Overflow that exogenously varies the recognition individuals receive. It focuses on new and less active users—a large but often overlooked segment of the platform's user base—where engagement strategies may have especially high leverage. The experiment tracks behavior over more than four months, enabling the analysis of both immediate and persistent effects. A structural model is estimated to quantify the roles of social and instrumental motivation, and counterfactual simulations are used to evaluate platform design strategies aimed at sustaining contribution.

The study offers several key contributions. First, it presents one of the largest randomized field experiments to date on a knowledge-sharing platform, focusing on a user segment—new and less active contributors—that has received limited attention in prior work. Second, it provides the first causal evidence distinguishing between social and instrumental motivations for online contribution behavior, using both reduced-form analysis and

structural estimation. Third, it is, to my knowledge, the only large-scale field experiment conducted after the widespread adoption of AI tools that have significantly reduced user contributions. Despite this broader decline in participation, the study shows that anonymous social recognition continues to meaningfully increase both the likelihood and volume of contributions. Finally, it offers actionable insights for platform design, demonstrating that low-cost recognition mechanisms can serve as effective and scalable alternatives to monetary or privilege-based incentives. These findings have broad relevance for digital platforms that rely on user-generated content, including review sites, educational forums, and social media environments.

The remainder of the paper is organized as follows. Section 2 introduces the data and institutional background of Stack Overflow. Section 3 develops the model of knowledge contribution. Section 4 describes the field experiment design, and Section 5 presents the experimental results. Section 6 discusses the identification and estimation strategy. Section 7 reports the structural estimates, and Section 8 presents simulated counterfactual analyses. Section 9 concludes.

2 Data and Setting

I collect data from Stack Overflow, a leading question-and-answer website for professional and enthusiast programmers. As of April 2023, it had over 20 million registered users, 24 million questions, and 35 million answers, with about 69% of questions answered. The site sees 5.9 million daily visits and around 3,900 new questions.[9] According to Stack Overflow's 2023 Developer Survey, 63% of respondents spend more than 30 minutes per day searching for answers or solutions on the platform, with 25% spending over 60 minutes daily.[10]

Stack Overflow and similar programming Q&A platforms have largely supplanted traditional programming books as day-to-day references since the 2000s, and today constitute a

[9]Data retrieved from https://stackexchange.com/sites?view=list#users, accessed April 14, 2023.

[10]Survey information available at https://survey.stackoverflow.co/2023/#section-productivity-impacts-daily-time-spent-searching-for-answers-solutions.

central resource in computer programming. Based on the tags assigned to questions, the ten most frequently discussed topics are JavaScript, Python, Java, C#, PHP, Android, HTML, jQuery, C++, and CSS.[11]

Figure 1a presents a screenshot of Stack Overflow. Similar to other online Q&A communities, users can post questions, and others voluntarily respond with answers. When users log into the platform, as illustrated in Figure 1a, questions are ranked and displayed based on their posting or update time. Specifically, questions that have been most recently posted, answered, or commented on appear at the top.[12] If users find a question or answer helpful, they can upvote it.[13]

The platform designs and maintains a user reputation system. Figure 1b displays a sample question and Figure 1c a sample answer, each showing the contributor's username and reputation score. Users earn reputation points when others upvote their contributions: one upvote yields 10 points.[14] Upvoting others does not grant points to the voter. Conversely, users can downvote content they find unhelpful, in which case both the recipient and the voter lose two reputation points. This penalty structure encourages users to be judicious with downvotes. Additionally, the platform monetizes user-generated content by displaying advertisements at the top and bottom of the page.

The platform grants users additional usage privileges as they surpass specific reputation thresholds. As shown in Figure 2a, users need at least 10 reputation points to lift new-user restrictions and create wiki posts; 15 points to upvote others' answers; and over 50 points to comment universally. The highest threshold, at 25,000 points, grants access to internal and Google site analytics of Stack Overflow. Privilege thresholds are more densely concentrated at lower reputation levels, with 12 thresholds at or below 200 points, and more sparsely

[11] This information is obtained from `https://stackoverflow.com/tags`, retrieved on April 14, 2023.

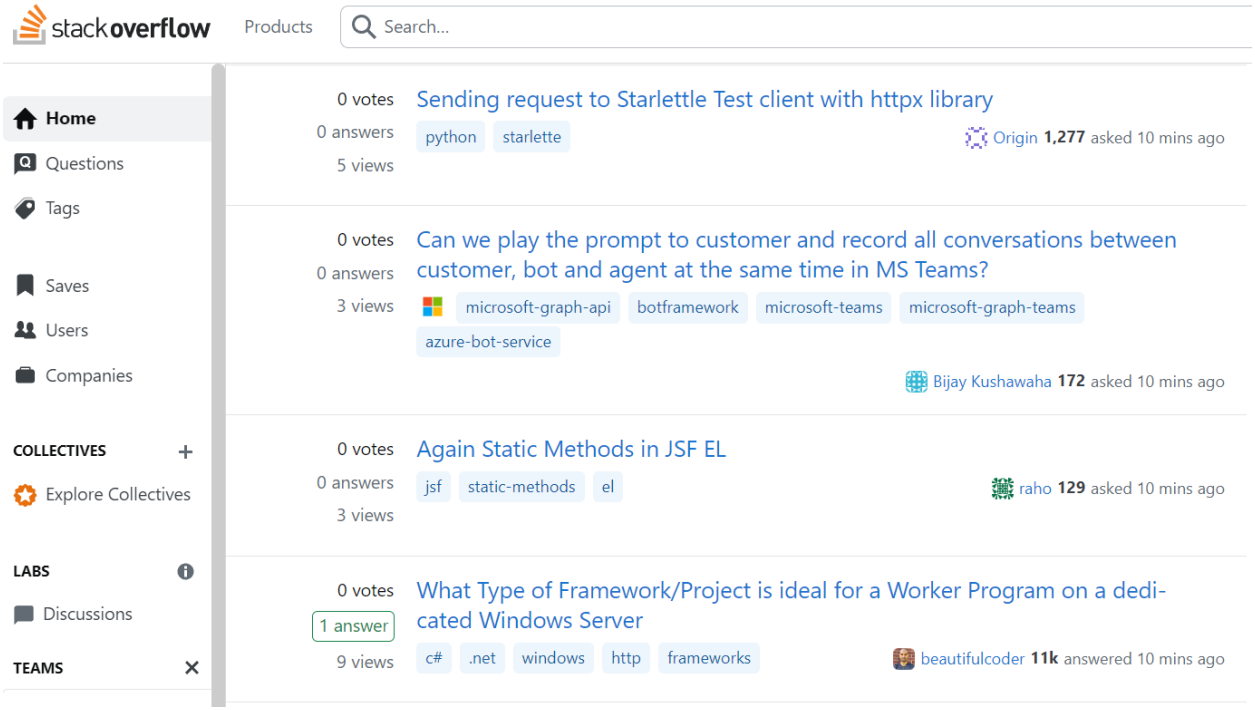
[12] This platform feature supports the experimental design, as the treatment—an additional upvote—is unlikely to meaningfully alter the front-page ranking of questions. This effectively rules out concerns that the intervention might artificially increase a post's visibility and attract additional user interactions.

[13] Additionally, the individual who posted the question can designate one of the responses as "accepted". An accepted answer is not necessarily the best but rather the one that the original asker found most helpful.

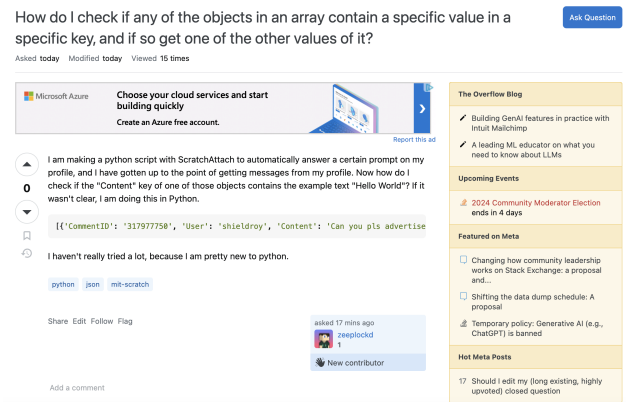
[14] An additional 15 reputation points are awarded if an answer is accepted by the original asker.

Figure 1: Stack Overflow and Sample Content

(a) Stack Overflow Front Page



(b) Stack Overflow - Sample Question



(c) Stack Overflow - Sample Answer

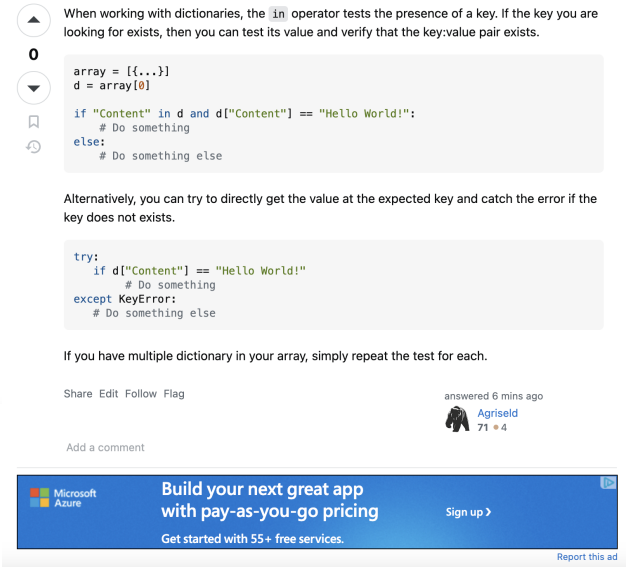


Figure 2: Privileges and Distances

(a) Thresholds for Different Privileges

(i) Below 200

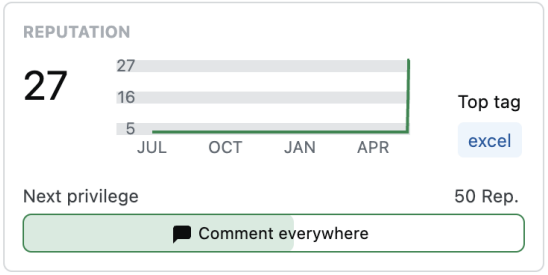


(ii) Above 200

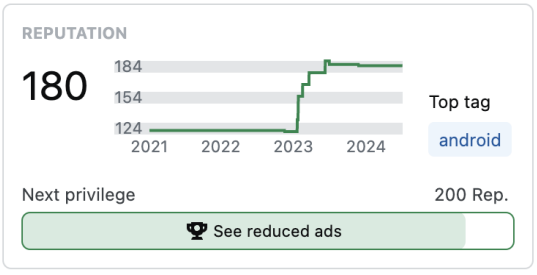


(b) Distance to the Next Threshold

(i) Example 1



(ii) Example 2



distributed at higher levels, with 13 thresholds between 250 and 25,000 points.

As users approach new privilege thresholds, the platform prominently displays their progress toward the next milestone at the top of their own profile pages—a feature intended to incentivize further contributions. Figure 2d illustrates two users at different reputation levels. The green bar at the bottom of each profile highlights the upcoming privilege and the remaining distance to unlock it.

3 A Model of Knowledge Contribution

To guide the experiment design, I specify a model that examines how individuals form expectations over time regarding the number of upvotes they will receive for contributing an additional answer on Stack Overflow, the online platform dedicated to technical programming-related questions.

Individual i values cumulative reputation points they already have on the platform at time t, denoted by R_{it}.

In addition, individual i values the number of privileges gained through crossing privilege thresholds, denoted by $q(R_{it})$. Note that $q(.)$ is a deterministic function, in the sense that given R_{it}, the individual i knows exactly how many privileges they have. More specifically, $q(R_{it})$ is a step function that increases by one as R_{it} reaches a certain threshold. Individual i's utility function follows:

$$U_i(R_{it}) = \gamma_i R_{it} + \varphi_i q(R_{it})$$

Intuitively, γ_i represents the social utility of gaining reputation, where it can be due to social status, altruism, warm glow, etc. And φ_i represents the instrumental utility, where additional reputation will help individuals gain additional privileges specified by the platform. Note that $\gamma_i = 0$ is a special case where individual i only values the instrumental utility of gaining additional privileges.

Let s_{it} denote the number of upvotes the individual i may get if they contribute an additional answer at time t. This will give individual i $10s_{it}$ additional reputation points, which means R_{it} evolves according to:

$$R_{it+1} = R_{it} + 10s_{it}$$

The number of upvotes per answer follows a Poisson distribution with parameter λ_{it}, which is also its expected value:

$$s_{it} \sim Poisson\left(\lambda_{it}\right)$$

$$E[s_{it}] = \lambda_{it}$$

Individual i at time period 0 has a prior belief about the distribution parameter λ_{i0} of the number of upvotes s_{i0} they will get if contributing an answer.

$$\lambda_{i0} \sim \Gamma\left(k_{i0}, \theta_{i0}\right)$$

Given the prior, the expected number of upvotes per answer is:

$$E[s_{i0}] = E[\lambda_{i0}] = k_{i0}\theta_{i0}$$

With each answer contribution, i will update their prior belief through Bayesian updating, with posterior:

$$\lambda_{it} \sim \Gamma\left(k_{i0} + \sum_{t=1}^{n} s_{it}, \frac{\theta_{i0}}{n\theta_{i0} + 1}\right)$$

where $\sum_{t=1}^{n} s_{it}$ is the sum of upvotes from historical answers, and n is the number of past answers up till time t.

Given the posterior, the expected number of upvotes per answer is:

$$E[s_{it}] = E[\lambda_{it}] = \left(k_{i0} + \sum_{t=1}^{n} s_{it}\right) \cdot \frac{\theta_{i0}}{n\theta_{i0} + 1}$$

Individual i has an effort cost c_i to contribute answers. ϵ_{it} follows a Type I Extreme Value distribution to account for random variations in the effort cost.

Let $I_{it} = \{\gamma_i, \varphi_i, c_i, R_{it}, k_{it}, \theta_{it}\}$ denotes the information set available to i at time t.

Consider individual i's decision at time $t + 1$.

Individual i will make the contribution decision based on the expected utility to gain from getting the reputation points:

$$\Delta U_{it} = E\left[U_i(R_{it} + 10s_{it}) \mid I_{it}\right] - U_i(R_{it}) - c_i + \epsilon_{it}$$

The individual i will choose to contribute an answer if the expected utility to gain is larger than 0.

$$D_{it} = \begin{cases} 1 & \Delta U_{it} > 0 \\ 0 & \text{otherwise} \end{cases}$$

In the model above, one additional upvote on one of individual i's past answers will shift individual i's posterior belief to:

$$\lambda_{it} \sim \Gamma\left(k_{i0} + \sum_{t=1}^{n} s_{it} + 1, \frac{\theta_{i0}}{n\theta_{i0} + 1}\right)$$

Given the shifted posterior, the expected number of upvotes per answer is:

$$E[s_{it}] = E[\lambda_{it}] = \left(k_{i0} + \sum_{t=1}^{n} s_{it} + 1\right) \cdot \frac{\theta_{i0}}{n\theta_{i0} + 1}$$

In addition, since one additional upvote will give individual i 10 additional reputation points, it will also change individual i's reputation R_{it} and the number of privileges gained

through reaching privilege thresholds, $q(R_{it})$. The additional reputation points will also change individual i's distance to the next privilege. D_{it} will therefore be changed.

Below are three model predictions about how an additional upvote will change individual i's decision.

Prediction 1. Given posterior belief, γ_i, φ_i and R_{it}, the lower the n, the increase in individual i's probability of contributing another answer is larger with an additional upvote.

Prediction 2. Given posterior belief, γ_i, φ_i and R_{it}, if the individual i expects to reach the next privilege threshold with one more contribution, then the increase in individual i's probability of contributing another answer is larger with an additional upvote.

Prediction 3. Given posterior belief, γ_i, φ_i and R_{it}, if the individual i expects to reach the next privilege threshold with one more contribution, the higher the φ_i, then, the increase in individual i's probability of contributing another answer is larger with an additional upvote.

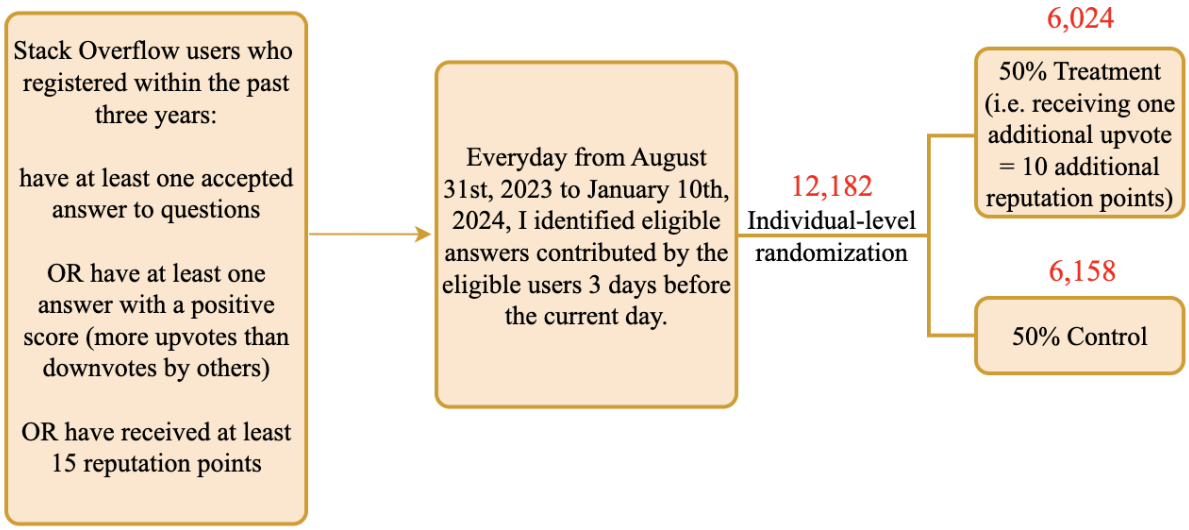
With the above predictions in mind, I designed and implemented a field experiment to measure the causal effect of receiving an additional upvote on individuals' online knowledge contribution decisions.

4 Experiment Design

The experiment design is summarized by Figure 3. I first identified the individuals who were eligible to enter my experimental sample. I focused on individuals who registered on Stack Overflow on or before August 31st, 2020, as this subgroup of individuals is relatively newly registered since Stack Overflow's foundation in 2008. I queried Stack Overflow's API to obtain detailed data of all of the 985,841 answers contributed by those 304,617 individuals on or before August 27th, 2023. With the data, I further restricted eligible individuals to those who have at least minimal experience in making meaningful contributions to the platform. More specifically, they had to satisfy at least one of the following criteria:

- Have contributed at least one accepted answer to questions
- Have contributed at least one answer with a positive score (more upvotes than downvotes by other users)
- Have received at least 15 reputation points

Figure 3: Experiment Design



By applying the minimal inclusion criterion above, I captured the vast majority of users capable of making meaningful knowledge contributions. This approach excluded individuals who had not contributed any helpful answers or whose reputation remained too low, as they were likely still learning how to navigate the platform, formulate effective responses, or identify appropriate questions.

To implement the experiment, I created Stack Overflow accounts and, over several months, actively answered technical programming questions to earn upvoting privileges for each account.

Between August 31, 2023, and January 10, 2024, I collected daily samples of answers posted exactly three days prior by the pre-determined pool of eligible users. The three-day lag ensured sufficient time for natural community evaluation—users could upvote or

downvote anonymously, delete their own answers, or have content removed by moderators. I restricted the sample to plausibly high-quality answers: those not deleted and with net non-negative scores after three days. These organic user responses served as a proxy for answer quality, helping ensure compliance with platform policies and the absence of deception.[15]

Each day, after identifying eligible users who posted high-quality answers, I retrieved their full posting histories and randomized them into treatment and control groups. Half were randomly assigned to receive one manual upvote—from one of my accounts—on one of their eligible answers.

To comply with Stack Overflow's moderation policies, I personally implemented all treatments between 12:00 a.m. and 5:00 a.m. ET, operating every day across my accounts for over four months.

In the case that an individual in the treatment group contributed multiple plausibly high-quality answers within that day, only the last qualified answer of the day would be treated. Essentially, the randomization is at the individual level, i.e. each treated user would only receive one upvote on one of the answers they posted three days before, and individuals who were included in either the control or the treatment group in previous days would not be sampled again in the future. This is to make sure that the treatment intensity is independent of the individuals' baseline activity level.

To ensure enough newly registered individuals and individuals with relatively low answering experience on the platform are included in the experiment, I expanded the eligible individuals to include those who newly satisfied the criteria during the past week, once per week during the four-month experiment period. This expansion allowed me to include individuals who registered before the experiment began and just satisfied the inclusion criteria in the past week, along with individuals who newly registered and satisfied the inclusion criteria in the past week.

The treatment would give 10 additional reputation points to the individual who con-

[15]The experiment received IRB approval prior to implementation.

tributes the answer and would indicate one additional anonymous individual finds one of their past answers helpful[16].

After implementing the treatment, I tracked each sampled individual's detailed subsequent answers, questions, comments, and profile information at a daily level for several months.

Table 1: Balance Table of Variables for the Treatment and Control Groups

Variable	Control Mean	Treatment Mean	t-Statistic	p-Value
Profile View Count	40.875	42.368	-0.186	0.852
Reputation	325.568	282.192	0.997	0.319
Reputation Change Last Week	9.865	10.043	-0.395	0.693
Reputation Change Last Month	17.034	16.262	0.643	0.520
Reputation Change Last Quarter	33.760	31.602	0.790	0.430
Reputation Change Last Year	115.916	102.839	1.010	0.313
Number of Past Answers	19.187	17.088	0.931	0.352
Number of Past Questions	3.008	3.045	-0.263	0.792
Profile Information Filled	0.551	0.542	1.051	0.293
Location Information Filled	0.335	0.334	0.195	0.845
Number of Users	6158	6024		

Notes: The t-stats reported follow the Welch's two-sample t-test.

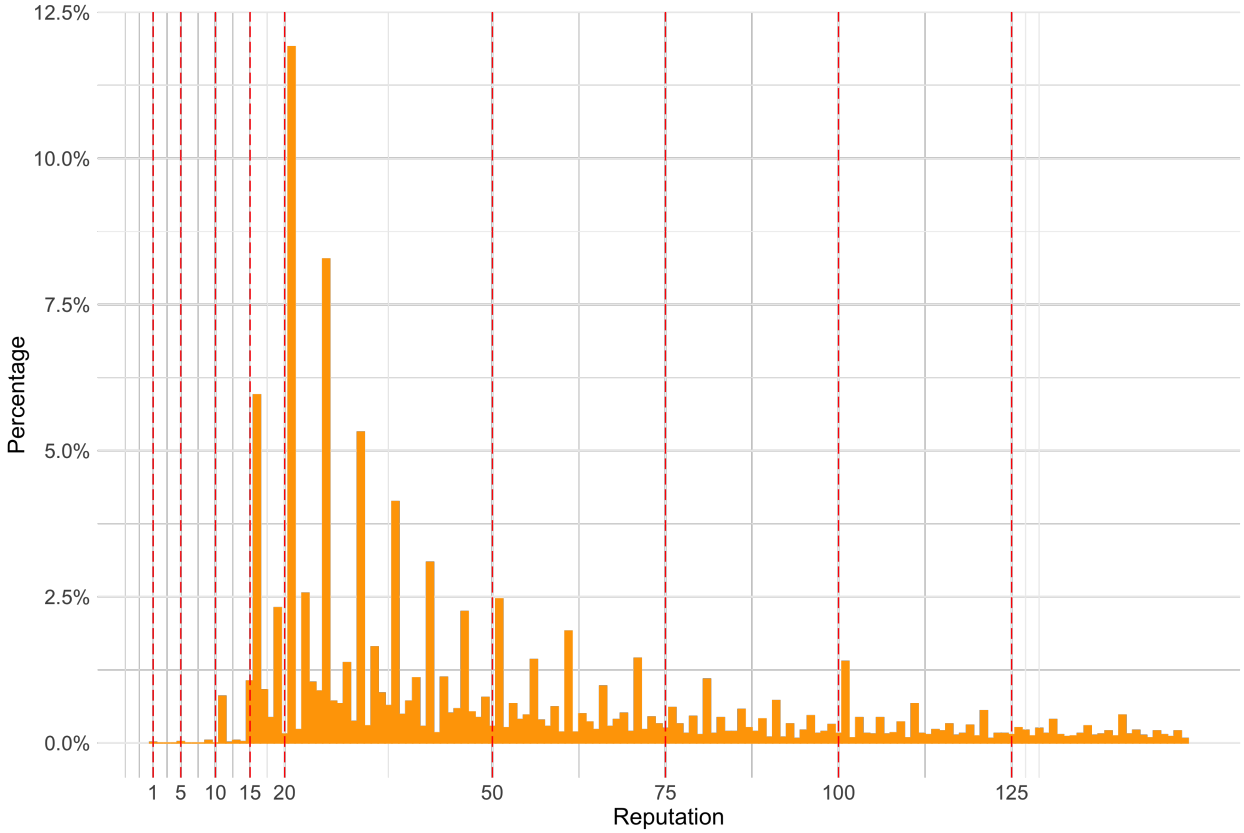
In total, I sampled 12,182 individuals during the experiment period of over four months. Among the sampled individuals, 6,024 are in the treatment group, and 6,158 are in the control group. I implemented the experiment for 132 consecutive days, and on average, I sampled 92 individuals per day. Table 1 shows the control group and treatment group are balanced on all key variables: profile view count by other users, reputation points, recent reputation changes, number of past answers, number of past questions, and availability of profile and location information.

Figure 4b shows the distribution of the total number of past answers by each individual at the time of being sampled. The distribution is right-skewed, with a majority of individuals

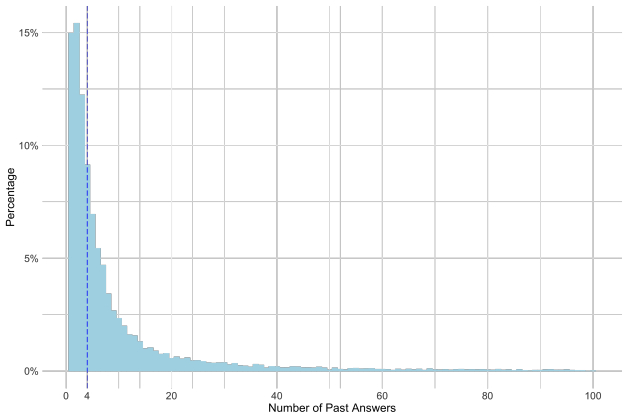
[16]The experiment carefully followed Stack Overflow's terms of usage and received IRB approval from the University of California, Berkeley before implementation. The IRB approval number is 2023-04-16245.

Figure 4: Distributions of Reputation Points, Number of Past Answers, and Average Score for the Experimental Sample

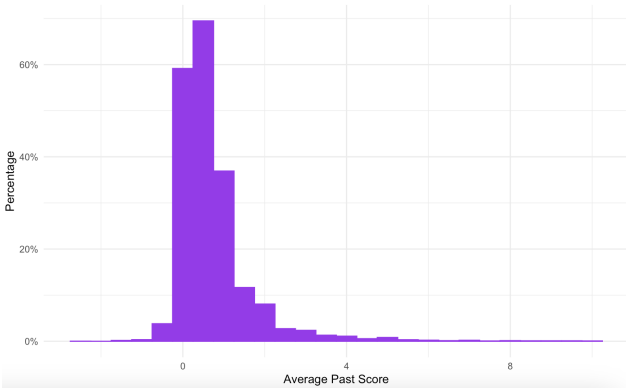
(a) The Percentage Distribution of Reputation Points at the Time of Being Sampled



(b) The Percentage Distribution of Number of Past Answers at the Time of Being Sampled



(c) The Percentage Distribution of Average Past Score at the Time of Being Sampled



having a low level of past answering experience, and a few individuals having a large number of past answers. The median individual in the sample contributed 4 answers in the past, and the average of the number of past answers is 18.15. Figure 4a shows the distribution of reputation for each individual at the time of being sampled. The red dashed vertical lines correspond to the privilege thresholds in Figure 2a. There is a mass of individuals with reputation points right above the thresholds. For example, there is a high proportion of individuals with reputation points right above 15, 20, and 50. However, there is also a substantial number of individuals who are less than 10 reputation points below the next privilege threshold, which means that one upvote from the experiment will make them cross the next privilege threshold and face another threshold - 3,044 individuals fall into this category. Overall, the distribution of reputation is right-skewed, with a median of 51 and a mean of 304.1. The additional 10 reputation points by the experiment treatment will move up the median individual's reputation by around 19.61%.

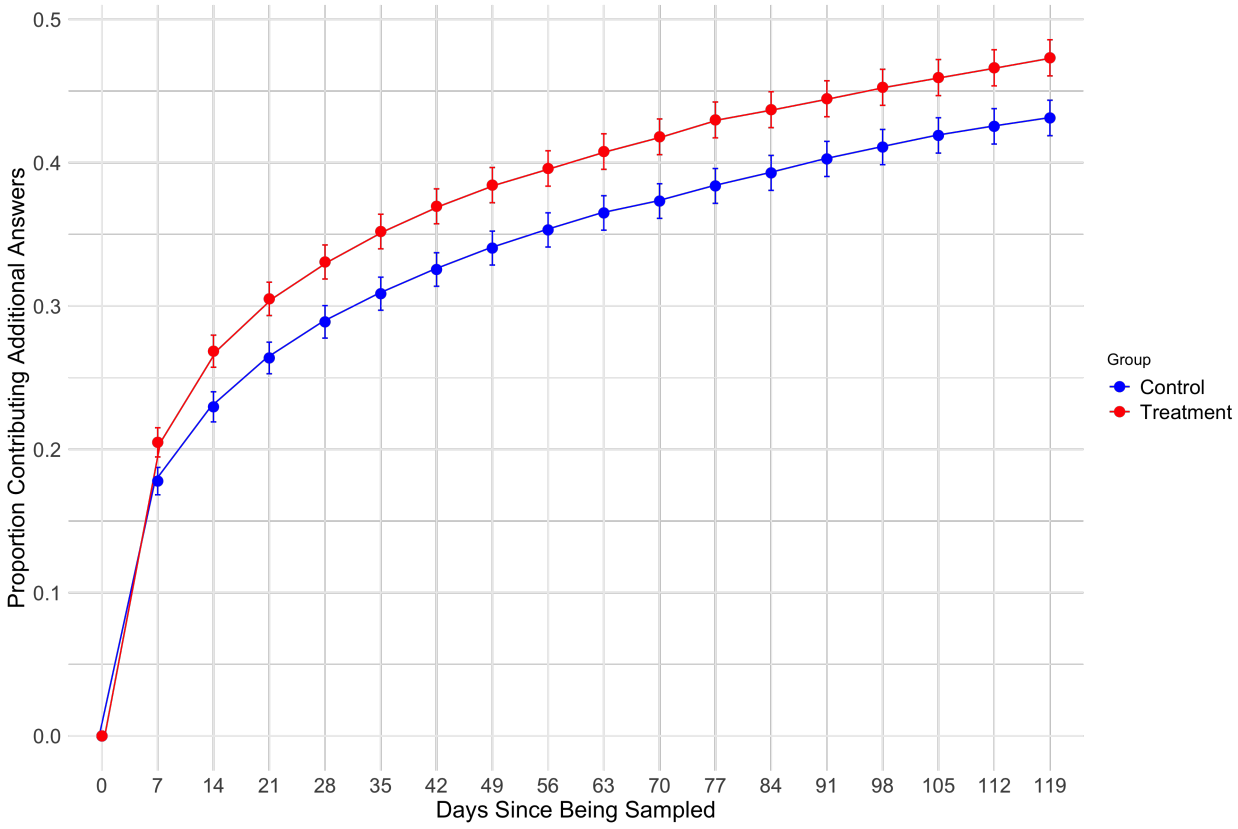
Figure 4c shows the distribution of the past average answer score (the number of upvotes minus the number of downvotes) for each individual at the time of being sampled. The distribution is also right-skewed, with a median of 0.50 and a mean of 0.78.

5 Results

5.1 Extensive Margin of Knowledge Contribution

Receiving one additional anonymous upvote significantly increases the likelihood of contributing another answer. Figure 5 presents the proportion of individuals in the treatment and control groups who contributed at least one additional answer after being sampled, measured over various time horizons. Contribution rates increase over time in both groups. In the control group, 17.8% of individuals contributed an additional answer within 7 days, rising to 26.1% within 21 days and 35.3% within 56 days. In the treatment group, the corresponding proportions are 20.5% within 7 days, 30.5% within 21 days, and 39.6% within 56

Figure 5: The Proportion of Individuals Contributing Additional Answers in Treatment and Control Groups Over Time



days.

The differences in contribution rates between the treatment and control groups are statistically significant. The gap increases from 2.7 percentage points within the first 7 days to 4.1 percentage points within 21 days, and stabilizes at approximately 4.3 percentage points thereafter—persisting for at least four months post-sampling. These patterns suggest that the treatment alters individuals' decisions to contribute, and the effect is unlikely to be driven by intertemporal substitution (i.e., individuals accelerating contributions they would have made later in the absence of the treatment).

Table 2: Contributing Additional Answers Over Time (Extensive Margin)

	Dependent variable: I(Additional Answers)			
	Within 7 Days	Within 14 Days	Within 21 Days	Within 120 Days
Treatment	0.027***	0.040***	0.042***	0.042***
	(0.007)	(0.008)	(0.008)	(0.009)
Observations	12,182	12,182	12,182	12,182
R^2	0.061	0.065	0.061	0.038
Adjusted R^2	0.047	0.051	0.048	0.024
Dep. Var. Mean	0.178	0.229	0.261	0.432
Percentage of Baseline	15.169	17.467	16.092	9.722
Sample Day FE	Yes	Yes	Yes	Yes
Reg. Month FE	Yes	Yes	Yes	Yes

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Note: The outcome variables are whether the individual contributed at least one additional answer within different numbers of days of being sampled. The standard error is clustered at the sample day level.

Table 2 presents linear probability regressions where the dependent variable indicates whether an individual contributed at least one additional answer within 7, 14, 21, or 120 days after being sampled. The regressions include sample-day fixed effects and registration-month fixed effects, with standard errors clustered at the sample-day level. Each of the 12,182 observations corresponds to an individual from either the treatment or control group.

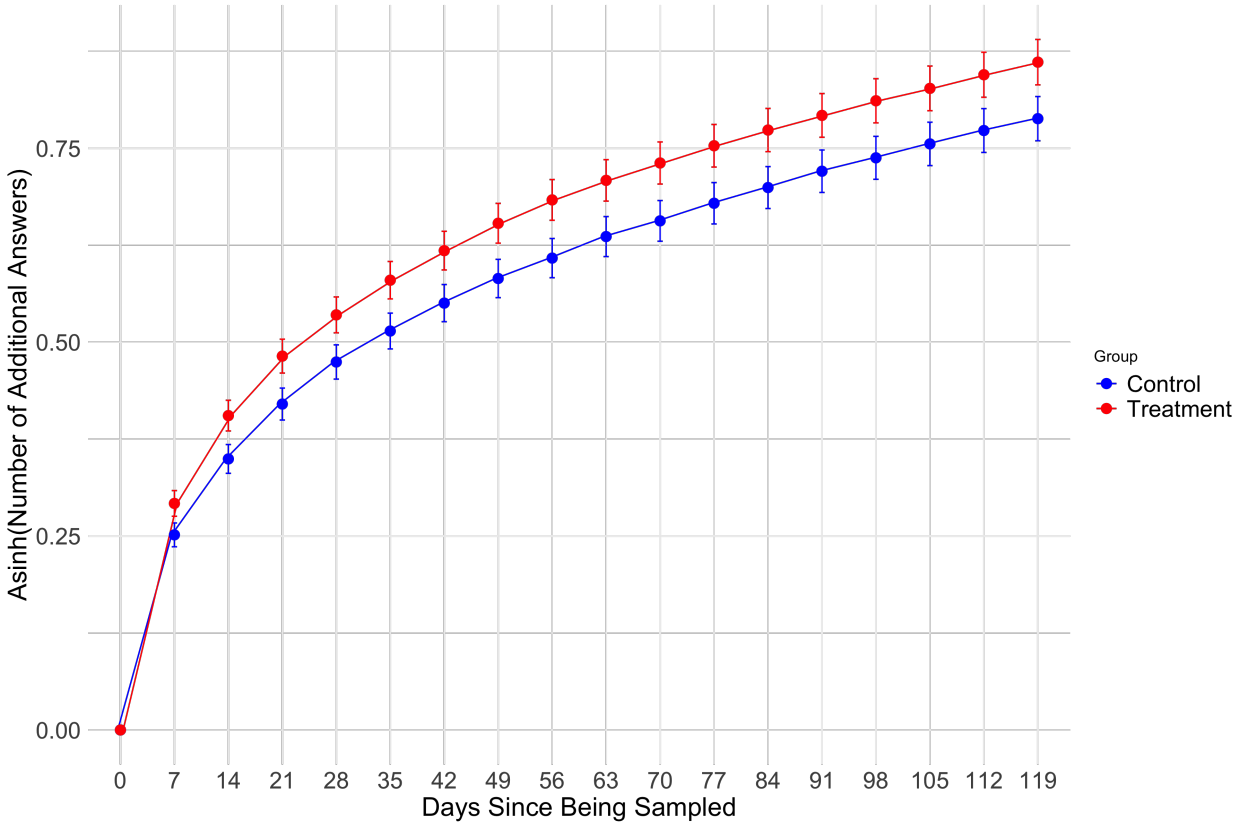
Receiving an additional upvote increases the probability of contributing within 7 days

by 2.7 percentage points, within 14 days by 4.0 percentage points, within 21 days by 4.2 percentage points, and within 120 days by 4.2 percentage points. These treatment effects represent 15.2%, 17.5%, 16.1%, and 12.2%, respectively, of the control group's contribution rates over the corresponding time windows.

Consistent with the pattern in Figure 5, the treatment effect gradually increases over time and stabilizes around 21 days after sampling, possibly because treated individuals need time to notice the upvote and identify relevant questions to answer. For this reason, the subsequent analysis uses a 21-day window to measure outcomes.

5.2 Intensive Margin of Knowledge Contribution

Figure 6: The Number of Additional Answers Per Individual in Treatment and Control Groups Over Time



Receiving one additional anonymous upvote significantly increases the number of answers

Table 3: The Number of Additional Answers Contributed within 21 Days (Intensive Margin)

	Dependent variable: Number of Additional Answers			
	Asinh(Number)	Poisson(Number)	Number	Number(Winsorized at 5)
Treatment	0.063***	0.141***	0.162**	0.093***
	(0.016)	(0.017)	(0.082)	(0.025)
Observations	12,182	12,182	12,182	12,182
R^2	0.070		0.094	0.087
Adjusted R^2	0.057		0.081	0.074
AIC		53,469.810		
Dep. Var. Mean	0.420		1.033	0.639
Sample Day FE	*Yes*	*Yes*	*Yes*	*Yes*
Reg. Month FE	*Yes*	*Yes*	*Yes*	*Yes*

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Note: The outcome variables are the number of answers, the number of answers winsorized at 5, and the arcsinh of the number of answers contributed by the individual within 21 days of being sampled. The standard error is clustered at the sample day level for the first, third, and fourth columns.

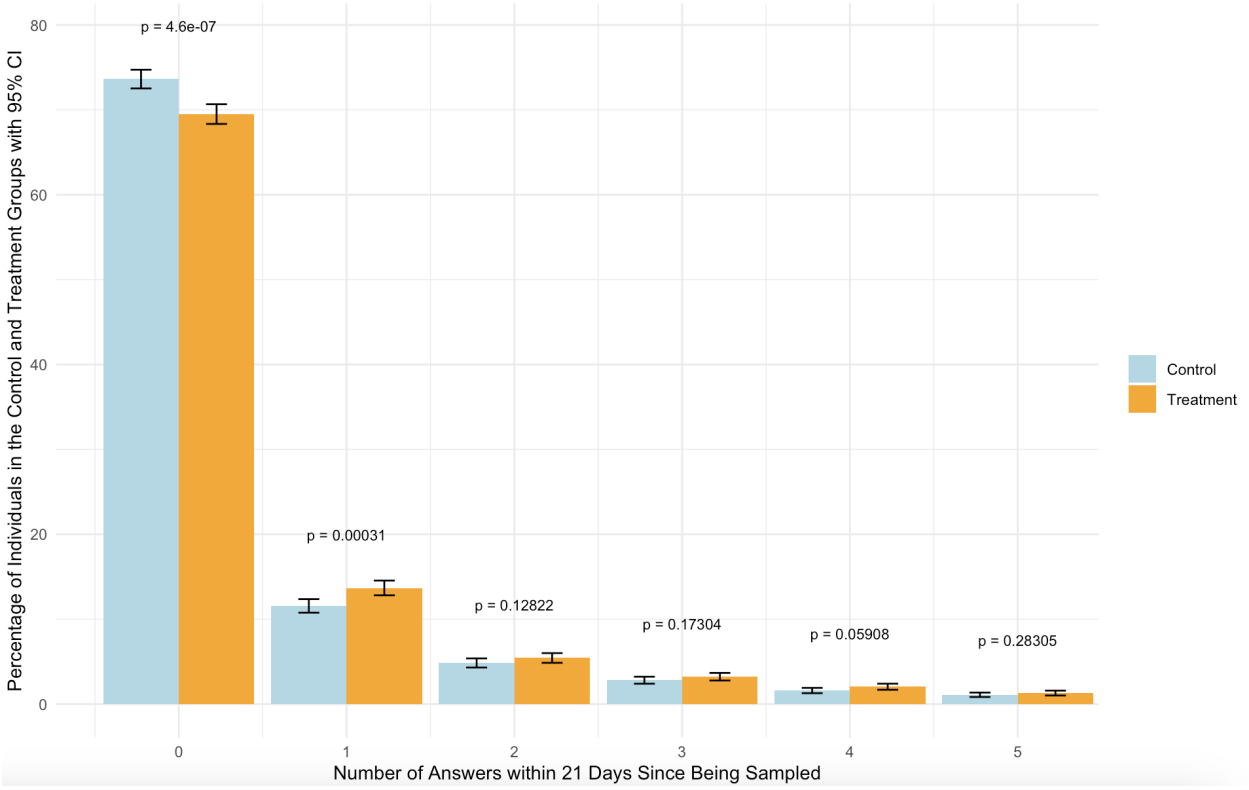
contributed. Figure 6 shows the average *asinh* of the number of additional answers per individual in the treatment and control groups over time. The gap between the two groups gradually widens during the first 21 days after sampling, and then remains stable for at least four months. Table 3 presents several regression specifications using the number of answers contributed within 21 days as the outcome variable. Overall, receiving one additional upvote increases the number of answers contributed per individual within 21 days by 0.171, corresponding to approximately 16.554% of the control group's average.

To further examine how the treatment affects the distribution of answer counts within the 21-day window, Figure 7 displays the percentage distribution of individuals by number of answers contributed, for both treatment and control groups. The x-axis is truncated at the 95th percentile (5 answers) to focus on the most common values. In the control group, 73.6% of individuals contribute no additional answers, 11.6% contribute one answer, and 14.8% contribute more than one. In the treatment group, the corresponding proportions are

69.5%, 13.7%, and 16.8%, respectively.

Comparing these distributions, the treatment group has a significantly smaller share of individuals who contribute nothing, and a significantly larger share who contribute exactly one answer. For individuals contributing more than one answer, the distributions are similar, though the treatment group has a slightly higher percentage. These results suggest that the treatment effect on the number of answers is primarily driven by shifting individuals from contributing nothing to contributing one answer.

Figure 7: The Percentage Distribution of Number of Answers within 21 Days Since Being Sampled



5.3 Secondary Outcomes: Login, Questions, and Voting

Receiving one additional anonymous upvote has no significant effect on whether an individual logs in at least once within 21 days of being sampled (Table 4a). This result is expected: there is no email or other external notification when an upvote is received, so

the treatment effect on the primary outcome—answer contribution—is unlikely to operate through a pure notification channel that prompts individuals to log in. However, the treatment does significantly increase the number of days individuals log in during the 21-day window. While the effect size is modest, the number of login days increases by 0.255, representing approximately a 3.9% increase relative to the baseline.

Receiving one additional anonymous upvote does not significantly affect the likelihood of posting additional questions or the number of questions posted within 21 days (Table 4b). If individuals were motivated to post content primarily to show off or signal existing reputation—such as using reputation to attract more attention—the treatment would have increased question posting as well. The null effect suggests this is not the case. Moreover, the absence of an effect on question posting implies that individuals do not update their beliefs about reputation to gain from questions based on reputation gained from their answers.

In contrast, the treatment significantly increases the probability of posting comments by 3.3 percentage points, and the number of comments within 21 days by approximately 5.8% relative to the control group (Table 4b). One plausible explanation is that commenting often accompanies answering: individuals may comment to request clarification on questions before answering, or to respond to follow-up comments on their own answers.

Receiving one additional anonymous upvote also significantly affects individuals' voting behavior within 21 days (Table 5). For upvoting, the treatment increases the probability of making at least one upvote by 2.7 percentage points, representing 7.30% of the baseline. The total number of upvotes also rises by approximately 5% relative to the control group.

For downvoting, the treatment significantly increases the probability of making at least one downvote by 1.1 percentage points. While the absolute effect is smaller than for upvoting, this is largely due to the low baseline rate—only 5.5% of individuals in the control group downvote at least once within 21 days. Thus, the treatment effect represents roughly a 20% increase relative to the baseline. The number of downvotes also increases, by approximately 2.1%.

Table 4: Summary of Login Activities and Additional Contributions

(a) Panel A: Login Activities

	Dependent variable:	
	I(Login)	Asinh(Num. Login Days)
Treatment	0.001	0.039**
	(0.004)	(0.018)
Observations	12,176	12,176
R^2	0.023	0.037
Adjusted R^2	0.009	0.023
Dep. Var. Mean	0.929	2.559

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

(b) Panel B: Contributing Additional Questions and Comments

	Dependent variable:			
	I(Questions)	Asinh(Num. Questions)	I(Comments)	Asinh(Num. Comments)
Treatment	0.008	0.007	0.033***	0.058***
	(0.005)	(0.006)	(0.008)	(0.020)
Observations	12,182	12,182	12,182	12,182
R^2	0.018	0.018	0.004	0.004
Adjusted R^2	0.049	0.055	0.028	0.057
Dep. Var. Mean	0.090	0.100	0.254	0.485

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Note: In panel A, the outcome variables are whether the individual login at least once and the arcsinh of the number of days the individual login within 21 days of being sampled. The standard error is clustered at the sample day level. In panel B, the outcome variables are whether the individual contributed at least one additional question, arcsihn of the number of additional questions contributed, whether the individual contributed at least one additional comment, and arcsinh of the number of additional comments contributed within 21 days of being sampled.

Table 5: Upvoting and Downvoting

	Dependent variable:			
	I(Upvotes)	Asinh(Num. Upvotes)	I(Downvotes)	Asinh(Num. Downvotes)
Treatment	0.027***	0.050**	0.011***	0.021**
	(0.008)	(0.020)	(0.004)	(0.009)
Observations	12,050	12,050	12,050	12,050
R^2	0.021	0.022	0.035	0.037
Adjusted R^2	0.007	0.008	0.021	0.023
Dep. Var. Mean	0.370	0.662	0.055	0.086
Sample Day FE	Yes	Yes	Yes	Yes
Reg. Month FE	Yes	Yes	Yes	Yes

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Note: The outcome variables are whether the individual did at least one upvote, arcsihn of the number of additional upvotes, whether the individual did at least one downvote, and arcsihn of the number of additional downvotes within 21 days of being sampled.

5.4 Heterogeneous Treatment Effect, Answer Experience

The effect of receiving an additional anonymous upvote varies with individuals' prior answering experience. As shown in Table 6a, I divide the sample into quartiles based on the number of past answers contributed at the time of treatment. For individuals with 1–2 prior answers, the estimated treatment effect is statistically significant at 2.5 percentage points (16.28% of the control mean). Among those with 3–4 prior answers, the effect is 4.61 percentage points (19.79%). For users with 5–10 prior answers, the effect is 6.30 percentage points (23.16%). For the most experienced users—those with more than 10 prior contributions—the effect declines to 3.77 percentage points (9.15%), though it remains statistically significant. This pattern is consistent with the belief-updating mechanism: as individuals accumulate experience, their beliefs about the reception of their contributions become more stable, reducing the informational weight of marginal social signals such as an additional anonymous upvote.

Table 6: Heterogeneity of Treatment Effect on Answer Contributions

(a) Panel A: Contributions by Previous Experience

	Dependent variable: I(Additional Answers)			
	1-2 Past Answers	3-4 Past Answers	5-10 Past Answers	Above 10 Past Answers
Treatment	0.025**	0.044***	0.063***	0.038**
	(0.012)	(0.015)	(0.015)	(0.018)
Observations	3,596	2,529	3,017	3,040
R^2	0.064	0.125	0.165	0.163
Adjusted R^2	0.017	0.061	0.115	0.113
Dep. Var. Mean	0.145	0.233	0.272	0.412
Percentage of Baseline	16.276	19.785	23.162	9.150
Sample Day FE	*Yes*	*Yes*	*Yes*	*Yes*
Reg. Month FE	*Yes*	*Yes*	*Yes*	*Yes*

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

(b) Panel B: Contributions by Reputation

	1-26 Reputation	27-51 Reputation	52-142 Reputation Reputation	Above 142 Reputation
Treatment	0.0332*	0.0569***	0.0684***	0.0327
	(0.0149)	(0.0166)	(0.0166)	(0.0172)
R^2	0.0724	0.1093	0.1122	0.1573
Adj. R^2	0.0215	0.0510	0.0552	0.1059
Num. obs.	3312	2799	2850	2993
Dep. Var. Mean	0.219	0.229	0.239	0.356
Percentage of Baseline	15.160	24.847	28.619	9.185
Sample Day FE	*Yes*	*Yes*	*Yes*	*Yes*
Reg. Month FE	*Yes*	*Yes*	*Yes*	*Yes*

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Note: In panel A, the outcome variables are whether the individual contributed at least one additional answer within 21 days of being sampled, split by the number of past answers the individual has contributed before being sampled. In panel B, the outcome variables are whether the individual contributed at least one additional answer within 21 days of being sampled, split by the number of reputation points the individual has before being sampled.

5.5 Heterogeneous Treatment Effect, Reputation

Receiving one additional anonymous upvote has a heterogeneous effect depending on individuals' reputation levels. As shown in Table 6b, I divide the sample into quartiles based on users' reputation points at the time of treatment. For individuals in the lowest quartile (1–26 reputation points), the treatment effect is statistically significant at 3.32 percentage points (15.16% of the control mean). Those in the second quartile (27–51 points) exhibit an effect of 5.69 percentage points (24.85%). In the third quartile (52–142 points), the estimated effect is 6.84 percentage points (28.62%). Among individuals in the highest quartile—those with more than 142 reputation points—the effect declines to 3.27 percentage points (9.19%) and is not statistically significant.

5.6 Heterogeneous Treatment Effect, Distance to Privilege Thresholds

Receiving one additional anonymous upvote has a heterogeneous effect depending on an individual's distance to the next privilege threshold. Table 7 presents the results of an analysis in which I group individuals based on their distance to the next threshold at the time of sampling, examining how the experimental upvotes change their subsequent motivation.

The first group includes individuals who are 1–10 reputation points below the next privilege threshold; receiving an upvote would allow them to reach the next threshold and shift their focus to a more distant one. The second group consists of those 11–20 points below the threshold; an upvote would bring them within 10 points of it. The third group includes individuals more than 20 points below the threshold, who would still remain over 10 points away even after receiving an upvote.

Experimental upvotes can influence an individual's future contribution decisions at the extensive margin through two primary channels: (1) by shifting beliefs related to social motivation, and (2) by providing additional reputation points, thereby altering the individual's

Figure 8: The Change of Social and Instrumental Motivations with the Treatment

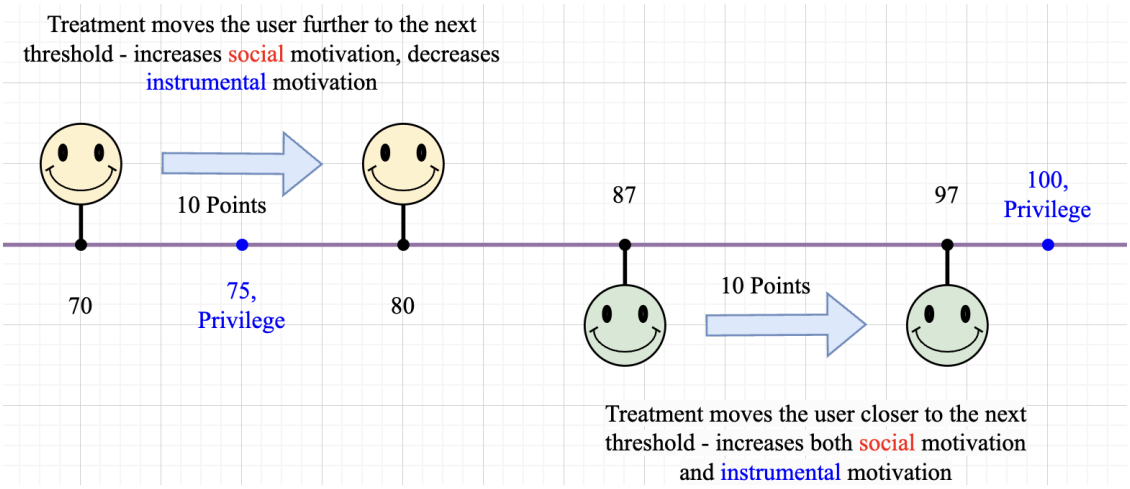


Table 7: Heterogeneity of Treatment Effect on Answer Contributions by Distance to the Next Privilege Threshold Before Being Treated

	Dependent variable: I(Additional Answers)		
	1-10 Distance	11-20 Distance	More than 20 Distance
Treatment	0.040**	0.055***	0.044***
	(0.020)	(0.017)	(0.011)
Observations	1,969	2,796	6,334
R^2	0.152	0.107	0.086
Adjusted R^2	0.070	0.049	0.060
Dep. Var. Mean	0.241	0.257	0.278
Percentage of Baseline	16.598	21.401	15.827
Sample Day FE	Yes	Yes	Yes
Reg. Month FE	Yes	Yes	Yes

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Note: The outcome variables are whether the individual contributed at least one additional answer within 21 days of being sampled, split by distance to the next threshold, before being treated/or not. Since the privilege thresholds are very concentrated below 20, the above results are for individuals with more than 20 reputation.

distance to the next privilege threshold—the instrumental component of motivation. To assess the importance of the instrumental channel, I estimate the treatment effect at the extensive margin separately for three groups. The first group consists of individuals who, due to the experimental upvote, cross the nearest reputation threshold and are thus further from the next threshold than they were before. If instrumental motivation plays a key role, this could reduce future contributions. The second and third groups include individuals who move closer to the next threshold as a result of the upvote, which should increase contributions if instrumental incentives are driving behavior.

Figure 8 presents a graphical illustration. The yellow user, with a reputation of 70, reaches a reputation of 80 after treatment—just surpassing the 75-point threshold for a new privilege. However, with the next threshold now farther away (at 100), their perceived likelihood of reaching it with the next contribution may decline, reducing instrumental motivation. In contrast, the green user starts at 87 and reaches 97—still below the threshold but now closer to it—so the treatment may increase instrumental motivation by raising the expected probability of obtaining the next privilege with an additional contribution. Thus, if instrumental motivation is the sole motivation, treatment may reduce contributions from users just below a threshold (like the yellow user) while increasing them for users farther away but approaching one (like the green user).

I find that the treatment has a positive and statistically significant effect on the extensive margin across all three groups. Furthermore, statistical tests fail to reject the null hypothesis of equal effects across groups, although the estimated effect is marginally larger for the second and third groups—those moved closer to the next privilege threshold. This pattern suggests that instrumental motivations to reach the next threshold cannot be the primary incentive to contribute. I revisit this point in the structural estimation analysis later in the paper.

5.7 Overall Future Answer Quality and Efforts

Conditional on contributing at least one answer within 21 days of being sampled, I compute each individual's average answer quality and effort. For answer quality, I consider whether an answer was accepted by the asker and the answer's score (total upvotes minus total downvotes) on the platform. For answer effort, I measure the number of words and the number of code examples included in the answer. As shown in Table 8, there are no significant differences between the treatment and control groups. This suggests that, conditional on contributing, the average quality and effort of answers are similar across the treatment and control groups.

Table 8: Quality and Efforts of Additional Answers (Conditional on Contributing)

	Accepted	Scores	Body Words	Codes
Treatment	−0.003	0.012	−4.658	0.030
	(0.009)	(0.019)	(4.825)	(0.115)
R^2	0.063	0.046	0.058	0.057
Adj. R^2	0.012	−0.006	0.007	0.006
Num. obs.	3327	3327	3327	3327
Dep. Var. Mean	0.118	0.150	164.063	2.683
Sample Day FE	Yes	Yes	Yes	Yes
Reg. Month FE	Yes	Yes	Yes	Yes

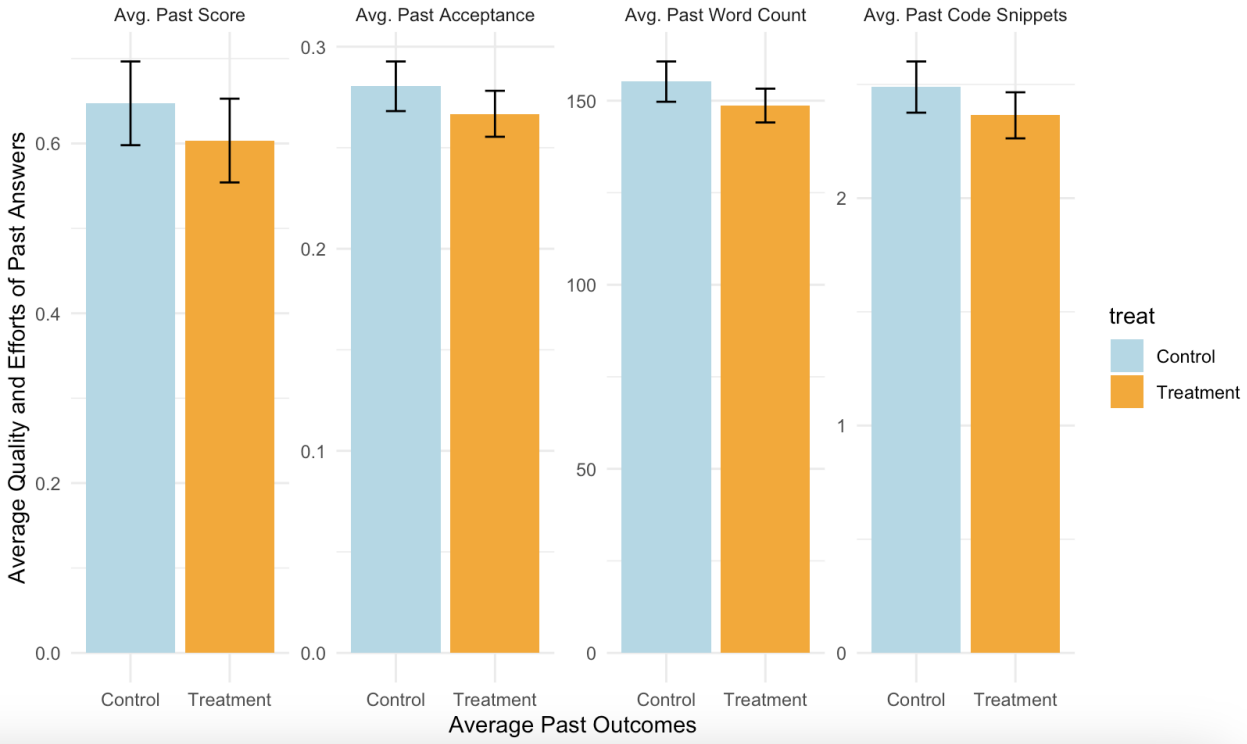
$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Note: The outcome variables are measured within 21 days of being sampled. Accepted is defined as the proportion of accepted answers, scores is defined as the average scores of answers, body words is defined as the average number of words in answers, and codes is defined as the average number of code pieces included in answers, for each individual who contribute at least one answer within 21 days of being sampled.

However, because this analysis is conditional on contributing additional answers, the observed effect may stem from two sources. First, the treatment may influence *who* chooses to contribute. Second, it may affect *how* individuals contribute—that is, whether they provide higher- or lower-quality answers or exert more or less effort compared to their own prior contributions. I examine these two channels in subsection 5.8 and subsection 5.9.

5.8 Who is Motivated to Contribute by the Treatment?

Figure 9: Individuals in the Control and Treatment Groups, Conditional on Contributing within 21 Days of Being Sampled



To decompose the treatment effect on future answer quality and effort, I track and analyze comprehensive data on all 221,093 answers posted by individuals prior to entering the sample. I compute each individual's average past answer quality and effort, conditional on contributing at least one answer within 21 days of being sampled. Figure 9 reports the average past answer score, the average acceptance rate, the average word count, and the average number of code examples per answer, aggregated by treatment and control groups. On average, individuals in the treatment group exhibit lower past answer quality and effort than those in the control group. While the differences are not statistically significant, this provides suggestive evidence that the treatment—an additional upvote increasing the answer score by one—encourages marginal contributors with lower historical answer quality to participate, consistent with the model.

5.9 Does the Treatment Lead to Lower Answer Quality and Efforts?

To further decompose the treatment effect on future answer quality and effort, I examine whether the treatment leads to within-individual changes in contribution behavior. Specifically, I compute the change in average answer quality and effort for each individual, defined as the difference between their post-treatment averages (within 21 days of being sampled) and their pre-treatment averages. This analysis is conditional on contributing at least one answer within the 21-day window. As shown in Table 9, conditional on contributing, individuals in the treatment group increased their answer quality and effort relative to their own past contributions more than those in the control group. However, the differences are not statistically significant.

Table 9: Change of Quality and Efforts Compare to Average Past Answers (Conditional on Contributing)

	ΔAccepted	ΔScores	ΔBody Words	ΔCodes
Treatment	0.014	0.063	1.390	0.115
	(0.011)	(0.039)	(4.698)	(0.104)
R^2	0.061	0.093	0.054	0.047
Adj. R^2	0.008	0.042	0.001	−0.006
Num. obs.	3327	3327	3304	3304
Dep. Var. Mean	−0.164	−0.493	8.593	0.187
Sample Day FE	*Yes*	*Yes*	*Yes*	*Yes*
Reg. Month FE	*Yes*	*Yes*	*Yes*	*Yes*

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Note: The outcome variables are measured within 21 days of being sampled. Accepted is defined as the proportion of accepted answers, scores is defined as the average scores of answers, body words is defined as the average number of words in answers, and codes is defined as the average number of code pieces included in answers, for each individual who contribute at least one answer within 21 days of being sampled.

6 Model Identification and Estimation

To quantify the relative importance of social motivation of gaining reputation (due to social status, altruism, warm glow, etc.) and instrumental motivation given by the platform (gaining more privileges with more reputation) in contribution decisions, I structurally estimate the model of individual decisions with heterogeneous utility parameters to match with moments from the experimental data.

The separate identification of utility parameters, γ_i and φ_i, is derived from comparing the contributing proportion of individuals within 10 reputation points of the next privilege threshold in the treatment and control groups. While the expected additional reputation points are independent of the distance to the next threshold, the expected additional privilege is not. The treatment group, which receives the experimental upvote, reaches the next privilege threshold and then faces a new, distant threshold. Thus, their short-term contribution decision is driven by the belief in gaining additional reputation rather than reaching the next threshold. In contrast, the control group, with less than 10 reputation points to the next threshold, is motivated by both gaining additional reputation and reaching the next privilege threshold.

The identification of the standard deviation of γ_i and φ_i arises from heterogeneous treatment effects across different subgroups of individuals. According to the individual learning model, holding γ_i and φ_i constant, if an individual does not expect to reach the next privilege threshold with their next contribution, their response to the treatment should decrease as the number of past contributions increases. The discrepancies between model predictions assuming homogeneous parameters and actual treatment effects indicate heterogeneous utility parameters, helping to identify the standard deviation of γ_i. Similarly, heterogeneous treatment effects among individuals within 10 reputation points of the next privilege threshold before the treatment help separately identify the standard deviation of φ_i.

I simplify the effort cost, c_i, to be heterogeneous across individuals but constant with

respect to the length of answers or the number of code examples included. This assumption is based on the finding that the estimated treatment effect on answer efforts is not statistically significant. It suggests that the primary effort cost of contributing an answer is the cognitive cost of understanding the question and devising a solution. Once an individual decides to contribute, the additional effort of typing more words or including code examples is relatively minor compared to the cognitive cost. Furthermore, I assume individuals have a rational prior when posting their first answer by calibrating k and θ based on the complete historical scores of all past answers by sampled individuals.

For estimation, I categorize individuals based on their reputation levels, proximity to the next privilege threshold (within 10 reputation points or not), and whether they are in the control or treatment group. I then compute the proportion of individuals in each subgroup who contributed within 21 days of sampling. Table 10 provides a complete list of the empirical moments used for the simulated method of moments, along with the corresponding number of observations for each moment. For each individual, I simulate 1,500 heterogeneous utility parameters, γ_i and φ_i, drawing from log-normal distributions. I predict individuals' decisions to contribute additional answers based on their history of answer scores and reputation at the time of sampling. Then, I predict their decisions to contribute after being treated with one additional upvote. I calculate the proportion of individuals contributing with and without the treatment. Finally, I compare the model-predicted proportions of contributors with and without treatment to the actual contributing proportions estimated from the treatment and control groups.

7 Structural Estimates

Table 11 shows structural estimates. γ_i is assumed to follow a log-normal distribution, with an estimated log-mean of -6.702 and an estimated log-standard of deviation 7.111. φ_i is assumed to follow a log-normal distribution, with an estimated log-mean of -19.640 and an estimated log-standard deviation of 14.895. The effort cost parameter c is estimated to

Table 10: Empirical Moments for the Simulated Method of Moments

	Proportion	Number Obs.
Reputation Below 19 - Control Group	0.231	532
Reputation Below 19 - Treatment Group	0.265	543
Reputation Between 20 and 39 - Control Group	0.229	2039
Reputation Between 20 and 39 - Treatment Group	0.266	1973
Reputation Between 40 and 49 - Control Group	0.267	450
Reputation Between 40 and 49 - Treatment Group	0.265	452
Reputation Between 50 and 64 - Control Group	0.240	504
Reputation Between 50 and 64 - Treatment Group	0.334	458
Reputation Between 65 and 74 - Control Group	0.219	215
Reputation Between 65 and 74 - Treatment Group	0.332	250
Reputation Between 75 and 89 - Control Group	0.221	253
Reputation Between 75 and 89 - Treatment Group	0.343	251
Reputation Between 90 and 99 - Control Group	0.282	110
Reputation Between 90 and 99 - Treatment Group	0.307	137
Reputation Between 100 and 114 - Control Group	0.270	222
Reputation Between 100 and 114 - Treatment Group	0.262	221
Reputation Between 115 and 124 - Control Group	0.274	95
Reputation Between 115 and 124 - Treatment Group	0.176	108
Reputation Between 125 and 189 - Control Group	0.231	459
Reputation Between 125 and 189 - Treatment Group	0.298	473
Reputation Between 190 and 199 - Control Group	0.211	38
Reputation Between 190 and 199 - Treatment Group	0.224	49
Reputation Between 200 and 239 - Control Group	0.257	74
Reputation Between 200 and 239 - Treatment Group	0.324	74
Reputation Between 240 and 249 - Control Group	0.235	17
Reputation Between 240 and 249 - Treatment Group	0.300	10

Table 11: Structural Estimates

	Estimate	Standard Error
$ln(\gamma)$ mean	-6.702	1.160
$ln(\gamma)$ sd	7.111	0.803
$ln(\varphi)$ mean	-19.640	9.962
$ln(\varphi)$ sd	14.895	7.423
c	2.840	0.733

be 2.840. All of the five parameters are precisely estimated. The estimates suggest that the instrumental preference parameter, φ_i, follows a more dispersed distribution compared to the distribution of the social preference parameter, γ_i. Figure A1 shows model fit.

8 Counterfactuals

8.1 Assessing the Relative Importance of Social Motivation and Instrumental Motivation

To assess the relative importance of social and instrumental motivations, I consider two counterfactual scenarios: one in which social motivation is removed, and another in which instrumental motivation is removed. First, I simulate the absence of social motivation. The results indicate that the proportion of individuals willing to contribute falls to 26% of the baseline control group. Next, I simulate the removal of instrumental motivation, in which case the proportion remains at 84% of the baseline control group. The substantially larger decline observed when social motivation is removed suggests that it plays a much more critical role in driving contribution decisions.

In subsection 8.2 and subsection 8.3, I consider two platform strategies aimed at enhancing social motivation: (i) amplifying the number of upvotes for answers that already have at least one upvote, and (ii) awarding an additional upvote to high-quality answers.

8.2 Changing the Probability of Gaining Upvotes by Highlighting Certain Answers

How do individuals' contribution decisions change when the probability of receiving upvotes is altered? In practice, platforms influence this probability by giving certain answers more or less exposure or by displaying notifications that nudge or discourage users from upvoting specific answers. In this counterfactual exercise, I simulate changes in the probability

of receiving upvotes for all past answers with a positive number of upvotes contributed by individuals in my sample.

Figure 10: Counterfactuals

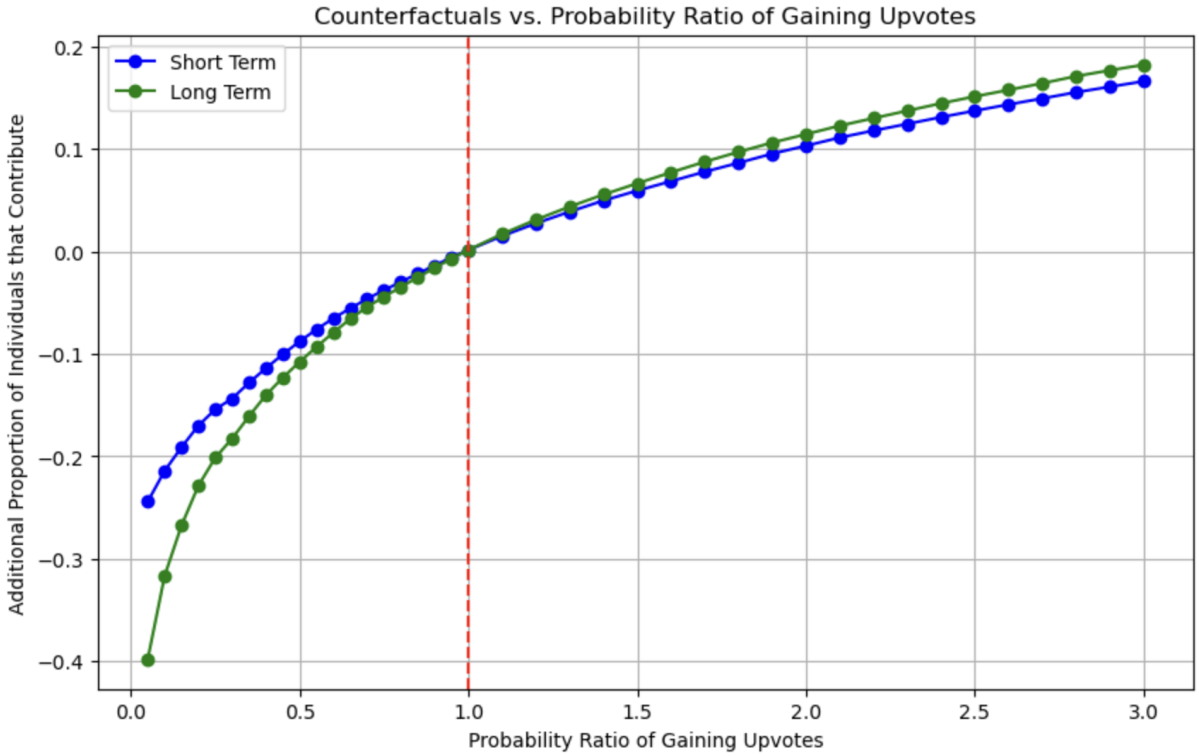
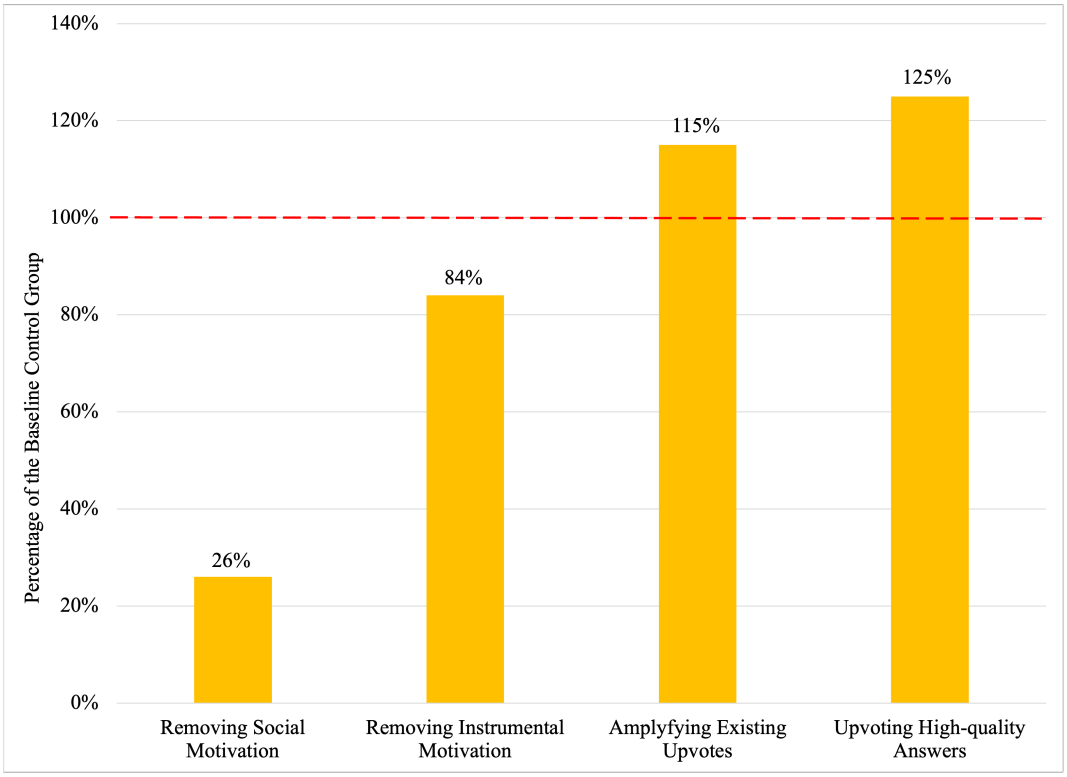


Figure 10 shows the change in the proportion of individuals that contribute, with counterfactual levels of probability of gaining upvotes relative to the baseline. The probability levels to the right of the vertical red correspond to increasing the probability of gaining upvotes, while the probability levels to the left of the vertical red line correspond to decreasing the probability. The blue line corresponds to the short-run, when individuals' prior has not adjusted, while the green line denotes the long-run, where individuals' prior has already adjusted according to the changed upvoting probability.

In the short run, increasing the probability of gaining upvotes on all past answers by 50% will induce an additional 5% of individuals to contribute. A 100% increase in this probability will result in about 9% more individuals contributing, and a 150% increase will lead to approximately 12% additional contributors. The efficacy of increasing the probability

Figure 11: Counterfactual Proportion in Percentage of Baseline



of upvotes shows diminishing returns. In the long run, as the prior adjusts, the proportion of contributing individuals increases slightly more.

However, decreasing the probability of gaining upvotes can have an increasingly negative marginal effect. As shown by the counterfactual probability levels to the left of the vertical line, a 10% decrease in the probability of gaining upvotes will reduce the proportion of contributing individuals by 2% in the short run, and a 50% decrease will reduce it by 8%. The slope becomes steeper as the probability further decreases. Additionally, in the long run, as the prior adjusts to the lower probability of gaining upvotes, the magnitude of the negative effect increases. The gap between the short-run (blue line) and long-run (green line) effects to the left of the vertical line widens as the probability further decreases.

8.3 Upvoting High-quality Answers by Introducing Expert Evaluations

When individuals decide which question to contribute an answer to, they can choose from a wide range of questions - every day, there are thousands of new questions posted, along with millions of existing questions. If an individual chooses to answer a less popular question or post an answer at a time when the visits to the site are low, their answers may be less likely to be noticed and upvoted by others, even though the quality of the answers may be high.

In this counterfactual, I consider giving one additional upvote to answers of high quality, defined by answers with at least one code example. This definition is consistent with existing research using data from Stack Overflow, such as Bregolin (2022). I show that giving one additional upvote to those answers will increase the proportion of contributing individuals by 25% of the baseline.

In the future, I plan to assess the quality of approximately 220,000 past answers contributed by individuals in my sample using advanced language models, such as ChatGPT. Subsequently, I will simulate a counterfactual scenario of upvoting answers deemed high-quality by the language model but which had not received any upvotes from other users at the time of sampling. This counterfactual analysis will provide insights into how the lack of recognition impacts contributions to online knowledge platforms.

In summary, Figure 11 illustrates the counterfactual proportion as a percentage of the baseline for four scenarios: removing social motivation, removing instrumental motivation, tripling the number of existing upvotes, and giving high-quality answers one additional upvote.

9 Conclusions

In this paper, I design and conduct a large-scale field experiment involving 12,182 individuals on one of the largest online question-and-answer platforms by collecting and treating daily samples for four and a half months. The treatment gives one additional upvote to eligible answers. I then track comprehensive data on individuals' subsequent daily posting behavior on the platform for four months.

I find that receiving one more upvote on an answer substantially impacts individuals' subsequent behavior. At the extensive margin, the probability of individuals making additional contributions increases by around 15% of the baseline. At the intensive margin, receiving one more upvote on an answer can increase the number of answers contributed within three weeks by around 6%. Both of the effects are sustainable over two months and, thus, are not driven by intertemporal substitutions.

The treatment effect is heterogeneous. It is larger on individuals with the number of past answers and reputation in the low-to-middle range and is slightly larger on individuals close to the next reputation threshold after being treated.

In terms of other outcomes, receiving one more upvote on an answer does not significantly impact whether individuals log in to the platform within 21 days of being sampled, but it slightly increases the number of days they log in by 0.55 days, around 3.9% of the baseline. The treatment has no effect on question contributions. However, it significantly increases both the probability of posting comments and the number of comments posted, with the effect magnitude on posting comments comparable to that on contributing answers. Additionally, the treatment significantly impacts individuals' voting behavior within 21 days. For upvoting, the treatment increases the probability by 7.30% of the baseline and the number of upvotes by around 5% of the baseline. For downvoting, the treatment increases the probability of making at least one downvote by around 20% of the baseline and the number of downvotes by 2.1%. The treatment does not seem to induce significantly less effort or worse

answers. Conditional on individuals deciding to contribute, there is suggestive evidence that individuals increase efforts.

To quantify how both motivators affect the primary outcome of interest - the extensive margin of whether or not to contribute additional online knowledge, I structurally estimate the model of individual decisions with heterogeneous utility parameters to explain my experimental findings. The estimates suggest that the extent to which individuals value social motivation and instrumental platform privileges varies widely across individuals.

Using estimates from the structural model, I explore a set of counterfactuals. I simulate scenarios where social motivation and instrumental motivation are removed, respectively. When social motivation is absent, the proportion of contributors drops to 26% of the baseline, whereas removing instrumental motivation retains 84% of the baseline. This larger reduction with the absence of social motivation suggests its more critical role in influencing contributions.

Then, I explore two platform strategies designed to enhance social motivation: amplifying the number of upvotes for answers that already have at least one upvote by highlighting those answers, and giving an additional upvote to high-quality answers by introducing expert evaluations. My findings indicate that tripling the number of upvotes for answers with at least one upvote can increase the proportion of contributing individuals by 15% compared to the baseline. Additionally, awarding one extra upvote to high-quality answers results in a 25% increase in contributions relative to the baseline.

Overall, the study suggests the important role of social motivation in motivating online content contributions and highlights the role of platform intervention in encouraging content contributions.

References

Ahn DY, Duan JA, Mela CF (2016) Managing user-generated content: A dynamic rational expectations equilibrium approach. *Marketing science* 35(2):284–303.

Andreoni J (2007) Giving gifts to groups: How altruism depends on the number of recipients. *Journal of public Economics* 91(9):1731–1749.

Athey S, Ellison G (2014) Dynamics of open source movements. *Journal of Economics & Management Strategy* 23(2):294–316.

Bregolin J (2022) Communication effort and the cost of language: Evidence from stack overflow. *Working Paper* .

Burtch G, He Q, Hong Y, Lee D (2022) How do peer awards motivate creative content? experimental evidence from reddit. *Management Science* 68(5):3488–3506.

Burtch G, Hong Y, Bapna R, Griskevicius V (2018) Stimulating online reviews by combining financial incentives and social norms. *Management Science* 64(5):2065–2082.

Burtch G, Lee D, Chen Z (2023) The consequences of generative ai for ugc and online community engagement. *Available at SSRN 4521754* .

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

Chen Y (forthcoming) *Digital Public Goods: Crowdsourcing, Crowdfunding, and User-Generated Content*.

Chen Y, Farzan R, Kraut R, YeckehZaare I, Zhang AF (2023) Motivating experts to contribute to digital public goods: A personalized field experiment on wikipedia. *Management Science* .

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

Conti A, Gupta V, Guzman J, Roche MP (2023) Incentivizing innovation in open source: Evidence from the github sponsors program. *Working Paper* .

Deolankar V, Fong J, Sriram S (2023) Content generation on social media: The role of negative peer feedback. *Available at SSRN 4522092* .

Eckles D, Kizilcec RF, Bakshy E (2016) Estimating peer effects in networks with peer encouragement designs. *Proceedings of the National Academy of Sciences* 113(27):7316–7322.

Gallus J (2017) Fostering public good contributions with symbolic awards: A large-scale natural field experiment at wikipedia. *Management Science* 63(12):3999–4015.

Goes PB, Guo C, Lin M (2016) Do incentive hierarchies induce user effort? evidence from an online knowledge exchange. *Information Systems Research* 27(3):497–516.

Guo M, Ni J, Shen Q, Xu Y (2023) Quantity vs diversity in online content production: Evidence from a knowledge sharing platform. *Johns Hopkins Carey Business School Research Paper* (22-12).

Huang JT, Kaul R, Narayanan S (2022) The causal effect of attention and recognition on the nature of user-generated content: Experimental results from an image-sharing social network. *Stanford University Graduate School of Business Research Paper* .

Jiménez Durán R (2021) The economics of content moderation: Theory and experimental evidence from hate speech on twitter. *Available at SSRN 4044098* .

Jin J, Li Y, Zhong X, Zhai L (2015) Why users contribute knowledge to online communities: An empirical study of an online social q&a community. *Information & management* 52(7):840–849.

Lacetera N, Macis M (2010) Social image concerns and prosocial behavior: Field evidence from a nonlinear incentive scheme. *Journal of Economic Behavior & Organization* 76(2):225–237.

Lerner J, Tirole J (2002) Some simple economics of open source. *The journal of industrial economics* 50(2):197–234.

Mummalaneni S, Yoganarasimhan H, Pathak V (2023) How do content producers respond to engagement on social media platforms? *Available at SSRN 4173537* .

Paridar M, Ameri M, Honka E (2024) More, faster, and better? effects of rewards on incentivizing the creation of user-generated content. *Available at SSRN 4580732* .

Srinivasan K (2023) Paying attention. *Working Paper* .

Sun Y, Dong X, McIntyre S (2017) Motivation of user-generated content: Social connectedness moderates the effects of monetary rewards. *Marketing Science* 36(3):329–337.

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

Wang Y, Goes P, Wei Z, Zeng D (2019) Production of online word-of-mouth: Peer effects and the moderation of user characteristics. *Production and Operations Management* 28(7):1621–1640.

Xu L, Nian T, Cabral L (2020) What makes geeks tick? a study of stack overflow careers. *Management Science* 66(2):587–604.

Zeng Z, Dai H, Zhang DJ, Zhang H, Zhang R, Xu Z, Shen ZJM (2023) The impact of social nudges on user-generated content for social network platforms. *Management Science* 69(9):5189–5208.

Zhang M, Luo L (2024) Effects of peer voting and followers on user contribution to online knowledge sharing platforms: Evidence from a field experiment. *Available at SSRN* .

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

Part I
Online Appendix

A Additional Tables and Figures

Table A1: Heterogeneity of Treatment Effect on Answer Contributions by Distance to the Next Privilege Threshold Before Being Treated

	Dependent variable: I(Additional Answers)		
	1-10 Distance	11-20 Distance	More than 20 Distance
Treatment	0.042***	0.055***	0.044***
	(0.014)	(0.017)	(0.011)
Observations	3,044	2,796	6,334
R^2	0.092	0.107	0.086
Adjusted R^2	0.038	0.049	0.060
Dep. Var. Mean	0.238	0.257	0.278
Percentage of Baseline	17.647	21.401	15.827
Sample Day FE	Yes	Yes	Yes
Reg. Month FE	Yes	Yes	Yes

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Note: The outcome variables are whether the individual contributed at least one additional answer within 21 days of being sampled, split by distance to the next threshold, before being treated/or not. This table includes all individuals, as a complement to Table 7.

Table A2: Average Quality and Efforts of Past Answers (Conditional on Contributing)

	Accepted	Scores	Body Words	Codes
Treatment	-0.018^{*}	-0.044	-5.801	-0.087
	(0.009)	(0.037)	(3.788)	(0.080)
R^2	0.069	0.085	0.065	0.069
Adj. R^2	0.018	0.035	0.014	0.019
Num. obs.	3327	3327	3304	3304
Dep. Var. Mean	0.282	0.643	155.556	2.495
Sample Day FE	Yes	Yes	Yes	Yes
Reg. Month FE	Yes	Yes	Yes	Yes

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Note: The outcome variables are measured within 21 days of being sampled. Accepted is defined as the proportion of accepted answers, scores is defined as the average scores of answers, body words is defined as the average number of words in answers, and codes is defined as the average number of code pieces included in answers, for each individual who contribute at least one answer within 21 days of being sampled.

Table A3: Ruling out an Alternative Mechanism - Treatment Effect on Answer Contributions by Reputation and Proportion of Reputation Gained from Answers

	Dependent variable: I(Additional Answers)		
	Group 1	Group 2	Combined
Treatment	0.067***	0.052***	0.050***
	(0.024)	(0.012)	(0.012)
High Reputation with Low Proportion			-0.030
			(0.020)
Treatment:High Reputation with Low Proportion			0.003
			(0.026)
Observations	1,185	4,625	5,810
R^2	0.172	0.064	0.063
Adjusted R^2	0.031	0.027	0.034
Dep. Var. Mean	0.179	0.250	0.235
Percentage of Baseline	37.430	20.800	21.277
Sample Day FE	Yes	Yes	Yes
Reg. Month FE	Yes	Yes	Yes

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Note: The outcome variables are whether the individual contributed at least one additional answer within 21 days of being sampled. The first column includes individuals who have an above-median reputation and with less than 50% of their reputation gained from answers, before being treated/or not. The second column includes individuals who have a below-median reputation and with more than 50% of their reputation gained from answers. If the heterogeneous treatment effect by answer experience/reputation in Table 6 is driven by satiation of reputation, the first column would have a smaller treatment effect. This is ruled out by the test in the third column, showing that the difference between the treatment effects of the two groups is statistically insignificant.

Table A4: Ruling out an Alternative Mechanism - Treatment Effect on Individuals Near Privilege Thresholds Unrelated to Contributing Answers

	Dependent variable: I(Additional Answers)		
	Within 10 Below Irrelevant Privileges	More than 10 Below Irrelevant Privileges	Combined
Treatment	0.098**	0.071***	0.075***
	(0.049)	(0.019)	(0.019)
Near			-0.041
			(0.026)
Treatment:Near			0.011
			(0.047)
Observations	518	2,074	2,592
R^2	0.312	0.143	0.132
Adjusted R^2	-0.010	0.066	0.070
Dep. Var. Mean	0.230	0.265	0.258
Percentage of Baseline	42.609	26.792	29.070
Sample Day FE	*Yes*	*Yes*	*Yes*
Reg. Month FE	*Yes*	*Yes*	*Yes*

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Note: The outcome variables are whether the individual contributed at least one additional answer within 21 days of being sampled. I focus on three privilege thresholds that are unrelated to contributing additional answers: setting bounties at 75 (offering some of reputation as bounty on a question), viewing close votes at 250 (viewing and casting close/reopen on ones own questions), and accessing review queues at 500 (accessing the first posts and late answers review queues). The first column includes individuals who have a reputation within 10 points below the three privilege thresholds, before being treated/or not. The second column includes individuals who have a reputation above the previous threshold, but more than 10 points below the three privilege thresholds. If the treatment effect in the first column of Table 7 is completely driven by additional privilege making it easier for people to contribute answers, then the first column in this table would have an insignificant treatment effect, when the additional privileges are unrelated to contributing answers. This is ruled out by the statistically significant positive treatment effect in the first column, and the test in the third column, showing that the difference between the treatment effects of the two groups is statistically insignificant.

Table A5: Ruling out an Alternative Mechanism - Treatment Effect on Individuals Contributing No More than One Additional Answer

	Dependent variable: I(Additional Answers)		
	1-10 Distance	11-20 Distance	More than 20 Distance
Treatment	0.026**	0.041**	0.029***
	(0.013)	(0.016)	(0.009)
Observations	2,628	2,344	5,284
R^2	0.070	0.088	0.040
Adjusted R^2	0.005	0.016	0.008
Dep. Var. Mean	0.126	0.125	0.139
Percentage of Baseline	20.635	32.800	20.863
Sample Day FE	Yes	Yes	Yes
Reg. Month FE	Yes	Yes	Yes

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Note: The outcome variables are whether the individual contributed at least one additional answer within 21 days of being sampled, split by distance to the next threshold, before being treated/or not. This table includes individuals who contributed no more than one additional answer within 21 days of being treated. This corresponds to 84.198% of the full sample.

Figure A1: Model Fit

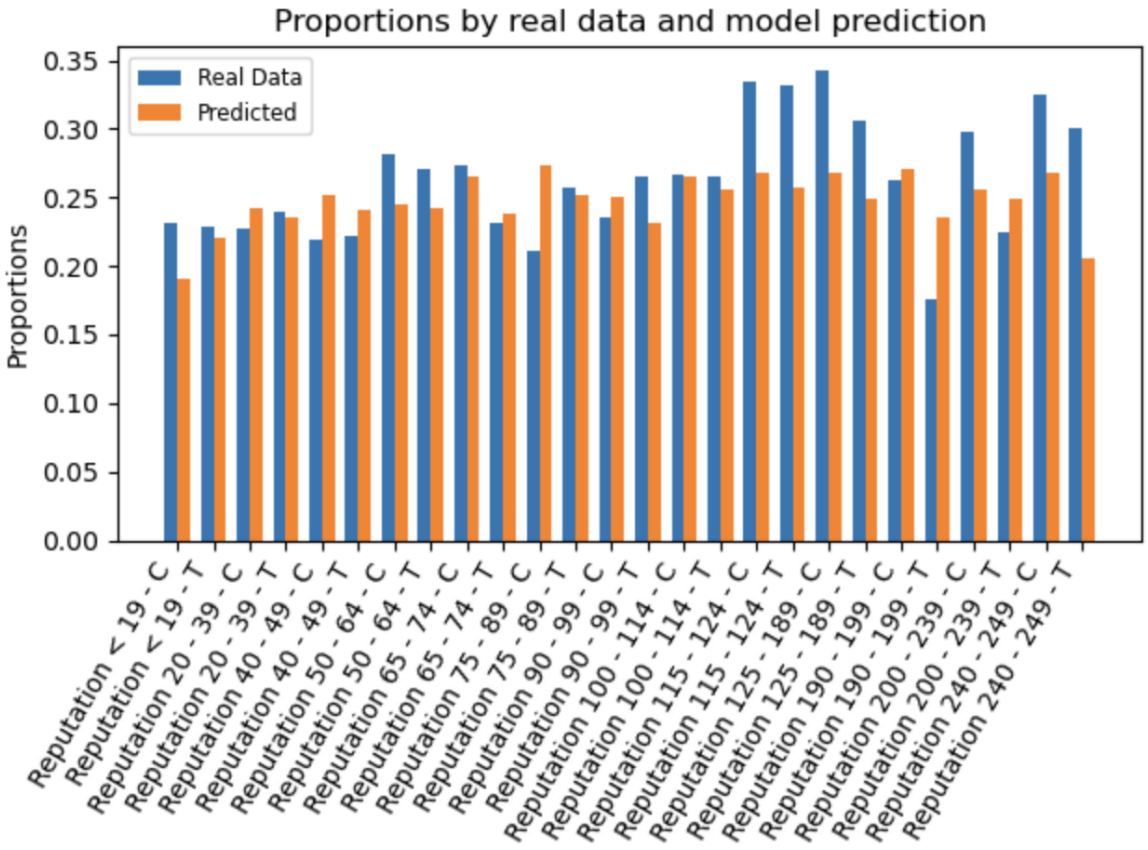
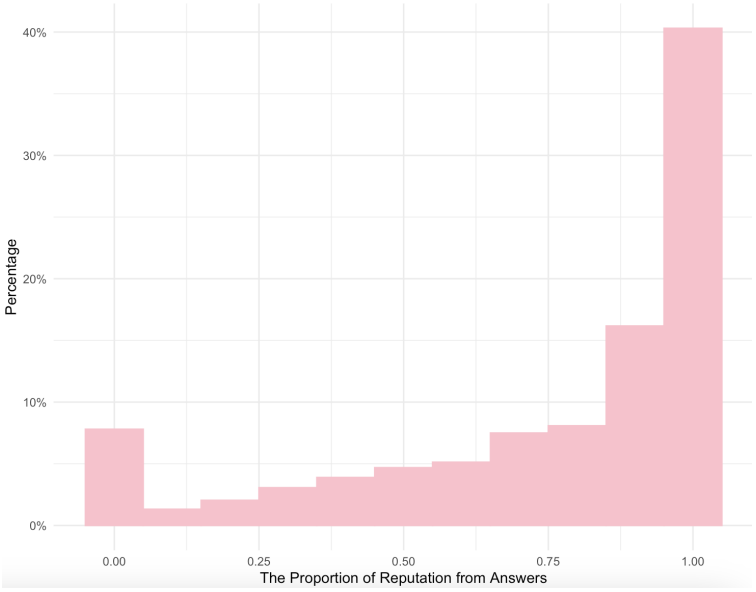
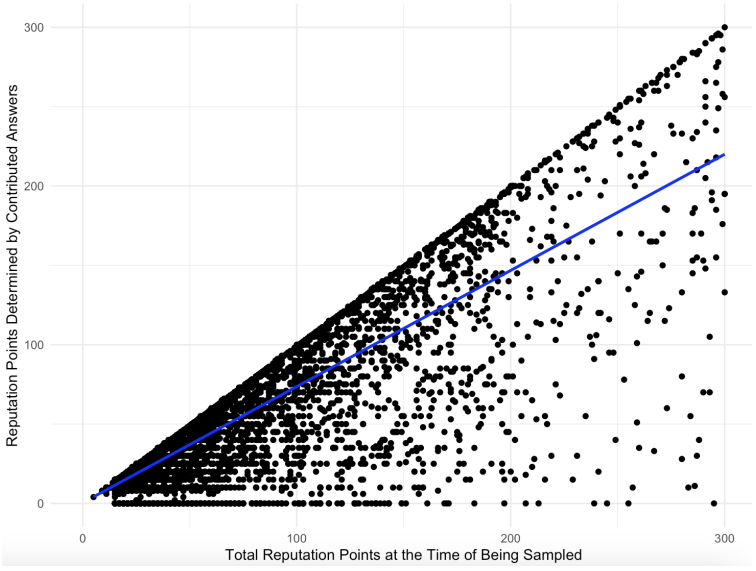


Figure A2: Reputation from Answers and Total Reputation at the Time of Being Sampled

(a) The Percentage Distribution of the Proportion of Reputation from Answers



(b) Reputation from Answers vs. Total Reputation



Note: Panel (a) of the figure shows the percentage distribution of the proportion of reputation from contributed answers for each individual at the time of being sampled. In addition to answering questions, individuals can gain reputation points by posting questions and suggesting edits, and they can lose reputation by downvoting others. Overall, 25% of individuals gain less than 59.52% of their reputation points through contributing answers. The median individual in the sample gains 90.91% of reputation through contributing answers. The mean proportion is 74.25%. Panel (b) of the figure shows a scatter plot of reputation points from answers versus total reputation points at the time of being sampled. Each point represents one individual.

B Proofs of Model Predictions

For completeness I restate the key objects (see Section 3 of the main paper for intuition).

Posterior belief. With n previous answers and total historical upvotes $S_n := \sum_{r=1}^{n} s_{ir}$, the conjugate-Gamma posterior is

$$\lambda_{it} \mid I_{it} \sim \Gamma\Big(k_{i0} + S_n,\ \frac{\theta_{i0}}{n\theta_{i0}+1}\Big),$$

so

$$\mu_{it} := \mathbb{E}[s_{it} \mid I_{it}] = \frac{\theta_{i0}}{n\theta_{i0}+1}\big(k_{i0} + S_n\big). \tag{1}$$

Deterministic part of contribution value.

$$V_{it} = \underbrace{\gamma_i\, 10\mu_{it}}_{\text{utility from extra reputation}} + \underbrace{\varphi_i\big[\mathbb{E}\, q(R_{it}+10s_{it}) - q(R_{it})\big]}_{\text{utility from extra chance of getting the additional privilege}} - c_i. \tag{2}$$

With an i.i.d. Type I extreme-value shock ϵ_{it}, the contribution probability equals $P_{it} = \frac{e^{V_{it}}}{1+e^{V_{it}}}$.

One extra up-vote on a past answer. The shock raises k_{i0} by 1 *and* R_{it} by 10. Post-shock objects carry "+" and I set

$$\mu_{it} = \frac{\theta_{i0}}{n\theta_{i0}+1}\big(k_{i0} + S_n\big), \tag{3}$$

$$\mu_{it}^{+} = \frac{\theta_{i0}}{n\theta_{i0}+1}\big(k_{i0} + S_n + 1\big), \tag{4}$$

$$\begin{aligned}\Delta\mu_n := \mu_{it}^{+} - \mu_{it} &= \frac{\theta_{i0}}{n\theta_{i0}+1}\big[(k_{i0}+S_n+1) - (k_{i0}+S_n)\big] \\ &= \frac{\theta_{i0}}{n\theta_{i0}+1}\,.\end{aligned} \tag{5}$$

$$\Delta\mu_n := \mu_{it}^{+} - \mu_{it}, \quad \Delta V_n := V_{it}^{+} - V_{it}, \quad \Delta P_n := \frac{e^{V_{it}^{+}}}{1+e^{V_{it}^{+}}} - \frac{e^{V_{it}}}{1+e^{V_{it}}}. \tag{6}$$

I list below two auxiliary lemmas that are used repeatedly in the proofs of Predictions 1-3.

Lemma 1. *For every integer $n \geq 0$,*

$$\delta\mu_n = \frac{\theta_{i0}}{n\theta_{i0}+1} \quad \textit{satisfies} \quad \delta\mu_n > \delta\mu_{n+1}.$$

Thus $\{\delta\mu_n\}$ is strictly decreasing in n.

Proof.

$$\delta\mu_n - \delta\mu_{n+1} = \frac{\theta_{i0}}{n\theta_{i0}+1} - \frac{\theta_{i0}}{(n+1)\theta_{i0}+1} = \frac{\theta_{i0}^2}{(n\theta_{i0}+1)\big((n+1)\theta_{i0}+1\big)} > 0.$$

□

Lemma 2. *For fixed $x \in \mathbb{R}$, the function $a \mapsto \frac{e^{x+a}}{1+e^{x+a}} - \frac{e^x}{1+e^x}$ is strictly increasing in a.*

Proof. Its derivative is $\frac{e'}{1+e'}x + a = e^{x+a}/(1+e^{x+a})^2 > 0$. □

Throughout, abbreviate

$$\Delta_q := \mathbb{E}\, q(R_{it}^+ + 10s_{it}) - \mathbb{E}\, q(R_{it} + 10s_{it}) \geq 0.$$

Prediction 1: Learning is stronger early on

Prediction EC.4 (Prediction 1)**.** Holding $\gamma_i, \varphi_i, R_{it}$ and the current posterior fixed, ΔP_n is strictly decreasing in the integer index n.

Proof. Lemma 1 gives $\delta\mu_{n_1} > \delta\mu_{n_2}$ for $n_1 < n_2$. From (2),

$$\Delta V_n = \gamma_i\, 10\, \delta\mu_n + \varphi_i\, \Delta_q$$

is therefore strictly decreasing in n. The logistic map in Lemma 2 preserves the inequality, so $\Delta P_{n_1} > \Delta P_{n_2}$. □

Prediction 2: Threshold proximity matters

Prediction EC.5 (Prediction 2)**.** Let τ be the smallest privilege threshold exceeding R_{it}.

(a) If $R_{it} + 10\mu_{it} \in (\tau - 10, \tau]$ (*near threshold*) then $\Delta_q > 0$ and $\Delta P_n > \gamma_i\, 10\, \delta\mu_n$.

(b) If $R_{it} + 10\mu_{it} \leq \tau - 10$ or $R_{it} + 10\mu_{it} \geq \tau + 10$ (*far from threshold*) then $\Delta_q = 0$ and $\Delta P_n = \gamma_i\, 10\, \delta\mu_n$.

Proof. Define $p = \mathbb{P}(R_{it} + 10s_{it} \geq \tau)$ and $p^+ = \mathbb{P}(R_{it} + 10 + 10s_{it} \geq \tau)$. Then

$$\Delta_q = p^+ - p = \mathbb{P}\Big(\frac{\tau - R_{it}}{10} - 1 \leq s_{it} < \frac{\tau - R_{it}}{10}\Big).$$

The bracket has positive length (hence $\Delta_q > 0$) exactly in the near-threshold region; otherwise it is empty, so $\Delta_q = 0$. Combine with (2) and use Lemma 2. □

Prediction 3: Higher instrumental value φ_i amplifies the effect

Prediction EC.6 (Prediction 3)**.** Within the near-threshold region of Proposition EC.5, $\partial\Delta P_n/\partial\varphi_i > 0$.

Proof. Near a threshold, $\Delta_q > 0$, so $\Delta V_n = \gamma_i \, 10 \, \delta\mu_n + \varphi_i \, \Delta_q$ is affine and strictly increasing in φ_i. Lemma 2 then gives the same sign for $\partial \Delta P_n / \partial \varphi_i$. □