

Authority and Delegation in Online Communities ^{*}

Jacopo Bregolin[†]
University of Liverpool

Latest version available [here](#)

November 2020

Abstract

Many online platforms rely on user-generated content and need to incentivize free effort. With data from Stack Exchange, I investigate whether users provide more and better quality contributions when endowed with more autonomy and authority over actions. Using a dynamic discrete choice model, I show that authority has positive marginal value that is heterogeneous across different types of users. I simulate counterfactuals with different designs. The results show that the platform would lose an important share of production and quality of content in the absence of delegation. The trade-off depends on the composition of the community, as the sensitivity to the incentives is heterogeneous.

^{*}This paper is part of my Phd dissertation at TSE. I am extremely indebted to my advisor, Milo Bianchi, co-advisor, Daniel L. Chen, and committee members Jacques Crémer and Daniel Ershov. Their support and guidance have been invaluable. I am grateful to Olivier de Groote and Heski Bar-Isaac for their extensive help. I also thank Pierre-André Chiappori, April Franco, Hulrich Hege, Astrid Hopfensitz, Aija Leiponen, Greg Lewis, Inna Lyubareva, Matthew Mitchell, Sophie Moinas, Elena Panova, Mark Schankerman, and Timothy Simcoe for helpful discussions and feedback. I thank members of the Finance, Industrial Organization, and Behavior, Institution and Development groups at TSE, as well as seminar participants at the IOEA Corsica, ZEW, Rotman School of Management, and VIDE brown bag for useful comments.

[†]jacopo.bregolin@liverpool.ac.uk

1 Introduction

Many companies rely on voluntary contributions by internet users. User-generated content is valuable because it provides information about quality (as for product reviews, [Luca 2011](#), [Lewis and Zervas 2019](#)), customer support at zero cost (as crowdsourced online forums¹), or platform enhancing features, like for Spotify and Google Maps. For some platforms, user content is the product itself. This is the case for social media platforms, information aggregators such as Wikipedia, and question-and-answer websites such as Stack Exchange. How can companies incentivize participation without remuneration? The literature has investigated the motives behind participation, identifying factors based on intrinsic utility ([Roberts, Hann, and Slaughter 2006](#)), peer effects ([Zhang and Zhu 2011](#), [Chen, Harper, Konstan, and Li 2010](#)), and virtual rewards ([Gallus and Frey 2016](#), [Goes, Guo, and Lin 2016](#)). The theoretical literature in personnel economics has identified another nonmonetary channel to incentivize participation, that is, the delegation of autonomy and decision rights ([Gibbons, Matouschek, and Roberts 2013](#), [Gambardella, Panico, and Valentini 2015](#)).² Empirical work on this channel is sparse, and to my knowledge, it has not been studied in the context of online communities.³

In this paper, I investigate whether the delegation of control rights and authority leads to an increase in online contributions. I identify whether and to what extent users are interested in obtaining more autonomy over tasks and study its role in contribution patterns. Using data from Stack Exchange, I show that people do value such autonomy. However, different types of users value obtaining and having such autonomy differently. Through counterfactual exercises, I explore organizational implications and find that delegation increases performance of highly extrinsically motivated participants, while all users are more willing to participate when endowed with more authority.

What does it mean to allocate authority in digital platforms? Every online com-

¹[Mozilla's Firefox](#) for instance.

²In this paper, I interchangeably use the terms of "delegation of decision rights", "delegation of control rights", and "autonomy over decisions". The term "authority" identifies autonomy over a decision whose outcome affects other individuals. For the definition of power and related concepts, I refer to [Sturm and Antonakis 2015](#)

³There is anyway theoretical work that studies the incentive effect of delegation, for instance [Rajan and Zingales \(1998\)](#), [Blanes I Vidal and Möller \(2007\)](#), and [Bester and Kräbmer \(2008\)](#). The literature has addressed as well other types of nonmonetary incentives, with similar dynamics to the delegation of authority. [Auriol and Renault \(2008\)](#) and [Besley and Ghatak \(2008\)](#) investigate status incentives, while the tournaments literature has studied promotion incentives ([Lazear and Rosen 1981](#)). These papers include rivalry between workers in obtaining status and promotions. In my work instead, delegation does not depend on other workers' actions. Finally, the literature has also identified several other non-monetary utilities. While conceptually farther from the delegation of authority, they are still relevant, as they may affect the users' decision process. Motivations that have been identified include intrinsic utility and firm recognition ([Roberts et al. 2006](#), [Nov 2007](#), [Ma and Agarwal 2007](#), [Jeppesen and Frederiksen \(2006\)](#)), the community size ([Zhang and Zhu 2011](#)), reference points on others' behavior ([Chen et al. 2010](#)), within-community reputation ([Chen, Ho, and Kim 2010](#)), peer recognition ([Jin, Li, Zhong, and Zhai 2015](#), [Chen, Wei, and Zhu 2017](#)), awards ([Gallus and Frey 2016](#)), sequential targets ([Goes et al. 2016](#)), and the signaling of skills ([Belenzon and Schankerman 2015](#), [Xu, Nian, and Cabral 2020](#)).

munity requires moderators (Gillespie 2018), but who has the authority to modify the community content differs across platforms. Facebook does not allow users to modify content and hires professional moderators. Users are only allowed to flag content that they believe violates Facebook’s rules. In contrast, Wikipedia allows every internet user to modify existing articles. Finally, Stack Exchange provides authority on moderation conditional on achieving given performance targets. What trade-offs affect this decision?

This paper focuses on the incentive effects of the allocation of authority based on performance and studies the trade-off that arises from conflicting incentives. It includes Facebook’s and Wikipedia’s strategies as limit cases, where the performance threshold required to obtain authority is set at either infinity or zero. I address two main incentive effects. First, if users value acquiring autonomy, delegation incentivizes effort until users reach the performance threshold (*dynamic incentive*). Second, if users value contributing when endowed with more autonomy, delegation relaxes the participation constraint, as it increases the value of participating (*static incentive*). A stronger *dynamic incentive* effect would suggest increasing the performance threshold, while a stronger *static incentive* would suggest decreasing it. The paper studies the platform’s trade-off by quantifying both incentive effects under different counterfactual performance thresholds.

To give an overview, the paper uses data from Stack Exchange, a family of websites in which registered users ask questions and provide answers on different topics. The website is moderated by experienced users who have full autonomy over editing questions and answers and who are not paid. New users’ edits need to be approved by the moderators. New users become moderators after reaching a performance threshold. In this context, I observe users’ contributions before and after they receive authority on editing. After providing evidence of the *static incentive* effect via a regression discontinuity design, I develop a dynamic discrete choice model to measure users’ preference for authority, allowing for heterogeneity across types. The paper finds that the incentives differently affect the different types. Specifically, the incentive responses depend on heterogeneous valuations for authority and different participation costs. The latter is an important since a potentially large incentive effect becomes negligible if the cost of participation is high. The final total number of contributions is strongly affected by the composition of the community, which is therefore a crucial factor in designing incentives.

The data I use include the contribution history (answers, questions, edits, and comments) of all participants in the English Language Learners website of Stack Exchange.⁴ The partition of users by type is data-driven and based on users’ profile pages. It aims to capture the heterogeneity of the broad motives behind participation. Three types emerge: *Anonymous* users provide very little information, *Identifiable* users provide community-relevant information, and *Informative* users provide a lot of information with links to external content (such as LinkedIn profiles).

The analysis proceeds in two steps. First, I test for the presence of *static incentives* by looking at both the acquisition and the loss of authority through a reduced-form analysis. To study the effect of the acquisition of authority, I use a regression discontinuity analysis in which the running variable is the distance from the threshold in points,

⁴<https://ell.stackexchange.com/>

i.e. the performance measure on which the threshold depends. I compare editing with commenting activity since the latter is not affected by the threshold. I find a significant and stable increase in the number of edits just after the acquisition of autonomy, while the number of comments follows an independent pattern. The effect is driven mainly by *Anonymous* users, while *Informative* users seem to respond in the long run. The study of the loss of autonomy exploits a variation in the platform design, which increases the performance required to obtain authority. I find that participants who anticipated the change and lost authority stopped making edits, while participation in answering questions did not change. The dynamic nature of the *dynamic incentive* effect does not allow clean identification in reduced form.

In the second step, to quantify the incentive effects and simulate counterfactuals, I use a structural dynamic discrete choice model. In each period, users decide their contribution in terms of the number of answers, the quality of answers, and the number of edits. The utility function includes a dummy variable equal to 1 if the user reaches the required performance threshold to obtain authority. Identification of the *dynamic incentive* relies on the effort that users make when approaching the performance threshold: higher effort allows them to reach the threshold more quickly. Systematic higher effort when approaching the threshold would identify a positive marginal utility of authority. Variation in the willingness to participate once endowed with authority identifies the *static incentive* effect. The utility function includes interactions of the dummy variable with variables capturing the intrinsic net benefit of participation, allowing for long-term changes in the net cost of contribution. In addition to variables that capture the cost of participation, the utility function includes other sources of motivation that are potentially correlated with the threshold: the number of points and the number of privileges accumulated. I estimate the flow utility parameters using *finite dependence* (Arcidiacono and Miller 2011), a methodological tool that allows the approximation of value functions without the full solution of the model.

The results show a positive marginal utility of authority and a significant increase in willingness to participate in editing once endowed with authority. *Anonymous* users show the highest marginal cost of contribution. Nevertheless, they have the highest value for authority, meaning that both the *dynamic incentive* and the *static incentive* effects are particularly relevant for them. *Informative* users are very sensitive to the *dynamic incentive*, while *Identifiable* users are not. The results suggest that authority on editing does not have spillover effects on answering, and only *Anonymous* users slightly substitute answering with editing.

With estimates from the model, I simulate counterfactual contribution histories under different performance requirements to obtain authority. In particular, I consider the case with a performance threshold equal to zero (full delegation), infinity (no delegation), or two intermediate levels. The results show that in the simplified context of the simulation, *Anonymous* users do not contribute due to their high costs of participation. Since *Identifiable* users are not sensitive to the *dynamic incentive*, the choice of the threshold level should focus on *Informative* users. Their participation in answering is maximized when they have to reach a performance threshold, but a threshold that is too high may

induce a smaller increase in participation. The optimal threshold level depends on the expected lifetime of participation and on whether the reputation points that can be accumulated are capped.

This paper has two main contributions. First, I show direct evidence of nonmonetary preferences and identify in real data the intrinsic value of authority. This result confirms experimental evidence showing that individuals value control rights and power (Fehr, Holger, and Wilkening 2013, Bartling, Fehr, and Herz 2014, Owens, Grossman, and Fackler 2014, Pikulina and Tergiman 2020).⁵

The second contribution relates to the organizational implications of these nonmonetary preferences. The paper shows that platform designers should take into account the incentive effects induced by the allocation of decision rights. In addition, the paper suggests that the platform will optimally target different users with different incentives. While the results of this paper are specific to the context of online communities, they may suggest implications for a broader set of environments, addressing puzzles that emerged in the literature on promotions. It can provide a plausible explanation for 1) the use of promotions rather than bonuses, even if bonuses are more flexible incentives (Baker, Jensen, and Murphy 1988, Gibbons and Waldman 1999) and 2) the commitment to promote employees on the grounds of observable measures not correlated to the skills required for the delegated tasks (Peter principle, Fairburn and Malcomson 2001, Benson, Li, and Shue 2019).

The paper proceeds as follows. Section 2 describes the website from which data are taken, while section 3 presents the data and the identification of user types within the online community. I then present the results from the reduced-form analyses in section 4, and the structural model in section 5. Finally, sections 6 and 7 report the results and the counterfactual simulations. Section 8 concludes.

2 Stack Exchange: “Self Managed” Platforms

I use data from Stack Exchange. Stack Exchange is a family of platforms founded in 2009 that provides users the opportunity to post questions and answers on a variety of topics. Each website in the group specializes in a particular topic: notably, *Stack Overflow*, the largest community, hosts questions and answers about programming languages, but there are 172 other websites, each focused on a different topic, from technology to the arts. These websites belong to the commercial company Stack Exchange Inc. which has raised 153 million dollars in venture capital, as of July 2020.⁶ To give a sense of the welfare produced to consumers, Stack Exchange receives 418.8 million monthly visits and 805.9 million monthly page views. It contains 3.3 million questions, which

⁵Nonexperimental work has identified a beneficial effect of delegation on performance but does not investigate whether a channel is an intrinsic value for authority. Bandiera, Best, Khan, and Prat (2020) use a field experiment, while Liberti (2018) use real data from a financial institution.

⁶<https://www.businesswire.com/news/home/20200728005330/en/Stack-Overflow-raises-85M-Series-funding-accelerate>

have received 3.6 million answers⁷. Instead of hiring experts to answer questions, Stack Exchange is crowd-based. Anyone can register and contribute to the platform: there are no registration fees, but contributions are not remunerated. Users do not need to register to browse the content. Its business strategy is similar to that of other websites, such as *Quora* or *Yahoo! Answers*, but differs from that of *Google Answers*, which was active between 2002 and 2006, where those responding to answers were paid.

The objective of the platform, as described by the creators, is to provide detailed and easily accessible solutions for specific questions.⁸ For instance, duplicate questions or questions on subjective topics are *closed* (i.e. they do not allow answers), and the answers to a question are ranked based on up-votes rather than the publication date.

Participation in Stack Exchange is subject to an incentive system based on virtual rewards, either reputation points or *badges*.

Badges.

Badges are comparable to medals and, to some degree, to firms' bonuses. Users obtain badges when they accomplish given performance targets, whereby performance depends on the quantity, quality, and timing of contributions. There are bronze, silver, and gold badges, based on performance required.

Reputation Points and Privileges.

Once the user publishes a question or an answer, other community members can up-vote it or down-vote it, allocating or removing reputation points from the author. More precisely, each up-vote provides 10 points, while each down-vote removes 2 points. To vote, users need to have accumulated at least 15 reputation points. The user can also receive points by suggesting a modification to existing content: if the suggestion is accepted and gets implemented, the user gains 2 points.⁹ The accumulation of points allows a user to obtain *privileges*. With few exceptions, privileges are rewards that give access to resources or actions. Users obtain them when reaching given threshold levels of reputation points, hierarchically. The higher the number of points accumulated is, the closer the user gets to having full administrative control of the website. Table 1 reports the list of privileges and the reputation points necessary to obtain each of them (rightmost columns).¹⁰ Privileges may be comparable to promotions in traditional companies: more experienced *employees* receive more information and authority from the company owners, as well as more responsibilities. A difference is that workers compete for a limited number of higher positions, while in Stack Exchange there is no competition. Users are guaranteed to obtain privileges if they reach the required performance.

⁷<https://stackexchange.com/about>

⁸<https://www.joelonsoftware.com/2008/09/15/stack-overflow-launches/>

⁹Figure 25 in the appendix provides the detailed rules to gain points

¹⁰Since these values have changed during the life of the website, the table provides two values for each privilege. In section 3, I provide more details on how the change happened and when each threshold has been applied.

2.1 Hierarchy in Stack Exchange and the Rationale for Delegation

In Stack Exchange, community members make most of the content decisions, but not all users have the same decision rights. Users can acquire decision rights and authority in two ways. The first way is to be elected. The platform organizes internal elections at an intermittent frequency. Elected users have a permanent mandate and authority to moderate the platform. Users can obtain the same control rights as elected moderators by accumulating points, which allows them to obtain privileges. Privileges provide either access to actions or more authority. Via the allocation of privileges, the platform commits delegating control to community members that achieve given performance measures. In this paper, I will focus specifically on the privilege that delegates authority in editing. Editing is the action of modifying existing content to improve or correct it. Before users reach that privilege, they are allowed to propose modifications, but their suggestions need to be approved by either the author of the modified content or by the voting of two users that already have the privilege. Once users obtain enough reputation points to obtain the privilege, their edits are directly implemented and do not require the approval of third parties. I consider this variation as an increase in editing authority.

Why would the platform want to delegate authority to community members? Compared to hiring professional moderators, delegating has the advantage that community members work for free. This induces important savings for the platform, as it needs many moderators.¹¹ There are two other important reasons. First, if users' willingness to make contributions to the platform is significantly higher when endowed with full authority, delegation relaxes a user's participation constraint. The intuition is equivalent to what [Gibbons et al. \(2013\)](#) refers to as "to pay the employee less": I define this effect as the *static incentive*, as it is independent of the dynamics of contribution. Second, to tie delegation to performance incentivizes participation if users value gaining authority. I call this effect the *dynamic incentive*. The incentive effects induced by delegation are particularly relevant in the context of voluntary work. As users are volunteers, their outside option from participating is high, and the absence of formal contracts reduces the cost of leaving.

If the platform wants to leverage both the *static incentive* and the *dynamic incentive*, it faces a trade-off. A positive *static incentive* effect would suggest delegating to every user, independent of performance, while a positive *dynamic incentive* effect would advise conditioning delegation on performance. Finally, if the platform delegates on performance, it needs to decide the level of performance required. A second trade-off emerges: a more demanding performance threshold incentivizes participants for a longer time. However, a higher performance threshold decreases the *static incentive* effect.

¹¹Facebook employs around 15000 moderators: [Charlotte Jee, MIT Technology Review, June 2020](#)

Privilege	type	Reputation Requirements	
		Graduated	Public Beta
access to site analytics	Milestone	25000	5000.0
trusted user	Milestone	20000	4000.0
protect questions	Moderation	15000	3500.0
access to moderator tools	Moderation	10000	2000.0
approve tag wiki edits	Moderation	5000	1500.0
cast close and reopen votes	Moderation	3000	500.0
create tag synonyms	Moderation	2500	1250.0
edit questions and answers	Moderation	2000	1000.0
established user	Milestone	1000	750.0
create gallery chat rooms	Communication	1000	
access review queues	Moderation	500	350.0
create tags	Creation	300	150.0
view close votes	Moderation	250	250.0
vote down	Moderation	125	125.0
edit community wiki	Creation	100	100.0
create chat rooms	Communication	100	
set bounties	Creation	75	75.0
comment everywhere	Communication	50	50.0
talk in chat	Communication	20	
flag posts	Moderation	15	15.0
vote up	Moderation	15	15.0
remove new user restrictions	Milestone	10	10.0
create wiki posts	Creation	10	10.0
participate in meta	Communication	5	5.0
create posts	Creation	1	1.0
vote in moderator elections		150	150.0
association bonus		200	200.0
shown in network reputation graph and flair		200	200.0
reputation leagues - top x% link in profile		201	201.0
qualify for first Yearling badge		201	201.0
run for moderator		300	300.0

Table 1: List of privileges that users can obtain when accumulating reputation points. The first column describes what the user obtains when achieving each privilege, the second shows the category of the privilege, while the third and the fourth show the number of reputation points required to obtain the privilege. The *Public Beta* column applies to the platform between January 2013 and February 2016, while the *Graduation* column applies from February 2016 onwards.

3 Data

Stack Exchange is composed of many websites that share the same structure, whose only difference is the main topic of the questions posted. In this paper, I use data from the website called *English Language Learners* (ELL), which focuses on questions and answers related to the use of English.¹² The creation of Stack Exchange websites follows a specific procedure. First, an initial community of users makes a proposal of creation in a specific platform called *Area 51* and starts contributing.¹³ When the website proves to have enough demand and a sustained amount of activity within *Area 51*, the platform administrators launch it with an independent URL. The website enters the beta period, which is divided into *private beta* and *public beta*. The *private beta* allows participation only to users who have contributed to the development phase. Normally, after a week, the website moves to the *public beta* phase, characterized by no restrictions on participation. Finally, once the platform administrators assess that the website can be sustainable over time, the site *graduates* to the final phase and receives a personalized design. Normally, the *graduation* and the new design would occur on the same date, but on the ELL website, the design occurred later.¹⁴ The timeline of these steps for the ELL website is reported in Figure 1. Once the website receives the new design, the reputation points required to obtain the privileges change.¹⁵




Figure 1: Timeline of the website

The data were retrieved on May 31st, 2020, and contain both information displayed in the user profile pages, as well as the content and modification history of posts (ques-

¹²The choice of this specific website is justified by several reasons. First, posts contain only text, and not equations or scripts, as it happens in more technical Q&A. As I will detail later, this is instrumental to measure quality of posts. Second, compared to other websites, it contains an additional source of variation, as the incentive structure was slightly modified in February 2016. Third, the topic of the English Language is not trivially linked to professional opportunities outside the community, which may otherwise drive most of participation.

¹³<https://area51.stackexchange.com/>

¹⁴The shift was due to a backlog of the designer team.

¹⁵Table 1 reports the number of reputation points required to obtain each privilege. The *Public Beta* column reports the requirements before the design, while the *Graduated* column reports the requirements after the design.

tions and answers). While the company makes most of the data available for free, I web-scraped the daily histories of reputation points obtained by users. At the time of downloading the data, the website counted 92,853 registered users, 121,633 published answers, and 77,357 questions. I constructed a panel of users’ participation on the website. I include users who have published at least one answer or edit while excluding those who have not gained a positive number of reputation points. Users are assumed to exit the platform after three months of inaction in answering or editing. The data are right-censored at the download date.

In the panel, there are 9,797 users who participated 713 days on average (with a range between 1 and 2,685 days). They published a total of 114,926 answers, on average 11.7 each, but with a very skewed distribution, ranging from 1 to 4,173. The edits I consider are edits to answers, either modifications of answer’s content or rollbacks, i.e. the recovery of a previous version. The users of the sample made 8,168 edits, of which 1,409 were suggested and the rest were directly implemented. Each user on average made 0.8 edits, with a range between 0 and 1,174. Figure 2 shows the cumulative production of users in both absolute terms and in shares. Users are ranked by the intensity of answering. It shows that activity in answers and edits is very concentrated, while questions are more homogeneously distributed.

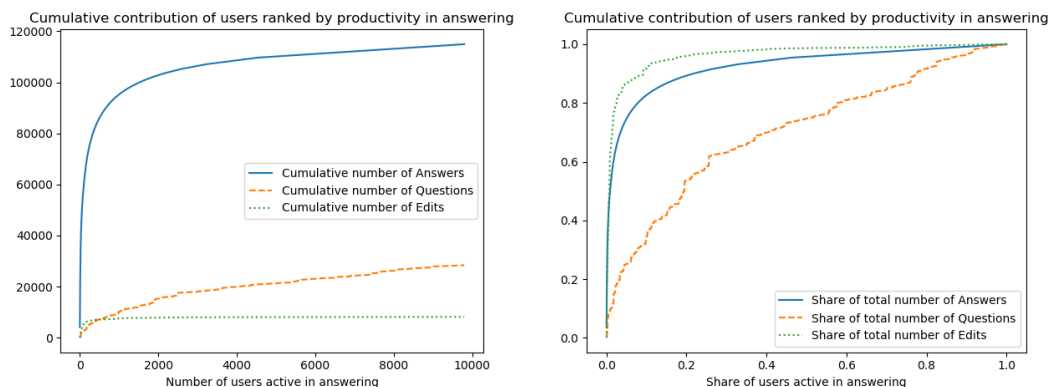


Figure 2: Production of answers, questions, and edits among active users. The x-axis reports the number of users (left) and the share of users (right) ranked by the number of answers published, while the y-axis reports the absolute number of answers, questions, and edits (left) and the respective shares (right).

In the sample, users reached on average 487 reputation points, with a range from 0 to 175,955, the 75th percentile being 208 (the zero is due to a particular case that got included in the sample). Figure 3 reports the number of users at each point in time who have reached the threshold to obtain more control over editing. Note that when the threshold value changed, some users lost their privilege.

I construct a proxy for answers’ quality using textual measures, including the number of words, links, and pictures. Details on this process can be found in appendix A.1. On average, users provided answers with a quality equal to 0.16, and the quality range

spanned from 0.004 to 14.107.

Finally, I construct a variable to measure the number of daily open questions in the topics of experience of the users. A question is open if it does not have an accepted answer. In summary, the variable is constructed by 1) clustering tags around topics, which are identified by exploiting the co-occurrence of tags in questions; 2) allocating open questions to topics; 3) recovering the user’s expertise on each topic based on her contributions; and 4) weighting the user’s available questions by her expertise. Appendix A.2 provides more details on the construction of this variable. On average, users have approximately 8,000 available questions on a given day, ranging from 22 to 21,744 (note that at the time of data retrieval, out of the 77,357 questions, 38,015 did not have an accepted answer). Table 2 reports some descriptive statistics.

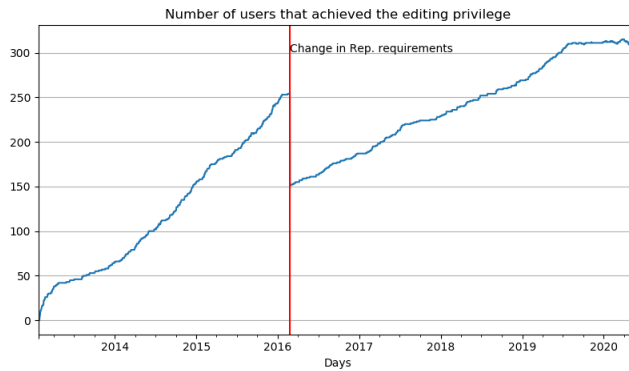


Figure 3: The number of users that have obtained authority over the editing task. In February 2016, an increase in the requirement of points to obtain this privilege induced the loss of the privilege for some users.

	Periods active	Questions	Answers	Edits	Avg Quality	Avg Availability of Q.
count	9797.00	9797.00	9797.00	9797.00	9797.00	9797.00
mean	713.16	2.90	11.73	0.83	14.18	7985.31
std	722.51	26.88	94.31	15.81	1.62	5168.64
min	0.00	0.00	1.00	0.00	6.01	22.73
25%	92.00	0.00	1.00	0.00	13.43	3770.96
50%	414.00	0.00	1.00	0.00	13.92	7532.39
75%	1183.00	0.00	4.00	0.00	14.47	11277.32
max	2685.00	906.00	4173.00	1174.00	44.96	21744.97

Table 2: Descriptive statistics across active users. The columns report, in order, the number of periods spent on the platform (it does not control for right censoring), the number of answers, questions, and edits made, the average answer’s quality, and the mean number of questions available, on average, in a given day.

3.1 User Characteristics

I construct a database to proxy for users' motivation. Users can choose what information to upload on their user pages and the type of information provided can be informative about the underlying motives for participation. The dataset contains information on whether users provide their location, a personal website, LinkedIn profile, and full name. In addition to these dummies, I include measures from the biographical note that users can include, in particular, the number of words and the number of links included. Tables 3 and 4 report summary statistics on these variables for the whole sample of registered users and the sample of users in the panel respectively. It emerges that most users do not have much information, but there is some heterogeneity.

Share of users		net AboutMe	AboutMe	links AboutMe
has full name	34.17 %	mean	21.92	31.12
has website	16.67 %	std	33.27	47.66
has location	31.93 %	min	1.00	1.00
has LinkedIn	1.54 %	25%	5.00	6.00
has bio note	25.82 %	50%	10.00	14.00
has links in bio	4.57 %	75%	25.00	36.00
Sample size	92,853	max	535.00	542.00
		Sample size	23,979	23,979
				4,239

Table 3: Statistics of user characteristics. **(Left)** Share of users that have the given characteristic. **(Right)** Distribution of the number of words in the biographical note (net of *stopwords*), the number of words in the biographical note (all), and the number of links in the biographical note, respectively. The statistics in the right table statistics are conditional on observing a positive value of each measure. The sample includes all users registered on the website on May 31st, 2020

Share of users		net AboutMe	AboutME	links AboutMe
has full name	26.06%	mean	25.4	37.74
has website	24.37%	std	36.8	55.89
has location	40.98%	min	1.0	1.0
has LinkedIn	1.18%	25%	5.0	7.0
has bio	39.33%	50%	13.0	18.0
has links	6.71%	75%	30.0	45.0
Sample size	9,797	max	340.0	510.0
		Sample size	3,853	3,853
				657

Table 4: Statistics of user characteristics, conditional on the user being part of the panel of active users. **(Left)** Share of users that have the given characteristic. **(Right)** Distribution of the number of words in the biographical note (net of *stopwords*), the number of words in the biographical note (all), and the number of links in the biographical note, respectively. The statistics in the right table statistics are conditional on observing a positive value of each measure. The sample includes all users registered on the website on May 31st, 2020

3.2 User Types

If the motives behind participation are strongly heterogeneous, then the incentive effects may differ across users. In this context, a platform incentive design that targets the average consumer may not maximize participation. To address this concern, I use the user characteristics to identify types so that the platform can 1) ex ante assess the composition of the participants and 2) combine different incentive strategies to target different types of users.

The identification of types relies on the assumption that the type of information provided is informative about the broad motives for participation.¹⁶ I use the data summarized in Table 3, which includes dummy variables taking a value equal to 1 if the given type of information is provided; the variables for the number of words and links in the biographical description, which I bin into three categories each; and the year of registration.

The approach implemented is to use a K-means clustering algorithm to cluster observations in a data-driven (or unsupervised) way. The challenges to address are twofold: first, the K-means algorithm does not apply to categorical variables. Second, the algorithm requires as input the number of clusters to be identified, which is not known ex ante. To address these challenges I adopt a Multiple Correspondence Analysis (MCA, which is similar to PCA but applies to categorical variables). The procedure transforms the data by exploiting cross-frequency tables of all variables, and produces new variables (*components*) that aggregate information in a hierarchical way from the first to the last *component*. The new data, while capturing the same information as the original data, provide continuous variables, addressing the first issue. I then decide how many clusters should be selected iteratively. Pick an arbitrary number of clusters k , run the K-means clustering, and plot the individual observations clustered in the k groups on the first two *components* plane. Repeat the procedure with the $k' \neq k$ cluster and evaluate by eye if, in the plot, the new distributions of groups better separate the observations into clusters.¹⁷ If observations do not separate into groups in the plot, then the evaluation must rely on some arbitrariness. This process leads to the identification of the three types. More details on the procedure and on why I adopted this specific approach rather than alternatives are given in appendix A.3.

Tables 5 and 6 provide summary statistics of individual characteristics by type for the whole sample of registered users and the sample of active users, respectively. It is possible to notice that the main discriminant of types is the number of pieces of information provided. The largest group, including more than half of the registered users, displays little if any information, and I define them as *Anonymous*. A second group includes

¹⁶The idea that users self-select in types by some of their choices on the platform is also adopted by Belenzon and Schankerman (2015), where they infer types from the choice of contributing to more or less open open-source software.

¹⁷This approach is borrowed from the literature on Marketing which focuses on *customer segmentation*. While the literature agrees that the K-Means algorithm is a generally accepted way to identify clusters, there is not a leading method to choose the number of clusters. The option adopted in this paper has been selected because, differently from other approaches, 1) it allows to deal with categorical variables, and 2) is ex ante agnostic on which variables are most relevant to describe the types.

information about themselves (location, biography) but not much information on their lives outside the community. This feature is in contrast to that of users in the last group, who generally provide a personal website and sometimes a LinkedIn profile. I refer to these two groups with the terms *Identifiable* and *Informative* respectively. It is relevant to note that some variables are not informative about the types. Providing a full name is quite homogeneous across types, and the year of registration is also orthogonal information.¹⁸

user type	Num.	Share of users who have ...					
	users	full name	website	location	LinkedIn	bio	links
<i>Anonymous</i>	65134	35.31%	1.65%	9.04%	0.00%	1.73%	0.00%
<i>Identifiable</i>	23260	29.25%	47.18%	85.09%	0.00%	79.82%	2.88%
<i>Informative</i>	4459	43.13%	76.74%	88.99%	32.05%	96.08%	80.06%

user type	stat.	net AboutMe	AboutMe	links AboutMe
		<i>Anonymous</i>	25%	9.00
	50%	15.00	22.00	
	75%	25.00	37.00	
	count	1129.00	1129.00	
	max	208.00	397.00	
	mean	21.80	33.38	
	min	1.00	1.00	
	std	25.03	40.68	
<i>Identifiable</i>	25%	4.00	5.00	1.00
	50%	8.00	11.00	1.00
	75%	17.00	25.00	1.00
	count	18566.00	18566.00	669.00
	max	361.00	505.00	3.00
	mean	15.48	22.51	1.11
	min	1.00	1.00	1.00
	std	23.59	35.79	0.35
<i>Informative</i>	25%	19.00	24.00	1.00
	50%	33.00	44.00	2.00
	75%	61.00	83.00	4.00
	count	4284.00	4284.00	3570.00
	max	535.00	542.00	55.00
	mean	49.90	67.83	3.16
	min	1.00	1.00	1.00
	std	51.56	71.11	4.20

Table 5: User characteristics by user type. The first table reports the share of users, for each type, to have the given information displayed. The second table reports the distribution of the number of words (without and with stop-words) and of the number of links contained in the biographical note (if any).

¹⁸The year of registration is used in the analysis but not reported in the table.

3.3 The Behavior of Different Types

While types are not directly based on behavior, participation patterns are very different across them.¹⁹ In this section, I present descriptive statistics of differences in behavior between types. I use the sample of active users from the panel, corresponding to the sample of table 6. In the sample, 5,414 of type *Anonymous*, 3,705 users are of type *Identifiable*, and 678 of type *Informative*.

Badges

Do types differ in the collection of badges? Badges are virtual medals rewarding the achievement of a performance target. The accumulation of badges may suggest sensitivity to short-term incentives. More challenging targets, silver and gold badges, point even more in this direction, as it is impossible to achieve them without tailoring behavior to the target. Figure 4 shows the average number of badges obtained by users in each group, where the vertical black bars are the standard errors of the means. It suggests that users who are more informative in their user pages are also react more to short-term incentives; *Informative* users obtain more badges than *Identifiable* users, and *Identifiable* users obtain more than *Anonymous* users.

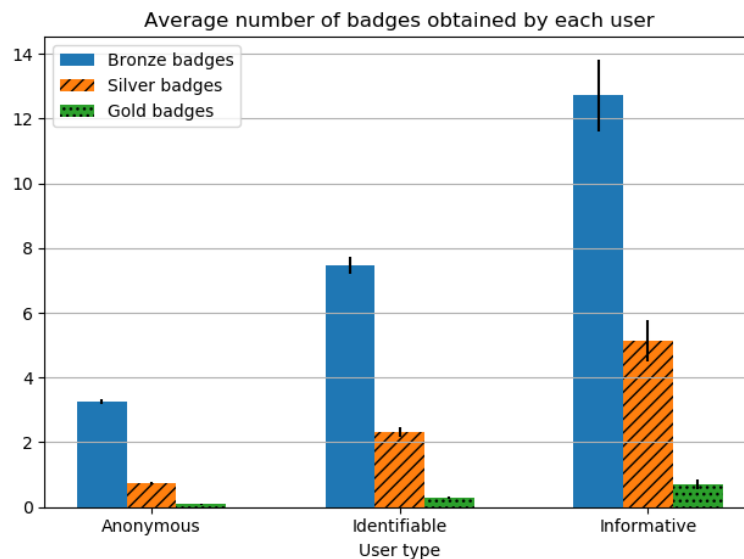


Figure 4: Average number of badges obtained by the active users of each type in the sample users included in the panel. Vertical bars are standard error of the mean.

¹⁹Some literature addresses unobserved heterogeneity by inferring types from observed actions (Arcidiacono and Miller 2011). In this paper, I do not adopt that approach for two reasons. First, I want the platform designer to be able to assess the composition of the community ex ante and in a simple manner. Second, since I observe an unbalanced panel of participation with censoring at the download date, inferring motives from observed actions may be biased by the selection of action I observe.

Time to Reach the Editing Threshold

The types also show heterogeneity in the probability, at each point in time, of having reached the delegation threshold.

I estimate the survival function, where the failing event is the achievement of the threshold number of points. Since in the data the value of the threshold changes, I estimate two survival functions, one for the users that registered before the change and one for users that registered when the threshold was already in its final value. Figure 5 shows the plot of the survival functions for each type. Users of the *Informative* type obtain authority the fastest.

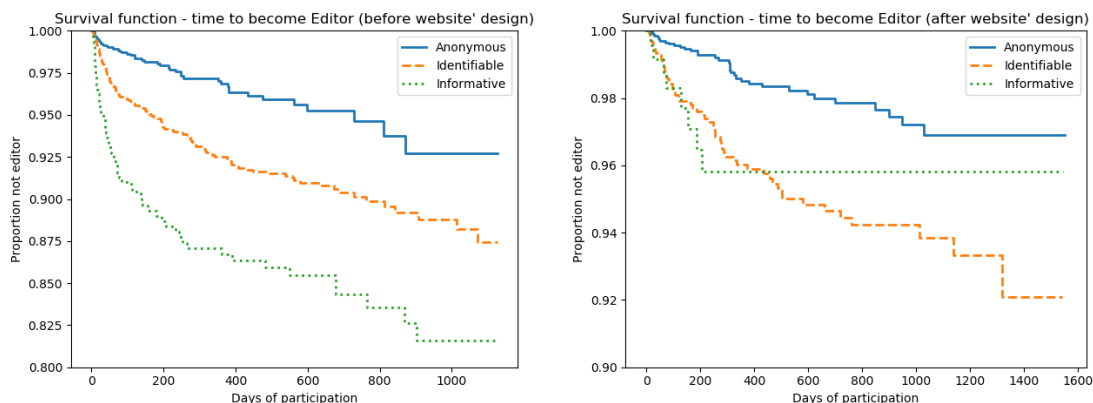


Figure 5: Survival function estimated for users who registered before the threshold change (left) and after the change (right). For the left graph, the data include time series cut at the date when the reputation threshold changed.

Share of Production

Table 7 reports instead the total and average production by type. The marginal contribution of *Informative* users is the highest, followed by that of the *Identifiable* type. The most relevant observation is on direct edits, where the 678 *Informative* users made nearly 60% of the total number of direct edits. Similar patterns can be identified in figure 6, where the share of total content produced each month is plotted by each type. It is possible to notice that *Informative* users reduced their contribution in answering while remaining the main editors of the website over time.

Participation in Elections

The platform seldom organizes election where community members can candidate. Elected users obtain full moderation authority permanently. Participation in an election and winning an election may also reveal information on the motives for participation. Participation in an election may signal that the user has a specific commitment to the



Figure 6: Time distribution of the number of active participants (top left) and of the share of content published by each type. Proceeding clockwise, the graphs report the share of answers, the share of suggested edits, and the share of direct edits.

community. To be a candidate in an election you need to have at least 300 reputation points, while to vote for candidates the requirement is of 150 points²⁰. Figure 7 reports the number of candidates by type, and the number of winners. It is possible to notice that candidates are generally of type *Identifiable* or *Informative*, while elected users are mostly *Informative*.

3.4 Summary of Types

Overall, the profile of each type emerges quite clearly from the descriptive evidence. The online community is in large part populated by *Anonymous* users, who are not particularly active in production. Low production implies a longer average time to achieve the delegation threshold. Nevertheless, the size of this group is such that it still contributes to nearly 30% of the total production of answers. In contrast, *Informative* users are the

²⁰For more details, see <https://stackoverflow.blog/2010/12/02/stack-exchange-moderator-elections-begin/>

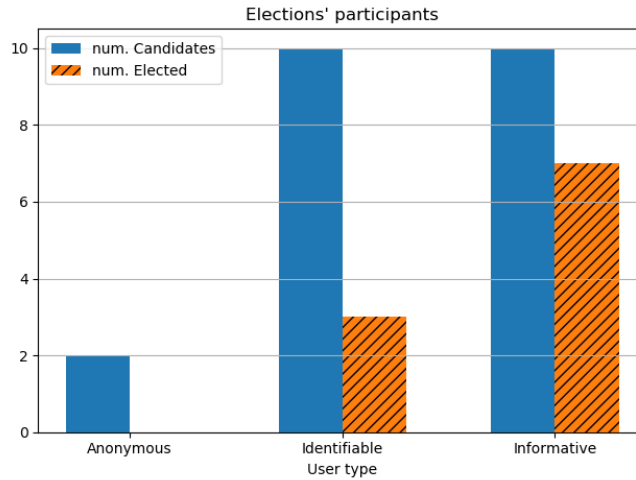


Figure 7: Number of candidates and number of winners of elections, by type.

smallest group of members but the most active. They provide a great deal of information on their profiles, suggesting important extrinsic motives. They produce the most and provide the majority of the editing activity. Their very high activity may also justify the higher likelihood of winning elections. Finally, *Identifiable* users are in between: they provide some information about themselves but no links or LinkedIn profiles, suggesting that they do not aim to signal outside of the platform. They contribute significantly but still take longer time to achieve the delegation threshold than *Informative* users.

user type	Num. users	Share of users who have ...					
		full name	website	location	LinkedIn	bio	links
<i>Anonymous</i>	5414	24.9%	2.4%	8.53%	0.0%	3.77%	0.0%
<i>Identifiable</i>	3705	25.32%	47.72%	80.08%	0.0%	80.54%	2.05%
<i>Informative</i>	678	39.38%	72.27%	86.43%	17.11%	98.08%	85.69%

user type	stat.	net AboutMe	AboutMe	links AboutMe
		<i>Anonymous</i>	25%	10.00
	50%	16.00	24.00	
	75%	28.00	43.00	
	count	204.00	204.00	
	max	208.00	397.00	
	mean	25.47	40.69	
	min	1.00	1.00	
	std	29.84	50.51	
<i>Identifiable</i>	25%	5.00	6.00	1.00
	50%	9.00	13.00	1.00
	75%	21.00	31.00	1.00
	count	2984.00	2984.00	76.00
	max	338.00	505.00	2.00
	mean	17.82	26.95	1.07
	min	1.00	1.00	1.00
	std	25.18	39.99	0.25
<i>Informative</i>	25%	23.00	32.00	1.00
	50%	40.00	56.00	2.00
	75%	71.00	104.00	3.00
	count	665.00	665.00	581.00
	max	340.00	510.00	55.00
	mean	59.39	85.23	3.18
	min	1.00	1.00	1.00
	std	57.62	85.84	4.92

Table 6: User characteristics by user type, for sample of active users in the panel. The first table reports the share of users, for each type, to have the given information displayed. The second table reports the distribution of the number of words (without and with stop-words) and of the number of links contained in the biographical note (if any).

Type	Num Users	num answers		num suggested edits		num direct edits	
		Total	Avg. per user	Total	Avg. per user	Total	Avg. per user
<i>Anonymous</i>	5414	32511.0	6.00	309.0	0.06	465.0	0.09
<i>Identifiable</i>	3705	63500.0	17.14	836.0	0.23	2272.0	0.61
<i>Informative</i>	678	18915.0	27.90	264.0	0.39	4022.0	5.93

Table 7: Total and average users' production by type.

4 Reduced-Form Analysis

In this section, I provide reduced-form evidence of a positive *static incentive* effect.²¹ In other words, I test the hypothesis that users are more willing to contribute when endowed with authority and find that users are significantly more willing to contribute when they obtain more authority over an action. This result is confirmed by the literature on power; one study states, “[...] Generally, research has shown that power increases an action orientation and, thus, leads directly to the taking of action for those who possess it [...]” (Sturm and Antonakis 2015). Only actions concerned with variation in authority see a significant change. A comparable action does not increase or decrease significantly, suggesting the absence of both complementarity and substitutability effects.

The analysis proceeds by initially presenting the effect of removing the *static incentive*, by looking at participation when users lose authority, and then describing the effect of introducing the *static incentive*, observing behavior when users gain authority.

4.1 Loss of Authority

Due to a change in reputation requirements, users could lose their editing privilege. In practice, before February 2016, authority on editing was allocated when the user reached 1000 points, but after February 2016, the requirement was 2000 points. Every user who had a number of points in between 1000 and 2000 on that date lost editing authority.²²

Users partly anticipated the change in reputation requirements. At a previous date, the *graduation* date, the platform stepped into its last phase of development, and users knew at that point that the reputation requirements were going to change. To account for this anticipation, I select users who participated on the website at the time of the *graduation* date and exclude those who registering later. I then select users at risk of being affected by the change in reputation requirements, i.e., users who had less than 2000 points on the *graduation* date. These users knew that if they had 1000 points or may have reached 1000 points but did not have 2000 points, they would lose their editing privilege.

Figures 8 and 9 show the share of users out of the considered sample making a positive number of edits and answers, respectively. Different colors and patterns separate participants based on the number of points accumulated in the given week. The orange and striped sections identify the share of users that, if the change in reputation requirements were to happen in that given week, would lose editing authority.

It is possible to see that users who lose the privilege stop making edits, as the orange and striped sections disappear after the change in reputation requirements. In other words, the share of users who did not reach the new threshold before the change but had

²¹The *dynamic incentive* effect cannot be identified in a reduced-form analysis with standard tools such as difference-in-difference and regression discontinuity. Reduced-form models cannot account for forward-looking behavior, and their identifying assumptions do not hold. Goes et al. (2016) claim to identify the dynamic incentive effect in reduced-form analysis relying on functional form assumptions and modifying the data to account for forward-looking behavior.

²²Please refer to section 3 for more details.

already reached the previous threshold do not participate in the next four months after the change.²³ The share of users who had reached the new threshold or had never reached the previous threshold remains positive after the change (except for the first week after the change). The same pattern is not observed for answering behavior. The users who lost the privilege maintain a similar participation level after the change in reputation requirements, independent of the reputation point bracket to which they belonged.




Figure 8: Number of users making a positive number of edits out of those having fewer than 2000 points at the graduation week.




Figure 9: Number of users making a positive number of answers out of those having fewer than 2000 points at the graduation week.

²³This phenomenon is not driven by the fact that everyone either reached 2000 points or never reached 1000 points. Figure 26 in the appendix shows that every week had a positive number of users with between 1000 and 2000 points, even when conditioning on the selected sample.

4.2 Gain of Authority

To test presence of a *static incentive* effect, I also look at contribution behavior when users obtain authority. I exploit the discontinuity created by the allocation of authority and implement a regression discontinuity design.²⁴ The outcome variable is either the number of edits published or the number of comments published. Editing and commenting are comparable actions because 1) neither are a main driver of the accumulation of points and 2) they both share the purpose of improving or helping improve existing content. They differ in the treatment status. Once users achieve the threshold number of points, they acquire editing authority. At that same threshold instead, commenting authority is not affected (users already have full authority to comment).

I estimate the following specification, in which the running variable is the number of reputation points:²⁵

$$Y_{it} = \alpha_i + \gamma_t + \beta_{r_{it}-\bar{R}} + a_{m_{it}} + b_{c_{it}} + \varepsilon_{it} \quad (1)$$

where Y is either the number of edits or the number of comments made, depending on the outcome of interest. α_i identifies the user fixed effect, γ_t the week fixed effect, r_{it} the number of reputation points that user i has in period t (binned in 50-point intervals), and \bar{R} the number of reputation points required to obtain editing authority. Note that this value depends on the calendar date since it is set at 1000 points before February 2016 and at 2000 points after. The parameters of interest are $\{\beta_{r-\bar{R}}\}_{\forall r}$, which identify the fixed effects of being $r - \bar{R}$ points away from the threshold \bar{R} . Note that while these fixed effects are not period-specific, the unit of observation is still a week. This means that the fixed effects capture a **weekly average** number of edits (or comments) at a given reputation point interval.²⁶ Finally, I include a dummy equal to 1 if the user is an elected moderator in time t and a dummy equal to 1 if the user is a candidate in a moderator election in time t (a_m and b_c). ε_{it} is an error term.

In Figure 10 I report the estimates for $\{\beta_{r-\bar{R}} | r - \bar{R} \in [-6, 6]\}$. The outcome variable is standardized to enable comparisons between editing and commenting. On the x-axis then, the value at 0 identifies the fixed effect estimate of having a number of points included in $[\bar{R}, \bar{R} + 50)$; the value at 1 identifies the fixed effect of having between $\bar{R} + 50$ and $\bar{R} + 100$ points; and so on.

Identification of the effect relies on the assumption that when users have a number of points in the neighborhood of the threshold, the only variable affecting behavior is

²⁴This approach is sometimes called an event study with two-way fixed effects.

²⁵I do not use time as a running variable for two main reasons. The first is that users are aware of the allocation rule, so they may adjust their behavior to receive the privilege sooner or later. The treatment date is then endogenous. The second reason is more technical: Sun and Abraham (2020) show that in an OLS regression with individual fixed effects, time fixed effects, and relative time fixed effects (i.e. fixed effects for the n^{th} period before or after the treatment), trends before and after the treatment date are not identified. If the treatment were completely unexpected, then the researcher would only be interested in the effect at the treatment time. In my context, there is anticipation, and it is therefore relevant to identify the trends.

²⁶If the observational units were the reputation point intervals themselves, the total amount of editing would depend on the time “spent” on each interval, and, as a consequence, on the rate of answering accumulating points.

the acquisition of the privilege. It is possible to see that the number of edits made increases significantly when users have a number of reputation points just above the threshold. The pattern of the number of comments made does not seem to depend on the threshold. This comparison suggests that authority increases the willingness to participate only for the action in which the users receive more authority, and there are no clear spillovers over the effort made in similar actions. In section A.4 of the appendix, I provide some robustness checks that look at effort levels around the precedent and the following privilege.




Figure 10: Estimates for the fixed effects of being in the n^{th} reputation point interval above or below the threshold, that is, the parameters $\{\hat{\beta}_{r-\bar{R}} | r - \bar{R} \in [-6, 6]\}$ in the regression specification 1. Vertical bars are confidence intervals. Outcome variables are the standardized number of edits (circles) and number of comments (squares).

Heterogeneity

Do these effects differ across the different types of users? Figure 11 (graph above) reports the estimates for the editing activity for the different types of users separately. Note that outcome variables are standardized within each user type. Estimates seems to suggest that the effect is strongest for *Anonymous* users, who increase the number of edits immediately after the threshold. *Informative* users increase their contribution significantly when they have more than 150 points above the threshold requirement, which suggest a long term effect of authority. The identifying assumptions are anyway stronger moving away from the threshold, and the causal interpretation is less reliable. Even when accounting for heterogeneity, the number of comments still does not seem to depend on the threshold. Figure 11 (graph below) reports the estimates by type when the outcome variable is the number of comments.

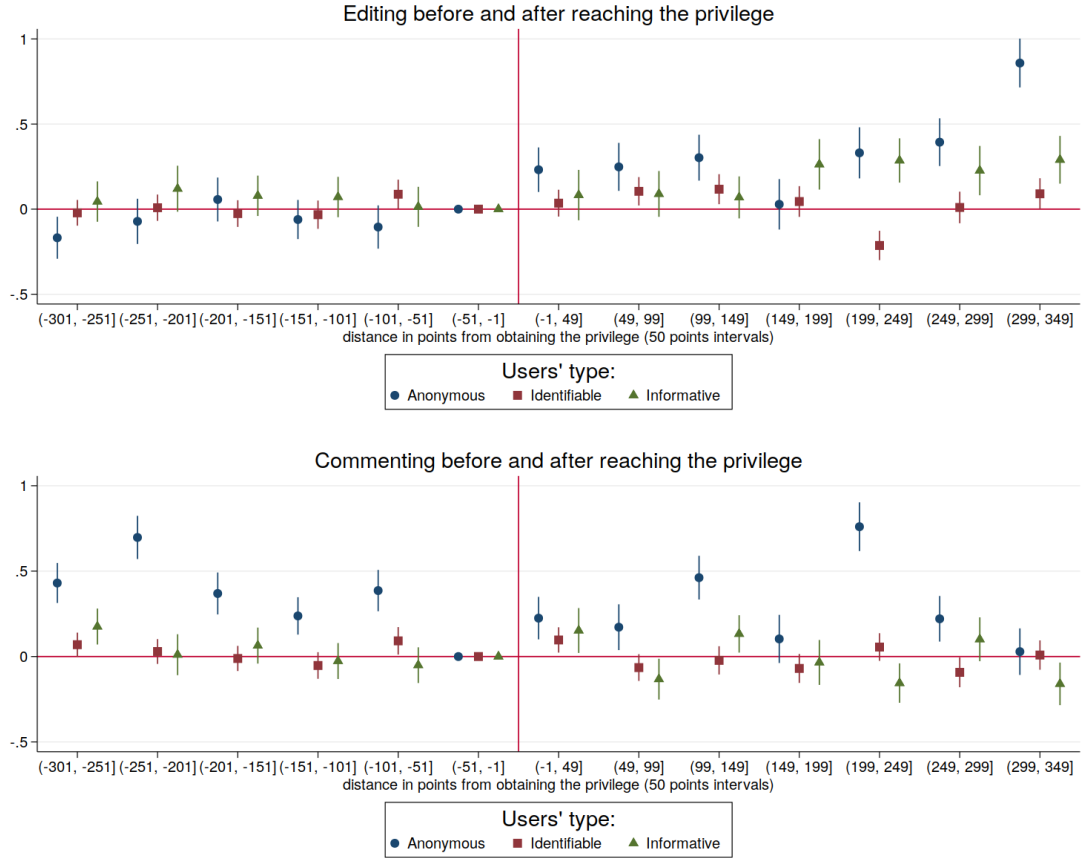


Figure 11: Estimates for the fixed effects of being in the n^{th} reputation point interval above or below the threshold, that is the parameters $\{\hat{\beta}_{r-\bar{R}}|r - \bar{R} \in [-6, 6]\}$ in the regression specification 1. The outcome variable is the standardized (by type) number of edits (graph above) or the standardized (by type) number of comments (graph below). Vertical bars are confidence intervals. Shapes differentiate the types *Anonymous* (circles), *Identifiable* (squares), and *Informative* (triangles).

5 Dynamic Discrete Choice Model

The reduced-form evidence tests for the presence of *static incentive* effects. Nevertheless, it has multiple limitations. First, it does not allow us to compare the incentive effect of allocating authority relative to other types of motives. Second, it does not test for or quantify the *dynamic incentive* effect. Finally, it does not allow us to simulate counterfactual behavior.

To overcome these limitations, I develop a dynamic discrete choice model that studies intertemporal choices and accounts for forward-looking behavior. Dynamic discrete choice models estimate preference parameters based on the concept of *revealed preferences*, that is, the assumption that choices are the outcome of (random) utility maximization and, as such, provide information on users' preferences. In the context of participation in online communities, users choose their effort to contribute to the platform. Their choice depends on the cost of effort, net of choice's intrinsic utilities, and expected future benefits. Benefits could be, for instance, a certain number of reputation points or the achievement of authority.²⁷

5.1 Model setup

In each period, the user decides whether to participate in the online community, and if she does, she decides her effort levels for two tasks, answering and editing. Effort is defined as a combination of the quantity and quality of answers and the quantity of edits. An action choice in period t is then a vector:

$$\boldsymbol{\alpha}_t = \begin{bmatrix} A_t \\ Q_t \\ E_t \end{bmatrix} \ni \mathcal{A}$$

where A indicates the quantity of answers, Q denotes the average quality of answers, and E indicates the quantity of edits. \mathcal{A} represents the choice set, including all possible combinations of effort levels in the two tasks.²⁸ Choices affect utility in two ways: first, the user pays the cost of effort, net of all the benefits that the actions provide in the given period. Second, choices affect the transition and future realizations of states.

The net cost of effort is specific to the action undertaken. The net cost of answering for user i in period t is defined as

$$C_{it}^A \equiv Q_{it} + A_{it}^{scarcity_{it}}$$

where $A^{scarcity}$ is the number of answers raised to a measure of the scarcity of questions to answer. The variable *scarcity* captures the inverse of the availability of questions to

²⁷Note that the application of dynamic discrete choice models to this context is conceptually similar to works that study dynamic investment decisions with discrete choice models. A typical application is the application to human capital investment decisions. Examples of this literature are [Arcidiacono, Aucejo, Maurel, and Ransom \(2016\)](#), [De Groote \(2019\)](#).

²⁸Effort levels are discretized to have 21 possible combinations of the quantity of answers, the quality of answers, and the quantity of edits. More details are provided in [appendix A.6](#).

answer. This variable is user-specific as it accounts for the topics the user can address. More details about the construction of this variable are available in appendix A.2. The *scarcity* variable takes values in $[1, \infty]$ so that when there are many questions to answer, the cost tends to be linear, while with fewer questions available, the cost becomes increasingly convex in the quantity of answers.

The net cost of editing is instead the number of edits that the user decides to make:

$$C_{it}^E \equiv E_{it}$$

The motives of participation directly included in the utility function are those directly related to the accumulation of reputation points: the number of accumulated reputation points themselves (R), a dummy equal to 1 if the individual has control over editing and 0 otherwise (*Authority*), and a variable with the cumulative number of privileges obtained by the user ($cumT$ for “cumulative thresholds”) to control for the possibility that users do not value authority per se but rather virtual vertical rewards (either perceiving them as virtual promotions or as sequential targets in a game). All other motives driving participation are assumed to not be correlated with the accumulation of points. They are captured by a choice-specific preference shock (ε). Note that unobserved motives may affect the cost of participation. The marginal cost of answering and editing is then net of all benefits from participation. These include the intrinsic value of participation (“having fun” in contributing per se), altruism, and reciprocity.

The per period flow utility of user i is then defined as

$$U_{it} = \beta_0 R_{it} + \beta_1 C_{it}^A + \beta_2 C_{it}^E + \beta_3 cumT_{it} + Authority_{it} (\beta_4 + \beta_5 C_{it}^A + \beta_6 C_{it}^E) + \varepsilon_{it}. \quad (2)$$

The parameters β_4 , β_5 , and β_6 together capture the user’s marginal utility from the acquisition of authority and, as a consequence, how the user responds to the dynamic incentive effect. The *static incentive* effect is instead captured by the coefficients β_5 and β_6 . They capture, respectively, a change in the willingness to make answers and edits after the user achieves authority in editing. β_5 corresponds to a spillover effect on answering, while β_6 corresponds to a direct effect of participation on the task endowed with authority. The latter coefficient captures the same incentive effect that was observed in the reduced-form analysis, with the difference that β_6 is an average for the change in participation for all reputation levels greater than the threshold.

The user chooses an optimal sequence of choices to maximize the total sum of the discounted utility from all her periods of participation. Let $\alpha^* \equiv \{\alpha_t\}_{t < T}$ be the sequence of optimal choices, where T is her last period of participation on the website. Then she chooses

$$\alpha^* = \arg \max_{\alpha} \mathbb{E} \left[\sum_{t=1}^T \delta^{t-1} U_{it}(\alpha_t) \right].$$

Timing of a Period

As represented in Figure 12, the timing is as follows. 1) The agent observes the values of the states realized at the end of the previous period, which includes the total number of reputation points she has obtained, the number of points she expects to receive from past efforts, how many privileges she has collected, whether she has already obtained editing authority or not, the availability of questions to answer, and her experience in terms of time spent on the website and the number of contributions. She then forms beliefs over the value of the states that may be realized in the next periods, conditional on the possible choices she could make. 2) She makes an effort decision over two tasks, maximizing her conditional value function. 3) The flow payoff is realized, and 4) at the end of the period, the new value of the states is realized.

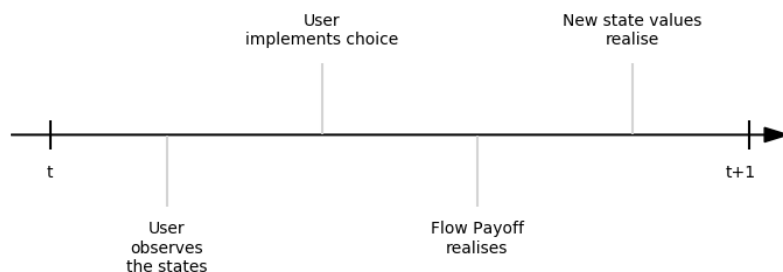


Figure 12: Timing of a period

5.2 Beliefs

Users form beliefs and expectations over the evolution of the state space, given the choices they make.

Evolution of Reputation Points

The points that the user expects to receive in the future depend on current and past actions, particularly on the choice of the quantity and quality of answers, as well as of edits in case they are suggested. Points also depend on edits that the user receives from other community members.²⁹

Consider for simplicity the beliefs that the user forms in the first period of participation t_0 . She considers choosing a triplet $\{A, Q, E\}$ of the number of answers, the average

²⁹For a detailed explanation of the rules to obtain points, please refer to figure 25 in the appendix

quality of answers, and the number of edits. The number of received edits on an answer on the publication day is modeled as a Poisson process in which the mean depends on the answer's quality and the user's experience. Similarly, the number of up-votes and down-votes arriving on the answer at the creation date are modeled as Poisson processes. Let j identify a given answer that the user published on publication day t_0 , which is also the first day of participation of the user. Then, the following random variables are, respectively, the number of modifications that answer j receives in period t_0 and the number of up-votes and the number of down-votes received by j in t_0 :

$$\begin{aligned}\text{Received Edits}_{j,t_0} &\sim \mathcal{P}(\lambda_{E,j,t_0}), \\ \text{Up-votes}_{j,t_0} &\sim \mathcal{P}(\lambda_{U,j,t_0}), \\ \text{Down-votes}_{j,t_0} &\sim \mathcal{P}(\lambda_{D,j,t_0}).\end{aligned}$$

The expected values of these random variables are given by

$$\lambda_{E,j,t_0} = \exp(\beta_0 + \beta_1 Q_{t_0} + \mathbf{EXP}_{t_0} \boldsymbol{\beta}_2) \quad (3)$$

$$\lambda_{U,j,t_0} = \exp(\gamma_0 + \gamma_1 Q_{t_0} + \gamma_2 \lambda_{E,j,t_0} + \mathbf{EXP}_{t_0} \boldsymbol{\gamma}_3) \quad (4)$$

$$\lambda_{D,j,t_0} = \exp(\delta_0 + \delta_1 Q_{t_0} + \delta_2 \lambda_{E,j,t_0} + \mathbf{EXP}_{t_0} \boldsymbol{\gamma}_3) \quad (5)$$

where \mathbf{EXP} is a vector of variables capturing the user's experience. Specifically, it includes the number of days in which the user has been participating on the website, and the cumulative number of answers that she has published.

If the user, in her first period of participation, published A_{t_0} answers, then she will expect to receive by the end of the period

$$\begin{aligned}\Lambda_{U,t_0} &= A_{t_0} \times \lambda_{U,j,t_0} \\ \text{and } \Lambda_{D,t_0} &= A_{t_0} \times \lambda_{D,j,t_0}\end{aligned}$$

which are the total expected amounts of up-votes and down-votes, respectively.

Finally, the number of approved suggested edits is modeled as a binomial distribution:

$$\text{ApprovedEdits}_{t_0} \sim \mathcal{B}(E_{t_0}, \pi)$$

The expected number of points that the user expects to receive at the end of period t_0 is given by

$$\mathbb{E}[\rho_{t_0} | \boldsymbol{\alpha}_{t_0}] = 10 \times \Lambda_{U,t_0} - 2 \times \Lambda_{D,t_0} + 2 \times \pi \times E_{t_0}.$$

The answers produced in period t_0 may also induce the arrival of up-votes and down-votes in the following periods. The process deterministic. Let Δt be the number of days passed from the publication day, such that if $t = t_0 + 1$, then $\Delta t = 1$. Then,

$$\begin{aligned}\lambda_{U,j,t_0+\Delta t} &= \lambda_{U,j,t_0} \times \exp\left(\frac{-\Delta t}{\tau_U}\right) \\ \lambda_{D,j,t_0+\Delta t} &= \lambda_{D,j,t_0} \times \exp\left(\frac{-\Delta t}{\tau_D}\right)\end{aligned}$$

$\lambda_{U,j,t_0+\Delta t}$ is the expected number of up-votes that the answer j , published in t_0 , receives in period $t_0 + \Delta t$, and similarly for down-votes. τ_U and τ_D are parameters.

Given these assumptions, effort induces the arrival of a number of points, which is decreasing over time. If the user chooses positive effort in several periods, these processes aggregate. The user expects to receive an amount of points in the future resulting from all present and past efforts. In general, the expected number of up-votes and down-votes arriving at the end of a given period y are

$$\begin{aligned}\Lambda_{U,t} &= \Lambda_{U,t-1} \times \exp\left(\frac{-1}{\tau_U}\right) + A_t \times \lambda_{U,j,t}(Q_t) \\ \Lambda_{D,t} &= \Lambda_{D,t-1} \times \exp\left(\frac{-1}{\tau_D}\right) + A_t \times \lambda_{D,j,t}(Q_t)\end{aligned}$$

and the expected number of points arriving at the end of the period is

$$\mathbb{E}[\rho_t | \{\alpha_{\tilde{t}}\}_{\tilde{t} \leq t}] = 10 \times \Lambda_{U,t} - 2 \times \Lambda_{D,t} + 2 \times \pi \times E_t.$$

To conclude, let R_t be the cumulative number of points that the user observes to have at the beginning of period t . Then, the user expects to have, at the end of the period

$$\mathbb{E}[R_{t+1} | R_t, \{\alpha_{\tilde{t}}\}_{\tilde{t} \leq t}] = R_t + \mathbb{E}[\rho_t | \{\alpha_{\tilde{t}}\}_{\tilde{t} \leq t}]$$

Evolution of the Experience Variables

The variables for users' experience evolve in a deterministic way. The number of days of participation on the platform increases by one unit each period, while the cumulative number of answers published increases based on the choice of the quantity of answers published.

Evolution of the Scarcity Variable

The availability (or scarcity) of questions to answer evolves in an exogenous way based on the general trend on the platform. Let *avail* be the variable capturing the number of available questions in the platform. Then

$$avail_{it} = avail_{it-1} + \nu_1$$

where ν_1 is identified in reduced form from the linear regression

$$avail_t = \nu_0 + \nu_1 t + \epsilon_t$$

As shown in the graphs in appendix A.2, the availability of questions to answer increases monotonically over time. Users then expect a steady increase, given by an estimated parameter. Note that the rate of increase in availability is not topic-specific.

5.3 Identification

The identification of preference parameters relies on *revealed preferences*. Since choices affect the value of the states, observed choices are informative on what the user cares about. On Stack Exchange, choices affect users' utility in two ways. First, they have a direct impact on the present utility. Direct effects on utility include the cost of effort and the intrinsic value of contributing. They are captured in the utility specification by C^A and C^E . Marginal utilities of these direct effects are identified as in a standard static conditional logit model. Everything else equal, a higher cost of effort implies that the user engages in the action less frequently. Beliefs do not play a role as the costs and intrinsic values do not affect future utilities. Second, choices induce returns in future periods that may affect future utilities, such as the arrival of points and the achievement of authority. These returns do not have any impact on the utility that the user receives in the same period that she makes her contribution choice. In this case, the choice affects only the discounted future payoff. The identification of the marginal utilities of these returns is strictly based on variation across value functions. If users value obtaining these returns, they will be more willing to choose an effort level that allows them to obtain them. That choice would then be observed with a higher frequency. It is relevant to note that the parameters β_0 , β_3 , and β_4 are not identified in a static model.

The computation of the value functions under the different possible choices relies on a technique called *finite dependence* (Arcidiacono and Miller (2011)).³⁰ This approach substantially reduces the computational burden since it allows the estimation without solving the model. The computation of the value functions requires the evaluation of the future expected utility for only a few periods ahead. This approach has a drawback. Whenever returns are not smooth but are step functions, which is the case for *cumT* and *Authority*, their marginal utility is identified only when the user can obtain the given return in the few periods ahead used for the computation of the value function. In practice, only the choices of users who are close to reaching the threshold(s) identify the marginal utility of *cumT* and *Authority*. For these users, certain levels of effort will allow them to obtain a privilege in the following periods, but others will not. The observation of a significant increase in effort when users have a number of points just below the threshold identifies a positive utility for the acquisition of the privilege.

5.4 Estimation

The estimation proceeds in several steps. First, I set the discount factor at 0.95. Second, I estimate in either reduced-form or nonparametrically all parameters not appearing in the utility function. This includes the probability that a suggested edit is approved, the rate of arrival of new questions, the parameters to the predict return for given levels of effort and experience, and the decay rate for the returns on effort. Third, I estimate the preference parameters following Arcidiacono and Miller (2011). The estimation relies on the conditional logit assumption, assuming that the idiosyncratic preference shocks

³⁰The work by Arcidiacono and Miller (2011) is rooted in a large econometric literature. Seminal papers are the works by Rust (1987), Hotz and Miller (1993), and Magnac and Thesmar (2002).

follow an extreme value Type 1 distribution. The derivation of the log-likelihood function is presented in appendix A.5. The algorithm preserves computational feasibility even without binning the state variables, i.e., without reducing the dimensionality of the state space.

6 Results

6.1 First-stage Reduced-Form Estimates

When users decide what action to take, they form beliefs on the arrival of points in the next period. For a given amount of answers and a given quality, they first predict the number of edits that they will receive on the publication date and then the number of up-votes and down-votes that their content will receive in the next period.

The expected number of edits made on an answer, excluding edits made by the author of the answer, is modeled as in equation 3 and table 8 reports the different specification estimates (the specification used in the structural model is the number 2). It is possible to notice that experience is negatively correlated with the arrival of edits. This finding implies that experience captures some skills that the quality variable is not able to measure: even controlling for quality, experienced users produce content that needs to be corrected less.

	(1)	(2)	(3)	(4)	(5)
Received Edits	Poisson	Poisson	Poisson	Poisson	OLS
Answer Quality	-0.0135* (-2.29)	-0.00178 (-0.30)	-0.00414 (-0.70)	-0.00178 (-0.20)	-0.000224 (-0.64)
Experience: num Answers		-0.000382*** (-9.59)	-0.000389*** (-9.79)	-0.000382** (-2.71)	-0.00000871*** (-4.94)
Experience: days in platform		-0.000609*** (-15.85)	-0.000584*** (-14.59)	-0.000609*** (-8.07)	-0.0000208*** (-9.30)
_cons	-2.946*** (-33.72)	-2.788*** (-32.12)	-2.504*** (-25.31)	-2.788*** (-20.59)	0.0598*** (10.40)
<i>N</i>	118552	118552	118552	118552	118552
Year FE	NO	NO	YES	NO	NO
std. err. clustered at author	NO	NO	NO	YES	YES

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 8: Estimates for beliefs on the arrival of edits on the publication day, given the answer’s publication quality and the user’s experience.

Once users have expectations regarding the number of edits they will receive, they predict how many up-votes and down-votes their content will receive. The parameter estimates of equations 4 and 5 are reported in table 9 and 10 respectively. As expected,

higher quality correlates with more up-votes and fewer down-votes. Received edits have a positive coefficient in both cases. One explanation is that edits improve quality, inducing more up-votes, but at the same time, users may want to penalize content of bad quality. Finally, more experienced users can expect more up-votes and fewer down-votes.

	(1)	(2)	(3)	(4)	(5)
Num Up-Votes	Poisson	Poisson	Poisson	Poisson	OLS
Answer Quality	0.0528*** (69.61)	0.0513*** (67.06)	0.0463*** (59.88)	0.0463*** (9.58)	0.0887*** (8.84)
Received Edits	0.457*** (55.61)	0.485*** (58.79)	0.472*** (57.26)	0.472*** (20.99)	0.935*** (15.46)
Experience: num Answers		0.0000465*** (11.79)	0.0000332*** (8.29)	0.0000332 (0.61)	0.0000563 (0.58)
Experience: days in platform		0.000112*** (24.04)	0.000271*** (52.35)	0.000271*** (7.17)	0.000372*** (6.34)
._cons	-0.414*** (-35.47)	-0.469*** (-39.59)	-0.101*** (-6.87)	-0.101 (-1.10)	0.532** (2.97)
<i>N</i>	118552	118552	118552	118552	118552
Year FE	NO	NO	YES	YES	YES
st. err. clustered at author	NO	NO	NO	YES	YES

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 9: Expected number of up-votes arriving on the publication day

When users choose effort, they form beliefs about the number of up-votes and down-votes they will receive not only in the next period but also in the following periods. Effort directly affects the number of up-votes and down-votes in the next period, while in the following periods the number of up-votes and down-votes decrease deterministically following an exponential function, as shown in figures 13 and 14.³¹ The model was estimated via a non-linear fit. Table 11 reports the estimates. In the figures, the red dots are the data values for λ_U and λ_D , and on the x-axis Δt is reported. The figures also compare the chosen model (exponential) with alternatives and present the decay at both daily and weekly levels.

Finally, the last parameter estimated in the first step is the rate of increase in the availability of questions over time. Since the community is increasing and some questions remain unanswered, users expect a steady increase in availability. Estimates are reported in table 12.

³¹The functional form is the standard function to model the amplitude of the cycles of a pendulum. Users' effort corresponds to the strength that starts the oscillation.

	(1)	(2)	(3)	(4)
Num Down-votes	Poisson	Poisson	Poisson	OLS
Answer Quality	-0.0517*** (-14.98)	-0.0488*** (-13.27)	-0.0488*** (-7.31)	-0.00434*** (-7.53)
Received Edits	0.732*** (27.00)	0.665*** (24.28)	0.665*** (18.98)	0.107*** (12.44)
Experience: num. Answers		-0.000236*** (-10.85)	-0.000236* (-2.44)	-0.0000156** (-3.10)
Experience: days in platform		-0.000239*** (-10.79)	-0.000239*** (-4.32)	-0.0000202*** (-4.78)
_cons	-1.664*** (-32.97)	-1.528*** (-28.44)	-1.528*** (-14.74)	0.169*** (17.97)
<i>N</i>	118552	118552	118552	118552
std. err. clustered at author	NO	NO	NO	YES

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 10: Expected number of down-votes arriving on the publication day

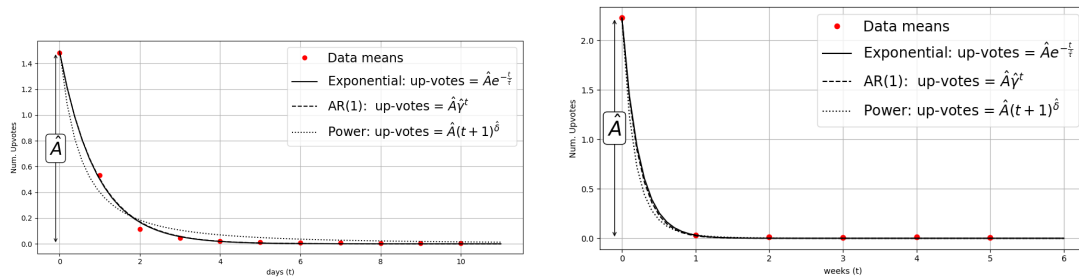


Figure 13: Returns in terms of up-votes from publishing content

6.2 Flow Payoff Parameters

Tables 14 and 15 present the flow payoff parameters obtained from the dynamic discrete choice model. The former presents parameter estimates when the utility function does not include the interaction terms, i.e., does not allow authority to affect the willingness to participate in answering and/or editing. This specification aims to omit the sensitivity to the *static incentive* to focus on the *dynamic incentive*. Under this specification, the coefficient of *Authority* captures the marginal utility from the acquisition of authority. A higher parameter implies a higher willingness to reach the threshold and, therefore, a stronger *dynamic incentive* effect. It is possible to notice that *Anonymous* and *Informative* users value obtaining authority, while the threshold does not motivate *Identifiable* users. Scaling the parameters by the marginal values of points provides a more concrete quantification of the value of authority. Table 13 reports the value of authority in terms of points for each user type separately. It also includes the average number of posts

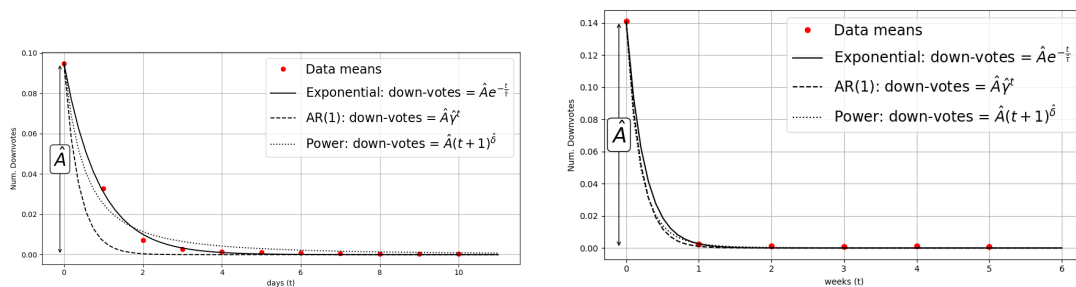


Figure 14: Returns in terms of down-votes from publishing content

period length	τ_U (up-votes)	τ_D (down-votes)
day	0.91	0.89
week	0.23	0.25

Table 11: Estimates of the parameters for the rate of decay on the arrival of up-votes and down-votes on past answers.

(including questions and answers) that a user of the given type had to create to achieve that number of points.

Estimates suggest that only *Anonymous* and *Informative* users are sensitive to the dynamic incentive.

Table 15 reports estimates for the flow utility parameters when the specification includes interaction terms of the variable authority with the net cost of participation. The coefficients of *CA x Authority* and *CE x Authority* capture the sensitivity to the *static incentive*. The former identifies possible changes in the willingness to answer questions, and the latter captures possible changes in the willingness to make edits. It is possible to notice that *Anonymous* users are less willing to contribute to answering after they reach the threshold. There are two possible interpretations of this result: either they lose interest in participating because their main motive was the achievement of authority or they substitute answering with editing. The effect is small: *Anonymous* users are 1% less willing to post answers. The answering behavior of other types of users is not affected by the achievement of the threshold. In contrast, all users are significantly more willing to make edits. *Anonymous*, *Identifiable*, and *Informative* users are 8%, 4%, and 5% more willing to make edits, respectively. In addition to this positive *static incentive* effect on editing, the cost of participation remains high for all users.

period length	Number of additional available questions in the next period
day	13.68
week	95.76

Table 12: Estimates of the increase in availability of answers in each period

User type	coef. authority	value in points	value in actions (avg)
<i>Anonymous</i>	1.5394	252 points	33 posts
<i>Identifiable</i>	0.1702	30 points	4 posts
<i>Informative</i>	1.4503	329 points	28 posts

Table 13: Marginal value of acquiring authority, by type. Proceeding from left to right, the table reports the parameter estimates for the marginal utility of authority, its counterpart value in terms of points (parameter scaled by the coefficient of R), and the average number of posts that a user of the given type would need to create to achieve those points. Posts include answers and questions.

Variables	(no Heterogeneity)	(<i>Anonymous</i>)	(<i>Identifiable</i>)	(<i>Informative</i>)
R	0.0074*** (0.0001)	0.0061*** (0.0005)	0.0056*** (0.0002)	0.0044*** (0.0003)
CA	0.0004* (0.0002)	-0.3669*** (0.0192)	0.00003 (0.0004)	0.0007** (0.0002)
CE	-0.6133*** (0.1661)	-3.3660*** (0.6046)	-4.4860*** (0.3161)	-2.0967*** (0.2319)
cumT	-0.8409*** (0.0205)	-0.4032*** (0.0310)	-0.7842*** (0.0276)	-0.8019*** (0.0548)
Authority	1.2052*** (0.1207)	1.5394*** (0.3577)	0.1702 (0.2536)	1.4503** (0.5118)
N. users	9,783	3,700	5,407	676
Sample size	991,657	471,837	407,098	112,722

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 14: Estimates for the flow payoff parameters considering the whole sample, or estimating separately for each type of user. Standard errors are in parentheses

Variables	(no Heterogeneity)	(<i>Anonymous</i>)	(<i>Identifiable</i>)	(<i>Informative</i>)
R	0.0069*** (0.0001)	0.0064*** (0.0005)	0.0057*** (0.0002)	0.0045*** (0.0004)
CA	-0.0001 (0.0008)	-0.3563*** (0.0196)	.00005 (.0006)	0.0007*** (0.0002)
CE	-10.3311*** (0.4979)	-7.9549*** (0.8927)	-6.1724*** (0.4051)	-5.7740*** (0.4757)
cumT	-0.7745*** (0.0206)	-0.4177*** (0.0322)	-0.7855*** (.028)	-0.7681*** (0.0563)
Authority	1.3162*** (0.1203)	1.5223*** (0.3577)	0.1713 (0.2535)	1.4709*** (0.5118)
CA x Authority	0.0609*** (0.0036)	-0.0048*** (0.0016)	-0.0018 (0.0011)	-0.0008 (0.0014)
CE x Authority	12.2064*** (0.5247)	0.6338*** (0.0593)	0.2507*** (0.0308)	0.2703*** (0.0274)
N. users	9,783	3,700	5,407	676
Sample size	991,657	471,837	407,098	112,722

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 15: Estimates for the flow payoff parameters for the whole sample or by type of user. The specification includes interaction terms of the net costs of actions with the control dummy. Standard errors are in parentheses

7 Counterfactual Analysis: Incentive effect of delegation on contributions

The estimated flow payoff parameters allow us to predict behavior under different delegation designs. Using a counterfactual analysis, I provide evidence of the trade-off faced by the platform when deciding how to allocate authority. In particular, I simulate counterfactual contributions in answering when the performance threshold is set to 1) zero so that everyone is endowed with authority; 2) infinity so that no one will ever obtain authority; and 3) two intermediate levels, where authority is allocated depending on the user reaching a pre-established positive but finite performance level. Note that all these scenarios are realistic. Wikipedia is a leading example of the case in which agents have full authority. On Wikipedia, every internet user is allowed to contribute by writing new articles and modifying existing content. On the other hand, most online retailers do not allow users to modify reviews provided by other contributors. In this case, there is no delegation. Users can sometimes rate existing reviews or flag inappropriate reviews but have no right to modify them. Stack Exchange instead represents an example of the intermediate case in which the allocation of authority depends on the achievement of a performance threshold.

To simulate contribution levels, I cannot rely on the results of [Arcidiacono and Miller \(2011\)](#). Some additional restrictions are therefore necessary to achieve computational feasibility. The approach used is to solve by backward induction the maximization problem, assuming that users participate in the website for a fixed amount of time.

The simulation proceeds in three steps. First, I compute the choice-specific transition probabilities. These are matrices mapping each possible combination of state values to future combinations of state values and providing the probability distribution of future state values, given a choice made. The state variables that I consider are the number of accumulated reputation points, the expected up-votes and down-votes realizing do to effort in the past, the availability of questions to answer, and the variables capturing experience: the number of answers already made and the number of days of participation on the platform. Details on the restrictions to the dimensionality of the state variables are in [appendix A.8.1](#). Second, I compute the value function backward, starting from the last period. I assume that users participate for 100 periods and then exit the platform definitively. Finally, in the third step, I forward-simulate the decisions in each period.

Each simulation is characterized by a different performance threshold.³² The considered threshold levels are 0, 500, 1000, and 99999. Since I set the maximum number of points that users can achieve to 1500, in the last scenario none of the users obtain authority.³³ [Figure 15](#) reports the simulated average number of answers made under the different delegation thresholds. The estimates used for the simulations are those of the utility function that includes the interaction terms. It reports the average number of answers posted by users of each type. On average, users reach the threshold at the vertical line of same color and pattern as the simulation, and reach 1500 points at the

³²In a given simulation, the performance threshold is fixed i.e. does not change across time.

³³Please refer to [appendix A.8.1](#) for more details on the accumulation of points in the simulations.

vertical lines with dots. Since users cannot accumulate more than 1500 points, after that line the marginal intrinsic utility from additional points is zero.

It is possible to see that the incentive effects induce very heterogeneous responses across types. *Anonymous* users participate very little, even though they should be the most sensitive to the incentives. The low production is caused by their high cost of participation, which, in the simplified context of the simulation, is not compensated by the incentive effects. The *Identifiable* users are instead not sensitive to the dynamic incentive effect. Their participation is not much affected by changes in the performance threshold. Finally, *Informative* users are instead very reactive to the incentive design. Their participation increases faster when approaching the threshold, while it is slacker in the case of full or no delegation.

To understand how the different incentive designs translate into final production on the platform, I sample users of each type following the proportion appearing in the real data. This corresponds to 55% of *Anonymous* users, 38% of *Identifiable* users, and 7% of *Informative* users. As shown in Table 16, the platform reaches the highest level of production when the performance threshold is set at 500 points. Most of the increase is attributable to the *Informative* users. Since they are a small share of the participants, the final change in production is limited. Figure 16 shows how the production of answers occurs during the lifetime of the platform.

Performance required	Answers	Change	<i>Anonymous</i>	<i>Identifiable</i>	<i>Informative</i>
0 Points	12562.0		92	10967	1503
500 Points	13374.0	+6.46%	+13.04%	+1.6%	+41.52%
1000 Points	13300.0	+5.87%	+13.04%	+2.43%	+30.54%
NO Delegation	12886.0	+2.58%	+13.04%	+2.01%	+6.12%

Table 16: Total number of answers produced on the platform under the different delegation designs. Columns report the number of answers produced and the relative change compared to the full delegation design, overall and by type.

These results show that the platform could exploit the *dynamic incentive* effect to increase the number of answers provided. Nevertheless, it faces a trade-off: since users are more willing to edit when endowed with authority, postponing delegation induces a lower production of edits. Table 17 reports the contributions made through editing under the different delegation designs.³⁴ It is relevant to note that any design different from the full delegation scenario produces fewer edits. The full delegation design is the setting that maximizes the *static incentive* effect.

³⁴Note that the sample used in the simulation is much smaller than the actual number of participants. In reality, the number of edits would be more substantial.

Performance required	Edits	Change	<i>Anonymous</i>	<i>Identifiable</i>	<i>Informative</i>
0 Points	16.0		2	11	3
500 Points	7.0	-56.25%	-50.0%	-54.55%	-66.67%
1000 Points	8.0	-50.0%	-50.0%	-54.55%	-33.33%
NO Delegation	7.0	-56.25%	-50.0%	-45.45%	-100.0%

Table 17: Total number of edits produced on the platform under the different delegation designs. Columns report the number of edits produced and the relative change compared to the full delegation design, overall and by type.




Figure 15: Average number of answers that users of each type make under different delegation designs. The x-axis reports weeks of participation on the platform. The vertical lines of the same pattern as the series identify the average period in which users achieve the threshold. The vertical lines with dots (named as re. cap. in the legend) identify the average period in which users reach 1500 reputation points, which is the maximum number of points that they can obtain in the simulation framework. After those lines, users cannot accumulate more points. Anonymous users never reach this limit. After 100 periods (weeks), the users exit the platform.

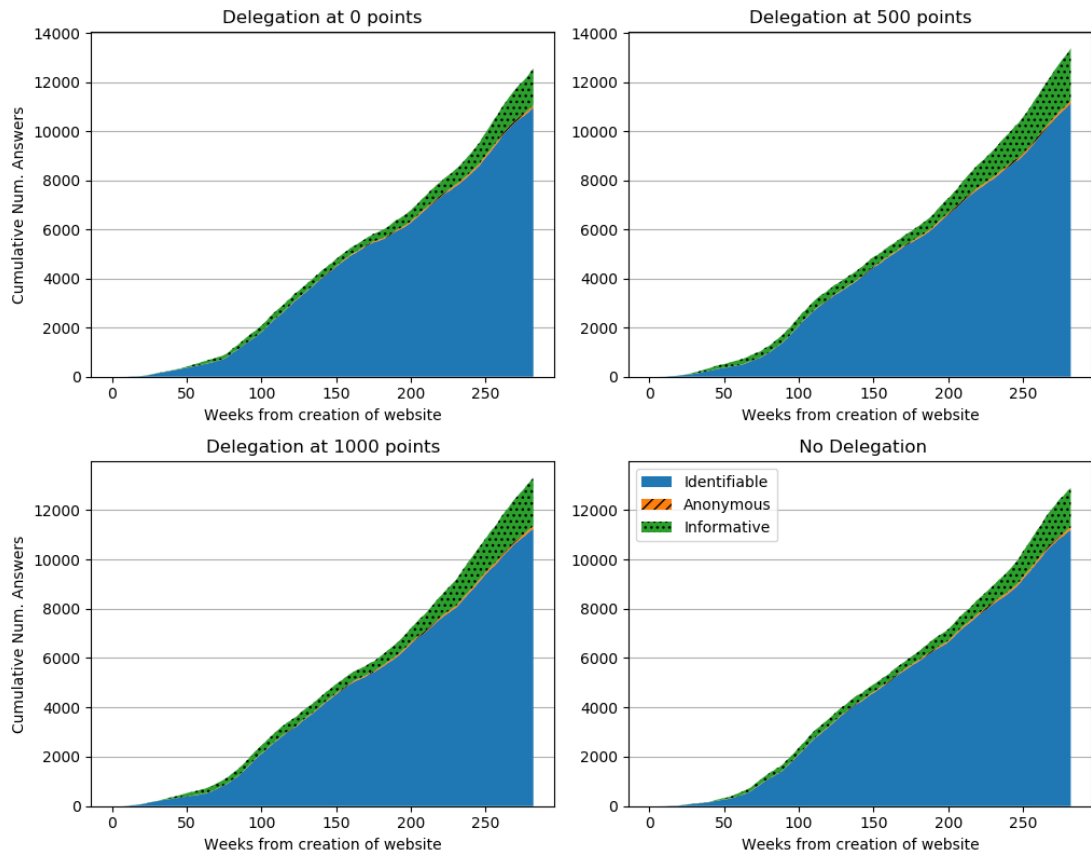


Figure 16: Cumulative number of answers made during the life of the platform, by type and delegation design.

8 Conclusion

In this paper, I show that users participating in online communities value the allocation of control rights and authority. I then study the implications for the platform design, investigating the incentive role of delegation.

First, the willingness to contribute to a given task depends on the level of autonomy and authority the user has with regard to the task. The paper indeed finds that users post significantly more edits if their edits are directly implemented and do not require third-party approval. To my knowledge, this is novel evidence in real data and contributes to the growing literature that studies the role of autonomy and authority for incentives and the optimal delegation structure (Liberti 2018, Bandiera et al. 2020). Interestingly, the allocation of authority on a task does not seem to affect contributions in other tasks. The paper finds evidence that the production of both comments and answers is not affected by the allocation of authority. These results contribute to the literature on multitasking (Holmstrom and Milgrom 1991), suggesting that incentives may not backfire in these contexts. In contrast to the others, *Anonymous* users slightly substitute answering with editing when they have more authority.

Second, the paper finds heterogeneity in the value of acquiring authority. *Anonymous* and *Informative* users are motivated by the acquisition of authority and increase their contribution when approaching the threshold needed to acquire it. In contrast, *Identifiable* users seem to be motivated by other factors.

The results regarding the preference for authority have important implications for platform design. Regarding the moderation task, the platform can incentivize participation via the *static incentive*. If the objective is to maximize contributions in the moderation task, the platform would need to provide authority to all participants from the registration date. The *dynamic incentive* has no impact on participation in editing and, as a consequence, there is no good reason to delegate authority based on performance. This is because suggested edits provide very few points and cannot be the main tool to reach the performance threshold. The scenario of full delegation would be comparable to the design adopted by Wikipedia. Nevertheless, delegation and commitment to allocating authority based on performance incentivize answering. The effect is driven mainly by *Informative* users, who increase their participation by 40% when incentivized via the *dynamic incentive*. The answering task is not much affected by the *static incentive* instead. If the platform aims to maximize the number of answers produced on the platform as much as possible, it should delay delegation and commit to providing authority based on performance in answering. The optimal performance threshold depends on users' cost of answering, and the average time users plan to stay on the platform. The optimal organizational design depends on 1) the type of action the platform needs to incentivize and 2) the composition of the community. On the first dimension, I show that participation in the different tasks (answering and editing) depends on different incentives. On the second dimension, I show that the sensitivity to the incentives is heterogeneous. If the community is not populated by *Informative* users, the platform's trade-off simplifies: the *dynamic incentive* becomes irrelevant and full delegation emerges

as leading strategy. Otherwise, the platform would be better off in targeting different types with different incentives. It would allocate authority based on performance for *Informative users*, and full authority to the other types. This paper provides a way for the platform to identify user types ex-ante, before observing their actions. It then allows the platform to assess the composition of the community and adopt the most suitable design.

References

- Arcidiacono, P., E. Aucejo, A. Maurel, and T. Ransom (2016). College attrition and the dynamics of information revelation. *Working Paper*. [25](#)
- Arcidiacono, P. and R. A. Miller (2011, November). Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. *Econometrica* 79(6), 1823–1867. [4](#), [15](#), [30](#), [37](#), [55](#), [56](#), [57](#)
- Auriol, E. and R. Renault (2008, Spring). Status and incentives. *The RAND Journal of Economics* 39(1), 305–326. [2](#)
- Baker, G. P., M. C. Jensen, and K. J. Murphy (1988, July). Compensation and incentives: Practice vs. theory. *The Journal of Finance* 43(3), 593–616. [5](#)
- Bandiera, O., M. C. Best, A. Q. Khan, and A. Prat (2020). The allocation of authority in organizations: A field experiment with bureaucrats. *Working paper*. [5](#), [42](#)
- Bartling, B., E. Fehr, and H. Herz (2014, November). The intrinsic value of decision rights. *Econometrica* 82(6), 2005–2039. [5](#)
- Belenzon, S. and M. Schankerman (2015). Motivation and sorting of human capital in open innovation. *Strategic Management Journal* 36, 795–820. [2](#), [13](#)
- Benson, A., D. Li, and K. Shue (2019, November). Promotions and the peter principle. *The Quarterly Journal of Economics* 134(4), 2085–2134. [5](#)
- Besley, T. and M. Ghatak (2008, May). Status incentives. *The American Economic Review: Papers & Proceedings* 98(2), 206–211. [2](#)
- Bester, H. and D. Kräbmer (2008, Autumn). Delegation and incentives. *RAND Journal of Economics* 39(3), 664–682. [2](#)
- Blanes I Vidal, J. and M. Möller (2007, Summer). When should leaders share information with their subordinates? *Journal of Economics & Management Strategy* 16(2), 251–283. [2](#)
- Bruneel-Zupanc, C. (2020). Discrete-continuous dynamic choice models: Identification and conditional choice probabilities estimation. *Working Paper*. [58](#)

- Chen, W., X. Wei, and K. Zhu (2017). Engaging voluntary contributions in online communities: A hidden markov model. *MIS Quarterly* 42(1), 83–100. 2
- Chen, Y., F. M. Harper, J. Konstan, and S. X. Li (2010, September). Social comparison and contributions to online communities: A field experiment on movielens. *American Economic Review* 100(4). 2
- Chen, Y., T.-H. Ho, and Y.-M. Kim (2010). Knowledge market design: A field experiment at google answers. *Journal of Public Economic Theory* 12(4). 2
- De Groote, O. (2019). A dynamic model of effort choice in high school. *Working Paper*. 25
- Fairburn, J. A. and J. M. Malcomson (2001, January). Performance, promotion, and the peter principle. *The Review of Economic Studies* 68(1), 45–66. 5
- Fehr, E., H. Holger, and T. Wilkening (2013, June). The lure of authority: Motivation and incentive effects of power. *The American Economic Review* 103(4), 1325–1359. 5
- Gallus, J. and B. S. Frey (2016, August). Awards: A strategic management perspective. *Strategic Management Journal* 37(8), 1699–1714. 2
- Gambardella, A., C. Panico, and G. Valentini (2015). Strategic incentives to human capital. *Strategic Management Journal* 36(1), 37–52. 2
- Gibbons, R., N. Matouschek, and J. Roberts (2013). Decisions in organizations. In R. Gobbons and J. Roberts (Eds.), *The Handbook of Organizational Economics*, Chapter 10, pp. 373–431. Princeton University Press. 2, 7
- Gibbons, R. and M. Waldman (1999). Careers in organizations: Theory and evidence. In O. C. Ashenfelter and D. Card (Eds.), *Handbook of Labor Economics*, Volume 3, part B, Chapter 36, pp. 2373–2437. North-Holland. 5
- Gillespie, T. (2018). *Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions that Shape Social Media*. Yale University Press, New Haven, CT. 3
- Goes, P., C. Guo, and M. Lin (2016, September). Do incentive hierarchies induce user effort? evidence from an online knowledge exchange. *Information Systems Research* 27(3), 497–516. 2, 20
- Greenacre, M. and J. Blasius (Eds.) (2006). *Multiple Correspondence Analysis and Related Methods*. Statistics in the Social and Behavioral Sciences Series. Chapman & Hall/CRC. 50
- Hagberg, A. A., D. A. Schult, and P. J. Swart (2008, 08). Exploring network structure, dynamics, and function using networkx. In G. Varoquaux, T. Vaught, and J. Millman (Eds.), *Proceedings of the 7th Python in Science Conference (SciPy2008)*, Pasadena, CA USA, pp. 11–15. 64

- Holmstrom, B. and P. Milgrom (1991). Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *Journal of Law, Economics, & Organization* 7, 24–52. [42](#)
- Hotz, V. J. and R. A. Miller (1993, July). Conditional choice probabilities and the estimation of dynamic models. *Review of Economic Studies* 60(3), 497–529. [30](#), [56](#)
- Hunter, J. D. (2007). Matplotlib: A 2d graphics environment. *Computing in Science & Engineering* 9(3), 90–95. [64](#)
- Jeppesen, L. B. and L. Frederiksen (2006, January-February). Why do users contribute to firm-hosted user communities? the case of computer-controlled music instruments. *Organization Science* 17(1), 45–63. [2](#)
- Jin, J., Y. Li, X. Zhong, and L. Zhai (2015). Why users contribute knowledge to online communities: An empirical study of an online social q&a community. *Information & Management* 52, 840–849. [2](#)
- Lazear, E. P. and S. Rosen (1981, October). Rank-order tournaments as optimum labor contracts. *Journal of Political Economy* 89(5), 841–864. [2](#)
- Lê, S., J. Josse, and F. Husson (2008). Factominer: An r package for multivariate analysis. *Journal of Statistical Software, Articles* 25(1), 1–18. [64](#)
- Lewis, G. and G. Zervas (2019). The supply and demand effects of review platforms. *Working Paper*. [2](#)
- Liberti, J. M. (2018, August). Initiative, incentives, and soft information. *Management Science* 64(8), 3469–3970. [5](#), [42](#)
- Luca, M. (2011). Reviews, reputation, and revenue: The case of yelpcom. *Harvard Business School Working Paper*. [2](#)
- Ma, M. and R. Agarwal (2007, March). Through a glass darkly: Information technology design, identity verification, and knowledge contribution in online communities. *Information Systems Research* 18(1), 42–67. [2](#)
- Magnac, T. and D. Thesmar (2002, March). Identifying dynamic discrete decision processes. *Econometrica* 70(2), 801–816. [30](#)
- McKinney, W. (2010). Data structures for statistical computing in python. In S. van der Walt and J. Millman (Eds.), *Proceedings of the 9th Python in Science Conference*, pp. 56–61. [64](#)
- Nov, O. (2007, November). What motivates wikipedians? *Communications of the ACM* 50(11), 60–64. [2](#)

- Owens, D., Z. Grossman, and R. Fackler (2014, November). The control premium: A preference for payoff autonomy. *American Economic Journal: Microeconomics* 4(4), 138–161. [5](#)
- Pedregosa, F., G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 12, 2825–2830. [58](#), [64](#)
- Pikulina, E. S. and C. Tergiman (2020, May). Preferences for power. *Journal of Public Economics* 185(104173). [5](#)
- Rajan, R. G. and L. Zingales (1998, May). Power in a theory of the firm. *The Quarterly Journal of Economics* 113(2), 387–432. [2](#)
- Roberts, J. A., I.-H. Hann, and S. A. Slaughter (2006, July). Understanding the motivations, participation, and performance of open source software developers: A longitudinal study of the apache projects. *Management Science* 52(7), 984–999. [2](#)
- Rust, J. (1987, September). Optimal replacement of gmc bus engines: an empirical model of harold zurcher. *Econometrica* 55(5), 999–1033. [30](#)
- Seabold, S. and J. Perktold (2010). Statsmodels: Econometric and statistical modeling with python. In *9th Python in Science Conference*. [64](#)
- Sturm, R. E. and J. Antonakis (2015, January). Interpersonal power: A review, critique, and research agenda. *Journal of Management* 41(1), 136–163. [2](#), [20](#)
- Sun, L. and S. Abraham (2020, December). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*. [22](#)
- Virtanen, P., R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. Jarrod Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and Contributors (2020). SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* 17, 261–272. [64](#)
- Xu, L., T. Nian, and L. Cabral (2020). What makes geeks tick? a study of stack overflow careers. *Management Science* 66(2), 587–604. [2](#)
- Zhang, X. M. and F. Zhu (2011, June). Group size and incentives to contribute: A natural experiment at chinese wikipedia. *The American Economic Review* 101(4), 1601–1615. [2](#)

Appendix A Details and Robustness

A.1 Construction of quality variable

The variable quality captures the variation of points received by an answer at its publication day explained by text characteristics.

Let \mathbf{X}_j be a vector of text characteristics of an answer j , right after publication, so before any modification occurs. Let \bar{t}_j be the publication date of answer j . I estimated the following linear model:

$$points_{j,\bar{t}_j} = \beta_0 + \mathbf{X}_j\boldsymbol{\beta}_1 + \mathbf{X}_j^2\boldsymbol{\beta}_2 + \varepsilon_j$$

with $points_{j,\bar{t}_j}$ being the points obtained by answer j 's author at the publication day. The quality of answer j is then defined as the predicted number of points from the above model.

The vector of text characteristics includes:

- number of words,
- precision, defined as the number of words excluding the stop-words, over total number of words,
- number of links,
- number of images.

Table 18 reports the estimates for the linear regression models used to predict the variable quality. The specification adopted corresponds to column (5).

A.2 Construction of scarcity variable

The construction of the variable scarcity follows several steps.

- Construct the variable *availability*, given by the cumulative number of questions appearing in the platform, that, each day, don't have yet an answer selected as best answer. This is equivalent for every users. The cumulative number of questions and the number of questions without an accepted answer are plotted in figure 17.
- Recover topics from question tags³⁵. To do this, I first construct a graph of tags, where a link between two tags exists if the two tags appear at least once in the same question. The intensity of the links are given by the number of times that the two tags have appeared in a same question. I then identify topics using the Page rank algorithm³⁶, i.e. a topic will be those tags that are connected to the

³⁵Questioners can add tags when posting a question

³⁶This is the Google search algorithm of the early times of the search engine

Dep. var: points	(1)	(2)	(3)	(4)	(5)
Length	0.00440*** (5.65)	0.00997*** (7.48)	0.00921*** (6.80)	0.00927*** (6.85)	0.00859*** (6.34)
Precision	9.219*** (8.81)	9.495*** (9.06)	32.20*** (4.38)	32.02*** (4.35)	30.91*** (4.20)
Num. figures	3.504*** (8.98)	3.554*** (9.11)	3.555*** (9.11)	6.915*** (10.42)	6.474*** (9.73)
Num. links	1.818*** (21.86)	1.807*** (21.73)	1.806*** (21.72)	1.784*** (21.44)	2.235*** (23.50)
Length ²		-0.00000991*** (-5.15)	-0.00000926*** (-4.78)	-0.00000932*** (-4.81)	-0.00000861*** (-4.44)
Precision ²			-22.54** (-3.12)	-22.39** (-3.10)	-21.81** (-3.02)
Num. figures ²				-1.231*** (-6.26)	-1.172*** (-5.95)
Num. Links					-0.0393*** (-9.78)
_cons	8.978*** (17.16)	8.437*** (15.81)	2.909 (1.57)	2.944 (1.59)	3.292 (1.78)
<i>N</i>	118552	118552	118552	118552	118552

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 18: Regressions to predict the quality variables. Model finally used is in column (5).

most other tags. I identify 6 topics, since the Page ranking value drops suddenly after those first 6 tags. The topics are: 'grammar', 'word-usage', 'meaning', 'sentence-construction', 'meaning-in-context', 'word-choice'. I then partition the graph around these 6 tags, using a Voronoi diagram. After this process I then have 6 topics, and a mapping from every tag to each of these topics. The word-clouds of tags related to each topic are plotted in figure 18.

- I allocate topics at the questions still to answer, recovered at the first bullet point: using the tags assigned to those questions, I obtain the share of each topic in each of the questions.
- I do a similar process for each user, on all questions he/she has answered, and recover the share of each topic in which he/she is expert about
- for each user i , I weight the available questions at each period t by his/her expertise, call this variable $avail_{it}$. Figure 19 shows the distribution of time of this variable, in average across users' lifetime in the website.

The variable scarcity is then defined as:

$$scarcity_{it} \equiv \frac{maxavail}{\log(avail_{it})}$$

where max_{avail} is the maximum value that $\log(avail_{it})$ takes in the data, across all i, t .

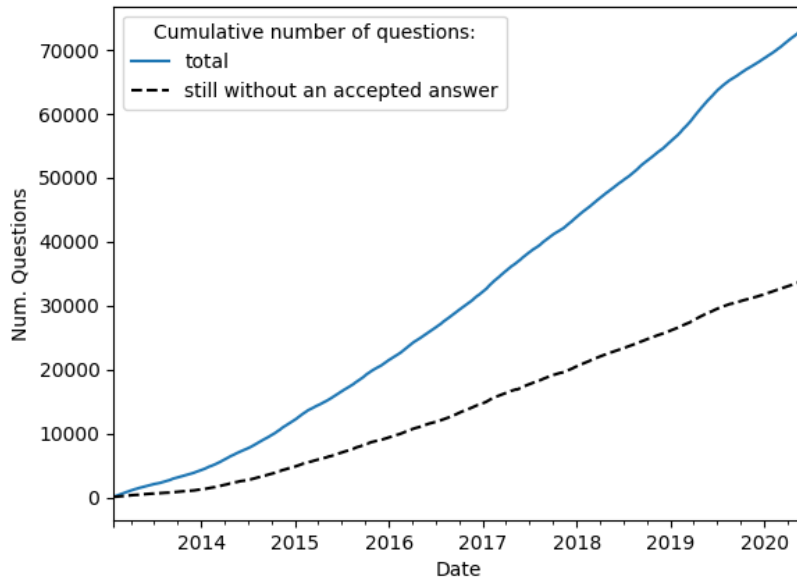


Figure 17: Cumulative number of questions in the platform, both total number and net of questions that have already selected an answer as *best answer*.

A.3 Construction of Types

The procedure I used to identify types follows few steps with the combination of quantitative assessment and interpretative assessment. First I aggregate information to reduce the dimensions of the individual characteristics. Then I employ an algorithm to identify clusters within the reduced space.

Information aggregation. The most simple approach to reduce dimensionality would be to aggregate the variables via, for example, sum. Since the individual characteristics include dummy variables taking value 1 if the user decided to display some given information, as well as the length of the biographical description, summing over them gives a measure of the amount of information displayed. This approach turned to not be a good solution, as behavior is not linearly correlated with the amount of information. The aggregated variable is then not informative on the different types of users.

A common alternative is to perform the Principal Component Analysis (PCA). This approach transforms the data by creating orthogonal vectors, each containing the largest possible variance of the original variables. The first vector will be the most representative of the original variance, the second will be the most representative of the residual variance, and so on. PCA anyway relies on quantitative continuous variables, as it relies

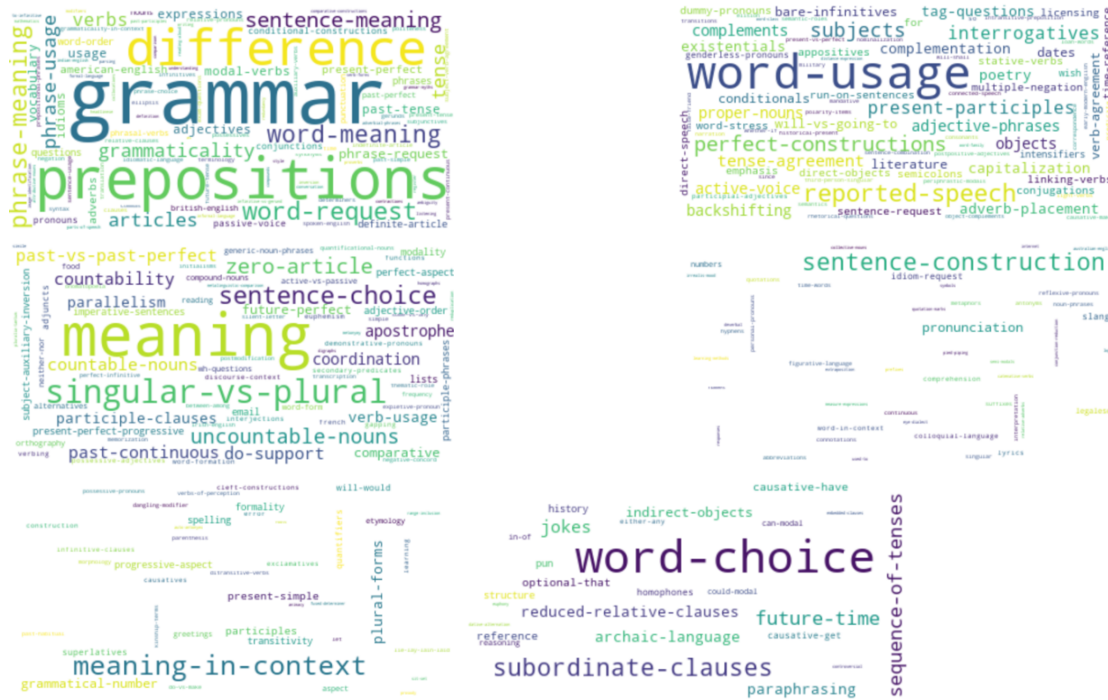


Figure 18: Word-clouds for each topic identified

on the computation of the variance, and it is not suitable to dummy variables.

In this work I adopt the Multiple Correspondence Analysis (MCA, [Greenacre and Blasius 2006](#)), a sort of PCA counterpart for categorical variables, which is a generalization of the Correspondence Analysis (CA). This method relies on the cross tabulation of each pair of variables, with the single categories being the rows and columns, and the joint frequency the measure in the cells.

As the PCA, the MCA algorithm outputs dimensions (or factors) that aggregate the information of the original variables. Individual users can then be plotted in the reduced bi-dimensional space formed by each pair of dimensions. In the discussion that follows I will focus on the plane formed by the first and second dimensions. Note that, since this algorithm is applied to categorical variables, I bin in three groups the variables representing the length of the biographical note and the variable with the number of links appearing in the biographical note.

Figure 20 shows the variable representation in the first two dimensions space. First it is possible to notice, on the axes, that the first dimension contains about 17% of the information of the individual characteristics, while the second dimension about 8%. The location of the variables on the plain tells the extent to which that dimension include information from the given variables. It is possible to see that the length of the biographical note is the most important source of information for both dimensions, while the inclusion of location and website in the user page is only captured by the first di-




Figure 19: Average expertise-weighted availability of questions per period of participation on the website, across users. Shadow areas identify the standard deviation.

mension.

Figure 21 instead represents on the same dimensions the individuals, i.e. the sample of users. This graph may help to understand if individuals cluster in groups, based on the information of the first two factors. It is possible to observe that clear clusters are not emerging. Nonetheless, points are not displayed in an uniform cloud with respect to the axis. While some are grouping around the (0,0) point, meaning that they have characteristics close to the average of the sample, others appear on the positive side of the first dimension. Users appearing in the upper right quadrant are more likely to have a LinkedIn profile, a website, and the location, compared to the average user, as well as longer biography with more links. Users in the bottom right quadrant are also more likely to have a website and the location, they tend to have a biography, but a short one.

Identification of groups. A typical clustering algorithm is the so called K-Means clustering. This algorithm requires the number k of groups that want to be identified, it picks k centroids (i.e. means of partitions of the observations) and updates the centroids so to minimize the within-cluster variance. This algorithm is also meant to work with continuous quantitative variables, so is not suitable to be directly applied on the original individual characteristics. I then apply the K-Means clustering procedure to the first 5 dimensions recovered after the application of the MCA procedure. These are continuous variable and still represent the information of the original data.




Figure 20: Variable representation on the first two dimensions plane, where the dimensions are obtained via the MCA of the individual characteristics.

By choosing three clusters (i.e. $k = 3$), the resulting individual representation is shown in figure 22, with individuals colored based on the allocated cluster.

A.4 Reduced form - robustness checks

A possible concern on the reduced form analysis is that the effect observed is not specific of the privilege allocating control on editing. In other words, we could observe a significant increase in the editing activity after each achievement of privileges. To check for this possibility, I estimate the exact same specification of section 4.2 around different thresholds.

In particular I consider the two privileges achieved just before and just after the allocation of authority. Figure 23 reports the estimates of the reputation-point intervals fixed effects, around the privilege “Established User”. This privilege does not allocate any resource, and it is just a recognition. It is obtained with 750 points during the beta phase of the site, and with 1000 points during the final phase. It is possible to notice that right around the threshold it is not observed a significant increase in editing. Moving further from the threshold shows instead an increase, but pre-treatment effects seems to suggest the presence of a trend, rather than a causal effect of the treatment. Finally, figure 24 shows again estimates for the same specification, but this time around



Figure 21: Representations of users on the first two dimensions plane, where the dimensions are obtained via the MCA of the individual characteristics.

the allocation of the “Creat Tag Synonyms” privilege. This privilege allows users to corrects tags. It is achieved either with 1250 points in the beta phase, or with 2500 points otherwise. Looking at the effects just in the neighborhood of the treatment, it is not really possible to identify a clear pattern.

A.5 Derivation of Likelihood function

Let $D \in \{1, 0\}$ be a binary variable that takes value equal to 1 when the user is given full ex-ante control over Edits. In addition, denote \mathbf{d}_t a vector of dummy variables, $d_{\alpha t}$, for each possible choice $\alpha \in \mathcal{A}$, such that $d_{\alpha t}$ is equal to 1 if in period t is selected choice α , and zero otherwise.

Choosing an action α^* in period t , the one period flow utility of user i is then given by:

$$U_{it}(d_{\alpha^*t} = 1) = \beta'_0 \mathbf{x}_{it}(d_{\alpha^*t} = 1) + \mathbf{1}\{D_t = 1\} \beta'_1 \mathbf{x}_{it}(d_{\alpha^*t} = 1) + \varepsilon_{i\alpha^*t}$$

Where the vector \mathbf{x}_t is described in section 5.1.

The term $\varepsilon_{i\alpha^*t}$ is instead a choice specific utility term not measurable by the econometrician.



Figure 22: Representations of users on the first two dimensions plane, where the dimensions are obtained via the MCA of the individual characteristics. Colors refer to cluster groups identified with k-means clustering on the MCA dimensions.

Individual problem

Define as \mathcal{Z} the set of all possible states z , i.e. all possible combinations of state variables, at t . This does not consider only the variables that enter the utility function (i.e. \mathbf{x}_t), but also variables that may affect users' beliefs on the probability distribution over future states.

A user selects a sequence of optimal decisions $\mathbf{d}^* \equiv \{\mathbf{d}_t^*\}_{t \leq T}$ that satisfies³⁷:

$$\mathbf{d}^* = \arg \max_{\mathbf{d}} \mathbb{E} \left[\sum_{t=1}^T \sum_{\alpha \in \mathcal{A}} \delta^{t-1} d_{\alpha,t} U_{\alpha t}(z_t) \right] = \mathbb{E} \left[\sum_{t=1}^T \sum_{\alpha \in \mathcal{A}} \delta^{t-1} d_{\alpha,t} (u_{\alpha t}(z_t) + \varepsilon_{\alpha t}) \right],$$

where δ is a discount factor and, at each period t , the expectation is taken with respect to z_τ and ε_τ , for $\tau \geq t + 1$.

In words, the agent, at each period, will choose whether to contribute in the platform

³⁷To make notation more readable, for any function f that depends on the agent's choice, I will use the following:

$$f_{\alpha t}() \equiv f_t(d_{\alpha t} = 1)$$

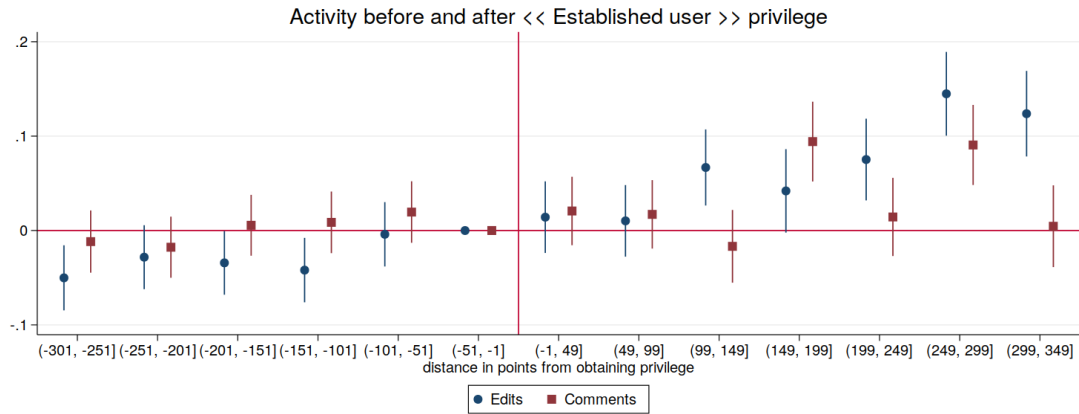


Figure 23: Estimates for reduced form effect around the *Establish User* privilege

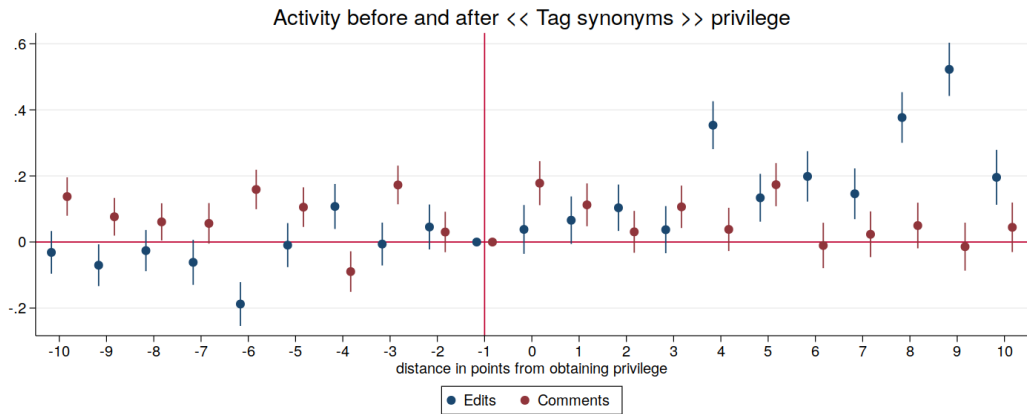


Figure 24: Estimates for reduced form effect around the *create tag synonyms* privilege

and eventually what type of contribution to make, between producing content (answers), performing moderation task (edits), or both.

Identification and estimation

For the characterization of the problem I follow [Arcidiacono and Miller \(2011\)](#).

Define the ex-ante value function at period t as the discounted sum of the expected future payoff under optimal behavior, and before the shock ε_t is realized³⁸. In other words, it is the continuation value of being in state z_t , before ε_t is realized and the decision at t

³⁸The reason why it is considered the ex-ante value function is because the shock is not observed by the researcher. Note nevertheless that at the time of the decision in period t , the shock is observed by the agent, who'll take it into account in her choice.

is taken. By applying Bellman's principle, it is then given by:

$$V_t(z_t) = \mathbb{E} \left[\sum_{\alpha \in \mathcal{A}} d_{\alpha,t}^* \left(u_{\alpha t}(z_t) + \varepsilon_{\alpha t} + \delta \sum_{z_{t+1} \in \mathcal{Z}} V_{t+1}(z_{t+1}) f_{\alpha t}(z_{t+1}|z_t) \right) \right]$$

where the expectation is taken with respect to $\varepsilon_{\alpha t}$, and $f_{\alpha t}(z_{t+1}|z_t)$ is the probability that the vector of states will take a certain value in the next period, given the choice made. This transition probability does not depend on all the history of past choices due to the assumptions made in the previous section.

Define then the conditional value function $\nu_{\alpha t}(z_t)$ as the value function $V_t(z_t)$ for a given choice α and net of the preference shock ε_t :

$$\nu_{\alpha t}(z_t) = u_{\alpha t}(z_t) + \delta \sum_{z_{t+1} \in \mathcal{Z}} V_{t+1}(z_{t+1}) f_{\alpha t}(z_{t+1}|z_t).$$

Finally, define the conditional choice probabilities $\mathbf{p}_t(z_t)$ as the vector that gives the probabilities of choosing option $\alpha \in \mathcal{A}$ given state z_t , taking expectations on the preference shock, so to explain different choices in the data given the same states:

$$p_{\alpha t}(z_t) = \int d_{\alpha t}^* g(\varepsilon_t) d\varepsilon_t,$$

with $g(\varepsilon_t)$ being the density of ε_t which is assumed to have continuous support. Building on [Hotz and Miller \(1993\)](#), [Arcidiacono and Miller \(2011\)](#) show that, under certain conditions, it exists a function ω for each $\mathbf{k} \in \mathcal{A}$ such that:

$$\omega_{\mathbf{k}}(\mathbf{p}_t(z_t)) = V_t(z_t) - \nu_{\mathbf{k}t}(z_t).$$

It follows that:

$$\nu_{\alpha t}(z_t) = u_{\alpha t}(z_t) + \delta \sum_{z_{t+1} \in \mathcal{Z}} (\nu_{\mathbf{k}t+1}(z_{t+1}) + \omega_{\mathbf{k}}(\mathbf{p}_{t+1}(z_{t+1}))) f_{\alpha t}(z_{t+1}|z_t),$$

which can be rewritten as:

$$\nu_{\alpha t}(z_t) = u_{\alpha t}(z_t) + \sum_{\tau=t+1}^T \sum_{\mathbf{k} \in \mathcal{A}} \sum_{z_{\tau} \in \mathcal{Z}} \delta^{\tau-t} (u_{\mathbf{k}\tau}(z_{\tau}) + \omega_{\mathbf{k}}(\mathbf{p}_{\tau}(z_{\tau}))) d_{\mathbf{k}\tau}^*(z_{\tau}, d_{\alpha t} = 1) \kappa_{\tau-1}^*(z_{\tau}|z_t, d_{\alpha t} = 1), \quad (6)$$

where the function $\kappa_{\tau}^*(z_{\tau+1}|z_t, d_{\alpha t} = 1)$ represents the cumulative probability of being in state $z_{\tau+1}$ in period $\tau + 1$ conditional on having been in state z_t and having chosen α in period t , i.e.

$$\kappa_{\tau}^*(z_{\tau+1}|z_t, d_{\alpha t} = 1) \equiv \begin{cases} f_{\alpha t}(z_{t+1}|z_t) & \text{for } \tau = t \\ \sum_{z_{\tau} \in \mathcal{Z}} \sum_{\mathbf{k} \in \mathcal{A}} d_{\mathbf{k}\tau}^* f_{\mathbf{k}\tau}(z_{\tau+1}|z_{\tau}) \kappa_{\tau-1}^*(z_{\tau}|z_t, d_{\alpha t} = 1) & \text{for } \tau = t + 1, \dots, T. \end{cases}$$

To write the conditional value function as in 6 is functional to implement the *Finite Dependence* property, generalized by Arcidiacono and Miller (2011). This property allows to rewrite the problem such that the agent considers only a subset of the future periods to make her decision.

The intuition behind the property goes as follows.

First of all the identification of the structural parameters will be based on the comparison of conditional value functions, since the likelihood of observing at t a choice α rather than α' given a specific state z_t corresponds to the probability that $\nu_{\alpha t}(z_t) - \nu_{\alpha' t}(z_t) > \varepsilon_{\alpha t} - \varepsilon_{\alpha' t}$.

Consider now two alternative choices, α and α' . If, by choosing either of the two, it is possible to follow sequences of decisions such that the probability distribution of the state variables is exactly equivalent, then, when substituting equation 6 into the difference $\nu_{\alpha t}(z_t) - \nu_{\alpha' t}(z_t)$, all future periods after the sequence of choices will cancel out.

Assumption over the distribution of the stochastic term.

Consider again two alternative choices, α and α' . Since we are interested in measuring the probability that $\nu_{\alpha t}(z_t) - \nu_{\alpha' t}(z_t) > \varepsilon_{\alpha t} - \varepsilon_{\alpha' t}$, we need to make assumptions on the distribution of the stochastic term $\varepsilon_{\alpha t}$. I will assume a Type I extreme value distribution. This allows to express the choice probabilities as:

$$p_{\tilde{\alpha} t}(z_t) = \frac{\exp(\nu_{\tilde{\alpha} t}(z_t))}{\sum_{\alpha \in \mathcal{A}} \exp(\nu_{\alpha t}(z_t))} = \frac{1}{\sum_{\alpha \in \mathcal{A}} \exp(\nu_{\alpha t}(z_t) - \nu_{\tilde{\alpha} t}(z_t))}$$

and the ex-ante value function as:

$$V_t(z_t) = \ln \left(\sum_{\alpha \in \mathcal{A}} \exp(\nu_{\alpha t}(z_t)) \right) + \gamma = -\ln(p_{\tilde{\alpha} t}(z_t)) + \nu_{\tilde{\alpha} t}(z_t) + \gamma$$

where γ is the Euler's constant and $\tilde{\alpha}$ is an arbitrary reference choice from \mathcal{A} . It follows that:

$$\omega_{\tilde{\alpha}}(\mathbf{p}_t(z_t)) = -\ln(p_{\tilde{\alpha} t}(z_t)) + \gamma.$$

Given a reference choice $\tilde{\alpha}$ then it is possible to write the difference of conditional value functions as:

$$\begin{aligned} \nu_{\alpha t}(z_t) - \nu_{\tilde{\alpha} t}(z_t) = & u_{\alpha t}(z_t) - u_{\tilde{\alpha} t}(z_t) + \\ & \sum_{\tau=t+1}^{t+\Delta_t} \sum_{\mathbf{k} \in \mathcal{A}} \sum_{z_\tau \in \mathcal{Z}} \delta^{\tau-t} (u_{\mathbf{k}\tau}(z_\tau) - \ln(p_{\mathbf{k}\tau}(z_\tau))) [d_{\mathbf{k}\tau}^*(z_\tau, d_{\alpha t} = 1) \kappa_{\tau-1}(z_\tau | z_t, d_{\alpha t} = 1) + \\ & - d_{\mathbf{k}\tau}^*(z_\tau, d_{\tilde{\alpha} t} = 1) \kappa_{\tau-1}(z_\tau | z_t, d_{\tilde{\alpha} t} = 1)] \end{aligned}$$

where Δ_t is the number of periods after which the agent faces the same probability distribution over the states, independently of having initially chosen α or $\tilde{\alpha}$.

The Log-likelihood function of the data is given by:

$$\begin{aligned}
L(\beta_0, \beta_1, \gamma) &= \sum_{i=1}^N \sum_{t=1}^T \sum_{\alpha \in \mathcal{A}} \log \left(\frac{\exp(\nu_{\alpha it}(z_{it}))}{\sum_{k \in \mathcal{A}} \exp(\nu_{kit}(z_{it}))} \right) \times d_{\alpha it} \\
&= \sum_{i=1}^N \sum_{t=1}^T \sum_{\alpha \in \mathcal{A}} \log \left(\frac{\exp(\nu_{\alpha it}(z_{it}) - \nu_{\bar{\alpha} it}(z_{it}))}{\sum_{k \in \mathcal{A}} \exp(\nu_{kit}(z_{it}) - \nu_{\bar{\alpha} it}(z_{it}))} \right) \times d_{\alpha it}
\end{aligned}$$

A.6 Details on Estimation of Structural model

A.6.1 Choice set

Because of computational time, the choice set must be constrained to a finite and limited number of options.³⁹ In my specification, users are allowed to make 21 possible choices of effort. They may not participate at all, make effort only in answering, only in editing, or in both. Answering effort is a combination of quantity and quality of answers, with two possible levels for quantity, and three possible levels of quality. Quantity of edits can take two possible levels. All options in the choice set are listed in the table 19.

The value of the possible levels are obtained by looking at the distribution of actions taken in the data by individuals at each week of participation. For what concerns the quantity of answers, I split the distribution at the 70th quantile, corresponding to three answers, so to categorize effort between low (1 to 3 answers) and high (4 or more). I then select, as possible option for the user, the median values of these two categories, so either 1 or 7 answers. A similar process is made for quality and edits. The distribution of quality is split in three categories at the 33th and 66th quantiles. The median values are 13.33, 14.12, and 15.97. Finally, the distribution of number of edits is split at the 75th quantile, leading to two categories: low effort, which includes 1 or 2 edits, and high effort, including 3 or more edits. The distribution of values within each category is plotted in figure 27 in the appendix. The choice of the quantile levels is arbitrary.

A.7 Conditional Choice Probabilities

Conditional choice probabilities are computed before estimation via a static logit⁴⁰ Before estimation, the data is scaled so that each variable would be in the range (0, 1). The scaling algorithm subtracts the minimum and divide by the difference between the maximum and the minimum. The multinomial logit model implemented is the following:

$$\begin{aligned}
\alpha_{it}^* &= \beta_0 R_{it-1} + \beta_1 \Lambda_{U,it-1} + \beta_2 \Lambda_{D,it-1} + \beta_3 \text{avail}_{it} + \beta_4 \text{AnswerNum}_{it} + \beta_5 \text{Seniority}_{it} + \\
&\quad + \beta_6 t + \beta_7 \text{date}_{it} + \text{cum}T_{it}
\end{aligned}$$

³⁹A more natural assumption would be that users make discrete choices of tasks, and continuous choices for effort levels. As of today, the econometric literature is not providing a way to do so. A first solution to this problem is provided in the recent work by Bruneel-Zupanc (2020).

⁴⁰Logistic regression in Scikit-learn (Pedregosa, Varoquaux, Gramfort, Michel, Thirion, Grisel, Blondel, Prettenhofer, Weiss, Dubourg, Vanderplas, Passos, Cournapeau, Brucher, Perrot, and Duchesnay 2011) with *saga* solver.

A	Q	E
0.0	0.00	0.0
0.0	0.00	1.0
0.0	0.00	4.0
1.0	13.33	0.0
1.0	13.33	1.0
1.0	13.33	4.0
1.0	14.12	0.0
1.0	14.12	1.0
1.0	14.12	4.0
1.0	15.97	0.0
1.0	15.97	1.0
1.0	15.97	4.0
7.0	13.33	0.0
7.0	13.33	1.0
7.0	13.33	4.0
7.0	14.12	0.0
7.0	14.12	1.0
7.0	14.12	4.0
7.0	15.97	0.0
7.0	15.97	1.0
7.0	15.97	4.0

Table 19: Possible effort levels that users are allowed to choose in estimation. Columns report, from left to right, the possible choice of effort in the number of answers, the average quality of answers, and the number of edits

where α_{it}^* is the choice made by user i in period of participation t , R is the number of reputation points, Λ_U and Λ_D are the expected number of up-votes and down-votes arriving from past effort, $avail$ is the number of available questions to answer, $AnswerNum$ is the number of answers already published up to period t , $Seniority$ the number of days passed since the registration day, $date$ is the calendar week, and $cumT$ the number of privileges obtained by the user. All parameters are choice specific.

A.8 Details on Simulation of Counterfactuals

A.8.1 Restrictions on the state values

Reputation points. It is assumed that users can accumulate at most 1500 reputation points. To adjust for this limit, which is not present in the real design, I scale the returns in points from up-votes / down-votes. Every up-votes provides 5 reputation points to the author, while every down-votes removes 1 point. The approval of suggested edits provide 1 point.

Expected number of points arriving from past actions. The variables Λ_U and Λ_D , which are normally continuous, are discretized. Λ_U can take value from zero to 0.2, with steps of 0.01, while Λ_D can take value from zero to 0.03, with steps of 0.01. The boundaries of these sets are generally never hit, and do not impose important restrictions. On the contrary, the discretization reduces the sensitivity of the model.

Availability of questions. I randomly allocate to users a registration date. Based on the dates of participation, I allocate the number of available questions to each user, as it appears to be in the real platform. To reduce dimensionality, I bin the variable so that the number of available question can be one of 5 unique values. Note that the number of available questions could still change across the time of a user's participation.

Experience variables. The number of answers already made and the days of participation are set to zero and are not allowed to increase. In other words, in the simulations I do not allow for learning while participating.

Appendix B Other figures

You can earn a maximum of 200 reputation per day from any combination of the activities below. [Bounty awards](#), [accepted answers](#), and [association bonuses](#) are not subject to the daily reputation limit.

You gain reputation when:

- question is voted up: +5
- answer is voted up: +10
- answer is marked "accepted": +15 (+2 to acceptor)
- suggested edit is accepted: +2 (up to +1000 total per user)
- bounty awarded to your answer: + full bounty amount
- one of your answers is awarded a bounty automatically: + half of the bounty amount ([see more details about how bounties work](#))
- site association bonus: +100 on each site (awarded a maximum of one time per site)
- example you contributed to is voted up: +5
- proposed change is approved: +2
- first time an answer that cites documentation you contributed to is upvoted: +5

If you are an experienced Stack Exchange network user with 200 or more reputation on at least one site, you will receive a starting +100 reputation bonus to get you past basic new user restrictions. This will happen automatically on all current Stack Exchange sites where you have an account, and on any other Stack Exchange sites at the time you log in.

You lose reputation when:

- your question is voted down: -2
- your answer is voted down: -2
- you vote down an answer: -1
- you place a bounty on a question: - full bounty amount
- one of your posts receives 6 spam or offensive flags: -100

All users start with one reputation point, and reputation can never drop below 1. Accepting your own answer does not increase your reputation. Deleted posts do not affect reputation, for voters, authors or anyone else involved, in [most cases](#). If a user reverses a vote, the corresponding reputation loss or gain will be reversed as well. Vote reversal as a result of voting fraud will also return lost or gained reputation.

At the high end of this reputation spectrum there is little difference between users with high reputation and ♦ moderators. That is intentional. We don't run this site. The community does.

Figure 25: Rules to obtain or lose reputation in Stackexchange(<https://stackoverflow.com/help/whats-reputation>)




Figure 26: Number of users that have accumulated different amount of reputation points, conditional on having less than 2000 points at the graduation week. The decreasing value is due to exiting of the platform. It is possible to see that some users are reaching the level of 2000 points and they will not lose the privilege at the design date, some never reached the privilege, and others, the orange ones, lose it.

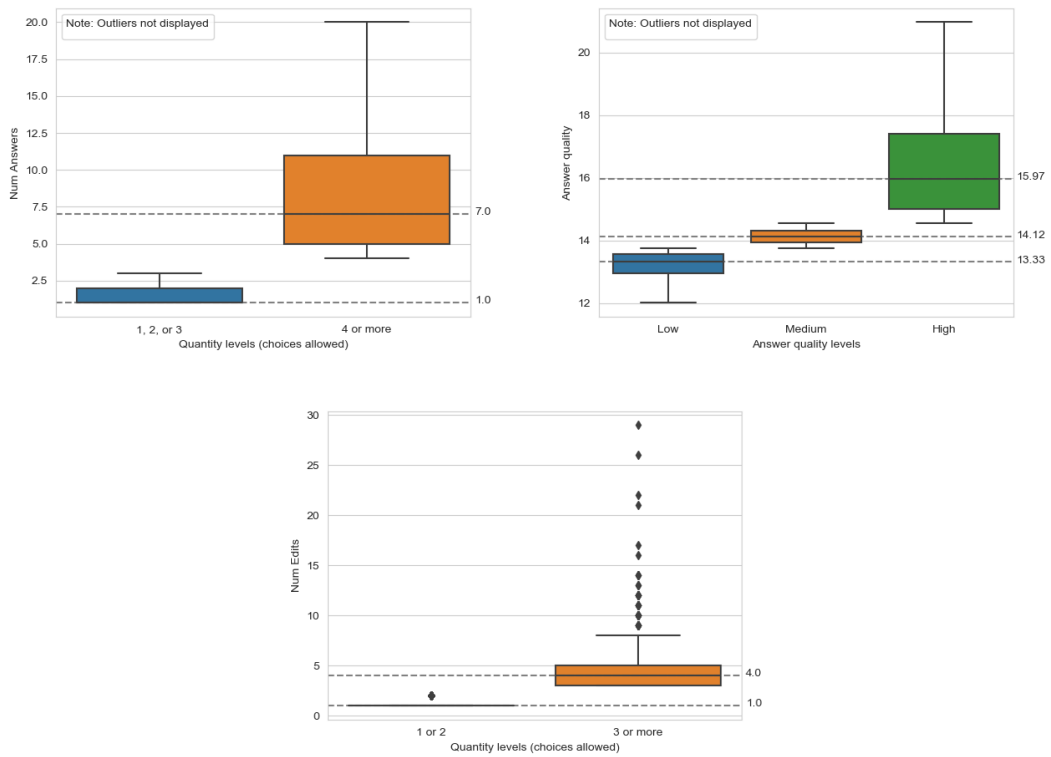


Figure 27: Categories of possible actions that users in the estimation are allowed to take, with the distribution of actual actions in each category. Values on the right vertical axis are the median value of each category, which make the set of options that users are allowed to choose.

Appendix C Credits for the software used

Pedregosa et al. (2011), Seabold and Perktold (2010), Hagberg, Schult, and Swart (2008), McKinney (2010), Lê, Josse, and Husson (2008), Virtanen, Gommers, Oliphant, Haberland, Reddy, Cournapeau, Burovski, Peterson, Weckesser, Bright, van der Walt, Brett, Wilson, Jarrod Millman, Mayorov, Nelson, Jones, Kern, Larson, Carey, Polat, Feng, Moore, Vand erPlas, Laxalde, Perktold, Cimrman, Henriksen, Quintero, Harris, Archibald, Ribeiro, Pedregosa, van Mulbregt, and Contributors (2020), Hunter (2007)

Other software used:

StataCorp. 2017. Stata Statistical Software: Release 15. College Station, TX: StataCorp LLC.