Delegation as Incentive for Public Good Provision: Evidence from an Online Community *

Jacopo Bregolin†
University of Liverpool

June 2024

Abstract

In many organisations, employees' learning and productivity rely on knowledge platforms' user-generated content, which has become a standard daily source of information for various tasks. As users contribute on a voluntary basis, platforms need to incentivise free effort. With data from Stack Exchange, I investigate whether users provide more and better quality contributions when endowed with more control over actions. Using a dynamic discrete choice model, I show that autonomy 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. When delegation is based on performance, the platform faces a trade-off, which depends on the composition of the community, and the tasks that the platform wants to incentivise.

*I am extremely grateful for comments and feedback to Heski Bar-Isaac, Milo Bianchi, Kevin Boudreau, Daniel L. Chen, Pierre-André Chiappori, Jacques Crémer, Olivier de Groote, Daniel Ershov, Chiara Farronato, April Franco, Guido Friebel, Hulrich Hege, Astrid Hopfensitz, Aija Leiponen, Greg Lewis, Inna Lyubareva, Matthew Mitchell, Sophie Moinas, Elena Panova, Mark Schankerman, and Timothy Simcoe. I thank for useful comments attendants at several internal seminars at TSE, and participants at the IOEA Corsica, ZEW Digitized Economy workshop, Rotman School of Management PhD seminar, VIDE brown bag, Economics of Platform Seminar, Platform Strategy Research Symposium, University of Liverpool, University of Minnesota Carlson School of Management, Católica Lisbon, SIOE, and ESEM. Earlier versions of this paper circulated with the title *Authority and Delegation in Online Communities.*

†jacopo.bregolin@liverpool.ac.uk

1 Introduction

Nowadays, productivity relies heavily on information crowdsourced from internet users and aggregated by knowledge platforms (e.g. Wikipedia and Stack Overflow) or large language models (e.g. OpenAI's GPT). Software programmers rely on online sample code, journalists write pieces based on social media content, and restaurants rely on online reviews, to mention a few. While user-generated content is a valuable public good, its production is not remunerated nor contractible. How can platforms that host such content incentivise high-quality contributions? The literature has identified several non-monetary drivers of effort.[1] Platforms may exploit those that depend on their design as crucial devices to incentivise participation. Typical organisational tools are non-monetary rewards, like awards, that leverage users' preferences for recognition and status (Gallus and Frey 2016). Nevertheless, awards may incentivise effort before they are assigned, but not necessarily after, as the allocation is generally permanent. The theoretical literature in personnel economics has identified another non-monetary channel, that is, the delegation of decision rights. In other words, platforms may incentivise participation through the strategic allocation of autonomy over decision-making (Gibbons, Matouschek, and Roberts 2013, Gambardella, Panico, and Valentini 2015). While autonomy could be used as a reward for productive users, its allocation may increase long-term task commitment (Beckmann and Kräkel 2022).[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 decision rights leads to an increase in quantity and quality of 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,

[1] Drivers of effort include intrinsic utility and firm recognition (Roberts, Hann, and Slaughter 2006, Nov 2007, Ma and Agarwal 2007, Jeppesen and Frederiksen (2006)), the community size (Zhang and Zhu 2011), reference points on others' behavior (Chen, Harper, Konstan, and Li 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, Guo, and Lin 2016), and the signaling of skills (Belenzon and Schankerman 2015, Xu, Nian, and Cabral 2020).

[2] In this paper, I interchangeably use the terms "decision rights", "control rights", and "autonomy".

[3] There 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 Krähmer (2008). The literature has also addressed other types of nonmonetary incentives, with similar dynamics to the delegation of control rights. 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.

I show that people value such autonomy. However, different types of users value obtaining and having such autonomy differently. Through counterfactual exercises, I explore organisational implications and find that delegation increases the number and quality of contributions, but less so when users are already committed to the site. In addition, delegation based on performance may increase productivity but backfire if the performance targets are too high.

What does it mean to allocate autonomy in digital platforms? Every online community 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 autonomy in editing content conditional on achieving given performance targets. What trade-offs affect this decision?

This paper focuses on the incentive effects of the allocation of control rights 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 autonomy is set at either infinity or zero. I address two main incentive effects. First, if users value acquiring autonomy, delegation incentivises effort until users reach the performance threshold (*dynamic incentive*). This is consistent with a theory of *career concerns* (Holmström 1999). Second, if users value contributing when endowed with more autonomy, delegation relaxes the participation constraint, as it increases the value of participating (*static incentive*). Possible justifications are higher task commitments (Beckmann and Kräkel 2022) and intrinsic value for autonomy (Bartling, Fehr, and Herz 2014). 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.

Stack Exchange is a family of websites where registered users ask questions and provide answers on different topics. The moderation of the website relies on moderators elected within the community and community members' edits. Edits are not directly implemented and require approval by moderators or the owner of the edited content. Nevertheless, when users collect enough reputation points, they gain control rights on the editing task and are able to implement their modifications directly. In this context, I observe users' contributions before and after they receive autonomy in editing. After providing reduced-form evidence of the *static incentive* effect, I develop a dynamic discrete choice model to measure users' preference for control rights, allowing for heterogeneity across types. The paper finds that the incentives differently affect the different types, and depend on heterogeneous valuations for autonomy and different participation costs. In particular, under stricter delegation designs, 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 (e.g. answers and edits) of all participants in the English Language Learners website of Stack Exchange.[4] This website includes questions on the use of the English language. 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, who provide very little information, *Informative* users, who provide partial information about their identity, including location, websites, and Linkedin profiles, and *Specialised* users, who provide more detailed information and, crucially, include the word "English" in the biography, suggesting more commitment to the topic of the site.

The analysis proceeds in two steps. First, I test for the presence of *static incentives* by looking at participation choices in editing before and after users gain control rights over the editing task. I use a regression discontinuity analysis in which the running variable is the distance from the threshold in points, i.e. the performance measure on which the threshold depends. The analysis exploits variation from both the gaining and losing of autonomy, as the performance threshold changed during the life of the site, with retroactive effects. I find that users contribute more edits and are more likely to participate in editing, once they acquire autonomy. The threshold does not affect contributions in other actions, like comments. The dynamic nature of the *dynamic incentive* effect does not allow clean identification in reduced form. Nevertheless, descriptive evidence shows an increase in answering when users approach the threshold.

In the second step, I use a dynamic discrete choice model (à la Rust 1987, but with continuous state space) to quantify the incentive effects and simulate counterfactuals. At each week of participation, 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 control rights, and allows preferences to depend on the degree of autonomy. 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 obtaining autonomy. Variation in the willingness to participate once endowed with autonomy identifies the *static incentive* effect. In addition to variables that capture the cost of participation, the utility function includes other sources of motivation

[4] https://ell.stackexchange.com/

potentially correlated with the threshold: the number of reputation 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 autonomy and a significant increase in willingness to participate in editing once endowed with control over the action. Interestingly, estimates suggest that the utility of answering increases as well when users gain control over editing. In particular, *Anonymous* users show the highest marginal cost of contribution in answering, but the cost is more than offset when they gain autonomy. *Informative* users are the most sensitive to the *dynamic incentive*, while *Specialised* users have the lowest costs in answering, and delegation does not affect much their answering activity.

With estimates from the model, I simulate counterfactual contribution histories under delegation designs that differ on the threshold reputation points requirement. In particular, I consider the case with a performance threshold equal to zero (full delegation), infinity (no delegation), or two intermediate levels, identifying more or less demanding targets. The results show that the platform maximises quality and quantity of contributions in answering by basing delegation on performance with a relatively low threshold. *Anonymous* users' participation is too costly to reach a more demanding threshold, and their contributions become sparse in that case. A design that excludes the possibility to delegate control removes any incentives for participation of *Anonymous* users and reduces participation of *Informative* users while not affecting much *Specialised* users.

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

The second contribution relates to the organisational implications of these non-monetary 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 identifies heterogeneity in the impact of incentives and shows that users' profile pages capture participation motives.

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

[5]Nonexperimental work has identified a beneficial effect of delegation on performance but does not investigate whether a channel is an intrinsic value for autonomy. Bandiera, Best, Khan, and Prat (2021) use a field experiment, while Liberti (2018) use real data from a financial institution.

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 setting and the rationale for delegation, 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. Section 6 reports the counterfactual analysis. Finally, section 7 concludes.

2 Stack Exchange: "Self Managed" Platforms

Stack Exchange is a family of 172 websites where users can freely and voluntarily ask and answer questions on a topic specific to each website. Participation does not involve any monetary transactions. The most well-known site is *Stack Overflow*, which hosts questions and answers about programming languages. These websites belong to a commercial company, which, as of July 2020, has raised 153 million dollars in venture capital and was sold in June 2021 for US1.8$ billion.[6] To give a sense of the welfare produced to consumers, in 2023 Stack Exchange received 418.8 million monthly visits and 806.3 million monthly page views. Users created 3.1 million questions, which received 3.5 million answers[7].

Participation in Stack Exchange is subject to an incentive system based on virtual rewards, either *badges* or *privileges*. Badges are comparable to medals or firms' bonuses and depend on the accomplishment of given performance targets. Privileges instead provide access to additional resources or actions, and, in general, to a more influential role in the community. Users achieve them sequentially, by accumulating reputation points. For instance, with 15 points, users achieve the possibility to upvote other users' contributions, while if they reach 20000 points, they achieve close to full administrative control of the site.[8] Users obtain reputation points in several ways, mostly from up-votes on their questions and answers, and from getting their answers *accepted* as the one solving the question. Users can also get a few points when they make suggested edits to other users' content, and the edits get approved. Edits are suggested until the user reaches the editing

[6] https://www.businesswire.com/news/home/20200728005330/en/Stack-Overflow-raises-85M-Series-funding-accelerate and https://www.prosus.com/news-insights/group-updates/2021/prosus-to-acquire-stack-overflow

[7] https://stackexchange.com/about

[8] The platform also delegates the website management through elections. At certain times, community members can vote to elect *moderators* who, once elected, jump at the top of the privilege hierarchy even if they do not satisfy the reputation requirement. Elected *moderators* keep their role permanently.

privilege, after which edits are directly implemented (i.e. do not require approval) and do not provide any point.[9]

2.1 Delegation of Control over Editing in Stack Exchange

The sequence of privileges that users can achieve by accumulating reputation points is comparable to a managerial hierarchy in the community. It allows the platform to delegate control over decision-making to volunteer users based on a performance indicator (reputation points). One particular example, which is the main focus of this paper, relates to decisions on the moderation of the platform through edits. Editing is the action of modifying existing questions and answers to improve or correct them. A user can always make edits, but these are not necessarily implemented. Users who have not achieved the *editing* privilege create *suggested edits*, which require approval by either the content creator or by users who already have the privilege. Users with the *editing* privilege, instead, can directly implement the edits. In other words, the platform uses the *editing* privilege to delegate control over the editing action. Delegation of such control is instrumental to the platform in two possible ways. On one side, the platform saves money as it does not need to hire personnel who review and approve suggested edits.[10] On the other side, it potentially create participation incentives. First, suppose users' willingness to make contributions to the platform is significantly higher when endowed with full control. In that case, delegation relaxes a user's participation constraint (Gibbons et al. (2013)).[11] I define this effect as the *static incentive*, as it is independent of the contribution dynamics. Second, tying delegation to performance incentivises participation if users value gaining control, as they would want to produce more to achieve it. I call this effect the *dynamic incentive*.

The platform faces a trade-off if it wants to leverage both the *static incentive* and the *dynamic incentive*. A positive *static incentive* effect would suggest delegating to every user to maximise edits, while a positive *dynamic incentive* effect would advise conditioning delegation on performance.

[9]In the appendix, figure 11 lists the rules to gain points, while table 10 reports the list of privileges and the reputation points necessary to obtain each of them.

[10]In 2020 Facebook employed about 15000 moderators who were considered insufficient: Charlotte Jee, MIT Technology Review, June 2020. More recently, a cost-saving reduction in the number of moderators of the platform X (previously Twitter) raised numerous concerns (https://www.theguardian.com/technology/2022/oct/28/twitter-takeover-fears-raised-over-disinformation-and-hate-speech)

[11]As noted by Sturm and Antonakis (2015), "*...research has shown that power increases an action orientation and, thus, leads directly to the taking of action for those who possess it...*".

3 Data

In this paper, I use data from the Stack Exchange website called *English Language Learners* (ELL), which focuses on questions and answers related to the use of English. This specific website is particularly suitable for the analysis for two reasons. First, posts contain only text, not equations or scripts, as in more technical Q&A. This allows us to measure the quality of the answers with text measures. Second, in the middle of the sample period, the site changed the reputation thresholds to achieve privileges, creating an additional variation on users' control over editing, as some users lost the *editing* privilege. Figure 1 reports the timing of the change, while figure 2 reports the number of users with the privilege over time.[12]

The data were retrieved on March 7^{th}, 2022, and contain the complete set of user profile pages, contributions (e.g. answers and edits), and users' reputation histories. I constructed a panel of users' weekly participation in the website by including users who contributed at least one answer or one edit.[13] Users' participation histories start with their first answer or edit and are assumed to end after three months of inaction in answering or editing. The panel includes the weekly number of contributed edits and answers, and the average quality of the published answers for 12141 users. I proxy for answer quality with the number of reputation points that an ordinary-least-squares model of textual characteristics predicts. Textual characteristics include the length of the answer, the number of links and images appearing in the answer, and their quadratic values.[14] The data are right-censored at the download date. Table 1 provides descriptive statics of user activity. As it is standard in online communities, participation is skewed, with a relatively small group of users contributing a substantial part of the site content. Consequently, users are heterogeneous in terms of the number of reputation points achieved and whether they have reached the threshold.

---

[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 site called Area 51 and starts contributing. When the website has enough demand and sustained activity within *Area 51*, the platform administrators launch it with an independent URL. The website enters the *Private Beta* period, where participation is limited to users who have contributed in the development stage and, soon after, the *Beta* period (initially called *Public Beta*), with open 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 personalised design. Normally, the *graduation* and the new design would occur on the same date, but on the ELL website, the design occurred later due to a backlog of the designer team. Once the website receives the new design, the reputation points required to obtain the privileges change. Figure 1 reports this timeline for the ELL site. Table 10 in the appendix reports the number of reputation points required to obtain each privilege and how that changed after the site design.

[13]I include all edits on questions' titles, questions' tags, and questions and answers' bodies.

[14]Section A.1 in the appendix provides details on the specification of the model and the estimated model parameters.

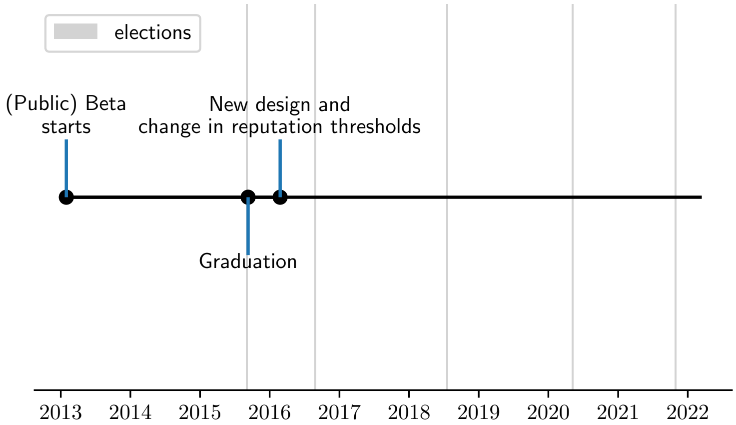

Notes. The site allows contributions from the broad public since a *Beta stage*. When it achieves sustained participation, it *graduates* and moves to more demanding reputation thresholds to achieve the privileges. The actual change in reputation thresholds happened when the site changed design. Elections allow a few users to be elected moderators and achieve most privileges without reputation requirements.

Figure 1: Timeline of the ELL website

3.1 User Types

The heterogeneity in user contribution patterns may suggest the existence of different reasons behind participation. In this context, a platform incentive design targeting the average consumer may not maximise participation and productivity, as different types of users may react to incentives differently. While the researcher does not directly observe why users participate, the information that users self-disclose on their profile pages may signal such motives. For instance, a user who contributes to Stack Exchange to show off her expertise as an English language expert has incentives to disclose personal information to be recognisable. At the same time, a user who answers questions for pure altruism may not mind leaving her profile empty.[15]

In practice, I identify user types with a data-driven approach that develops sequentially. First, I extract from user profile pages a set of categorical variables that describe what information users have disclosed about themselves. These include whether the user provides a full name, a biographical text and its length, a website, a location, and a LinkedIn profile and whether she includes the word *English* in her biography.[16] Second, I transform these categorical variables into

[15] Belenzon and Schankerman (2015) also exploit user choices to identify types. In their case, they rely on whether programmers contribute to more or less "open" open-source software.

[16] The inclusion of the term *English* suggests that the user possibly justifies her ability to

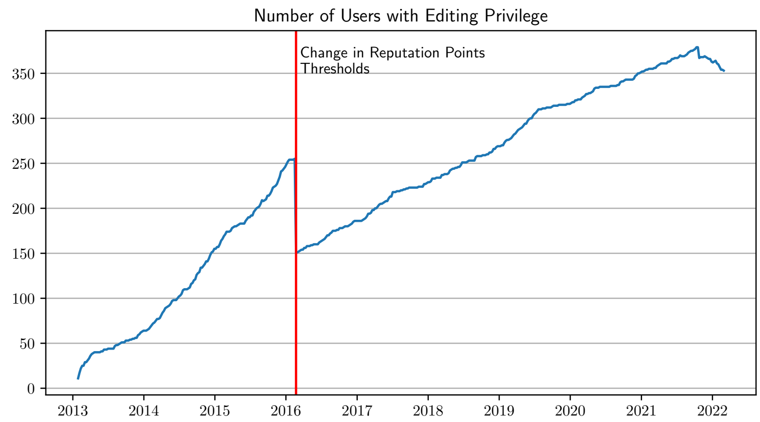

Notes. Number of users that achieved the editing privilege, net of users who left the platform. In February 2016, an increase in the requirement of points to obtain this privilege induced the loss of the privilege for some users. Data are available since the *Beta* period.

Figure 2: Users with editing privilege

quantitative measures and reduce their dimensionality through a Multiple Correspondence Analysis (MCA, Greenacre and Blasius 2006). Finally, I use a K-Means clustering algorithm to cluster users in groups.[17]

The procedure leads to the identification of three types of users, which I call *Anonymous*, *Informative*, and *Specialised*. *Anonymous* users have empty profile pages or provide very scarce information. *Informative* users provide partial information: mostly their location, some biographical notes, and a website, like a LinkedIn profile. *Specialised* users, instead, tend to have longer biographical text. They crucially differ from *Informative* users as they always include the term *English* in their self-description, suggesting more commitment to the specific topic of the site. Table 2 reports to what extent each type of user provided specific pieces of information.

Different types of users also behave differently. Table 3 reports descriptive statistics on users' contributions by user type. It emerges that *Anonymous* users are the least productive, while *Specialised* users are the most productive and likely to reach the editing threshold, confirming the hypothesis that they are the most committed to the platform. *Informative* users appear to be in-between.

answer questions on the use of the language (e.g. "*...language learner (C2 in English[...])...*" [user 740]).

[17]Section A.3 in the appendix provides details on the implementation of this procedure.

	mean	std	min	median	max
Number of Weeks of Activity	138.45	137.47	1.00	86.00	476.00
Amount of Reputation Points Reached	480.42	3516.62	0.00	110.00	179616.00
Number of Answers	11.55	101.30	0.00	1.00	4850.00
Number of Edits	5.68	106.86	0.00	0.00	6167.00
Average Answer Quality	13.50	1.53	5.52	13.26	41.85
Reached Editing Privilege	0.04	0.19	0.00	0.00	1.00

Notes. Descriptive statistics of user activity in the website, for those users included in the panel. Statistics are at the user level. The variable *Reached Privilege* takes the value 1 if the user reaches the threshold to achieve the *editing* privilege, so its mean value corresponds to the share of users who achieved the privilege. Statistics on answer quality are conditional on a positive number of contributed answers.

Table 1: Descriptive Statistics

	All Sample	Anonymous	Informative	Specialised
Share of Users with Full Name	26.55%	23.89%	31.35%	31.03%
Share of Users with Website	23.83%	4.4%	57.54%	67.98%
Share of Users with Location	40.22%	11.61%	93.74%	75.69%
Share of Users with LinkedIn	1.32%	0.0%	4.07%	0.79%
Share of Users without Bio	61.67%	90.95%	10.22%	0.0%
Share of Users with Short Bio	17.68%	6.06%	40.22%	26.09%
Share of Users with Long Bio	20.65%	2.99%	49.56%	73.91%
Share of Users with Term *English* in Bio	4.17%	0.0%	0.0%	100.0%
Sample Size	12141	7801	3834	506

Notes. Share of users who have specific information on their profile pages, in the whole sample, and in each user type-specific sub-sample. *Bio* stands for Biography, i.e. the *About Me* section of the user profile.

Table 2: User Types Characteristics

		mean	std	min	median	max
Anonymous	Weeks of Activity	118.92	126.95	1.00	59.00	476.00
	Reputation Points Reached	256.30	1284.85	0.00	80.00	65197.00
	Number of Answers	6.30	48.22	0.00	1.00	2681.00
	Number of Edits	1.10	15.52	0.00	0.00	609.00
	Avg. Answer Quality	13.39	1.42	6.77	13.22	41.63
	Reached Privilege	0.02	0.14	0.00	0.00	1.00
Informative	Weeks of Activity	170.04	147.18	1.00	135.50	476.00
	Reputation Points Reached	661.67	4705.15	0.00	138.00	179616.00
	Number of Answers	15.06	122.72	0.00	1.00	4171.00
	Number of Edits	9.78	158.57	0.00	0.00	6167.00
	Avg. Answer Quality	13.69	1.70	5.52	13.34	41.85
	Reached Privilege	0.05	0.22	0.00	0.00	1.00
Specialised	Weeks of Activity	200.34	153.75	3.00	180.00	476.00
	Reputation Points Reached	2562.29	9919.16	2.00	237.50	139682.00
	Number of Answers	65.82	304.92	0.00	5.00	4850.00
	Number of Edits	45.17	279.15	0.00	0.00	4474.00
	Avg. Answer Quality	13.85	1.68	11.02	13.41	29.53
	Reached Privilege	0.18	0.39	0.00	0.00	1.00

Notes. User-level descriptive statistics on platform contributions by user type. The variable *Reached Privilege* takes value 1 if the user reaches the threshold to achieve the *editing* privilege. It follows that the mean value corresponds to the share of users who achieved the privilege.

Table 3: Descriptive Statistics by Type

4 Preference for Control and the Static Incentive Effect

Reduced-form evidence shows that users are more willing to contribute edits after achieving the editing threshold, suggesting they are sensitive to the static incentive effect. This is possible to see in figure 3, which reports estimates of the following regression:

$$Y_{it} = \alpha_i + \gamma_t + \beta_{r_{it}-\bar{R}} + \boldsymbol{W}'_{it}\boldsymbol{\rho} + \varepsilon_{it} \tag{1}$$

where i indexes users, and t indexes week. Y is either the number of edits (left graph in figure 3) or a dummy equal to 1 if the user contributed any edits (right graph in figure 3). 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 week t, and $\bar{R}$ the number of reputation points required to obtain the editing privilege (i.e. 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 $\bar{R}$.[18] Finally, I include control variables. One set of variables aims to control for other drivers of editing activity. It includes a dummy equal to 1 when the user is an elected moderator, a dummy equal to 1 when the user is a candidate in a moderator election, and dummies equal to one in the week the user has achieved editing-related badges.[19] A second set of variables aims to control for the user's time availability and includes the number of answers and comments produced. The sample includes only users who achieved the editing privilege. This prevents selection effects, as not all users have reached the threshold.

Figure 3 reports estimates of the parameters $\{\beta_{r-\bar{R}}\}_{\forall r}$. They show that users are more likely to participate in editing content and edit more when they have the *editing* privilege. These results suggest that users prefer contributing edits if they have full control over their implementation.[20]

Concerning the *dynamic incentives*, a reduced-form analysis with standard tools cannot identify the effect of the threshold as it cannot account for forward-

[18]In practice, I bin the reputation points missing to reach the thresholds in 100-points intervals, with the first and last bins defined such that $r - \bar{R} < -600$ and $r - \bar{R} \geq 700$ respectively.

[19]Badges are sort of virtual medals that users receive when they accomplish specific activity targets. The editing-related badges are the *Copy Editor* and *Strunk & White* badges (https://ell.stackexchange.com/help/badges).

[20]Table 8 in the appendix provides the complete set of estimates. It also provides estimates for a placebo regression where the outcome variable is the number of weekly comments written by the user. The *editing* privilege does not affect the degree of control over the action of commenting. Estimates show no change in the intensity and willingness to contribute comments around the achievement of the privilege. This result suggests that the increase in editing activity is not driven by a more general increase in commitment to the platform.

looking behaviour.[21] Nevertheless, an estimation of model 1 with the number of contributed answers as an outcome variable suggests that users increase contributions and are more likely to participate in answering before they approach the threshold, as shown in section A.5 in the appendix.

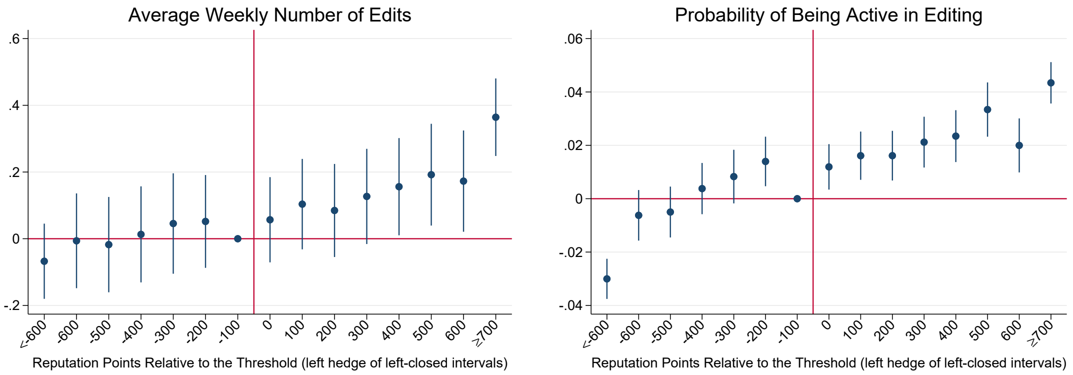

Notes. Reputation points fixed effects before and after achieving the editing privilege. Sample of users who reached the threshold.

Figure 3: Editing Contributions Relative to Achieving the Privilege

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 behaviour.

To overcome these limitations, I develop a dynamic discrete choice model that studies intertemporal choices and accounts for forward-looking behaviour. 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 maximisation and, as such, provide information on users' preferences. In the context of participation in online communities, users choose their efforts 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 privileges.[22]

[21] Goes et al. (2016) address this problem relying on functional form assumptions and modifying the data to account for forward-looking behaviour.

[22] Note that the application of dynamic discrete choice models to this context is conceptually

5.1 Users' Participation Choices

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:

$$\boldsymbol{\alpha_t} = \begin{bmatrix} NA_t \\ QA_t \\ NE_t \end{bmatrix} \ni \mathcal{A}$$

where NA indicates the number of answers, QA denotes the average quality of answers, and NE indicates the number of edits. $\mathcal{A}$ represents the choice set, including all possible combinations of effort levels in the two tasks.[23]

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

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

where δ is a discount factor and $U_{it}(\boldsymbol{\alpha_t})$ is the flow utility that user i receive at each period of participation.

Every period of participation proceeds as follows:

1. the user observes the values of the states realised at the end of the previous period, which includes the total number of reputation points she has obtained in the past, how many privileges she has collected, whether she has already achieved the *editing* privilege, the number of questions available to answer, and her experience in terms of time spent on the website and the number of contributions. Note that, excluding the user's experience, all states are a function of the reputation points achieved,

2. she forms beliefs over the value of the states that may be realised in the next periods, conditional on past choices and the possible new contribution choice she could make,

similar to works that study dynamic investment decisions with discrete choice models. A typical application is human capital investment decisions. Examples of this literature are Arcidiacono, Aucejo, Maurel, and Ransom (2016) and De Groote (2024).

[23] Effort levels are discretised to limit the computational burden. In practice, $NA \in \{0, 1, 7\}$, $QA \in \{0, 12.69, 13.4, 15.13\}$, and $NE \in \{0, 1, 9\}$, resulting in 21 possible combinations of efforts (note that while it is allowed to have effort in only edits or only answers, a positive number of answers requires a positive answer quality, and vice-versa). More details are provided in section A.6.2 in the appendix.

3. she makes a contribution decision in editing and answering,

4. the flow payoff realises,

5. the states update to their new values.

The per-period flow utility is defined as:

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

where R_{it} are user's reputation points realised at the end of period $t-1$, $cumT_{it}(R_{it})$ is the total number of privileges achieved, and $Control_{it}(R_{it})$ is a dummy equal to 1 if the user achieved the *editing* privilege.[24] The variables C_{it}^A and C_{it}^E are direct net utilities (or net costs) of answering and editing respectively. They are defined as:

$$\begin{aligned} C_{it}^A &\equiv QA_{it} + NA_{it}^{scarsity_{it}}, \\ C_{it}^E &\equiv NE_{it}, \end{aligned}$$

where $NA^{scarcity}$ is the number of answers raised to a measure of the scarcity of questions to answer.[25] Finally, ε_{it} is an idiosyncratic choice-specific preference shock.

The parameter β_0 captures the marginal utility of accumulating reputation points. A positive estimate would suggest that users either enjoy collecting points per se, as it would happen if they treat the site as a video game, or benefit from reputation points externalities (e.g. if points signal ability to employers). The parameters associated with C_{it}^A and C_{it}^E capture instead direct utility from making a certain contribution choice, including a cost of effort and some intrinsic benefit from the action (e.g. if the user is altruistic or enjoys participating per se). The parameters β_4, β_5, and β_6 together capture the user's marginal utility from the acquisition of control and, as a consequence, how the user responds to the *dynamic incentives*. β_6 informs on the *static incentive* effect, as it captures changes in the willingness to make edits after the user achieves the *editing* privilege.

[24] The values of these variables do not depend on period t choices, but on past choices only. A static model would not be able to identify their respective parameters.

[25] The variable *scarcity* captures the inverse of the availability of questions to answer. It takes values in $[1, \infty]$ so that when there are many questions to answer, the cost of answering tends to be linear in the number of answers, while with fewer questions available, the cost becomes increasingly convex. Details on the construction of the *scarcity* variable are in section A.2 in the appendix.

5.2 Beliefs

Users form beliefs and expectations over the evolution of the state space, given the contribution choices they make. In this section, I make assumptions on how users form such expectations. Section 5.5.1 provides estimates of the parameters involved with these processes.

5.2.1 Evolution of Reputation Points

Users gain reputation points mostly from up-votes received on their answers and lose them when they receive down-votes. Votes may arrive the same week the user publishes the answer or later. Suggested edits also provide reputation points if and when approved. I make assumptions on the process of arrival of votes and edit approval as a function of the user's effort and community edits of the user's answers (which can affect the answers' quality).

Consider the beliefs that the user forms in the first period of participation t_0. Each answer j that the user publishes in period t_0 receives, at the end of the period, a number of community edits and votes that follow a Poisson process:

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

The expected values of these random variables are:

$$\lambda_{E,j,t_0} = \exp\left(\beta_0 + \beta_1 QA_{t_0} + \beta_2 Seniority_{t_0} + \beta_3 Practice_{t_0}\right), \tag{3}$$

$$\lambda_{U,j,t_0} = \exp\left(\gamma_0 + \gamma_1 QA_{t_0} + \gamma_2 \lambda_{E,j,t_0} + \gamma_3 Seniority_{t_0} + \gamma_4 Practice_{t_0}\right), \tag{4}$$

$$\lambda_{D,j,t_0} = \exp\left(\delta_0 + \delta_1 QA_{t_0} + \delta_2 \lambda_{E,j,t_0} + \delta_3 Seniority_{t_0} + \delta_4 Practice_{t_0}\right), \tag{5}$$

where QA_{t_0} is the average quality of the user's answers published in period t_0, and *Seniority* and *Practice* are measures of the user's experience. *Seniority* is the number of days the user has been participating on the website, and *Practice* is the cumulative number of answers she has published (both zero if $t = t_0$).

If the user, in her first period of participation, published NA_{t_0} answers, then she will expect to receive by the end of the period a number of up-votes and down-votes as follows:

$$\begin{aligned}
\Lambda_{U,t_0} &= NA_{t_0} \lambda_{U,j,t_0}, \\
\Lambda_{D,t_0} &= NA_{t_0} \lambda_{D,j,t_0}.
\end{aligned}$$

I model the number of approved edits, out of NE_{t_0} contributed suggested edits, as a binomial process, such that:

$$\text{ApprovedEdits}_{t_0} \sim \mathscr{B}(NE_{t_0}, \pi).$$

It follows that the expected number of reputation points (ρ) that the user expects to receive at the end of period t_0 is given by[26]:

$$\mathbb{E}[\rho_{t_0}|\boldsymbol{\alpha}_{t_0}] = 10\Lambda_{U,t_0}(NA_{t0}, QA_{t0}) - 2\Lambda_{D,t_0}(NA_{t0}, QA_{t0}) + 2\pi NE_{t_0}.$$

The answers produced in period t_0 may also induce the arrival of up-votes and down-votes in the following periods. The process is deterministic and follows an exponential decay.[27] Let Δt be the number of days passed from the publication day, such that if $t = t_0 + 1$, then $\Delta t = 1$. Then,

$$\lambda_{U,j,t_0+\Delta t} = \lambda_{U,j,t_0} \exp\left(\frac{-\Delta t}{\tau_U}\right),$$
$$\lambda_{D,j,t_0+\Delta t} = \lambda_{D,j,t_0} \exp\left(\frac{-\Delta t}{\tau_D}\right).$$

$\lambda_{U,j,t_0+\Delta t}$ 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.

If the user chooses positive effort in several periods, these processes aggregate. In general, the expected number of up-votes and down-votes arriving at the end of a given period t is, respectively,

$$\Lambda_{U,t} = \Lambda_{U,t-1} \exp\left(\frac{-1}{\tau_U}\right) + NA_t\lambda_{U,j,t}(QA_t),$$
$$\Lambda_{D,t} = \Lambda_{D,t-1} \exp\left(\frac{-1}{\tau_D}\right) + NA_t\lambda_{D,j,t}(QA_t),$$

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

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

To conclude, let R_{t-1} 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

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

[26] One up-vote gives 10 points, one down-vote removes 2 points, and the approval of a suggested edit gives 2 points.

[27] Alternative assumptions give similar results. Section A.6.3 in the appendix shows how different functional forms fit the data.

5.2.2 Evolution of Questions' Availability

The availability 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

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

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. As stated in equation 2, effort choices affect users' utility directly through C^A and C^E. The identification of marginal utilities of these direct effects relies on a standard static conditional logit model (McFadden 1974). At the same time, positive effort leads to the expectation of obtaining more reputation points in the future and, as a consequence, of a higher probability of achieving the *editing* privilege. In other words, choices affect users' value function. Such variations allow the identification of reputation-points-related marginal utilities, which would not be identified in a static model. Most importantly, they allow to identify users' value in achieving the *editing* privilege.

5.4 Estimation

The estimation proceeds in several steps. First, I set the discount factor at 0.95. Second, I estimate the parameters that affect users' beliefs over the evolution of the state space, either in reduced form or nonparametrically.[28] Third, I estimate the preference parameters following Arcidiacono and Miller (2011).[29] The estimation relies on the conditional logit assumptions and assumes that the idiosyncratic preference shocks follow an extreme value Type 1 distribution.

The estimation technique allows us to maintain computational feasibility while not restricting the state space (which includes continuous variables) and allowing for a relatively large choice set (21 choices). I achieve this in the following way. First, I exploit the Conditional Choice Probabilities (CCP) estimator (Hotz and Miller 1993) and compute, in a first step, the reduced-form probability of making

[28] Section section 5.2 describe those processes.

[29] The work by Arcidiacono and Miller (2011) builds on an extensive econometric literature, e.g. Rust (1987), Hotz and Miller (1993), and Magnac and Thesmar (2002).

contribution choices given certain values of the states. I estimate such probabilities with an l2-penalised logit model on scaled data between 0 and 1.[30] Second, I compute value functions exploiting *finite dependence* (Arcidiacono and Miller 2011). This is possible because, given the assumptions on the evolution of the state space, positive effort in a period t followed by no contributions leads to the arrival of reputation points with positive probability only for a few periods ahead. For an intuition of how *finite dependence* works, consider the example where a user contributes two answers in period t and expects these answers to receive up-votes and down-votes in the next three periods. Suppose the user stops contributing in the next four periods. In that case, the expected number of reputation points that she will have at $t+4$ is equivalent to those she would have had by contributing nothing in period t, two answers in period $t+1$, and nothing again up to period $t+4$ (i.e. to a relative path of contributions where the positive effort is shifted by one period). Consequently, by differentiating the value functions of the two choice paths, the utility coming from actions after period $t+4$ cancels out. Exploiting this feature, a full solution of the model is not necessary, and I can compute the value function using only the expected utilities along the selected choice paths.[31] Section A.6.1 in the appendix provides a formal definition of *finite dependence* and the derivation of the log-likelihood function. Finally, I exploit the assumption of linearity of the utility function to compute expected values of the states outside of the estimation algorithm. This implies that I do not need to compute state-transition probability matrices or reduce the state space's dimensionality.

5.5 Results

5.5.1 First-stage Estimates of Beliefs Parameters

Table 4 reports estimates of the parameters that drive users' expectations on the evolution of the state space, given their choices. The first set of estimates relates to the number of edits a user expects to receive from the community on an answer she publishes. Estimation exploits a Generalised Least Squares (GLS) Poisson

[30] Section A.6.4 in the appendix provides additional details.

[31] This approach has advantages and disadvantages. Apart from reducing the computational burden, it relaxes some assumptions as it does not require fixing a terminal period and does not strictly impose rational and perfect anticipation of utility in the far future. On the downside, whenever returns are not smooth but are step functions, which is the case for $cumT$ and $Control$, only data of users who may experience variation in returns in those few periods ahead can identify their marginal utilities. For users' choices to be informative of users' marginal value of $Control$, users need to have already enough reputation points such that different levels of effort have different impacts on the probability of achieving the privilege in the next few periods. In practice, this implies that the estimation of the coefficient of $Control$ relies on relatively small variation.

model. Naturally, a higher-quality answer requires fewer edits. Estimates confirm this intuition, as higher quality and higher user experience lead to fewer edits. The second and third sets of estimates identify predictors of up-votes and down-votes. Again, estimates reflect what one would expect. Higher answer quality and higher user experience lead to more up-votes and less down-votes. Community edits instead increase the chances of both up-votes and downvotes. A possible justification is that community edits improve the answer, attracting more up-votes, but, at the same time, are more likely to occur on low-quality answers, which attract down-votes. These models are also estimated through a GLS Poisson model.

The parameter π reports the rate at which suggested edits get approved. The estimate, which is simply the average approval rate in the data, indicates that most suggested edits get approved. The parameters τ_U and τ_D characterise the exponential decay rate at which up-votes and down-votes continue to arrive in the following weeks after the publication of the answer. The small value of these parameters suggests a very steep decrease in votes with time. Finally, the rate of question availability indicates the availability of questions has substantially increased over time.

5.5.2 Flow Payoff Parameters

Table 5 reports the estimates of the flow payoff parameters. The first column reports estimates for the full sample, while the other columns report estimates by type of users.[32] Users have a positive marginal utility from accumulating reputation points and experience a direct cost from contributing in answering and editing. Nevertheless, the achievement of the editing privilege offsets such costs. The acquisition of control over edits substantially increases the utility of contributing edits, and has a positive externality on the utility of contributing answers. In addition, the positive coefficient of *Control* suggests that users have a positive marginal utility from crossing the threshold and gaining the *editing* privilege. These effects, which are consistent across user types, suggest that the delegation of control over the editing task substantially impacts users' participation preferences.

Estimates suggest heterogeneity across different types of users. *Specialised* users have the lowest direct cost from answering, and the achievement of the *editing* privilege seems to affect them relatively less. Since they are users that include the word *English* in their biography, they possibly have higher expertise and may be interested in participating independently of the incentive designs. On the contrary, the *editing* privilege most affects the *Anonymous* users, who also

[32]State-transition parameters are assumed to be the same across user types, while *CCP* and flow payoff estimates are specific to each user type data.

have the highest preference for accumulating reputation points, suggesting high sensitivity to the platform design. Informative users are somehow in-between.

	Estimate	Std. Error
$\lambda_{E,j,t} = \exp\left(\beta_0 + \beta_1 QA_t + \beta_2 Seniority_t + \beta_3 Practice_t\right)$		
β_0	-2.7154	0.0835
β_1	-0.0089	0.006
β_2	-0.0006	0.00003
β_3	-0.0002	0.00003
$\lambda_{U,j,t} = \exp\left(\gamma_0 + \gamma_1 QA_t + \gamma_2 \lambda_{E,j,t} + \gamma_3 Seniority_t + \gamma_4 Practice_t\right)$		
γ_0	-0.5276	0.0114
γ_1	0.0553	0.0008
γ_2	0.476	0.0079
γ_3	0.0001	0.000004
γ_4	0.00004	0.000003
$\lambda_{D,j,t} = \exp\left(\delta_0 + \delta_1 QA_t + \delta_2 \lambda_{E,j,t} + \delta_3 Seniority_t + \delta_4 Practice_t\right)$		
δ_0	-1.6716	0.0523
δ_1	-0.0468	0.0038
δ_2	0.6565	0.0268
δ_3	-0.0002	0.00002
δ_4	-0.0003	0.00002
$\text{ApprovedEdits}_t \sim \mathscr{B}(NE_t, \pi)$		
π	0.8115	
$\Lambda_{U,t} = \Lambda_{U,t-1} e^{-\frac{1}{\tau_U}} + NA_t \lambda_{U,j,t}(QA_t)$		
τ_U	0.2297	
$\Lambda_{D,t} = \Lambda_{D,t-1} e^{-\frac{1}{\tau_D}} + NA_t \lambda_{D,j,t}(QA_t)$		
τ_D	0.2463	
$avail_t = \nu_0 + \nu_1 t + \epsilon_t$		
ν_1	108.6329	0.5489

Notes. Estimates of parameters governing belief formation processes over the state space's evolution. $\{\beta\}$, $\{\gamma\}$, and $\{\delta\}$ parameters are estimated with Generalised Least Squares Poisson models and capture, respectively, determinants of the number of edits on own answers, the number of up-votes, and the number of down-votes. π is estimated non-parametrically and captures the probability a suggested edit gets approved. τ_U and τ_D are the decay time of up-votes and down-votes arriving on answers since answers' publication and are estimated via non-linear least squares. ν_1 is the rate at which questions' availability grows over time and is estimated via ordinary least squares.

Table 4: First-Stage Estimates

	(All Sample)	Anonymous	Informative	Specialised
R	0.00880	0.00949	0.00865	0.00695
	(0.000167)	(0.000309)	(0.000246)	(0.000336)
C^A	-0.234	-0.304	-0.172	-0.0331
	(0.00931)	(0.0140)	(0.0143)	(0.0219)
C^E	-2.046	-2.345	-1.977	-2.403
	(0.109)	(0.245)	(0.139)	(0.258)
Tcum	-0.541	-0.615	-0.509	-0.224
	(0.0159)	(0.0224)	(0.0273)	(0.0396)
Control	1.285	1.128	1.489	0.549*
	(0.111)	(0.210)	(0.157)	(0.244)
$C^A \times$ Control	0.370	0.469	0.307	0.134
	(0.0106)	(0.0189)	(0.0158)	(0.0242)
$C^E \times$ Control	2.575	3.501	2.248	3.162
	(0.117)	(0.278)	(0.152)	(0.266)
N	1680904	927651	651890	101363
N. Users	12141	7801	3834	506

Standard errors in parentheses

Notes. Estimates of structural parameters on the whole sample and by type of user.

Table 5: Dynamic discrete choice model estimates

6 Counterfactual Delegation Designs

As the allocation of control over the editing task affects users' participation preferences, the quality and quantity of contributions depend on the delegation design. In this section, I simulate counterfactual contribution histories under alternative delegation thresholds. In particular, I compare four designs: one that provides full delegation to every user unconditionally on effort, a design that never delegates, and two designs that allocate control at different levels of productivity (measured in reputation points). Note that all these scenarios are realistic. Wikipedia is a leading example of the case in which agents have full control. On Wikipedia, every 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 achieving 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 assume a fixed amount of periods of participation (60) and solve by backward induction users' maximisation problem. In addition, I add restrictions on the dimensionality of the state space by limiting the possible values of the state variables and reducing the returns of effort. Most importantly, the maximum number of reputation points that users can accumulate is 2000, with up-votes giving 5 reputation points and down-votes removing 1 point. Details on the restrictions to the dimensionality of the state variables are in section A.7.1 in the appendix.

The simulation proceeds in three steps. First, I compute the choice-specific transition probabilities. These matrices provide the probability of transitioning from each possible combination of state values to all possible future combinations of 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 that may realise due 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 in the platform. Second, I compute the value function backwards, starting from the last period. Finally, in the third step, I forward-simulate the decisions in each period for 100 users who share the same preferences but experience different choice-specific shocks.

The simulations differ on the number of reputation points necessary to achieve the *editing* privilege.[33] The threshold levels are 0, 750, 1500, and 10000. The first scenario represents a context of full delegation where users have control over the editing task since the beginning of their participation history. The second and third scenarios condition delegation to a, respectively, easier and harder performance target. Finally, the threshold cannot be reached in the last scenario, representing an environment with no delegation.

6.1 Counterfactual Results

Figures 4 and 5 report, respectively, the average number and quality of answers that users would produce across their history of participation under different delegation designs. The red (light) shade identifies periods in which users have the privilege, while the blue (dark) shade identifies periods in which users reached the maximum possible amount of reputation points.

Concerning the number of answers contributed, the *dynamic incentive* is not very effective. Under a full delegation scenario, users already sustain a substantially high number of contributions. This is because they receive positive direct utility when contributing with control rights on editing. By delaying delegation to a relatively reachable threshold, participation remains substantially similar as users expect to obtain the privilege relatively soon. Nevertheless, if the threshold is very demanding, users may not be able to reach it. Interestingly, this differs across different types of users. *Anonymous* users find it too costly to reach the threshold and keep a relatively low intensity of contributions. On the contrary, *Informative* and *Specialised* users contribute a lot to achieve it. Finally, with no prospects of receiving control rights, contributions are much lower, and only *Specialised* users sustain a more significant contribution.

The quality choice shows a different pattern, suggesting that *dynamic incentives* are effective on answer quality. Indeed, it is possible to notice that with the low-threshold scenario, users increase the quality of their contributions just before they reach the threshold.

Table 6 instead reports the total quantity of answers, their average quality, and the total number of edits that users would contribute overall in each scenario. It shows that the platform would maximise the quantity and quality of answers by setting a positive delegation threshold, but not too large. A positive threshold allows the exploitation of the dynamic incentive effect, inducing users to sustain more significant and higher-quality contribution patterns until they obtain the privilege. Nevertheless, any delay in allocating control over editing reduces the production of edits. The platform then faces a tradeoff between incentivising higher quality answering or incentivising editing.

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

The scenario with no delegation is inferior to any other option. Indeed, it eliminates the beneficial effects of both the *dynamic incentive* and the *static incentive*.

User Type	Delegation at:	Answers	Δ Answers	Quality	Δ Quality	Edits	Δ Edits
All Users	0 Points	37669		14.11		52834	
	750 Points	37793	+0.33%	14.12	+0.11%	39429	-25.37%
	1500 Points	35746	-5.1%	14.08	-0.2%	24170	-54.25%
	No Delegation	309	-99.18%	13.51	-4.24%	662	-98.75%
Anonymous	0 Points	38342		14.16		54000	
	750 Points	38239	-0.27%	14.20	+0.27%	40244	-25.47%
	1500 Points	22735	-40.7%	14.16	-0.04%	15984	-70.4%
	No Delegation	226	-99.41%	13.51	-4.56%	585	-98.92%
Informative	0 Points	37603		14.09		45438	
	750 Points	37778	+0.47%	14.10	+0.06%	33534	-26.2%
	1500 Points	37719	+0.31%	14.07	-0.2%	21961	-51.67%
	No Delegation	16228	-56.84%	13.82	-1.91%	689	-98.48%
Specialised	0 Points	36790		14.03		53822	
	750 Points	36725	-0.18%	14.10	+0.51%	39989	-25.7%
	1500 Points	36861	+0.19%	14.13	+0.7%	26562	-50.65%
	No Delegation	32667	-11.21%	13.90	-0.95%	511	-99.05%

Notes. Total number of contributions in answers and edits, and average quality of answers predicted by the simulations. Columns starting with Δ report changes in contributions with respect to a design with full delegation at zero reputation points.

Table 6: Overall Contributions in the Site, by Delegation Design

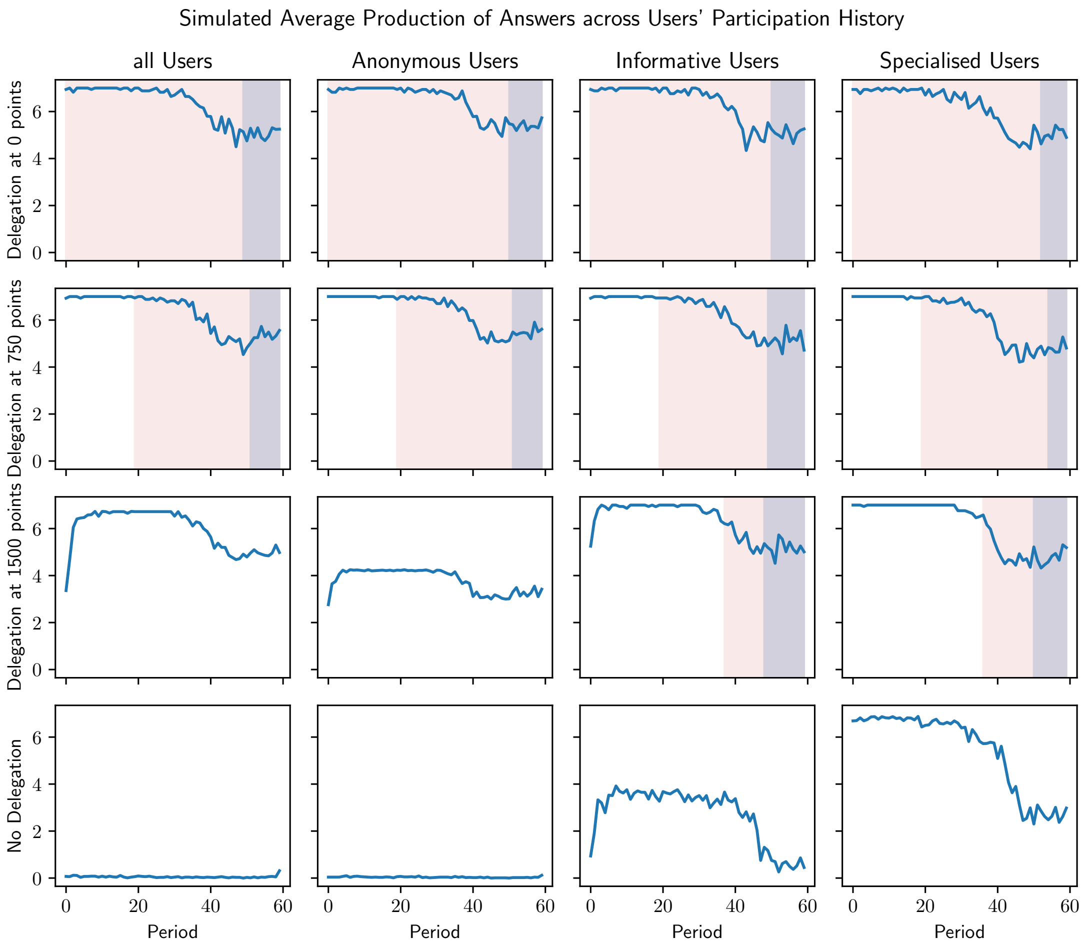

Notes. Each panel reports the average number of answers that users would contribute, given the preferences specific to their type and the delegation threshold. The red (light) shaded area corresponds to periods when all users have achieved the *editing* privilege, while the blue (dark) shaded area identifies periods when users have reached the limit amount of reputation points (i.e. 2000). Users are allowed to publish 0, 1, or 7 answers.

Figure 4: Simulated Contribution in Answering

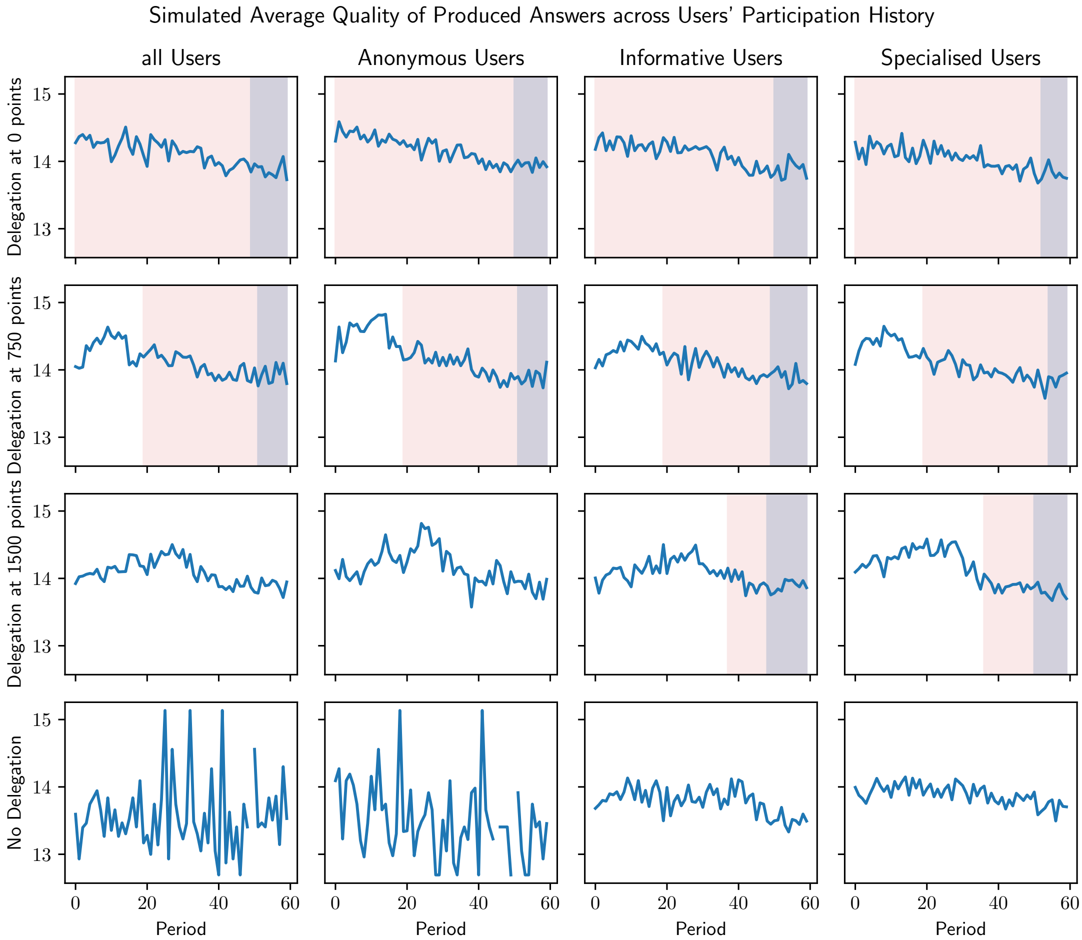

Notes. Each panel reports the average quality of answers that users would contribute, given the preferences specific to their type and the delegation threshold. The red (light) shaded area corresponds to periods when all users have achieved the *editing* privilege, while the blue (dark) shaded area identifies periods when users have reached the limit amount of reputation points (i.e. 2000). Users are allowed to publish answers with quality levels equal to 12.69, 13.41, or 15.13.

Figure 5: Simulated Average Quality of Contributed Answers

7 Conclusion

In this paper, I show that, in online communities, users value the allocation of control rights on actions. 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 about the task. The paper finds that users post significantly more edits if they 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 for incentives and the optimal delegation structure (Liberti 2018, Bandiera et al. 2021). Interestingly, allocating autonomy on a task has a positive externality on other tasks. The paper finds evidence that the cost of answering reduces when users gain autonomy over editing. These results contribute to the literature on multitasking (Holmstrom and Milgrom 1991), suggesting that incentives may not backfire in these contexts.

Second, the paper finds heterogeneity in the value of acquiring control rights. *Anonymous* and *Informative* users are motivated by the acquisition of autonomy and contribute little in scenarios where they would never be granted full control. In contrast, *Specialised* users appear to be more committed to the platform and less sensitive to incentives.

These results have important implications for platform design and the creation of user-generated public goods. Platforms can exploit the *static incentive* effect to maximise the contribution of edits. In an environment where peer moderation is essential (e.g. because of the crucial importance of quality rather than quantity of content), platforms should delegate autonomy to all users independently of performance measures. The *dynamic incentive* has no impact on participation in editing, and, as a consequence, there is no good reason to delegate based on performance. A full delegation design may be optimal in settings where, like in Wikipedia, it is important to maximise content quality, which is not necessarily objective, and peer evaluations are fundamental. Nevertheless, when a second task suffers from free-riding (e.g. answering in Stack Exchange), delegation and commitment to allocating authority based on performance can help incentivise contributions in that task and its quality. In that case, the design should identify an optimal threshold that is positive but reachable, given users' contributing costs.

In the setting analysed, a no-delegation design is never optimal. Nevertheless, this paper abstracts from environments where users are in competition or conflict with each other while focusing on platforms where users are generally aligned on the objective of creating useful knowledge. In the presence of conflict, like on social media platforms (e.g., Facebook or Twitter/X), edits may not be constructive, leading the platform to retain control over moderation.

References

Arcidiacono, P., E. Aucejo, A. Maurel, and T. Ransom (2016). College attrition and the dynamics of information revelation. *Working Paper*. 15

Arcidiacono, P. and P. Ellickson (2011, September). Practical methods for estimation of dynamic discrete choice models. *Annual Review of Economics 3*, 363–394. 43

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. 5, 19, 20, 25, 41, 42, 43

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. 6

Bandiera, O., M. C. Best, A. Q. Khan, and A. Prat (2021, November). The allocation of authority in organizations: A field experiment with bureaucrats. *The Quarterly Journal of Economics 136*(4), 2195–2242. 5, 30

Bartling, B., E. Fehr, and H. Herz (2014, November). The instrinsic value of decision rights. *Econometrica 82*(6), 2005–2039. 3, 5

Beckmann, M. and M. Kräkel (2022, October). Empowerment, Task Commitment, and Performance Pay. *Journal of Labor Economics 40*(4), 889–938. 2, 3

Belenzon, S. and M. Schankerman (2015). Motivation and sorting of human capital in open innovation. *Strategic Management Journal 36*, 795–820. 2, 9

Benson, A., D. Li, and K. Shue (2019, November). Promotions and the peter principle. *The Quarterly Journal of Economics 134*(4), 2085–2134. 6

Besley, T. and M. Ghatak (2008, May). Status incentives. *The American Economic Review: Papers & Proceedings 98*(2), 206–211. 2

Bester, H. and D. Krähmer (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*. 44

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. (2024). Dynamic effort choice in high school: Costs and benefits of an academic track. *Journal of Labor Economics*. 15

Fairburn, J. A. and J. M. Malcomson (2001, January). Performance, promotion, and the peter principle. *The Review of Economic Studies 68*(1), 45–66. 6

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. 6

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, 14

Greenacre, M. and J. Blasius (Eds.) (2006). *Multiple Correspondence Analysis and Related Methods*. Statistics in the Social and Behavioral Sciences Series. Chapman & Hall/CRC. 10, 36

Holmström, B. (1999, January). Managerial incentive problems: A dynamic perspective. *Review of Economic Studies 66*(1), 169–182. 3

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. 30

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. 19, 42

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

Liberti, J. M. (2018, August). Initiative, incentives, and soft information. *Management Science 64*(8), 3469–3970. 5, 30

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. 19

McFadden, D. (1974). *Conditional Logit Analysis of Qualitative Choice Behavior*. Academic Press. 19

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

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. 4, 19

Sturm, R. E. and J. Antonakis (2015, January). Interpersonal power: A review, critique, and research agenda. *Journal of Management 41*(1), 136–163. 7

Xu, L., T. Nian, and L. Cabral (2020, February). 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

A.1 Construction of a Proxy for Answer Quality

I define the *quality* of an answer as the predicted number of reputation points that the answer is expected to receive, given its text characteristics. Let $\boldsymbol{X}_j$ be a vector of text characteristics of an answer j, based on the content of the answer right after publication (i.e. before any edits occur). Let also $points_j$ be the number of reputation points allocated to the author of answer j on the day of publication of answer j and because of answer j only. Using the whole sample of answers published in the site (as of March $7^t h$ 2022), I estimate the following linear model:

$$points_j = \beta_0 + \boldsymbol{X}_j \boldsymbol{\beta}_1 + \boldsymbol{X}_j^2 \boldsymbol{\beta}_2 + \varepsilon_j. \tag{6}$$

it follows that the measure of quality of an answer k is:

$$quality_k = \hat{\beta}_0 + \boldsymbol{X}_k \hat{\boldsymbol{\beta}}_1 + \boldsymbol{X}_k^2 \hat{\boldsymbol{\beta}}_2.$$

The vector of text characteristics includes: 1) the number of words (*length*); 2) The share of meaningful words, that is, the number of words that do not appear in the list of *stopwords* over the total number of words (*precision*); 3) the number of pictures included in the text; and 4) the number of links included in the text. Table 7 reports the estimates for equation 6.

A.2 Construction of scarcity variable

The scarcity variable is a measure of the quantity of available questions in the site at each point in time. It is defined as:

$$scarcity_t \equiv \frac{maxavail}{\log(avail_t)},$$

where $avail_t$ is the number of unanswered questions in the website in week t, and $maxavail$ is the $\max\{\log(avail)\}$.

A.3 Construction of Types

The procedure I used to identify types follows the steps below:

1. **Extraction of data from user profile pages**. Using data from the profile pages of all users registered in the site, I construct the following variables: 1) A dummy equal to 1 if the user has a full name, i.e. the user name is composed by two words which start with an upper-case letter followed by

	Reputation Points
Precision	31.21
	(6.48)
Precision^2	-22.34
	(6.37)
Length	0.01
	(0.00)
Length^2	-0.00
	(0.00)
Number of Pictures	5.25
	(0.59)
(Number of Pictures)^2	-1.03
	(0.17)
Number of links	2.20
	(0.09)
(Number of Links)^2	-0.04
	(0.00)
Constant	2.80
	(1.63)
Obs	143342
Adj. R^2	0.01
F-stat	157.31

Notes. Regression to predict the number of points received by answers on the day of publication, as a function of their text characteristics. This model is used to predict a proxy for answer quality.

Table 7: Model Estimates to Predict Answer Quality

lower-case letters; 2) a dummy equal to 1 if the user has a LinkedIn profile link, i.e. if the word *linkedin* appears in the website url or in the biography section; 3) a dummy equal to 1 if the user populates the website field; 4) a dummy equal to 1 if the user populates the location field; 5) a categorical variable taking value 0 if the user does not provide any biographical text, 1 if the user's biography is shorter than the median, and 2 if it is longer than the median (in terms of number of words).

2. **Quantification of the variables and reduction of dimensionality**. I implement A Multiple Correspondence Analysis (MCA, Greenacre and Blasius 2006), which works in a similar way as the Principal Component Analysis (PCA) but for categorical variables, and 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 informa-

tion of the original variables. I set the algorithm to compute 5 dimensions. 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. Figure 6 shows the variable representation in the first two dimensions space. First it is possible to notice that the first dimension contains about 36% of the information of the individual characteristics, while the second dimension about 18%. The location of the variables on the plain tells the extent to which each dimension include information from the given variables. It is possible to see that the length of the biographical text is the most important source of information for both dimensions, the inclusion of the location and a website in the user page is only captured by the first dimension, while the presence of a full name only by the second dimension. Figure 7 instead represents, on the same two dimensions, the individuals, i.e. the sample of users. This graph helps understanding 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. The interpretation of the graph suggests that the average consumer has relatively little information displayed in the profile page. Users clustered in the top-right of the graph are more likely than the average user to have a LinkedIn profile, a long biography, a full name, a website, and the location, while users clustered in the bottom right are more likely to have a small biography, the location, and a website.

3. **Clustering of users in groups**. In the third step, I implement the K-means clustering algorithm on the 5 MCA dimensions, and for arbitrary number of clusters k. For a given number k, the algorithm 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 meant to work with continuous quantitative variables, so it is not suitable to be directly applied on the original individual characteristics. Figure 8 shows the same individual representation as in figure 7, but with different colours for each cluster, in the case of $k = 3$. To select the best number of cluster k, I compute clusters for many possible values of k and plot the resulting clusters on the individual graph of the MCA. I then select the value of k that seems to create the most separate clusters (by eye).[34]

[34] This approach is borrowed from the marketing literature on *customer segmentation*. While the literature considers the K-Means algorithm a standard tool to identify clusters, there is no leading method to select the number of clusters.

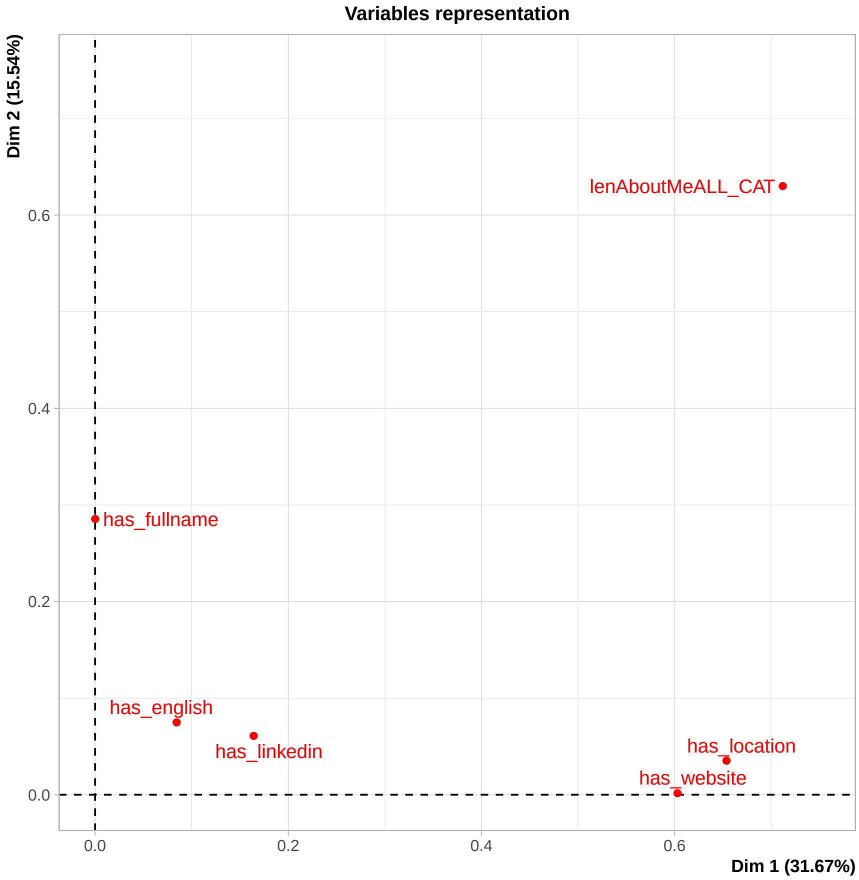

Figure 6: Variable representation on the first two dimensions plane, where the dimensions are obtained via the MCA of the individual characteristics.

A.4 Estimates table for reduced form parameters

Table 8 reports the full list of estimates of the reduced form model discussed in section 4. In model 1 the dependent variable is the number of weekly edits, in model 2 is a dummy equal to 1 if the user made at least one edit. Columns 3 and 4 report estimates for a different action (comments) which should not be affected by the achievement of the privilege. Model 3 has the number of comments as the dependent variable, while module 4's dependent variable is a dummy equal to 1 if the user made at least one comment. Estimates show that editing increases after users achieve the privilege, while contributing via comments is not affected.

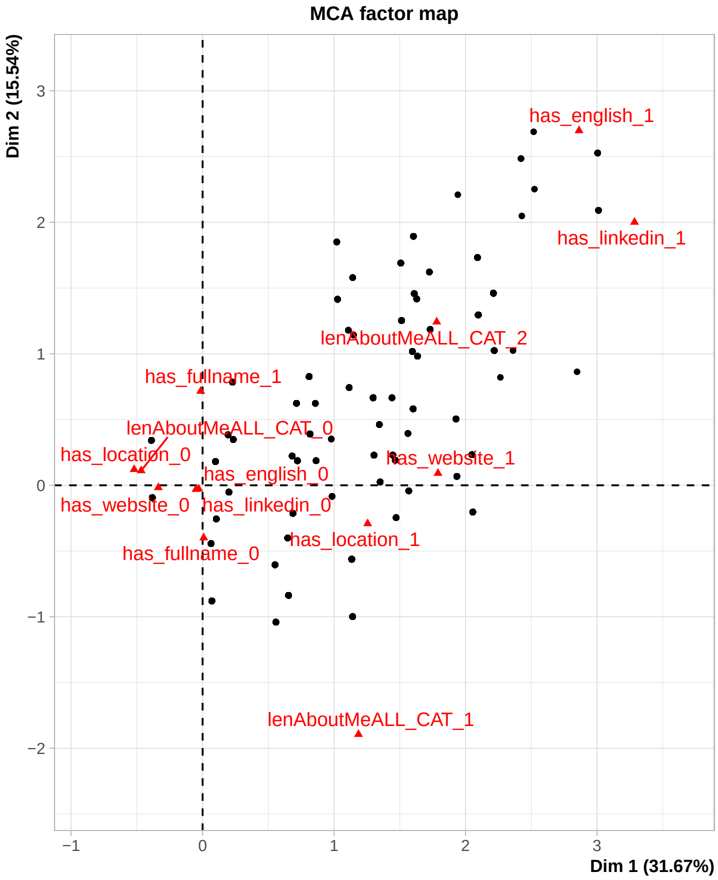

Figure 7: Representations of users on the first two dimensions plane, where the dimensions are obtained via the MCA of the individual characteristics.

A.5 Reduced-form Estimates of Answering

Figure 9 reports estimates for the empirical model presented in section 4 with, as outcome variable, the number of answers produced (left in the figure) or a dummy equal to one if the user made any answer (right in the figure).

A.6 Details on the Structural Model

A.6.1 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 $\boldsymbol{d}_t$ a vector of dummy variables, $d_{\alpha t}$, for each possible choice $\boldsymbol{\alpha} \in \mathcal{A}$, such that $d_{\alpha t}$ is equal to 1 if in period t is selected choice $\boldsymbol{\alpha}$, and zero otherwise.

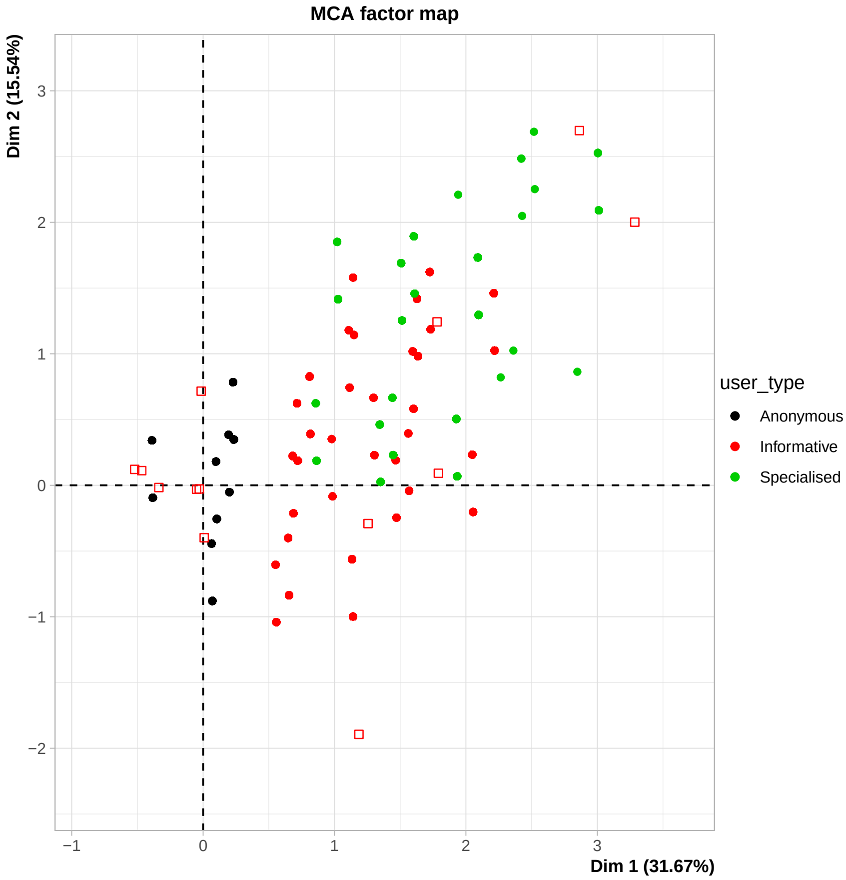

Figure 8: Representations of users on the first two dimensions plane, where the dimensions are obtained via the MCA of the individual characteristics. Colors refer to cluster groups identified with k-means clustering on the MCA dimensions.

Choosing an action $\boldsymbol{\alpha}^*$ in period t, the one period flow utility of user i is then given by:

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

Where the vector $\boldsymbol{x_t}$ is described in section 5.1.

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

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. $\boldsymbol{x_t}$), but also variables that may affect users' beliefs on the probability

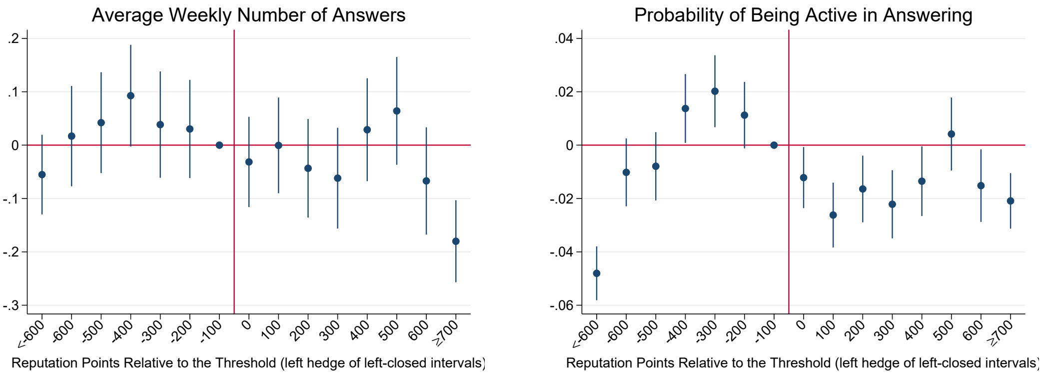

Notes. Reputation-points fixed effects before and after achieving the editing privilege. Sample of users who reached the threshold.

Figure 9: Answering Contributions Relative to Achieving the Privilege

distribution over future states.

A user selects a sequence of optimal decisions $\boldsymbol{d}^* \equiv \{\boldsymbol{d}_t^*\}_{t \leq T}$ that satisfies[35]:

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

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 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 realized[36]. In other words, it is the continuation value of being in state z_t, before ε_t is realized and the

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

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

[36]The 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.

decision at t is taken. By applying Bellman's principle, it is then given by:

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

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 $\boldsymbol{\alpha}$ 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 $\boldsymbol{p}_t(z_t)$ as the vector that gives the probabilities of choosing option $\boldsymbol{\alpha} \in \mathcal{A}$ 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 $\boldsymbol{k} \in \mathcal{A}$ such that:

$$\omega_k(\boldsymbol{p}_t(z_t)) = V_t(z_t) - \nu_{kt}(z_t).$$

It follows that:

$$\nu_{\alpha t}(z_t) = u_{\alpha t}(z_t) + \delta \sum_{z_{t+1}\in\mathcal{Z}} \left(\nu_{kt+1}(z_{t+1}) + \omega_k(\boldsymbol{p}_{t+1}(z_{t+1}))\right) f_{\alpha t}(z_{t+1}|z_t),$$

which can be rewritten as:

$$\nu_{\alpha t}(z_t) = u_{\alpha t}(z_t) + \sum_{\tau=t+1}^{T} \sum_{\boldsymbol{k}\in\mathcal{A}} \sum_{z_\tau\in\mathcal{Z}} \delta^{\tau-t}(u_{k\tau}(z_\tau) + \omega_k(\boldsymbol{p}_\tau(z_\tau)))d^*_{k\tau}(z_\tau, d_{\alpha t}=1)\kappa^*_{\tau-1}(z_\tau|z_t, d_{\alpha t}=1), \tag{7}$$

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.

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

To write the conditional value function as in 7 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 7 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 $\varepsilon_{\alpha_t t}$. I will assume a Type I extreme value distribution.
This allows to express the choice probabilities as:

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

and the ex-ante value function as:

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

where γ is the Euler's constant and $\tilde{\boldsymbol{\alpha}}$ is an arbitrary reference choice from $\mathcal{A}$. The interpretation of the term $-\ln(p_{k\tau}(z_\tau))$ is that it compensates for the possibility that the choice $\tilde{\boldsymbol{\alpha}}$ may not be optimal, given the draws of the errors (Arcidiacono and Ellickson 2011). It follows that:

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

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

$$\nu_{\alpha t}(z_t) - \nu_{\tilde{\alpha}t}(z_t) = u_{\alpha t}(z_t) - u_{\tilde{\alpha}t}(z_t) + \sum_{\tau=t+1}^{t+\Delta_t} \sum_{\boldsymbol{k}\in\mathcal{A}} \sum_{z_\tau \in \mathcal{Z}} \delta^{\tau-t} \left(u_{k\tau}(z_\tau) - \ln(p_{k\tau}(z_\tau))\right) [d^*_{k\tau}(z_\tau, d_{\alpha t} = 1)\kappa_{\tau-1}(z_\tau|z_t, d_{\alpha t} = 1) + \\ - d^*_{k\tau}(z_\tau, d_{\tilde{\alpha}t} = 1)\kappa_{\tau-1}(z_\tau|z_t, d_{\tilde{\alpha}t} = 1)]$$

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 $\boldsymbol{\alpha}$ or $\tilde{\boldsymbol{\alpha}}$.

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

$$\begin{aligned} L(\boldsymbol{\beta}_0, \boldsymbol{\beta}_1, \gamma) &= \sum_{i=1}^{N} \sum_{t=1}^{T} \sum_{\alpha\in\mathcal{A}} \log\left(\frac{\exp(\nu_{\alpha it}(z_{it}))}{\sum_{k\in\mathcal{A}} \exp(\nu_{kit}(z_{it}))}\right) \times d_{\alpha it} \\ &= \sum_{i=1}^{N} \sum_{t=1}^{T} \sum_{\alpha\in\mathcal{A}} \log\left(\frac{\exp(\nu_{\alpha it}(z_{it}) - \nu_{\tilde{\alpha}it}(z_{it}))}{\sum_{k\in\mathcal{A}} \exp(\nu_{kit}(z_{it}) - \nu_{\tilde{\alpha}it}(z_{it}))}\right) \times d_{\alpha it} \end{aligned}$$

A.6.2 Choice set

Because of computational time, the choice set must be constrained to a finite and limited number of options.[37] 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 9.

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 70^{th} 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 33^{th} and 66^{th} quantiles. The median values are 12.69, 13.41, and 15.13. Finally, the distribution of number of edits is split at the 75^{th} quantile, leading to two

[37] A 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).

categories: low effort, which includes 1 to 4 edits, and high effort, including 5 or more edits. The choice of the quantile levels is arbitrary.

A.6.3 Arrival of Votes on Answers

Section 5.2 describes functional form assumptions on the process of arrival of votes on answers. In particular, I assume that the arrival of up-votes and down-votes on an answer published on a given week follow an exponential decay process. This assumption reflects the fact that, on average, most up-votes and down-votes arrive on the same week the user publishes the answer, and very few votes arrive in the following weeks. Alternative functional forms that describe this sharp decrease in the number of votes across time would give similar predictions. Figure 10 shows how different alternative functional form would fit the data.

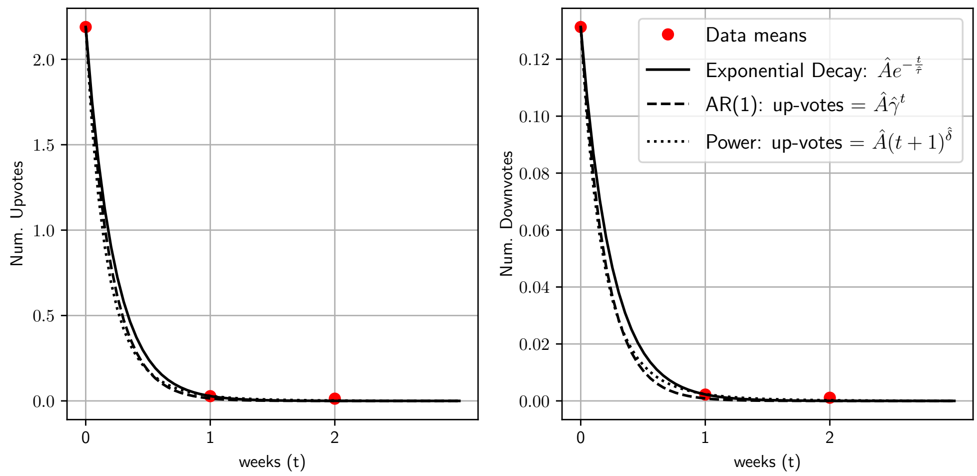

Notes. Comparison of different possible functional form assumptions to describe the process of arrival of votes on answers, since publication week (i.e. $t = t_0$). Dots report data means across answers.

Figure 10: Arrival of Votes on Answers

A.6.4 Conditional Choice Probabilities

Conditional choice probabilities are computed before estimation via a static logit[38] 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

[38] Logistic regression in Scikit-learn with *saga* solver.

between the maximum and the minimum. The multinomial logit model implemented is the following:

$$\begin{aligned}\alpha^*_{it} =&\beta_0 R_{it-1} + \beta_1 \Lambda_{U,it-1} + \beta_2 \Lambda_{D,it-1} + \beta_3 avail_{it} + \beta_4 AnswerNum_{it} + \beta_5 Seniority_{it} + \\ &+ \beta_6 t + \beta_7 date_{it} + cumT_{it}\end{aligned}$$

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.7 Details on Simulation of Counterfactuals

A.7.1 Restrictions on the state values

Reputation points. It is assumed that users can accumulate at most 2000 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 4 unique values (6024, 18072, 30120, 42168). Note that the number of available questions could still change across the time of a user's participation.

Experience variables. Experience variables are set to the median value (computed in the full dataset) and are not allowed to change. In practice, the number of answers already made is set to 2 for all users, and the days of participation are set to 812. In other words, in the simulations I do not allow for learning while participating.

A.8 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 11: Rules to obtain or loose reputation in Stackexchange (https://stackoverflow.com/help/whats-reputation)

A.9 Credits for the software used

Pedregosa, Varoquaux, Gramfort, Michel, Thirion, Grisel, Blondel, Prettenhofer, Weiss, Dubourg, Vanderplas, Passos, Cournapeau, Brucher, Perrot, and Duchesnay (2011), Seabold and Perktold (2010), Hagberg, Schult, and Swart (2008), McKinney (2010), Lï¿½, 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.

	(1) NumEdits	(2) At least one edit	(3) NumComments	(4) At least one comment
-700	-0.0675	-0.0300***	-0.443***	-0.0697***
	(0.0575)	(0.00384)	(0.0829)	(0.00616)
-600	-0.00621	-0.00623	-0.00129	-0.0196*
	(0.0725)	(0.00484)	(0.104)	(0.00777)
-500	-0.0177	-0.00500	0.104	-0.00739
	(0.0729)	(0.00486)	(0.105)	(0.00781)
-400	0.0130	0.00379	0.109	0.0176*
	(0.0735)	(0.00490)	(0.106)	(0.00787)
-300	0.0456	0.00829	0.205	0.0316***
	(0.0768)	(0.00512)	(0.111)	(0.00823)
-200	0.0518	0.0140**	0.0635	0.0245**
	(0.0710)	(0.00473)	(0.102)	(0.00761)
-100	0	0	0	0
	(.)	(.)	(.)	(.)
0	0.0569	0.0119**	0.0214	-0.00932
	(0.0652)	(0.00435)	(0.0939)	(0.00699)
100	0.104	0.0161***	-0.0462	-0.0251***
	(0.0691)	(0.00461)	(0.0996)	(0.00741)
200	0.0847	0.0161***	0.0531	0.00132
	(0.0712)	(0.00475)	(0.103)	(0.00763)
300	0.127	0.0212***	0.0376	-0.00576
	(0.0728)	(0.00486)	(0.105)	(0.00781)
400	0.156*	0.0234***	-0.0670	-0.0102
	(0.0744)	(0.00496)	(0.107)	(0.00797)
500	0.192*	0.0334***	0.0321	0.00813
	(0.0778)	(0.00519)	(0.112)	(0.00834)
600	0.173*	0.0200***	-0.0196	-0.00418
	(0.0775)	(0.00517)	(0.112)	(0.00830)
700	0.364***	0.0434***	0.137	0.000954
	(0.0593)	(0.00396)	(0.0854)	(0.00636)
NumAnswers	0.0707***	0.0133***	1.339***	0.0439***
	(0.00385)	(0.000257)	(0.00437)	(0.000325)
NumComments	0.149***	0.0124***		
	(0.00177)	(0.000118)		
is_candidate	4.169***	0.304***	3.079***	0.243***
	(0.410)	(0.0273)	(0.591)	(0.0439)
is_moderator	4.233***	-0.00521	-3.031***	0.0350**
	(0.122)	(0.00815)	(0.176)	(0.0131)
CopyEditor	30.70***	0.192***	15.40***	0.225***
	(0.622)	(0.0415)	(0.895)	(0.0666)
StrunkWhite	7.767***	0.401***	7.048***	0.213***
	(0.303)	(0.0202)	(0.436)	(0.0325)
Observations	154520	154520	154520	154520
User and Week FE	Yes	Yes	Yes	Yes

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 8: Estimates of reduced form model.

NA	0.00	0.00	0.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
QA	0.00	0.00	0.00	12.69	12.69	12.69	13.41	13.41	13.41	15.13	15.13
NE	0.00	1.00	9.00	0.00	1.00	9.00	0.00	1.00	9.00	0.00	1.00

NA	1.00	7.00	7.00	7.00	7.00	7.00	7.00	7.00	7.00	7.00
QA	15.13	12.69	12.69	12.69	13.41	13.41	13.41	15.13	15.13	15.13
NE	9.00	0.00	1.00	9.00	0.00	1.00	9.00	0.00	1.00	9.00

Notes. NA is number of answers (in a week), QA is the average answer quality for the answers made, and NE is the number of edits (in a week).

Table 9: Possible Combinations of Effort Levels

Name	Private Beta	(Public) Beta	Designed	Description
create posts	1	1	1	Ask a question or contribute an answer
participate in meta	5	5	5	Discuss the site itself: bugs- feedback- and governance
skip lecture on how to ask	-	-	10	
create community-wiki answers	10	10	10	Create answers that can be easily edited by most users
remove new-user restrictions	1	10	10	Post more links- answer protected questions
vote up	1	15	15	Indicate when questions and answers are useful
flag posts	15	15	15	Bring content to the attention of the community via flags
post instantly self-answered questions	15	15	15	
comment everywhere	1	50	50	Leave comments on other people's posts
set bounties	75	75	75	Offer some of your reputation as bounty on a question
edit community wikis	1	100	100	Collaborate on the editing and improvement of wiki posts
vote down	1	125	125	Indicate when questions and answers are not useful
create tags	1	150	300	Add new tags to the site
vote in moderator elections	-	150	150	
association bonus	200	200	200	
shown in network reputation graph and flair	200	200	200	
shown as "beta user" on area 51	200	200	-	
reduced advertisements	-	-	200	
reputation leagues, top x% link in profile	201	201	201	
qualify for first yearling badge	201	201	201	
view close votes	1	250	250	View and cast close/reopen votes on your own questions
run for moderator	-	300	300	
access review queues	350	350	500	Access the First posts and Late answers review queues
see vote counts	100	750	1000	ESTABLISHED USER- You've been around for a while- see vote counts
edit freely, se and lqp/a queue*	500	1000	2000	edit posts of others without review; access the Suggested edits and the Low quality posts or Low quality answers review queues
no popup asking to comment when downvoting	2000	2000	2000	
non-nofollow link in user profile	2000	2000	2000	
suggest tag synonyms	1250	1250	2500	Decide which tags have the same meaning as others
vote to close and reopen	15	500	3000	Help decide whether posts are off-topic or duplicates
review tag wiki edits	750	1500	5000	Approve edits to tag wikis made by regular users
moderator tools	1000	2000	10000	Access reports- delete questions- review reviews
reduce captchas	1000?	2000	10000	
protect questions	1750	3500	15000	Mark questions as protected
trusted user	2000	4000	20000	Expanded editing- deletion and undeletion privileges
access to site analytics	2500	5000	25000	Access to internal site analytics

Notes. List of privileges that users can achieve and associated amount of reputation points required. The *(Public) Beta* reports the required reputation points between January 2013 and February 2016, while the *Designed* column applies from February 2016 onwards.

Table 10: Privileges and associated reputation points requirements