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

Jacopo Bregolin†
University of Liverpool

November 2025

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 increases the marginal value of contributions, which is heterogeneous across different types of users. I simulate counterfactuals with alternative delegation designs. The results show that the platform would lose an important share of production and quality of content in the absence of delegation. However, with delegation based on performance, the platform influences the relative contribution of the user types to the platform content.

*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, Nick de Roos, Daniel Ershov, Chiara Farronato, April Franco, Guido Friebel, Hulrich Hege, Astrid Hopfensitz, Aija Leiponen, Greg Lewis, Inna Lyubareva, Noemi Mantovan, 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. The theoretical literature in personnel economics has identified another non-monetary channel, that is, the delegation of decision rights.[2] 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)). Users may value autonomy when contributing because of either private benefits or because it gives them more control over outcomes. If that is the case, the platform may prefer to provide autonomy unconditionally to incentivise contributions (Aghion and Tirole (1997)).[3] However, organizational designs where the allocation of autonomy is conditional on achieving performance targets may allow the platform to 1) focus users on specific tasks rather than others (Holmstrom and Milgrom (1991)), 2) create virtual *promotions* and the associated *career* incentives (Holmström (1999)), and 3) screen the best users across time (Gershkov and Perry (2012)).

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, a question and answer (Q&A) website, I show that people value such autonomy.

[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]This implication requires sufficiently close objectives of the platform and the users.

However, different types of users value having such autonomy differently. Through counterfactual exercises, I explore organisational implications and find that delegation increases the number and quality of contributions, suggesting that more autonomy to users boosts content production. However, the delegation of autonomy conditional on achieving performance targets allows the platform to change the relative contribution of the different user types. The online community splits into a smaller, more committed group of users, who are more likely to sustain contributions across time, and a larger group who are initially more engaged, but who are more likely to become inactive (*enthusiasts*). Conditioning delegation on performance, the platform increases the relative contribution of *enthusiasts*, as they achieve the performance target faster.

What does it mean to allocate autonomy in digital platforms? Knowledge platforms need to incentivise effort on two main tasks: content production and the moderation of existing content (Gillespie (2018)). While the production of content is generally unrestricted, platforms differ in how much autonomy they grant users for moderation. 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?

To study this question, I use data from Stack Exchange. Stack Exchange is a family of websites where registered users ask questions and provide answers on different topics. Users can collect *reputation points*, mostly through community upvotes on their contributed answers. The moderation of the website relies on moderators elected within the community and community members' edits. Users have control over the implementation of their edits if they have a number of reputation points larger than a threshold $\bar{R}$. If users have fewer reputation points than that, their edits are *suggested*. To get implemented, they require the approval of either the owner of the edited content or by users with more than the threshold number of points.

The data I use include the contribution histories (e.g. answers and edits) of participants in the English Language Learners website.[4] This website is part of the Stack Exchange family and focuses on questions concerning the use of the English language.

The analysis proceeds in two steps. First, I study the impact of variation in autonomy on users' propensity and intensity in editing. On the one hand, I estimate a regression discontinuity with staggered treatment around the reputation threshold $\bar{R}$ to show that the number of edits and the likelihood of editing increase

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

once users obtain autonomy on editing. On the other hand, I exploit a variation of the reputation threshold with retroactive effect, which led to some users losing autonomy. Consistent with the previous result, users who lose control are less likely to edit and produce fewer edits.

Second, I use a dynamic discrete choice model (à la Rust (1987), but with continuous state space) to quantify users' marginal utility from participation 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 allows preferences to depend on the degree of autonomy by interacting the drivers of the utility with a dummy equal to 1 if the user gained autonomy. The identification relies on revealed preference and an inter-temporal tradeoff generated by time discounting. Indeed, users trade off contributions in the current week versus the following, and prefer contributing in the current week if the discounted expected returns from contributing (e.g. reputation points and autonomy) outweigh the costs. Methodologically, I estimate the utility parameters using a conditional-choice-probability (CCP) based estimator with *finite dependence*, a tool that allows substantial computational gains as it does not require a full solution of the model (Arcidiacono and Miller (2011)).

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, users' utility of answering increases as well when users gain control over editing, suggesting positive spillovers of autonomy across tasks. The estimation identifies two unobserved types, to whom I will refer as *committed* and *enthusiasts*. *Committed* users are relatively few (about 20% of the sample), they are more likely to sustain contributions across time, and are possibly more extrinsically motivated, as they tend to disclose more information in their profile page. *Enthusiasts* instead form a larger group that is more active in the short run, but also more likely to become inactive with time. They differ on two key dimensions. First, *committed* users have higher initial costs in editing, but which are marginally decreasing faster. Second *committed* users value autonomy more when editing.

With estimates from the model, I simulate counterfactual contribution histories under delegation designs that differ on the threshold of reputation points $\bar{R}$. In particular, I consider the case with a performance threshold equal to zero (full delegation), 500, 1000, and 2000.

The results show that the platform would maximise quality and quantity of contributions with full delegation and without conditioning autonomy on performance. This is because both types of users make more effort if they have autonomy. However, over 150 weeks of participation, a positive performance threshold increases the relative contribution of *enthusiasts* with respect to *committed* users.

Indeed, *enthusiasts* reach the threshold faster, meaning that, with a positive reputation threshold, they have autonomy for a longer share of their contribution history. *Committed* users take more time to accumulate enough points to gain autonomy. However, they are less likely to become inactive, sustaining contribution levels much longer. With more demanding thresholds, this effect partly offset the previous one as, in the long run, *enthusiasts* are likely to become inactive before gaining autonomy.

Finally, I do not find that the quality of edits is significantly different across user types, suggesting that the dynamic selection across time does not necessarily help in screening low-ability types.

This paper contributes to the literature in several ways. First, I show direct evidence of non-monetary preferences for control and identify in real data the value of autonomy. This result confirms experimental and survey-based evidence showing that individuals value control rights and power (Fehr, Holger, and Wilkening (2013); Bartling, Fehr, and Herz (2014); Owens, Grossman, and Fackler (2014); Pikulina and Tergiman (2020); Meagher and Wait (2021)).

Second, the paper contributes to the empirical literature on the beneficial effects of providing autonomy on workers' performance. Bandiera, Best, Khan, and Prat (2021) finds that autonomy increases the efficiency of procurement officers in India, while Liberti (2018) shows that bank workers put more effort into production and the use of soft information.[5]

Third, the paper contributes to the literature that studies user behaviour in online communities, in particular in Stack Exchange. Aaltonen and Wattal (2025) study how the type of rejection notices affects user retention; Xu et al. (2020) and Forderer and Burtch (2024) investigate career concerns, Bregolin (2025) studies how the use of a foreign language affects communication quality, and Burtch, Lee, and Chen (2024) study how the popularisation of large language models is affecting user contributions. However, none of these papers use data from the ELL website.

Finally, while the results of this paper are specific to the context of online

[5]A substantial empirical literature has studied patterns of delegation in firms. De Varo and Prasad (2015) and Hong, Kueng, and Yang (2019) study the complementarity between delegation of control and performance-based monetary incentives. They show that it depends on job complexity (De Varo and Prasad (2015)) and that is inversely related to the delegation of control from managerial to non-managerial roles (Hong et al. (2019)). McElheran (2014) shows that delegation is more likely if the firm's adaptation to local problems is more important than within-firm coordination. Katayama, Meagher, and Wait (2018) find evidence that firms jointly optimise communication protocols and delegation of control decisions, and that delegation is more likely when valuable human capital is dispersed across the firm. Bloom, Sadun, and Van Reenen (2010), Meagher and Wait (2014), and Liu, Meagher, and Wait (2022) show that market conditions affect firms' incentives to delegate by making local knowledge more salient and affecting employees' trust of managers, which is associated with higher levels of delegation (Bloom, Sadun, and Van Reenen (2012); Meagher and Wait (2020)).

communities, they may suggest implications for a broader set of environments, addressing puzzles that emerged in the literature on promotions. They can provide a plausible explanation for the use of promotions rather than bonuses, even if bonuses are more flexible incentives (Baker, Jensen, and Murphy (1988), Gibbons and Waldman (1999)). In addition, they contribute to the investigation of the fact that often firms commit 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, and section 3 presents the data. 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 2025, 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, *badges* and *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 upvotes on their questions and answers,

[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.

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.[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, is the delegation of autonomy on editing content. Editing is the action of modifying existing questions and answers to improve them. Users can always make edits. However, if users have not collected enough reputation points to unlock the editing privilege, their edits are not directly implemented. They are rather suggestions, and need to be approved by the author of the modified content or by other users who already have the editing 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 implementation of the edits. 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.

The theoretical and experimental literature suggests that individuals may value autonomy either for its intrinsic value and private benefits, or because it gives them control over outcomes that would otherwise differ from their preferred ones.[11] Under either alternatives, users are more willing to contribute if endowed with

[9]In appendix D, figure 13 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]For example, Bartling et al. (2014), Fehr et al. (2013), and Pikulina and Tergiman (2020) find in experimental settings that individuals enjoy the intrinsic value of control. Beckmann and Kräkel (2022) find evidence of higher task commitment if workers have autonomy. Meagher and Wait (2021) find evidence of preference for power in particular when the individual background particularly values the individual level of power. Aghion and Bolton (1992) assume that entrepreneurs receive private benefits from control. The theoretical literature that studies delegation with principal-agent models relies on the assumption that more autonomy allows the agent to affect the outcome in their preferred way and, as a consequence, incentivise their effort (Aghion and Tirole (1997); Mookherjee (2006); Bester and Krähmer (2008); Hart (2017). Finally, Sturm and Antonakis (2015) underline how power affects the direct inclination to take an action.

control. The platform can provide users with autonomy to incentivise their effort. Under this consideration, and unless the platform would believe that users have preferences conflicting with the success of the community, the theory would suggest optimality of full delegation of control.

However, two features of online knowledge platforms may lead to alternative optimal designs. First, users contribute on two tasks, answering and editing. If the platform believes that one is more important than the other, it may optimally design stronger incentives on that task than the other (Holmstrom and Milgrom (1991); Dewatripont, Jewitt, and Tirole (2000)). For instance, in Q&A websites like Stack Exchange, answering contributions may be first order. From this perspective, full autonomy on answering and autonomy based on performance on editing create asymmetry of incentive and may lead users to focus more on answering than editing. This effect is expected to be even stronger if the performance indicators are based on successful answering.[12]

Second, there is no formal contract that ties users to participate across time, nor that could be rescinded if contributions are of low quality. Platforms need to retain users, while possibly learning users' abilities and selecting better types across time. In this context, the delegation of autonomy based on performance can be beneficial for two reasons. One the one hand it creates *career* incentives (Holmström (1999)). Users may want to contribute on the platform to show they are worth getting autonomy, or to signal their skills outside of the community (Xu et al. (2020); Forderer and Burtch (2024)).[13] On the other hand, it may create incentive for only high-ability users to stay active on the platform, as the reward of autonomy will be granted only conditional to a path of successful achievements feeding into the performance measure (Gershkov and Perry (2012)).[14]

Overall, the optimal delegation decision for the platform is not obvious. The provision of autonomy on both tasks should incentivise participation overall. However, the alternative design where users have autonomy on answering, but need to achieve performance targets to obtain autonomy on editing, may induce users 1) to

[12]Asymmetric incentives are instead suboptimal if the platform cares at both tasks equally. Since performance on editing is harder to measure than for answering - it is easier to observe if the answer satisfied the questioner rather than observing the degree of improvement after an edit - incentives based on performance should be avoided at all (Holmstrom and Milgrom (1994))

[13]If the signalling outside of the paltform is the major driver, it may still be preferable to delegate control unconditional on performance, as outcomes would be more informative of users' skills (Blanes i Vidal (2007))

[14]The literature has also addressed other types of nonmonetary incentives, that have similar features 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 (Lazear and Rosen (1981)) and rank based incentives (Ehrenberg and Bognanno (1990)). These papers include rivalry between workers in obtaining status and promotions. In my work instead, delegation does not depend on other workers' actions.

focus more on answering rather than editing, 2) create additional *career* incentives before autonomy is awarded, and 3) select the best users in the long run.

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 websites. 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 number of users with the privilege over time and shows that, in February 2016, about 90 users lost the privilege.[15]

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 a total of 5 answers and/or edits.[16] Users' participation histories start with their first answer or edit and continue until the data retrieval date. The panel includes the weekly number of contributed edits and answers, and the average quality of the published answers for 2625 users.

The measure of answer quality is an aggregated measure of several proxies for quality, namely the speed of response, measures of clarity, completeness, and informativeness assessed by a large language model (LLM), and text characteristics, including length and number of links.[17] In practice, the quality variable corre-

[15]Figure 12 in appendix D reports the timing of the change within the sample period. 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 12 reports this timeline for the ELL site. Table 10 in appendix D reports the number of reputation points required to obtain each privilege and how that changed after the new design.

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

[17]Using different data, Harper, Raban, Rafaeli, and Konstan (2008) find that length and links

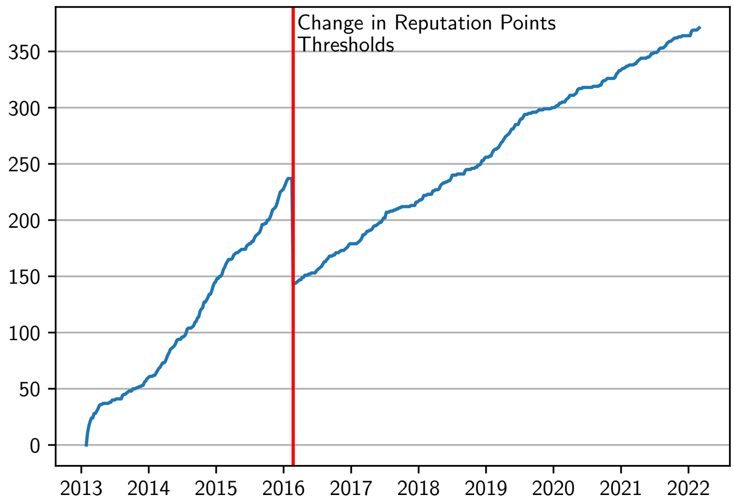

In February 2016, an increase in the requirement of points to obtain this privilege induced the loss of the privilege for some users.

Figure 1: Number of users with the editing privilege

sponds to the first component of the standardised and PCA-transformed set of proxy measures.[18]

The data are right-censored at the download date. Table 1 provides descriptive statistics 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.

are important predictors of answers' quality. Section A.1 in appendix A provides details on the construction of the proxy variables and the aggregation method.

[18]The first component explains 34% of the variance.

	mean	std	min	median	max
Number of weeks in the sample	288.77	121.69	3.00	303.00	476.00
Amount of reputation points reached	1801.98	7402.57	6.00	414.00	179616.00
Number of answers	48.41	213.86	0.00	10.00	4850.00
Number of edits	25.84	228.71	0.00	0.00	6167.00
Average answer quality	6.11	0.92	2.78	6.17	8.35
Reached editing privilege	0.16	0.37	0.00	0.00	1.00

Sample size: 2625 users. The variable *Reached editing privilege* takes the value 1 if the user reaches the threshold to achieve the *editing* privilege within the sample period, and zero otherwise. Statistics on answer quality are conditional on a positive number of contributed answers.

Table 1: User-level descriptive statistics

4 Preference for Control

Reduced-form evidence shows that users are more willing to contribute edits after achieving the editing threshold. This is possible to see in figure 2, which reports estimates of the following regression:

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

where i indexes users, and t indexes weeks. Y is either the number of edits (left graph in figure 2) or a dummy equal to 1 if the user contributed any edits (right graph in figure 2). α_i identifies the user fixed effect, α_t the calendar 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 $\bar{R}$.[19] 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.[20] A second set of variables aims to control for the user's time availability and includes the number of answers and comments produced.

For the reduced form analysis only, I exclude users who did not reach the threshold to prevent selection effects.

[19] 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.

[20] Badges are sort of virtual medals. The editing-related badges are the *Copy Editor* and *Strunk & White* badges (https://ell.stackexchange.com/help/badges).

Figure 2 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.[21]

A reduced-form analysis with standard tools cannot identify the effect of the threshold on answering behaviour as it cannot account for forward-looking behaviour, and the number of reputation points are endogenous to answering activity.[22]

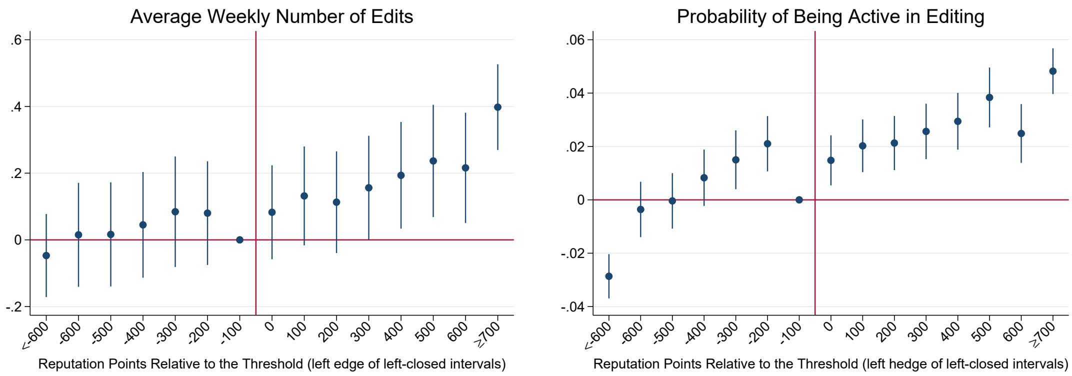

Reputation points fixed effects before and after achieving the editing privilege. Sample of users who reached the threshold.

Figure 2: Editing Contributions Relative to Achieving the Privilege

4.1 Loss of Control

An alternative way to observe the role of control on the propensity to edit is to focus on users who lost control once the reputation threshold increased. Indeed, on February 24th, 2016, the editing threshold increased from 1000 to 2000. On that day, users with more than 1000 points but less than 2000 lost the editing privilege.

[21]Table 8 in appendix B 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. It provides evidence that contributions in a task unaffected by the editing privilege do not increase after users obtain autonomy in editing, suggesting that the increase in edits is attributable to the privilege.

[22]Goes et al. (2016) address this problem relying on functional form assumptions and modifying the data to account for forward-looking behaviour. 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 B.2 in appendix B.

Figure 3 reports the editing activity of this specific group of users. It shows the total number of edits they made each week (left panel) and the number of users who made them (right panel). The colour identifies the number of reputation points the users had when they participated and contributed the edits, with the dark grey identifying users with control under the initial threshold but without under the new threshold. It is possible to notice that users stopped contributing edits right after they lost control. Editing activity then recovers, but it is driven by users who reached the new threshold (black bars).[23]

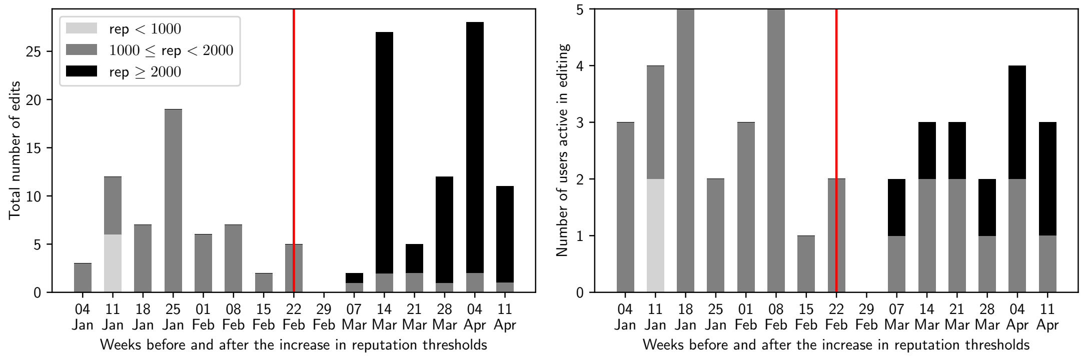

Sample of 94 users who, on the week of February 22nd, 2016, had a number of reputation points ≥ 1000 and < 2000. The left panel reports the number of edits they made each week (intensive margin), while the right panel the number of them who made the edits (extensive margin). The colour identifies the number of reputation points the users had when they participated and contributed the edits: light grey refers users with less than 1000 points, dark grey users with at least 1000 points, but less than 2000, and black users with at least 2000 points. The vertical line identifies the week when the change in reputation thresholds took place.

Figure 3: Editing Contributions Around the Change of Reputation Threshold

5 Dynamic Discrete Choice Model

The reduced-form evidence has multiple limitations. First, it does not allow us to compare the incentive effect of allocating control relative to other types of motives. Second, it does not test for or quantify dynamic responses to the delegation of control. Third, it does not allow us to simulate counterfactual behaviour. Finally, it does not allow for investigating selection at the threshold.

[23] In this context, a regression discontinuity design would not measure the parameter of interest, as users may reach the new threshold soon after the change in reputation thresholds. In addition, any more formal regression analysis would suffer from power issues due to the small sample size of the affected users (94 users).

To overcome these limitations, I develop a dynamic discrete choice model that accounts for forward-looking behaviour and unobserved heterogeneity on user types. 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 the choice's intrinsic utilities, and expected future benefits. Benefits could be, for instance, a certain number of reputation points or the achievement of privileges.[24]

5.1 Users' Participation Choices

Users join the platform and start contributing on random dates. Their participation history starts when they first contribute an answer, a question, or an edit. Once on the platform, users cannot exit, but can decide not to contribute. For the sake of exposition, let T identify their last week of participation, which could be interpreted as their death.[25]

In each week of participation $t \in \{0, ..., T\}$, users decide their 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 week t is then a vector:

$$\alpha_t = \begin{bmatrix} NA_t \\ QA_t \\ NE_t \end{bmatrix} \in \mathscr{A}$$

where NA indicates the number of answers, QA denotes the average quality of answers, and NE indicates the number of edits. $\mathscr{A}$ represents the choice set, including all possible combinations of effort levels in the two tasks.[26]

[24] Note that the application of dynamic discrete choice models to this context is conceptually similar to works that study dynamic investment decisions with discrete choice models. A typical application is human capital investment decisions. Examples of this literature are Arcidiacono, Aucejo, Maurel, and Ransom (2025) and De Groote (2024). Other examples of papers using this estimation strategy are Yoganarasimhan (2013) and Khorunzhina (2013), even though their context and exact application differ from this paper.

[25] Exit from the site is observational equivalent to zero contributions in all following periods. For identification, users do not need to anticipate T. However, they need to believe that $T > t+3$ for any week t of participation. Note that, as I do not observe T, I assume $T > \mathscr{T}$, where $\mathscr{T}$ is the last week observed in the data, i.e. the week ending on March 6th, 2022

[26] I discretise the effort levels to limit the computational burden. In practice, $NA \in \{0, 1, 7\}$, $QA \in \{0, 4.68, 6.61, 7.76\}$, and $NE \in \{0, 1, 9\}$, resulting in 21 possible combinations of efforts (with $QA > 0 \iff NA > 0$). More details are provided in section C.2 in appendix C.

Each user i chooses an optimal sequence of choices to maximise the total sum of the discounted utility from all her periods of participation. Let $\alpha^* \equiv \{\alpha_{\mathbf{t}}\}_{t<T}$ be such sequence of optimal choices. Then,

$$\alpha^* = \arg\max_{\alpha} \mathbb{E}\left[\sum_{t=t_0}^{T} \delta^{t-1} U_{\alpha t}(z_t)\right], \tag{2}$$

where δ is a discount factor and $U_{\alpha t}(z_t)$ is the flow utility that user i in states z_t receives in week t by choosing choice α.

Every week of participation proceeds as follows:

1. First, the user observes the values of the states realised at the end of the previous period, which include the total number of reputation points she has obtained, the number of questions available to answer, and her experience in terms of time spent on the website and the number of contributions. The number of reputation points implies how many privileges she has collected and whether she has already achieved the *editing* privilege.

2. Second, she forms beliefs over the value of the states that may be realised in the future, conditional on past choices and the possible new contribution(s) she could make. (Section 5.2 describes the assumptions on how users form beliefs.)

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

4. Finally, the flow payoff is realised and the states update to their new values.

User i per-period flow utility from choosing choice α in week t is defined as:

$$\begin{aligned} U_{\alpha t} &= u_{\alpha t}(z_t) + \varepsilon_{\alpha t}; \\ U_{\alpha t} &= \beta_{0\theta} R_t + \beta_{1\theta} C^A_{\alpha t} + \beta_{2\theta} (C^A)^2_{\alpha t} + \beta_{3\theta} C^E_{\alpha t} + \beta_{4\theta} (C^E)^2_{\alpha t} + \beta_{5\theta} cumT_t \\ &\quad + Control_t \left(\beta_{6\theta} + \beta_{7\theta} R_t + \beta_{8\theta} C^A_{\alpha t} + \beta_{9\theta} (C^A)^2_{\alpha t} + \beta_{10\theta} C^E_{\alpha t} + \beta_{11\theta} (C^E)^2_{\alpha t}\right) + \varepsilon_{\alpha t}, \end{aligned} \tag{3}$$

where R_t are user's reputation points realised at the end of period $t-1$, $cumT_t$ is the total number of privileges achieved, and $Control_t$ is a dummy equal to 1 if the user achieved the *editing* privilege. $cumT_t$ and $Control_t$ are deterministic functions of R_t and the privilege-thresholds system. The variables $C^A_{\alpha t}$ and $C^E_{\alpha t}$ are proxies for the amount of effort in answering and editing, respectively. They are defined as:

$$\begin{aligned} C^A_{\alpha t} &\equiv QA_{\alpha t} + NA^{scarcity_t}_{\alpha t}, \\ C^E_{\alpha t} &\equiv NE_{\alpha t}, \end{aligned}$$

where $NA^{scarcity}$ is the number of answers raised to a measure of the scarcity of questions to answer.[27] Finally, ε_{it} is an idiosyncratic choice-specific preference shock.

The parameter values are specific to an unobserved user type θ, which captures unobserved heterogeneity. The parameter β_0 captures the marginal utility of accumulating reputation points (R). 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^A and C^E capture instead direct utility from making a certain contribution choice, including a cost of effort and the intrinsic benefit from the action (e.g. if the user is altruistic or enjoys participating per se). The variables R, C^A and C^E are interacted with $Control$ to allow for a change in user preferences once they obtain autonomy.

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.

5.2.1 Evolution of Reputation Points

Users gain and lose reputation points in several ways, but mostly through upvotes and downvotes on their content.[28] Votes may arrive the same week the user publishes the answer or later. Suggested edits also provide reputation points if and when approved. In this section, I make assumptions on the processes of arrival of votes and edit approvals. In a nutshell, users will expect to receive more reputation points in the future if they produce more and better-quality content.

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

[27] The variable *scarcity* is the inverse of the number of unanswered questions. 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.3 in the appendix.

[28] A detailed breakdown of the different ways is in figure 13 in appendix D. Here, I focus on those associated with the strategic choices studied, i.e. posting answers and edits, which are the main sources of reputation points.

period, a number of community edits and votes that follow a Poisson process:

$$\text{Received Edits}_{j,t_0+1} \sim \mathscr{P}(\lambda_{E,j,t_0}),$$
$$\text{Up-votes}_{j,t_0+1} \sim \mathscr{P}(\lambda_{U,j,t_0}),$$
$$\text{Down-votes}_{j,t_0+1} \sim \mathscr{P}(\lambda_{D,j,t_0}).$$

The expected values of these random variables are:

$$\lambda_{E,j,t_0} = \exp\left(\gamma_0 + \gamma_1 QA_{t_0} + \gamma_2 Seniority_{t_0} + \gamma_3 Practice_{t_0}\right), \tag{4}$$
$$\lambda_{U,j,t_0} = \exp\left(\gamma_4 + \gamma_5 QA_{t_0} + \gamma_6 \lambda_{E,j,t_0} + \gamma_7 Seniority_{t_0} + \gamma_8 Practice_{t_0}\right), \tag{5}$$
$$\lambda_{D,j,t_0} = \exp\left(\gamma_9 + \gamma_{10} QA_{t_0} + \gamma_{11} \lambda_{E,j,t_0} + \gamma_{12} Seniority_{t_0} + \gamma_{13} Practice_{t_0}\right), \tag{6}$$

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:

$$\Lambda_{U,t_0} = NA_{t_0} \lambda_{U,j,t_0},$$
$$\Lambda_{D,t_0} = NA_{t_0} \lambda_{D,j,t_0}.$$

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

$$\text{ApprovedEdits}_{t_0+1} \sim \mathscr{B}(NE_{t_0}, \pi).$$

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[29]:

$$\mathbb{E}[\rho_{t_0+1} | \alpha_{t_0}] = 10\Lambda_{U,t_0}(NA_{t0}, QA_{t0}) - 2\Lambda_{D,t_0}(NA_{t0}, QA_{t0}) + 2\pi NE_{t_0}.$$

The answers produced in period t_0 may also induce the arrival of up-votes and down-votes in the following periods. I assume the process is deterministic and follows an exponential decay.[30] Let Δt be the number of weeks passed from the publication week, such that if $t = t_0 + 1$, then $\Delta t = 1$. Then,

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

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

[30] Alternative assumptions give similar results. Section C.3 in the appendix shows how different functional forms fit the data.

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

$$\mathbb{E}[\rho_{t+1}|\{\alpha_{\tilde{t}}\}_{\tilde{t}\leq t}] = 10\Lambda_{U,t} - 2\Lambda_{D,t} + 2\pi NE_t.$$

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

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

5.2.2 Evolution of Questions' Availability

The availability of questions to answer evolves linearly: $avail_{it} = avail_{it-1} + \nu_1$.

5.3 Identification

The identification of the parameters affecting the beliefs on upvotes, downvotes, and edits arrival on own posts exploits cross-sectional variation at the post level, in the first week of publication. The panel dimension of post histories instead identifies the rates at which upvotes and downvotes arrive in the weeks following the publication. Indeed, the data include how many upvotes and downvotes arrive on a post each week.

The identification of the flow payoff parameters relies on the concept of *revealed preferences*, that is, the observation of user effort choices allows us to infer users' preferences. In practice, it follows identification results in the literature on dynamic discrete choice models, including, but not exclusively, Rust (1987), Hotz and Miller (1993), Magnac and Thesmar (2002), Arcidiacono and Miller (2011), and Arcidiacono and Miller (2020).

5.3.1 Identification of flow payoff parameters - theory

Let the conditional value function $\nu_{\alpha t}(z_t)$ be the user's value function of making choice α in week t, net of the preference shock $\varepsilon_{\alpha t}$. In addition, assume, as it is standard in the literature, that the preference shock follows a Type 1 Extreme Value distribution. Given the maximization problem stated in equation 2 and based on results from the literature (Arcidiacono and Miller (2019)), it follows that:

$$\begin{aligned}\nu_{\alpha t}(z_t) &= u_{\alpha t}(z_t)\\ &+ \sum_{\tau=t+1}^{\mathscr{T}} \sum_{k\in\mathscr{A}} \sum_{z_\tau\in\mathscr{Z}} \delta^{\tau-t}(u_{k\tau}(z_\tau) - \ln(p_{k\tau}(z_\tau)) + \gamma))d_{k\tau}(z_\tau, d_{\alpha t}=1)\kappa_\tau(z_\tau|z_t, d_{\alpha t}=1)\\ &+ \mathscr{V}_\alpha,\end{aligned} \tag{7}$$

where $\mathscr{T}$ is the last period in the sample, $p_{\alpha\tau}(z_\tau)$ is the probability of choosing α when in state z_τ (so called *conditional choice probability*, CCP hereafter), $d_{\alpha t}$ is a dummy equal to 1 if the user chooses α in week t, and $\kappa_\tau(z_\tau|z_t, d_{\alpha t}=1)$ is the cumulative probability of being in state z_τ in week $\tau > t$ conditional on having been in state z_t and having chosen α in week t:

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

Finally, γ is Euler's constant and $\mathscr{V}_\alpha$ is the remaining utility that comes from participating after week $\mathscr{T}$ and until T. This latter part of the conditional value function is not identified as the sample ends at $\mathscr{T}$. However, identification is recovered using a property called *finite dependence*, as explained below.

The result in equation 7 show that for any choice α in week t, the future value term of the conditional value function can be expressed relative to any sequence of arbitrary choices. Consider two possible sequences of arbitrary choices that start in week t with choice α and α' respectively. As defined in Arcidiacono and Miller (2019), the pair of choices $\{\alpha, \alpha'\}$ exhibits r-period [finite] dependence if there exists sequences of decision[s] [...] from α and α' such that:

$$\kappa_{t+r+1}(z_{t+r+1}|z_t, d_{\alpha t}=1) = \kappa_{t+r+1}(z_{t+r+1}|z_t, d_{\alpha' t}=1)$$

In other words, *finite dependence* holds if a sequence of arbitrary contribution choices starting with α and one starting with α' would lead the user to given state values with the same probability. In that case, it is possible to write the differenced value functions of those choices as:

$$
\begin{aligned}
&\nu_{\alpha t}(z_t) - \nu_{\alpha' t}(z_t) \\
&\quad = u_{\alpha t}(z_t) - u_{\alpha' t}(z_t) \\
&\quad + \sum_{\tau=t+1}^{t+r} \sum_{k\in\mathscr{A}} \sum_{z_\tau\in\mathscr{Z}} \delta^{\tau-t} \left(u_{k\tau}(z_\tau) - \ln(p_{k\tau}(z_\tau))\right) [d_{k\tau}(z_\tau, d_{\alpha t}=1)\kappa_\tau(z_\tau|z_t, d_{\alpha t}=1) \\
&\quad - d_{k\tau}(z_\tau, d_{\alpha' t}=1)\kappa_\tau(z_\tau|z_t, d_{\alpha' t}=1)]
\end{aligned} \tag{8}
$$

Note that, as long as $\mathscr{T} > t+r$, $\mathscr{V}_\alpha = \mathscr{V}_{\alpha'}$. The unidentified part of the future utility drops out of the difference.[31] Given the assumptions on the flow utility function and transition probabilities, and sufficiently long data to be able to estimate the CCPs in a first step, all components of equation 8 are identified.

The exploitation of this result has two major advantages. First, it is not necessary to specify a terminal period, allowing for lighter assumptions on users' exit from the site. Second, the computational requirements are considerably lower than with alternative methods.[32]

5.3.2 Identification of flow payoff parameters - application

In the specific context of this paper, *finite dependence* holds because of the decreasing arrival of reputation points in the weeks following a contribution (see section 5.2). As a consequence, when users make a contribution, they expect it to produce reputation points only for a few weeks. As an example, let $\tilde{\alpha}$ be a possible contribution choice that leads, in expectation, to more reputation points in the following 3 weeks. Let also $\alpha^0 \equiv [0,0,0]$ identify the choice of no participation. Consider now the following two sequences of arbitrary choices: 1) $\{\tilde{\alpha}_{t_0}, \alpha^0_{t_0+1}, \alpha^0_{t_0+2}, \alpha^0_{t_0+3}, \alpha^0_{t_0+4}\}$; 2) $\{\alpha^0_{t_0}, \tilde{\alpha}_{t_0+1}, \alpha^0_{t_0+2}, \alpha^0_{t_0+3}, \alpha^0_{t_0+4}\}$. Following either sequence, the user would have, in expectation, the same amount of reputation points at $t+4$. Indeed, the second sequence just shifts of one week the arrival of the reputation points.[33] In this case, we would say that $\{\tilde{\alpha}, \alpha^0\}$ exhibits 3-period dependence. Given this example and

[31] In the current paper, $r = 3$.

[32] Some traditional alternative methods, like solving for the value functions via backward recursion, would be unfeasible in this setting, where the panel dimension extends up to 470 weeks. To my knowledge, this is the first paper that implements a dynamic discrete choice model with such a long time dimension.

[33] In practice, this equivalence is only approximate because some parts of the utility flows are time dependent: the return of effort depends on the number of weeks spent on the website, the cost of answering depends on the number of available questions, and the CCPs are allowed to be time-dependent. However, the change of these time-dependent variables across two consecutive weeks is minimal. In this example, I abstract from these time dependences and remove the time indices on effort levels and CCPs.

the flow utility specification (eq. 3), equation 8 would rewrite as (omitting the squared terms of the flow utility for simplicity):

$$
\begin{aligned}
&\nu_{\tilde{\alpha}t_0}(z_t) - \nu_{\alpha^0 t_0}(z_t) \\
&\qquad \approx \beta_1(1-\delta)C^A_{\tilde{\alpha}} + \beta_3(1-\delta)C^E_{\tilde{\alpha}} + \beta_8 Control_{t_0}(1-\delta)C^A_{\tilde{\alpha}} + \beta_9 Control_{t_0}(1-\delta)C^E_{\tilde{\alpha}} \\
&\qquad + \beta_0 \sum_{\tau=t_0+1}^{t_0+3} \delta^{\tau-t_0}\mathbb{E}^{\tilde{\alpha}t_0}[R_\tau - R_{\tau-1}] + \beta_5 \sum_{\tau=t_0+1}^{t_0+3} \delta^{\tau-t_0}\mathbb{E}^{\tilde{\alpha}t_0}[cumT_\tau - cumT_{\tau-1}] \\
&\qquad + \beta_6 \sum_{\tau=t_0+1}^{t_0+3} \delta^{\tau-t_0}\mathbb{E}^{\tilde{\alpha}t_0}[Control_\tau - Control_{\tau-1}] \\
&\qquad + \beta_7 \sum_{\tau=t_0+1}^{t_0+3} \delta^{\tau-t_0}\mathbb{E}^{\tilde{\alpha}t_0}[Control_\tau R_\tau - Control_{\tau-1}R_{\tau-1}] \\
&\qquad - \delta\left(\ln(p_{\alpha^0}(\mathbb{E}^{\tilde{\alpha}t_0}R_{t_0+1})) - \ln(p_{\tilde{\alpha}}(\mathbb{E}^{\tilde{\alpha}t_0}R_{t_0}))\right) \\
&\qquad - \sum_{\tau=t_0+2}^{t_0+3} \delta^{\tau-t_0}\left(\ln(p_{\alpha^0}(\mathbb{E}^{\tilde{\alpha}t_0}R_\tau)) - \ln(p_{\alpha^0}(\mathbb{E}^{\tilde{\alpha}t_0}R_{\tau-1}))\right)
\end{aligned} \tag{9}
$$

where $\mathbb{E}^{\tilde{\alpha}t_0}[z_t]$ captures the user's expected value of the state in week t having contributed $\tilde{\alpha}$ in week t_0 and not participated after (i.e. having chosen the first arbitrary sequence of choices). The last two lines of the equation are offset terms that compensate for the chance that the arbitrary choices in the sequence are not optimal.

Based on this intuition, it is possible to construct similar differenced value functions for any pair of choices $\{\alpha, \alpha^0\}_{\alpha \in \mathscr{A}}$, where higher effort choices require longer sequences of no participation.

5.3.3 Identification of flow payoff parameters - identifying variation

Equation 9 clarifies the source of variation that identifies each flow utility parameter. Since the differenced conditional value functions (CVFs) capture the additional utility of making a contribution in a week rather than the next, identification relies on the existence of an inter-temporal trade-off. The identification of the direct net utility of contributions (i.e. C^A and C^E and their interaction with $Control$) requires imposing $\delta < 1$ and uses intuitive variation, with larger effort choices inducing larger values for those variables. The other variables are all functions of the total number of reputation points collected by the user (i.e. R). Since contribution choices impact them only in the future, their identification relies on variation in the future value terms. The variation comes from the total amount of points that users expect to receive, which depends on the levels of effort and experience. Since $cumT$ and $Control$ are discrete variables, in the differenced CVFs, there is identifying variation only if the choice allows the user to reach one of the reputation thresholds for which those variables vary in the periods of the finite dependence path. As a consequence, only observations with initial levels of reputation points not too far away from those thresholds will help in identifying their marginal value.[34] Note that the coefficients of R, $cumT$, and $Control$ would not be identified in a static model.

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. Third, I estimate the preference parameters.

5.4.1 Estimation of beliefs' parameters

The estimation of the parameters that drive user beliefs exploits data at the answer level. The parameters $\{\gamma_1, ..., \gamma_{13}\}$ drive beliefs on the expected number of edits, upvotes, and downvotes that an answer receives given its quality and the author's experience. Their estimation uses Poisson models on cross-sectional data, where each observation includes the number of edits, upvotes, and downvotes received by the answer on the publication week (dependent variables), the measure of quality of the answer, and the measures of experience of the author (see section 5.2). The probability π that a suggested edit is approved is the simple average approval rate. The rate of change in questions' availability is the OLS estimate of a trend. Finally, I estimate the decay parameters for the arrival of upvotes and downvotes on answers with a non-linear least squares estimation on panel data of the answer's histories in weeks from the publication date. Standard errors assume independence of each estimation.

5.4.2 Estimation of flow utility parameters without unobserved heterogeneity

For clarity, I first discuss estimation in the absence of unobserved heterogeneity. The estimation of the flow utility parameters proceeds with the following steps. First, I estimate the CCPs in reduced form. Because the state space is continuous, I cannot use a bin estimator to estimate the CCPs nonparametrically. Instead, I

[34] Consider the example of section 5.3.2 with the $Control$ threshold set at 2000 and $\tilde{\alpha}$ such that $R_{t_0+3} - R_{t_0} = 200$. Then, line 3 of eq. 9 (i.e. the variation identifying β_6) is different from zero only for observations with $1800 \leq R_{t_0} \leq 1999$.

use a flexible L2-penalised multinomial logit model on scaled data between 0 and 1 and with saga solver.[35] Second, using the estimated CCPs and belief parameters, I compute the terms of the conditional value functions. More precisely, for each observation in the sample and each choice $\alpha \in \mathscr{A}$, I compute the terms of the sum in the differenced CVFs (see eq. 9, excluding the parameters). This step leads to the construction of a matrix with each column vector associated with one of the variables in the flow utility function. Third, I use such a matrix to estimate the flow utility parameters with a standard conditional logit model (McFadden (1974)).

5.4.3 Estimation of flow utility parameters with unobserved user types

I allow flow utility parameters to differ across two unobserved user types. For this purpose, I adapt the EM (Expectation Maximisation) algorithm proposed by Arcidiacono and Miller (2011) to my context. The algorithm maintains computational feasibility by searching for the values of the unobserved types and the flow utility parameters in sequential steps.

The algorithm proceeds with the following steps. First, I initialise the population probabilities for users to belong to an unobserved type, the differenced CVFs (computed as in section 5.4.2), and the flow utility parameters. To characterise unobserved user types, I vary the initial flow utility parameters such that one user type has higher costs of participation (i.e. lower marginal utility of C^A and C^E) than the other. Second, given the flow utility parameters, I compute the users' individual probability of being of a certain type. Third, with such individual probabilities, I update the population probabilities. Fourth, I estimate new CCPs for each user type, weighting the flexible multinomial logit with the individual probabilities. Fifth, I compute new differenced CVFs for each user type using the updated CCPs. Finally, I update the value of the flow utility parameters, maximising the following log-likelihood function:

$$\sum_{i=1}^{N}\sum_{\theta\in\Theta}\sum_{t=1}^{T}\sum_{\alpha\in\mathscr{A}} q_{i\theta}\ln\left(\frac{\exp(\nu_{\alpha it}(z_{it},\theta)-\nu_{\alpha^0 it}(z_{it},\theta))}{\sum_{k\in\mathscr{A}}\exp(\nu_{kit}(z_{it},\theta)-\nu_{\alpha^0 it}(z_{it},\theta))}\right)\times d_{\alpha it} \tag{10}$$

where θ indexes user types and $q_{i\theta}$ is the probability that user i has type θ. With updated flow utility parameters, the algorithm restarts from step one and repeats until convergence.

[35] As explanatory variables, I include R, R^2, $cumT$, $Control$, the number of upvotes and downvotes expected to arrive from past contributions, the number of available questions, the experience (number of weeks spent on the site and number of answers posted), the period of participation and its squared value, the date (week) and its squared value, and the number of reputation points missing to reach the editing threshold (and its squared value). Section C.4 in appendix C provides details.

5.4.4 Standard errors

I compute standard errors for the flow payoff parameters using a bootstrap procedure. In practice, given N users in the data, I sample with replacement N users and build 150 alternative datasets with their participation histories.[36] I then estimate the flow payoff parameters in each of these samples following the procedure described in section 5.4.3, and compute their standard deviation. This bootstrap procedure allows me to take into account the noise in both the reduced-form CCP estimation and the conditional logit estimation on the expected value terms. However, the standard errors do not incorporate the error in the belief parameters' estimates (used for the construction of the expected value terms) and, as a result, are possibly undervalued. The bootstrap procedure cannot incorporate the belief parameters' estimation, as it relies on different data from that of the flow payoff parameter estimation.

5.5 Results

5.5.1 First-stage Estimates of Beliefs Parameters

Table 2 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. 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 upvotes and fewer downvotes. Community edits instead increase the chances of both upvotes and downvotes. A possible justification is that community edits improve the answer, attracting more upvotes, but, at the same time, are more likely to occur on low-quality answers, which attract downvotes.

The parameter π reports the rate at which suggested edits get approved and shows that more than 80% of suggested edits get implemented. 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 that the availability of questions has substantially increased over time.

[36] In practice, the standard errors are computed with 146 estimations of the model.

	Estimate	Std. Error
$\lambda_{E,j,t} = \exp\left(\gamma_0 + \gamma_1 QA_t + \gamma_2 Seniority_t + \gamma_3 Practice_t\right)$		
γ_0	-1.7871	0.04
γ_1	-0.1351	0.0066
γ_2	-0.0005	3e-05
γ_3	-0.0002	3e-05
$\lambda_{U,j,t} = \exp\left(\gamma_4 + \gamma_5 QA_t + \gamma_6 \lambda_{E,j,t} + \gamma_7 Seniority_t + \gamma_8 Practice_t\right)$		
γ_4	-0.5711	0.0083
γ_5	0.184	0.0012
γ_6	0.665	0.0046
γ_7	0.0002	3e-06
γ_8	-4e-05	3e-06
$\lambda_{D,j,t} = \exp\left(\gamma_9 + \gamma_{10} QA_t + \gamma_{11} \lambda_{E,j,t} + \gamma_{12} Seniority_t + \gamma_{13} Practice_t\right)$		
γ_9	-1.0762	0.0263
γ_{10}	-0.1528	0.0043
γ_{11}	0.5786	0.0187
γ_{12}	2e-05	1e-05
γ_{13}	-0.0003	1e-05
$\text{ApprovedEdits}_t \sim \mathscr{B}(NE_t, \pi)$		
π	0.8115	0.0033
$\Lambda_{U,t} = \Lambda_{U,t-1} e^{-\frac{1}{\tau_U}} + NA_t \lambda_{U,j,t}(QA_t)$		
τ_U	0.2297	0.0067
$\Lambda_{D,t} = \Lambda_{D,t-1} e^{-\frac{1}{\tau_D}} + NA_t \lambda_{D,j,t}(QA_t)$		
τ_D	0.2463	0.0092
$avail_t = \nu_0 + \nu_1 t + \epsilon_t$		
ν_1	108.6329	0.5489

Estimates of parameters governing belief formation processes over the state space's evolution. The $\{\gamma\}$ 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 (on the whole sample of suggested edits and not only for users in the estimating sample). τ_U and τ_D are the decay times 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 2: First-Stage Estimates

5.5.2 Flow Payoff Parameters

Table 3 reports the estimates of the flow payoff parameters for each unobserved user type. Overall, all users have a positive marginal utility from accumulating reputation points and experience a direct cost (even though marginally decreasing) from contributing to answering and editing. In addition, users gain value in contributing both answers and edits after they achieve control on editing, particularly for edits, which is consistent with the reduced form evidence.

The estimation identifies two unobserved types: a smaller group of 478 users (about 19%), whom I call *committed*, and a larger group of 2147 users who I call *enthusiasts*. User types differ mostly in how they value participating in their contribution histories. *Committed* users have higher initial costs of editing than *enthusiasts*, but such costs decrease faster with more contributions, and are highly offset by extra marginal utility of editing after they gain control, which doubles that of *enthusiasts*. Preferences for answering are more similar across user types. However, *committed* users' marginal utility from answering is substantially higher once they gain autonomy on editing.

5.5.3 Unobserved user types characteristics

Table 4 reports descriptive statistics by user type. On average, *committed* users have consistently more information on their user profiles than *enthusiasts*. In particular, they have longer biographical descriptions, and they are more likely to display a full name (of the format *Name Surname*), a website, their location, and their LinkedIn profile. This suggests that A users are possibly more extrinsically motivated, as more information makes them more recognisable both within and outside the online community.

Conditionally on editing, *committed* users provide lower quality edits, as measured by a large language model (LLM)'s assessment.[37] However, the rate at which their edits are rolled back (i.e. reverted to a previous version) and the rate at which their suggested edits are approved are not statistically different.

Finally, the distribution of the year of registration on the platform is not different across types.

5.6 Model Fit

To test the fit of the model, I forward simulate the data exploiting the same CCP-based representation of the value function used in estimation. For the exact

[37] The measure of edit quality captures the number of issues in the posts that have been fixed with the edit. I use Mistral AI. For each edit, I propose to the algorithm the versions of the post before and after the edit, and ask it to assess how many issues each has. I then compute the difference. Section A.2 in appendix A provides additional details.

	Type: Committed	Type: Enthusiast	T-test Difference
R	0.0022 (0.0005)	0.0031 (0.0004)	-1.49
C^A	-0.2976 (0.0507)	-0.3091 (0.0547)	0.16
C^E	-2.9029 (0.2318)	-2.0027 (0.4069)	-1.94
$Tcum$	-0.1123 (0.0458)	-0.0779 (0.0237)	-0.72
$Control$	-0.1615 (0.8982)	-0.1439 (0.5949)	-0.02
$R \times Control$	0.0007 (0.0005)	0.0009 (0.0004)	-0.40
$C^A \times Control$	0.2331 (0.121)	0.2355 (0.0886)	-0.02
$C^E \times Control$	4.2768 (0.5368)	2.0908 (0.5196)	3.25
$(C^A)^2$	0.007 (0.0013)	0.0059 (0.0007)	0.69
$(C^E)^2$	0.2883 (0.0247)	0.2116 (0.0408)	1.74
$(C^A \times Control)^2$	0.0337 (0.0138)	-0.0015 (0.0021)	2.54
$(C^E \times Control)^2$	-0.4076 (0.0555)	-0.1879 (0.0544)	-3.13
Pop. probability of type	0.1883 (0.0667)	0.8117 (0.0667)	

Estimates are reported separately for each unobserved type. The column *T-test difference* reports the value of the t-statistic for the null hypothesis $H_0 : \beta_X^A = \beta_X^B$, where β_X^θ is the marginal utility of variable X for user type θ. Bootstrapped standard errors are in parentheses.

Table 3: Flow utility estimates

same users in the sample and their exogenous states, the procedure sequentially computes the structural probabilities of making given contribution choices. For each period, it computes the structural probabilities, draws contribution choices from such probabilities, and updates the endogenous states of the following period. The structural probabilities correspond to the likelihood function (i.e. the expression within the round brackets in equation 10) evaluated with the type-specific flow-utility parameters for each choice $\alpha \in \mathscr{A}$. Such computation exploits the expression of the likelihood function in terms of the differenced conditional value functions (equation 8). This approach allows for the simulation of the data without making additional assumptions on the state space, and without having to set a terminal period.

While the model approximates contribution patterns quite well, it under-predicts the total number of reputation points that users gain. Two main factors play a role in this misprediction. First, due to the discretisation of the choice set, the model is not able to replicate the highly skewed user contribution patterns.[38] It follows that a few users obtain many more reputation points than the model predicts. Second, users may receive points from sources not incorporated in the model, in particular from votes on questions. The unobserved types partly address the first issue, allowing for capturing more active contributors. To address the second issue, at

[38] As a reminder, the largest number of weekly answers allowed in the model is 7. Whenever, in the data, users contribute between 4 and 122 answers, the model assumes a choice of 7 answers, which is the observed median value within that range.

	Type	mean	std. err.	median	n. users
Number of words in bio	Committed	30.481	2.852	4.000	478
	Enthusiasts	23.396	1.097	3.000	2147
Share of users with full name	Committed	0.272	0.020	0.000	478
	Enthusiasts	0.258	0.009	0.000	2147
Share of users with website	Committed	0.366	0.022	0.000	478
	Enthusiasts	0.350	0.010	0.000	2147
Share of users with location	Committed	0.559	0.023	1.000	478
	Enthusiasts	0.517	0.011	1.000	2147
Share of users with linkedIn	Committed	0.019	0.006	0.000	478
	Enthusiasts	0.012	0.002	0.000	2147
Share of users with term *english* in bio	Committed	0.111	0.014	0.000	478
	Enthusiasts	0.106	0.007	0.000	2147
Year of registration	Committed	2016.349	0.115	2016.000	478
	Enthusiasts	2016.128	0.049	2016.000	2147
Share of edits rolled back	Committed	0.008	0.003	0.000	312
	Enthusiasts	0.009	0.002	0.000	817
Share of edits rolled back (min. 3)	Committed	0.008	0.002	0.000	205
	Enthusiasts	0.009	0.002	0.000	415
Edit quality with LLM	Committed	1.081	0.035	1.039	312
	Enthusiasts	1.189	0.032	1.000	817
Edit quality with LLM (min. 3)	Committed	1.132	0.027	1.143	205
	Enthusiasts	1.212	0.024	1.200	415
Approval rate of suggested edits	Committed	0.853	0.013	0.950	307
	Enthusiasts	0.844	0.009	1.000	893
Approval rate of suggested edits (min. 3)	Committed	0.843	0.012	0.890	206
	Enthusiasts	0.851	0.009	0.923	433

Descriptive statistics on observable characteristics. The first set of variables (up to the registration year) relates to information users disclosed in their profile pages at the time of the data download. The rest of the measures are proxies for user-level edit quality, including the rate at which users saw their edit reverted (*rolled back*), the rate at which users saw their suggested edits being approved, and a measure of edit quality computed with a large language model. *(min 3)* indicates statistics that were computed excluding users with less than 3 edits/suggested edits.

Table 4: Characteristics of unobserved types

	sample	mean	std	min	median	max
Amount of Reputation Points Reached	data	1295.15	4562.84	0.00	341.00	110252.00
	model	410.27	970.24	0.00	244.00	19434.00
Average Answer Quality	data	6.09	0.90	4.68	6.22	7.76
	model	6.09	0.69	4.68	6.11	7.76
Number of Answers	data	25.95	76.57	0.00	8.00	1038.00
	model	14.25	14.06	0.00	12.00	210.00
Number of Edits	data	10.76	59.88	0.00	0.00	1105.00
	model	3.15	9.65	0.00	1.00	171.00
Reached Editing Privilege	data	0.14	0.35	0.00	0.00	1.00
	model	0.03	0.16	0.00	0.00	1.00

Sample size: 2625 users. Statistics on participation histories are capped at 150 weeks, with censoring at the download date as in the data. The *model* corresponds to the forward simulated data with two unobserved types.

Table 5: Model fit of descriptive statistics

each period in the simulation, I increase user accumulated reputation points with the points they received from their own questions.[39]

Table 5 reports the model fit of the descriptive statistics, with user participation histories capped at 150 weeks (which is the length of the simulated histories). Besides the intervention mentioned above, the model under-predicts the total number of reputation points reached by some users and, as a consequence, the share of users reaching the threshold for control on editing. However, the qualitative insights are preserved. The simulated data reproduces well the reduced form results of section 4. Figure 10 in appendix C replicates figures 2 and 3 showing similar patterns.

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.

To simulate contribution levels, I use same the CCP-based method exploited for the fit of the model, by only varying the value of the editing threshold.[40] This

[39]This operation makes the simplifying assumption that users ask questions to receive answers, and not strategically to accumulate reputation points. As a consequence, they do not anticipate points arriving from questions.

[40]For details, see section 5.6

approach allows substantial flexibility such that I do not need to constrain the state space, the arbitrary length of the simulated history of participation does not affect individual choices, and the reduced computational burden allows for longer simulations. However, it requires observing the CCPs in the counterfactual environments. In other words, it requires knowing the probability that the users make certain contribution decisions in the counterfactual state space. I predicted such probabilities using a flexible multinomial logit model, which includes the number of missing reputation points to the threshold. The model is estimated with the observed data.[41] In the counterfactuals, the state space differs only in the number of points that the users need to accumulate to reach the editing threshold. It follows that the simulations rely on two main assumptions:

1. *The number of points required to achieve the editing threshold affects users' decisions only by affecting the probability that users reach such a threshold, given their contribution choices.*

2. *Conditional on the number of reputation points missing to achieve the threshold, the user's choices would be identical under a different threshold.*

In practice, I simulate counterfactual participation histories for the exact same users and state values as in the data, but for different values of the reputation threshold. The threshold is set, in different counterfactuals, to zero points, 500 points, 1000 points, and 2000 points. The above assumptions imply that the estimated reduced-form choice probability model predicts the correct choice probabilities under a counterfactual number of points missing to achieve the threshold.[42]

6.1 Counterfactual Results

Figure 4 reports the simulated average contribution histories of the users in the sample for 150 weeks of participation and by user type.[43]

[41] The change in threshold observed in the data allows us to separately identify the effect of the reputation points and of the reputation points missing from the threshold. Details are in appendix C.4.

[42] The assumptions would be violated if the users' behaviour depends on the number of users who have autonomy. Indeed, one may think that power has particular intrinsic value if it is not shared. Under this hypothesis, it would not be realistic to compare a scenario where everyone has control and a scenario where control is delegated based on performance without accounting for peer effects. However, I show that in this setting, such peer effects do not affect behaviour. By focusing on users who did not lose the editing privilege when the website changed the threshold, it is possible to see that the sudden drop in the number of users with the privilege did not affect their contribution patterns. This result is reported in figure 11 in appendix C.6.

[43] The values reported are the weekly average of 10 separate simulations for the 2216 users in the sample who were observed for at least 150 weeks.

The simulations remark on the importance of the delegation design. Users have a general preference to contribute when they have control over editing. Autonomy in moderation increases the utility of participation on both editing and answering. As a consequence, users contribute more with lower delegation thresholds. Strictly positive delegation thresholds reduce the probability that users obtain autonomy on editing, leading to lower users' average contributions and a higher likelihood of no participation.

These patterns are, however, very different across user types. The majority of users (*enthusiasts*) are particularly active at the beginning of their contribution histories, suggesting an initial excitement or temporary motivation. As such, they reach the threshold faster and contribute more and higher-quality answers at the beginning. However, their contribution rate reduces substantially with time. With full delegation, more than 50% of them do not participate at all after almost three years of participation, and this number grows to 90% with delegation at 2000 points. Note that this behaviour is not driven by anticipation of exit, but by a decreasing utility of participation across time. On the contrary, *committed* users (about 19% of the sample) contribute fewer answers at the beginning, but maintain more sustained contribution levels across time. In addition, their rate of inactivity is almost constant across time.

A major difference across types is also in the propensity for editing. *Committed* users are more likely to contribute edits across their participation history. As this propensity is much higher when users have autonomy on editing, the platform gives up on a substantial amount of community moderation by tying delegation of autonomy to performance.

Overall, delegation on performance allows the platform to change the relative weight that different user types have on content production. More demanding performance thresholds decrease the average contribution of all users, but relatively more that of *committed* users, who take longer time to reach the threshold. This relative effect is, anyway, partly compensated by higher inactivity of *enthusiasts*. Table 6 illustrates this insight by reporting how much users of each type contribute along the simulated 150 weeks. Conditional on full autonomy over editing (i.e. 0-point delegation threshold), *committed* users contribute about 30% more answers than *enthusiasts*.[44] However, with a strictly positive threshold the relation reverses and *enthusiasts* contribute 5-8% more answers than *committed* users. User's average answer quality follows a similar pattern.

For what concerns edits, the relative effect of delegation on performance on user types is different, as the value of autonomy on editing is much higher for *Committed* users. Indeed, they always contribute more edits than *enthusiasts*.

[44]Note that in absolute terms *enthusiasts* users contribute more as they compose the majority of the community

	Type	0 points	500 points	1000 points	2000 points
Answers	Committed	122.93	20.55	15.42	13.56
	Enthusiast	92.56	21.87	16.38	14.62
Quality	Committed	1.77	0.83	0.79	0.77
	Enthusiast	1.46	0.84	0.82	0.80
Edits	Committed	69.90	9.05	6.07	4.72
	Enthusiast	36.94	6.28	3.55	2.78
Reached priv.	Committed	1.00	0.19	0.07	0.02
	Enthusiast	1.00	0.22	0.05	0.02
No part.	Committed	0.07	0.52	0.56	0.57
	Enthusiast	0.24	0.61	0.68	0.70

Average user-level contributions for a simulated participation of 150 weeks, by delegation threshold. Includes the total contribution in answering (*Answers*), average answer quality conditional on answering (*Quality*), total contribution in editing (*Edits*), whether they reached the editing privilege (*Reached priv.*), and the rate of inactivity (average number of weeks of no participation - *no part.*)

Table 6: Simulated total contributions of a user

While the gap in edits halves (in percentage terms), shifting the threshold from 0 to 500, more demanding thresholds increase the gap again. This is because with higher thresholds, the rate of inactivity of *enthusiasts* users increases much more than for *committed* users. With a threshold set at 500, 15% more *enthusiasts* obtain autonomy, but this gap disappears with a 2000-point threshold.

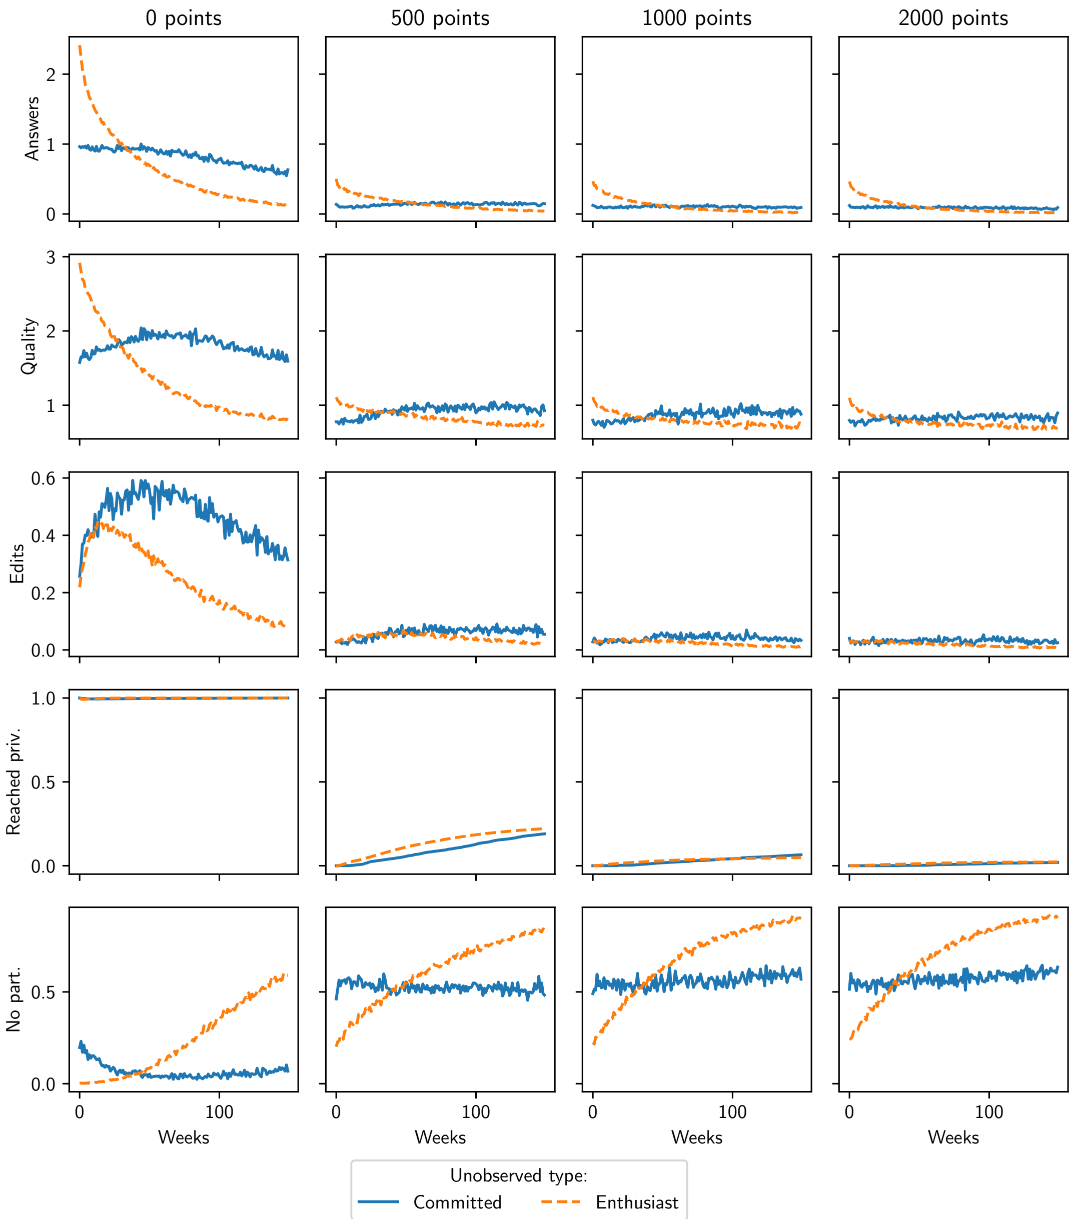

Simulated average weakly contributions by user type for different delegation designs (columns). The editing threshold is set to either 0, 500, 1000, or 2000 reputation points.

Figure 4: Simulated contribution paths by delegation threshold

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 digital platforms' 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. Indeed, the paper finds evidence that the utility from answering increases when users gain autonomy over editing. These results contribute to the literature on multitasking (Holmstrom and Milgrom (1991)), suggesting that when performance is less measurable in one task (editing) than in the other (answering), the allocation of autonomy proves to be an incentive that would not create biases towards specific tasks. This would not be the case for performance-based incentives, as discussed by the literature. In addition, results confirm that the asymmetric allocation of autonomy across tasks leads to relatively more contributions in more autonomous tasks. However, the overall level of contribution remains lower than by providing autonomy across tasks.

Second, the paper finds heterogeneity in the value of autonomy and participation more generally. The small group of *committed* users value editing with autonomy much more than *enthusiast* users. In addition, *Committed* users sustain more contributions across time under all delegation designs, while *enthusiast* users are more likely to become inactive. At the same time, when delegation is based on performance, *enthusiasts* reach the delegation threshold faster and experience autonomy from an earlier stage of their participation history.

These results have important implications for platform design and the creation of user-generated public goods. They suggest that if platforms want to maximise user-generated content, they should provide as much autonomy as possible to users. This may partly explain Wikipedia's success across the years, and goes in line with more recent policies of Stack Exchange, which, since February 2023, is fixing lower performance thresholds for future sites.

However, the platform may find optimal to delegate editing autonomy on performance if it wants to give more weights to certain types of users in the overall content production. *Committed* users tend to contribute more on average, but take a longer time to achieve the reputation points necessary for the acquisition of autonomy. Performance-based delegation would lead *enthusiasts* to contribute more answers than *committed* users. This effect, however, disappears if the threshold is too high, since *enthusiasts* would then be likely to become inactive before they are able to obtain autonomy.

References

Aaltonen, A. and S. Wattal (2025, July). Not Good Enough, but Try Again! The Impact of Improved Rejection Communications on Contributor Retention and Performance in Open Knowledge Collaboration. *Management Science*. 5

Aghion, P. and P. Bolton (1992, 07). An incomplete contracts approach to financial contracting. *The Review of Economic Studies 59*(3), 473–494. 7

Aghion, P. and J. Tirole (1997, February). Formal and real authority in organizations. *The Journal of Political Economy 105*(1), 1–29. 2, 7

Arcidiacono, P., E. Aucejo, A. Maurel, and T. Ransom (2025, January). College attrition and the dynamics of information revelation. *Journal of Political Economy 133*(1). 14

Arcidiacono, P. and P. Ellickson (2011, September). Practical methods for estimation of dynamic discrete choice models. *Annual Review of Economics 3*, 363–394. 49

Arcidiacono, P. and R. A. Miller (2011, November). Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. *Econometrica 79*(6), 1823–1867. 4, 18, 23, 47, 48

Arcidiacono, P. and R. A. Miller (2019). Nonstationary dynamic models with finite dependence. *Quantitative Economics 10*(3), 853–890. 19

Arcidiacono, P. and R. A. Miller (2020, April). Identifying dynamic discrete choice models off short panels. *Journal of Econometrics 215*(2), 473–485. 18

Auriol, E. and R. Renault (2008, Spring). Status and incentives. *The RAND Journal of Economics 39*(1), 305–326. 8

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, 34

Bartling, B., E. Fehr, and H. Herz (2014, November). The instrinsic value of decision rights. *Econometrica 82*(6), 2005–2039. 5, 7

Beckmann, M. and M. Kräkel (2022, October). Empowerment, Task Commitment, and Performance Pay. *Journal of Labor Economics 40*(4), 889–938. 7

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

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. 8

Bester, H. and D. Krähmer (2008, Autumn). Delegation and incentives. *RAND Journal of Economics 39*(3), 664–682. 7

Blanes i Vidal, J. (2007, February). Delegation of decision rights and the winner's curse. *Economics Letters 94*(2), 163–169. 8

Bloom, N., R. Sadun, and J. Van Reenen (2010). Does Product Market Competition Lead Firms to Decentralize? *The American Economic Review 100*(2), 434–438. 5

Bloom, N., R. Sadun, and J. Van Reenen (2012, November). The Organization of Firms Across Countries*. *The Quarterly Journal of Economics 127*(4), 1663–1705. 5

Bregolin, J. (2025, June). Communication Quality and the Cost of Language: Evidence from Stack Overflow. *Management Science 71*(6), 4669–4687. 5

Bruneel-Zupanc, C. (2020). Discrete-continuous dynamic choice models: Identification and conditional choice probabilities estimation. *Working Paper*. 49

Burtch, G., D. Lee, and Z. Chen (2024, May). The consequences of generative AI for online knowledge communities. *Scientific Reports 14*(1), 10413. 5

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*. 14

De Varo, J. and S. Prasad (2015). The Relationship between Delegation and Incentives Across Occupations: Evidence and Theory. *The Journal of Industrial Economics 63*(2), 279–312. 5

Dewatripont, M., I. Jewitt, and J. Tirole (2000, May). Multitask agency problems: Focus and task clustering. *European Economic Review 44*(4), 869–877. 8

Ehrenberg, R. G. and M. L. Bognanno (1990, December). Do Tournaments Have Incentive Effects? *Journal of Political Economy 98*(6), 1307–1324. 8

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, 7

Forderer, J. and G. Burtch (2024, June). Estimating Career Benefits from Online Community Leadership: Evidence from Stack Exchange Moderators. *Management Science*. 5, 8

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

Gershkov, A. and M. Perry (2012, January). Dynamic Contracts with Moral Hazard and Adverse Selection. *The Review of Economic Studies 79*(1), 268–306. 2, 8

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

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, 12

Harper, F. M., D. Raban, S. Rafaeli, and J. Konstan (2008). Predictors of answer quality in online q&a sites. In *CHI, Exploring Web Content.* 9

Hart, O. (2017, July). Incomplete Contracts and Control. *American Economic Review 107*(7), 1731–1752. 7

Holmström, B. (1999, January). Managerial incentive problems: A dynamic perspective. *Review of Economic Studies 66*(1), 169–182. 2, 8

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. 2, 8, 34

Holmstrom, B. and P. Milgrom (1994). The Firm as an Incentive System. *The American Economic Review 84*(4), 972–991. 8

Hong, B., L. Kueng, and M.-J. Yang (2019, November). Complementarity of Performance Pay and Task Allocation. *Management Science 65*(11), 5152–5170. 5

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. 18, 47

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

Katayama, H., K. J. Meagher, and A. Wait (2018, November). Authority and communication in firms. *Journal of Economic Behavior & Organization 155*, 315–348. 5

Khorunzhina, N. (2013, December). Structural estimation of stock market participation costs. *Journal of Economic Dynamics and Control 37*(12), 2928–2942. 14

Lazear, E. P. and S. Rosen (1981, October). Rank-order tournaments as optimum labor contracts. *Journal of Political Economy 89*(5), 841–864. 8

Liberti, J. M. (2018, August). Initiative, incentives, and soft information. *Management Science 64*(8), 3469–3970. 5, 34

Liu, D., K. J. Meagher, and A. Wait (2022, December). Market conditions and firm morality: Employee trust in the honesty of their managers. *Journal of Economic Behavior & Organization 204*, 89–106. 5

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. 18

McElheran, K. (2014). Delegation in Multi-Establishment Firms: Evidence from I.T. Purchasing. *Journal of Economics & Management Strategy 23*(2), 225–258. 5

McFadden, D. (1974). *Conditional Logit Analysis of Qualitative Choice Behavior*. Academic Press. 23

Meagher, K. and A. Wait (2021). Power up: A Cognitive Explanation of Heterogeneity in Preferences for Power. *Working paper*. 5, 7

Meagher, K. J. and A. Wait (2014, November). Delegation of Decisions About Change in Organizations: The Roles of Competition, Trade, Uncertainty, and Scale. *The Journal of Law, Economics, and Organization 30*(4), 709–733. 5

Meagher, K. J. and A. Wait (2020, November). Worker Trust in Management and Delegation in Organizations. *The Journal of Law, Economics, and Organization 36*(3), 495–536. 5

Mookherjee, D. (2006, June). Decentralization, Hierarchies, and Incentives: A Mechanism Design Perspective. *Journal of Economic Literature 44*(2), 367–390. 7

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, 7

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, 18

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, 5, 8

Yoganarasimhan, H. (2013, December). The Value of Reputation in an Online Freelance Marketplace. *Marketing Science 32*(6), 860–891. 14

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 on Data Construction

A.1 Answer Quality

Since the paper models the quality of an answer as a choice of the user, the associated empirical measure aims to capture the amount of effort the user has chosen, rather than an ex-post evaluation of the answer quality such as the number of upvotes the answer receives.

To construct this measure, I proceed with the following steps. First, I identify several possible proxies that capture different dimension of the answer quality. The proxies are calculated for each answer and include:

- The number of days separating the question's publication date and the answer's publication date. Because higher effort implies a quicker response, I make this measure negative, so that higher quality means less days.

- Text characteristics: 1) the number of words in the answer; 2) the share of informative words in the answer, computed as the number of words not included in the *stopwords* list over the total number of words; 3) the number of links in the answer; and 4) the number of pictures included in the answer.[45]

- LLM-based measures of clarity, completeness, and informativeness. More precisely, I ask Mistral AI (version *mistral-small-2503*) to allocate to the answers one category for each aspect of quality. The option given are, respectively, 1) *Unclear*, *Partially clear*, and *Clear*; 2) *Unrelated*, *Partial answer*, and *Complete answer*; and 3) *Scarce*, *Informative*, *Deepened*. I mark as zero the lowest category in each dimension, and 2 the highest. Figure 5 displays the exact prompt.

Table 7 reports descriptive statistics of the proxies.

Second, I standardise each proxy variable and apply a Principal Component Analysis (PCA) transformation. The standardisation ensures that none of the variables will mechanically dominate in the computed components, as they originally take heterogeneous ranges of values. Third, I extract the first component as the measure of quality. The first component explains 34.21% of the variation. While this number is relatively small, there is a substantial drop in explained variation moving to the following components, as shown in figure 6.

[45]The are several established lists of *stopwords* and I use the one available in the Python package *Natural Language Processing Toolkit*. Words included are, for instance, personal pronouns, prepositions, and conjunctions.

```
You will be presented with questions and answers on the use
of the English language, extracted from a crowdsourced Q&A website.
You will need to assess the quality of the answer.
You will output ONLY a json object containing the following information:
{
clarity: string // degree of clarity of the answer.
Choose only from the following list:
1) Unclear; 2) Partially clear; 3) Clear.
completeness: string // degree to which the answer adresses the question.
Choose only from the following list:
1) Unrelated; 2) Partial answer; 3) Complete answer.
informativeness: string // degree of informativeness of the answer.
Choose only from the following list: 1) Scarce; 2) Informative; 3) Deepened.
explanation: string // justify your assessments with in a sentence.
}
```

Figure 5: Mistral AI classification prompt for answers' quality

	mean	std	min	median	max
Days From Question (negative)	-48.06	238.35	-3297.00	0.00	0.00
LLM-based Clarity [0-2]	1.58	0.51	0.00	2.00	2.00
LLM-based Completeness [0-2]	1.55	0.51	0.00	2.00	2.00
LLM-based Informativeness [0-2]	1.44	0.62	0.00	2.00	2.00
Text-based Precision	0.49	0.07	0.00	0.48	1.00
Length	111.54	95.20	1.00	87.00	2855.00
Num. of Pictures	0.02	0.19	0.00	0.00	9.00
Num. of Links	0.34	0.92	0.00	0.00	82.00

Sample size: 147888 answers. The speed or response is the number of days between the question and the answer's publication dates. It is made negative as less days suggest higher quality. The measure of *precision* is the number of informative words over the total number of words. The LLM measure were computed with Mistral AI.

Table 7: Proxy measures of answer quality

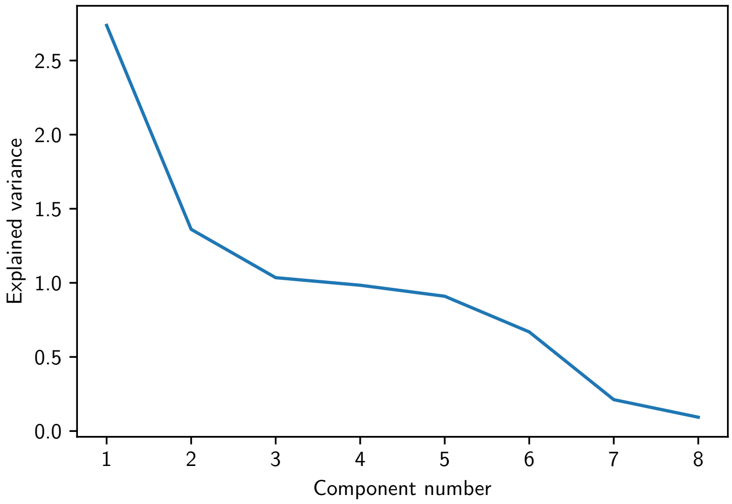

The variance explained by each component shows a substantial kink after the first component.

Figure 6: Explained variance of answer quality proxies by PCA components

A.2 Edit Quality

For each edit, I recover the version of the post (question or answer) before and after the edit. For a given edit, I then prompt an LLM (Mistral AI) to assess how many issues the two versions have. The algorithm then outputs a number for each version. I take the difference between the number of issues found in the version before the edit and those found in the version after the edit, and I use this difference as a measure of edit quality. This procedure involves only edits on posts' bodies, as edits on tags and titles are too small and would lead to too noisy measurements. Figure reports the prompt used in querying the algorithm.

A.3 Construction of scarcity variable

The scarcity variable is a measure of the quantity of available questions on 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 on the website in week t, and $maxavail$ is the $\max\{\log(avail)\}$.

```
You will be presented with two versions of an answer[question]
extracted from a crowdsourced Q&A website.
Identify all the issues in both versions. Issues include grammar mistakes,
unclarity of the meaning, or lack of references and hyperlinks.
You will output ONLY a json object containing the following information:

{
version1: int // search and count the total number of issues in version 1.
version2: int // search and count the total number of issues in version 2.
explanation: string // explain which issues in a sentence without commas.
}
```

Figure 7: Mistral AI classification prompt for edit quality

Appendix B Details on Reduced Form Results

B.1 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.

B.2 Reduced-form Estimates of Answering

Figure 8 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).

	(1)	(2)	(3)	(4)
	NumEdits	At least one edit	NumComments	At least one comment
-700	-0.0470	-0.0286***	-0.449***	-0.0742***
	(0.0635)	(0.00423)	(0.0912)	(0.00668)
-600	0.0151	-0.00360	0.00432	-0.0229**
	(0.0795)	(0.00530)	(0.114)	(0.00836)
-500	0.0164	-0.000393	0.129	-0.00787
	(0.0796)	(0.00530)	(0.114)	(0.00837)
-400	0.0450	0.00828	0.149	0.0181*
	(0.0809)	(0.00539)	(0.116)	(0.00851)
-300	0.0844	0.0150**	0.250*	0.0331***
	(0.0845)	(0.00563)	(0.121)	(0.00889)
-200	0.0802	0.0210***	0.117	0.0308***
	(0.0793)	(0.00528)	(0.114)	(0.00834)
-100	0	0	0	0
	(.)	(.)	(.)	(.)
0	0.0827	0.0148**	0.0338	-0.0108
	(0.0718)	(0.00478)	(0.103)	(0.00756)
100	0.132	0.0202***	-0.0378	-0.0281***
	(0.0757)	(0.00504)	(0.109)	(0.00796)
200	0.113	0.0213***	0.0575	-0.00240
	(0.0778)	(0.00518)	(0.112)	(0.00818)
300	0.156*	0.0256***	0.0544	-0.00683
	(0.0796)	(0.00530)	(0.114)	(0.00837)
400	0.194*	0.0294***	-0.0543	-0.00802
	(0.0815)	(0.00543)	(0.117)	(0.00857)
500	0.237**	0.0384***	0.0429	0.00816
	(0.0858)	(0.00572)	(0.123)	(0.00903)
600	0.216*	0.0249***	0.00867	-0.00248
	(0.0843)	(0.00562)	(0.121)	(0.00887)
700	0.398***	0.0482***	0.218*	0.00569
	(0.0656)	(0.00437)	(0.0941)	(0.00690)
NumAnswers	0.0688***	0.0131***	1.339***	0.0439***
	(0.00396)	(0.000264)	(0.00448)	(0.000328)
NumComments	0.150***	0.0125***		
	(0.00183)	(0.000122)		
is_candidate	4.162***	0.304***	3.057***	0.242***
	(0.421)	(0.0280)	(0.605)	(0.0443)
is_moderator	4.246***	-0.00252	-3.032***	0.0330*
	(0.125)	(0.00835)	(0.180)	(0.0132)
CopyEditor	30.68***	0.190***	15.39***	0.223***
	(0.639)	(0.0425)	(0.916)	(0.0671)
StrunkWhite	7.756***	0.399***	7.033***	0.213***
	(0.311)	(0.0207)	(0.447)	(0.0327)
Observations	146542	146542	146542	146542
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.

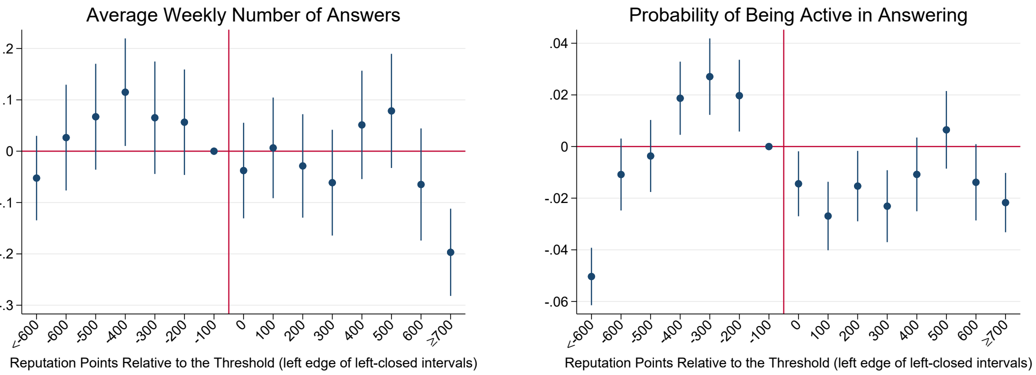

Reputation-points fixed effects before and after achieving the editing privilege. Sample of users who reached the threshold.

Figure 8: Answering Contributions Relative to Achieving the Privilege

Appendix C Details on the Structural model

C.1 Derivation of Likelihood function

Individual problem

Define as $\mathscr{Z}$ 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, but also variables that may affect users' beliefs on the probability distribution over future states.

A user selects a sequence of optimal decisions $\mathbf{d}^* \equiv \{\mathbf{d}_t^*\}_{t\leq T}$ that satisfies[46]:

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

where δ is a discount factor, $d_{\alpha t}$ is equal to 1 if in period t is selected choice α, and zero otherwise, 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

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

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

(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[47]. In other words, it is the continuation value of being in state z_t, before ε_t is realized and the decision at t is taken. By applying Bellman's principle, it is then given by:

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

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

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

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

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

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

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

It follows that:

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

[47] 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.

which can be rewritten as:

$$\nu_{\alpha t}(z_t) = u_{\alpha t}(z_t) + \tag{11}$$

$$\sum_{\tau=t+1}^{T} \sum_{\mathbf{k}\in\mathscr{A}} \sum_{z_\tau\in\mathscr{Z}} \delta^{\tau-t}(u_{k\tau}(z_\tau) + \omega_k(\mathbf{p}_\tau(z_\tau)))d^*_{k\tau}(z_\tau, d_{\alpha t}=1)\kappa_\tau(z_\tau|z_t, d_{\alpha t}=1), \tag{12}$$

where the function $\kappa_{\tau+1}(z_{\tau+1}|z_t, d_{\alpha t}=1)$ 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.

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

To write the conditional value function as in 12 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 12 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, as it is standard in the literature.
This allows to express the choice probabilities as:

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

and the ex-ante value function as:

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

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

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

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

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

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

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

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

C.2 Choice set

To reduce computational time for the estimation, I discretise the choice set to 21 possible combinations of effort, including a no participation decision.[48] All options in the choice set are listed in the table 9. The set of possible levels of contributions in answering is $\{1, 7\}$. A contribution of one answer corresponds to contributions

[48] A more natural assumption would be that users make discrete choices of task (answering and/or editing), and continuous choices for effort intensity. Bruneel-Zupanc (2020) provide a possible approach to estimate discrete-continuous choices.

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	4.68	4.68	4.68	6.61	6.61	6.61	7.76	7.76
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	7.76	4.68	4.68	4.68	6.61	6.61	6.61	7.76	7.76	7.76
NE	9.00	0.00	1.00	9.00	0.00	1.00	9.00	0.00	1.00	9.00

NA is the 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). Values in the tables are rounded to the second decimal.

Table 9: Possible Combinations of Effort Levels

between 1 and 3 answers in the data (which account for 70% of the weeks with positive contributions), while a choice of 7 corresponds to actual choices between 4 and 122 answers. 1 and 7 are the median values in the respective ranges. Similarly, for answer quality the possible choices are $\{4.68, 6.61, 7.76\}$, which correspond to the median values of the ranges $(0.517, 5.789]$, $(5.789, 7.386]$, and $(7.386, 10.819]$. The values 5.789 and 7.386 are the 33rd and 66th quantiles of the distribution of positive quality choices. Finally, users are allowed to contribute one or 9 edits, where a choice of one corresponds, in the data, to contributions between 1 and 4 edits (accounting for 75% of weeks with positive contributions), while the choice of 9 corresponds to contributions between 5 and 406 edits.

C.3 Details on Beliefs Assumptions

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 9 shows how different alternative functional forms would fit the data.

C.4 Reduced Form Conditional Choice Probabilities

Conditional choice probabilities are computed before estimation via a static logit[49] 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

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

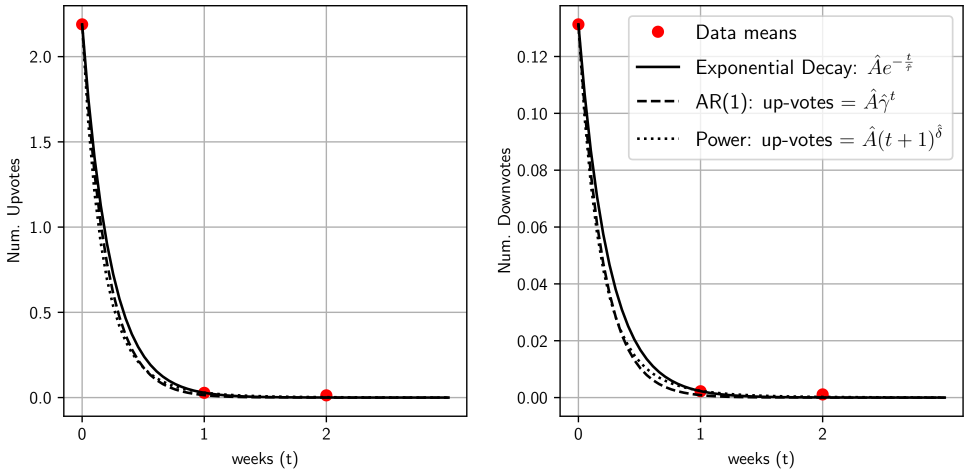

Comparison of different possible functional form assumptions to describe the process of arrival of votes on answers, since publication week (i.e. $t = t_0$). Dots report data means across answers.

Figure 9: Arrival of Votes on Answers

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

$$\begin{aligned}\alpha^*_{it} =&\beta_0 R_{it-1} + \beta_1 \Lambda_{U,it-1} + \beta_2 \Lambda_{D,it-1} + \beta_3 avail_{it} + \beta_4 AnswerNum_{it} + \beta_5 Seniority_{it} \\ &+ \beta_6 t + \beta_7 date_{it} + \beta_8 cumT_{it} + \beta_9 Control_{it} + \beta_{10} rep2Control_{it} + \beta_{11} R^2_{it-1} \\ &+ \beta_{12} t^2 + \beta_{13} date^2_{it} + \beta_{14} rep2Control^2_{it}\end{aligned}$$

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, $Seniority$ is the number of days passed since the registration day, $date$ is the calendar week, $cumT$ is the number of privileges obtained by the user, $Control$ is a dummy equal to 1 if the user has autonomy and zero otherwise, and $rep2Control = R - \bar{R}$, where $\bar{R}$ is the reputation threshold for autonomy. All parameters are choice specific.

C.5 Details on the Model Fit

Figure 10 reproduces the reduced form evidence of section 4 with the simulated data using the model.

C.6 Details on Simulation of Counterfactuals

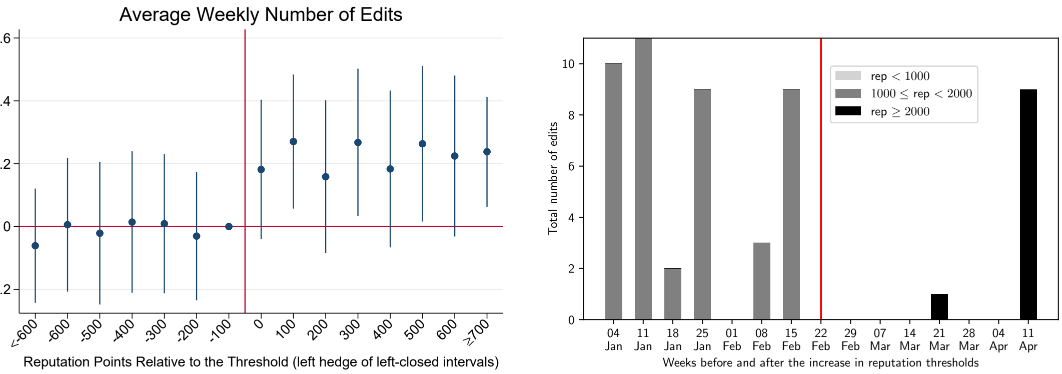

The model correctly predicts the reduced form evidence. Users increase participation in editing once they gain autonomy, while the reduce it when they lose autonomy.

Figure 10: Replication of the reduced form evidence with simulated data

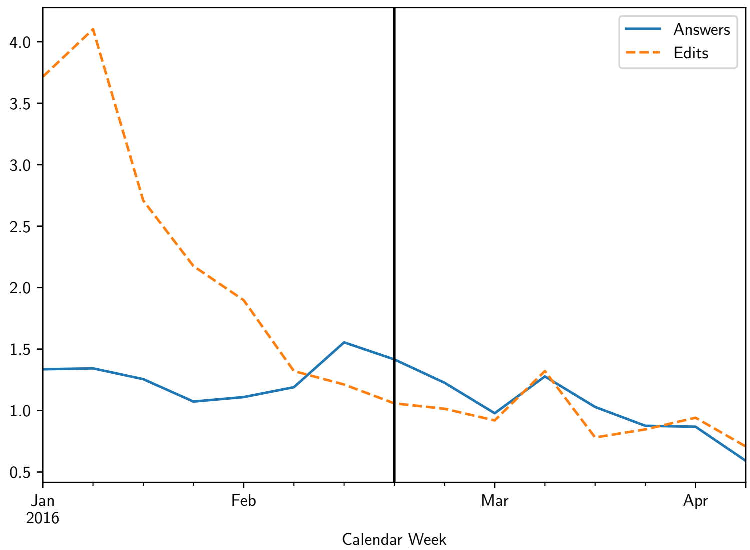

The graph reports the average number of contributions in answering and editing made by users who did not lose autonomy when the threshold changed. It shows that a sudden drop in the number of users with autonomy does not affect contribution patterns of users with autonomy.

Figure 11: Peer effect of autonomy

Appendix D Details on the Website Design

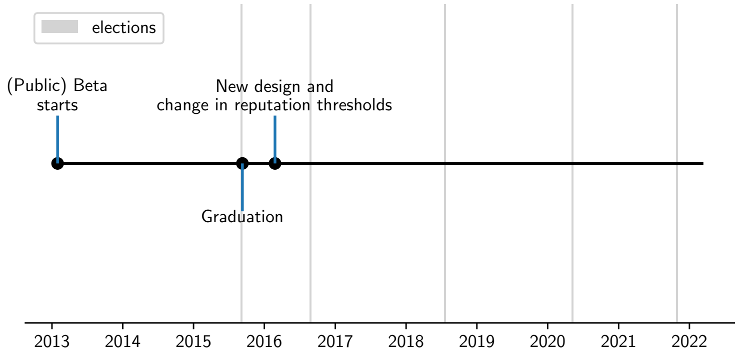

The site allows contributions from the broad public since a *Beta stage*. When it achieves sustained participation, it *graduates* and moves to more demanding reputation thresholds to achieve the privileges. The actual change in reputation thresholds happened when the site changed design. Elections allow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

Figure 12: Timeline of the ELL website

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

List of privileges that users can achieve and the 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

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

Appendix E 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), Le, 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.