Delegation and Authority in Online Communities

Jacopo Bregolin
Toulouse School of Economics

September 2020
VERY PRELIMINARY

Abstract

Many online platforms rely on user-generated content and need to incentivize free effort. In this paper, I investigate if users provide more and better quality contributions when endowed with more autonomy over actions. Using a dynamic discrete choice model, I show that control rights have positive marginal value that is heterogeneous across different types of users, where types are ex-ante identifiable by the platform. I simulate counterfactuals with different delegation designs. Results show that the platform would lose an important share of production and quality of content in absence of delegation.

1 Introduction

Many companies rely on the voluntary contribution of internet users to raise revenues. Some firms exploit user-generated content to enhance and differentiate their product, like online retailers via product reviews[1]. For other businesses instead, users' content is the product itself. This is the case of social media platforms, or question-and-answer websites like Stack Exchange. How companies can incentivise participation without remuneration?
In this paper I investigate if the delegation of control rights and authority could induce more and better contributions in an online community. I address the question of whether and to what extent users are interested in reaching more autonomy over tasks, and study its role in contribution patterns.

I exploit data from Stack Exchange, a website where registered users can ask questions and provide answers on different topics. The identifying variation comes from a characteristic of the website's rules. The moderation of the site is made by the community itself, but if a user make an edit of some content, the implementation is subject to a third party decision. Nevertheless, at the reach of a given performance measure, users are delegated full autonomy, becoming able to implement edits without third-party approval.

This work brings two main contributions. The first relates to a positive analysis of non-monetary preferences. My work aims to provide evidence that individuals value authority and control rights over actions. Behavioral Economics theory has investigated how several factors, like status and power, may affect behavior and optimal organizational design, but little has been shown with non-experimental data. I use real data of observed choices, and a reveled preference approach.
This investigation brings to the literature an additional dimension of non-monetary preferences, where reciprocity, reference dependence, and signaling, between others, have already been identified as driving factors of voluntary contributions.
The second contribution relates to the organizational implications of these non-monetary preferences. I study the productivity trade-offs faced by the company when deciding the delegation mechanism, allowing for different responses across users' types. As suggested by theoretical work in fact, companies can allocate authority to community members as an incentive device, to either lower remuneration or induce more participation.

For the analysis, I implement a mix of statistical tools and econometric approaches. I identify types of users with a Multiple Correspondence Analysis, combined to a K-Means clustering algorithm. This technique let me aggregate information from what users display in their user pages, and discriminate users without relying on ad-hoc assumptions.

The inference is then split in two methodological approaches. In a first part I carry some reduced form analysis which tests behavioral hypothesis. I use a regression discontinuity approach and compare behavior across treated and untreated actions. In the second part I develop a dynamic discrete choice structural model, which implements a dynamic version of the conditional logit model. The choice of this model is crucial for two reasons. From one side, it addresses the intrinsic dynamic nature of participation. If delegation is conditional on performance outcomes, individuals choose effort based on the expected returns over the performance measure, i.e. they decide their contribution level based on today's cost of effort, and future discounted benefits. A static model would not be able to identify marginal utilities from this type of investment decision. On the other side, a model of inter-temporal choices allows me to simulate counterfactual production histories, and provide policy implications for designers of online communities.

The identification of types led to three groups of users. About half of active contributors are part of one and are characterized by little or no information displayed in their user pages. A second quite large group provides more information: they are more likely to have a website link, their location, and a small biographical note. Finally, the users part of the last group, the smallest, are the most extrinsically motivated. They tend to have a Linkedin profile, a relatively long biographical note with at least some links inside, location, and website.
While the types are identified only by the information they display in the user page, they have specifically

[1] Other examples are applications that let users share content and reviews, like Google Maps or Spotify, or companies with online forums for questions and answers related to their product, Mozilla's Firefox for instance.

different behavioral patterns. The group with most information displayed is the most active, with the highest per-user number of answers and edits. They make most of the moderation activity, and they collect most of the badges (target based virtual rewards). On the contrary, the group with the least information displayed is the one with the lowest marginal contribution.
The reduced form analysis suggests that the intensity with witch users make actions is correlated to the degree of autonomy they have over them. Interestingly, the group of users with the least information is the most sensitive to the allocation of authority, and they are substantially driving the result. This heterogeneity is also identified in the structural model. This same group shows to have a large net cost of effort, but the highest marginal value for control. The allocation of control increases their willingness to make effort more than for the others, even though it has a positive effect on all contributors.

Counterfactual simulations show that, in absence of a pre-announced delegation system based on performance, the website would lose an important share of the answers published. This effect is mainly driven by users displaying little information, who drop their participation. The quantity of edits would not be much affected. If instead the website design would not implement any delegation at all, the amount of editing would also drop, as users have an important preference for editing autonomy.

The paper proceeds as follow. Section 2 presents an overview of the related literature, section 3 describes the website from which data is taken from, while sections 4 and 5 present the data and the identification of user types within the online community. I then present results from the reduced form analyses in section 6, and the structural model in section 7. Finally, sections 8 and 9 report the results and the counterfactual simulation. Section 10 conclude.

2 Related Literature

This paper relates to different strands of literature. It contributes to the Experimental Economics literature which studies and measure individual preferences in laboratory experiments. On the policy side, this paper addresses the large theoretical literature in Organizational Economics, studying delegation systems and optimal allocation of decision rights. Finally, it relates to a growing body of multidisciplinary work which studies online communities, and individual participation in digital platforms.

2.1 Delegation and Authority

What does it means to delegate a task inside an organization, and what are the implications?
Generally, the literature in Organizational Economics identifies this action in the strategic choice of a principal, normally located in some upper position within a formal hierarchy, who needs to decide whether to delegate control rights over resources to an agent. In other words, delegation is the empowering of a subordinate to take an action, which initially was of strict competence of an higher role in the organizational hierarchy.
This allocation may or may not grant power[2] and authority[3]. Power and authority, on top of autonomy over decision-making, imply an additional component, that is the ability to influence other people's actions and wills. As a consequence, the delegation of control rights over a decision process provides authority if that decision is meant to impact other people's actions and/or payoffs.

These definitions are instrumental to underline that the behavioral implications of delegation may be two-fold. Individuals could be affected by both the allocation of control per se, as well as the allocation of relative power over peers. If delegation provides more authority, these effects are generally confounded and not separately identifiable. This is the case for instance in most promotions, where the higher degree of autonomy goes hand in hand with more authority over other workers.
In the context of the present work, these two effects are confounded. The delegation occurs on the editing task, which directly affects production of peer contributors. The potential value users derive by this allocation per se, may be induced by a preference for autonomy in determining the editing outcome, a preference for being able to directly affect other people actions, or a combination of the two.

[2] I refer to the definition of power given in Sturm and Antonakis (2015): *Power is having the discretion and the means to asymmetrically enforce one's will over entities*, [...] involving "power to," or the "production of intended effects" [...] and "power over," which involves compelling others to do what one wants them to do [...]

[3] I consider authority as institutionalized power.

While this paper still suffers of this under-identification issue, it is still relevant to point out this difference, to make clear what the existing literature has studied[4].

Theoretical work in Organizational Economics (Aghion and Tirole 1997, Dessein 2002, Aghion et al. 2004, Bester and Krähmer 2008, Bester 2009) has specifically focused on the delegation of control right over a decision process, whose outcome affects both parties. The trade off rises as principal and agent have different preferences, and the information crucial to take the decision may not be homogeneously distributed.
Empirically, these trade offs specifically induced by incentive misalignment are studied by Bandiera et al. (2020) in a field experiment.
Other work studies how delegation can incentivize effort. Rajan and Zingales (1998) study the strategic allocation of access to critical resources, while Blanes I Vidal and Möller (2007) shows that the allocation of control allows the agent to signal her/his ability, incentivizing effort via career concerns.

This literature assumes that preferences are strictly based on the final payoff induced by the decision: control of the decision process has no value per se, but only to the extent the other party may not choose the preferred option.

An prolific literature in Management and Psychology has anyway identified that *individual power holder might change as a result of possesing power*, as reported by Sturm and Antonakis (2015). More recently, the Experimental Economics literature has addressed the question of whether individuals value control rights per se. Following the standard framework proposed by Aghion and Tirole (1997), Fehr et al. (2013) develop an experiment where a principal endowed with control needs to decide whether to delegate it or not. They observe that individuals tend to retain control even when delegation would be optimal, suggesting a positive intrinsic value of control and autonomy. In another experiment, Bartling et al. (2014) develop a framework to separately identify the size, if any, of the intrinsic value of decision rights[5]. They find that, on average, the principal value the delegation lotteries 16.7% more than the control lotteries.
The experimental design by Owens et al. (2014) has a similar purpose, but it does not adopt the framework of incentive misalignment in a principal-agent relationship. The intrinsic value of control over decisions is observed on the betting on themselves, rather than on a partner, in answering a quiz, even when beliefs on outcomes would induce delegation.
While these experiments measure the value of decision rights, the work by Pikulina and Tergiman (2020) investigates the value of being able to affect other people's payoff, separately identifying it from the preference for control and autonomy.

These results suggests that individuals have intrinsic value for control and authority. This has important implications on the optimal organizational design. As pointed by Gibbons et al. (2013) in section 5.1.2, if agents perceive this intrinsic values, then the commitment to delegation can be an incentive device, and the allocation of control can be used as a reward alternative to monetary remuneration. The same type of incentive is discussed in Besley and Ghatak (2008), where individuals perceive intrinsic value from status.
This may give solutions to the puzzles related to promotions. In fact, it can provide a plausible explanation for the use of promotions 1) rather than bonuses, even if bonuses are more flexible incentives (Baker et al. 1988, Gibbons and Waldman 1999) 2) tied to observable performance measure rather than management skills (Peter principle, Fairburn and Malcomson 2001, Benson et al. 2019).

2.2 Participation in online communities

The implementation of incentive systems that do not rely on monetary rewards is even more relevant in the context of online communities. The literature in Business Economics, Information Science, and some brunches of Computer Science has been investigating the motives behind voluntary provision of effort in digital environments.
Important parts of this work relied on the elicitation of preferences via surveys (Roberts et al. 2006, Nov 2007, Ma and Agarwal 2007). Jeppesen and Frederiksen (2006), exploiting sources like interviews and

[4] The Economics literature has not coordinated on specific definitions. I describe the existing work through the mentioned definitions, and this may lead to incoherence with the titles of the papers cited.
[5] Note that I consider synonyms *the control rights over decisions*, *the decision rights*, and *autonomy over decisions*.

questionnaires, find that users may contribute as a hobby, and they value recognition from the firm that their contribution is helping.
Another approach has been to conduct field experiments, or exploit natural experiments. Zhang and Zhu (2011) identify the size of the community as a crucial determinant of participation, exploiting the shut downs of Chinese Wikipedia. Chen et al. (2010) implement an experiment in a movie review crowd sourced platform, and finds that users show reference dependence: the amount of contribution is affected by information on the distance from the median contributor. The experiment carried in Google Answers by Chen et al. (2010) instead identifies the reputation system as an important driver of effort, even more than monetary remuneration.
Moving away from the experimental approach, Chen et al. (2017) develop a dynamic model (but with no rational expectations) with a latent motivation state that can change across time. They find an important role for peer recognition.
Closer to the incentive system I look in this paper, Goes et al. (2016) studies the role of a hierarchy within the community, and how sequential rewards providing more and more resources affect participation. Differently from my approach, they adopt a static model.
Peer-recognition is also identified as a relevant motive by Jin et al. (2015). In the context of Stack Overflow, a website for question-and-answers about programming languages, a source of motivation is to signal to the labor market skills in programming. This role of signalling has been identified by Xu et al. (2020).
Belenzon and Schankerman (2015) also find that labor-market signaling is a driver, together with intrinsic motivation,and reputation. This paper relates to my work also on a different dimension. To address the heterogeneity of users, I identify user types from individual characteristics. Belenzon and Schankerman (2015) also discriminate users across types. Their approach is anyway different, as they use participation choices as discriminating factors.

Theoretically, the literature that studies contribution patterns on online communities is more scarce. Roberts et al. (2006) provides a theoretical framework based on an organizational and social psychology model, while Jain et al. (2014) use a game-theoretic approach. The latter is specifically relevant for my work, as it specifically models optimal design of a question-and-answers website.

3 Stack Exchange: "self managed" platforms

Stack Exchange (SE) is a family of platforms born in 2009 that provides to registered users the possibility to post question and answers on a variety of topics[6]. Each website of the group specializes on a particular topic: notably *Stack Overflow*, the largest community, hosts Questions and Answers (Q&A) about programming languages, but there are some other 172 website, each focused on a different topic, from technology to arts.
These websites belong to the commercial company Stack Exchange Inc. which has[7], at December 2019, raised 70 million dollars in venture capital.
To give a sense of the welfare produced to consumers, Stack Exchange (SE) receives 418.8 million monthly visits and 805.9 million monthly page views. 3.3 million questions have been asked, and they received 3.6 million answers[8]. Instead of hiring experts to answer questions[9], SE is crowd based. Anyone can freely register and contribute in the platform without receiving any type of monetary compensation. Content produced can then be accessed by any Internet user.
Stack Exchange sophisticated incentive system based on reputation points and *badges* that reward contribution, both per se and via measures of its quality.

Badges. Badges are sort of medals assigned to the registered user that accomplishes a target set of actions and outcomes. They can be compared to bonuses in firms. There are bronze, silver and golden badges, based on how demanding the target is (generally, moving from bronze to gold, there is an exponential increase in requirements).
Targets relate to all types of possible activities: questions, answers, and moderation between the principal ones.

[6]For what follows, with the noun *post* I refer to a question or an answer.
[7]Please refer to: https://stackoverflow.com/company
[8]https://stackexchange.com/about
[9]This was the case for Google's service *Google Answers*, which have been active between 2002 and 2006.

Reputation points and Privileges. After the publication of content, users may be rewarded with reputation points. The main channel is publishing questions and answers. Figure 29 in the appendix provides the detail list of ways to gain or loose points. Community members (with already 15 points at least) can give positive or negative votes to posts and each vote turns in reputation points for the author. The accumulation of point allow the authors to obtain *privileges*. With few exceptions, privileges are rewards that give access to resources or actions. They are awarded when users reach given threshold levels of reputation points, in a hierarchical way. The higher the amount of points accumulated, the closer the user gets to have full administrative control of the website.

Privileges could be compared to promotions in traditional companies: more experienced *employees* are allocated more information and authority from the company owners, as well as more responsibilities.

The actions that the users get the possibility to implement relate to creation of content (to vote up and down, create tags, etc.), communication (chat with fellow users) and moderation (flag posts, modify questions and answers, cast close posts, etc.).

The objective of the platform owners is to provide detailed solutions for very specific issues and potentially cover most aspects people may need help for. Information should be easy to find and not noisy. For these reasons quality is at the core of the business strategy. Joel Spolsky, the inventor and co-founder, writes on his blog in reference to an example question on fizzbars[10]:

> After all, for the next 20 years, this question will be the canonical place on the web where programmers will come to find out about enlarging fizzbars without overwriting snibbits

Moderation is then a crucial task, and it is relevant for the company to implement a system sustainable in time, able to address an increase in size in contributors and content.

3.1 Hierarchy and delegation in Stack Exchange

In Stack Exchange, the communities created around each website are administered by hired employees. Nevertheless, most of the routinely moderation and administration is made by non-remunerated community members. Any registered user can acquire administrative rights by accumulating the privileges listed in table 1. The reputation points necessary to obtain each privilege are reported in the rightmost columns[11]. In addition, each website runs few internal elections to elect moderators. Elected users obtain all moderation tools without fulfilling the reputation requirements. Elections are infrequent, leading to few users elected with perpetual mandate.

In this paper, I identify the allocation of resources and authority via privileges as a delegation system. Through this rule, the platform administrators commit ex-ante to provide access and more autonomy on the use of resources, conditional on the achievement of performance targets.

Note anyway an important difference with the setting adopted in organization economics models. In Stack Exchange there is no rivalry in obtaining privileges. Authority is shared by all users that achieved the privileges. Rights that a user X obtain over other people's content can be obtained by others with respect to X's content.

3.2 The allocation of control and authority on editing

In the website, users can mainly contribute providing content (question/answers) and editing, that is improving of content produced by someone else.

Once they register, users have full authority on content production. They can publish questions and/or answers, they can modify them after publication and erase them[12]. This is not the case instead for moderation. Users are allowed (and incentivized) to suggest modifications to the content of posts published by fellow community members, but the suggested modification is not directly implemented. A suggested edit (as they will be called herein) to a post needs approval from either the owner of the post, or two positive votes from users that have the authority to do so. If the suggested edit receives two negative votes instead, then it is rejected and not implemented.

When users obtain the *edit questions and answers* privilege, they are now able to directly implement the edits to other people's posts, as well as to vote approval or rejection of suggested edits. In this paper the

[10] from the blog Joel on Software.

[11] Since these values changed during the life of the website, each privilege is matched to two threshold value. In section 4 I will provide more details on how the change happened, and when each threshold has applied.

[12] There is one exception: if the answer and/or question is considered by some moderator of long term interest for the community, moderators can remove the full authority of the owner on the given question/answer.

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

Table 1: Reputation point requirements to obtain the privileges

edit questions and answers privilege is then capturing the delegation of control rights over editing, since it provides more autonomy over the action, and authority over fellow users.

3.3 What is the rationale for delegation?

Several arguments can be addressed to justify different level of delegation. Gibbons et al. (2013) (section 5.1) propose three main reasons, namely to economize on limited resources, to make better use of the employee's information, and to pay the employee less.
The first and second reasons are clearly very relevant for Stack Exchange. Financial resources may be scarce as most of the services provided are free of charge, and the possibility to outsource tasks to voluntary contributors induce important savings for the company. Facebook, who relies on paid moderators, counts about 15,000 content reviewers, but many more would be required to assure healthy working conditions[13]. At the same time, to hire moderators in the labor market may be harder for Stack Exchange. The platform hosts questions and answers covering very specialized topics, so moderators need very specific competences. Contributors anyway have this knowledge, so delegation may profit of a self-selected pool of competent users.
The last reason does not apply in strict terms, as users are not remunerated. From a principal-agent perspective anyway, it means that the platform could use delegation as an incentive device. Under the conditions that 1) users value the allocation of control and authority, and 2) they anticipate they can achieve more control and authority, delegation can induce more participation.

In this paper I will focus on the third reason, which will be at the center of the counterfactual analysis.

4 Data

Stack Exchange is composed by many websites sharing the same structure, whose only difference is the main topic of question posted. In this paper I use the website called *English Language Learners* (ELL), which focuses on questions and answers related to the use of English.

Each website of Stack Exchange is created following a specific procedure. The proposal of creation is opened in a specific platform called *Area 51*[14]. The community of people interested to have a Q&A website on that topic starts to interact asking and answering questions. When the website, within the *Area 51*, proves to be "healthy", that is to have enough demand and a sustained amount of activity, it is launched with an independent url, and steps to the beta phase. The beta period is divided in private beta and public beta. The private beta allows participation only to users that have contributed in the development phase. Normally, after a week the website moves to the public beta phase, characterized by no restrictions to participation.

A relevant difference across beta phases and final phase is that the amount of reputation needed to reach the privileges changes[15]. Normally the variation in reputation thresholds would occur at the *graduation* of the site, i.e. when the site steps in the last phase, and at this same date the site gets its definitive graphical design. Nevertheless, for technical reasons in the ELL community the dates were shifted. First the site graduated, and later the reputation thresholds were changed, as shown in the timeline in figure 1.

The data was retrieved on May 31^{st}, 2020, and contains both information displayed in the user profile pages, as well as content and modification history of posts (questions and answers). The history of amount of reputation points that users had at each point in time is not available, and I obtained it by web-scraping.
At the time of download of the data, the website counted 92,853 registered users, 121,633 published answers, and 77,357 published questions.

From this data, I constructed a panel of users' participation in the website.
As the main purpose of the website design is to incentivize answering and editing, I include only users that have done both or either of these actions. Users that have not gained a positive number of reputation points are also excluded. The history of participation of every user is cut after 3 months of no answering

[13] Charlotte Jee, MIT Technology Review, June 2020
[14] https://area51.stackexchange.com/
[15] Table 1 reports the amount of reputation points required to obtain each privilege.

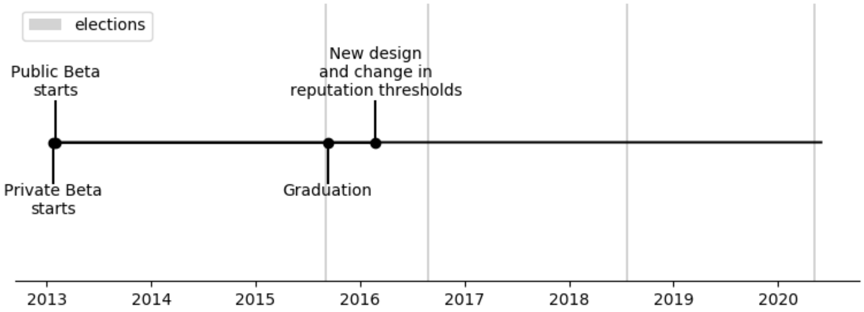

Figure 1: Time line of the website history

or editing activity, or by the download date.
In the panel there are 9797 users, who participated in average 713 days (with a range between 1 and 2685 days). They published a total of 114,926 answers, in average 11.7 each, but with a very skewed distribution, ranging from 1 to 4173.
The edits I consider are edits to answers, either modifications of answer's content or rollbacks, i.e. the recovery of a previous version of the answer. The users of the sample made 8168 edits, of which 1409 were suggested and the rest directly implemented. Each user in average made 0.8 edits, with a range between 0 and 1174.
In the sample, users reached in average 487 reputation points, with a range going from 0 to 175,955, the 75^th percentile being 208 (the zero is due to a particular case that got included in the sample). Figure 2 reports the number of users, at each point in time, who have reached the threshold to obtain more control over editing. Note that when the threshold value changed, this made some users loose the privilege.
Using textual characteristics of answers at the creation time, I construct a measure of answer quality. Details on this process can be found in the appendix A.1. In average, users made answers with quality equal to 0.16, and the quality range spans between 0.004 and 14.107.
Finally, I construct a variable to measure the number of questions available to be answered at each day, given the topic of experience of the users. A question is available if it does not have an accepted answer. In a nutshell, the process to construct the variable is to allocate each not-answered question to a topic, and wight the size of available topics with users' expertise on each of them. Details on the construction of this variable can be found in appendix A.2. In average, users have about 8,000 available questions in a given day, ranging from 22 to 21744 (note that at the time of data retrieval, out of the total 77,357 questions, 38,015 do not have an accepted answer).

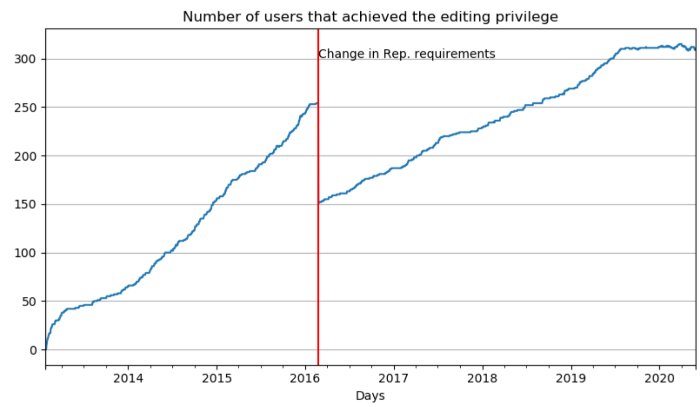

Figure 2: Number of users that have obtained control over the editing task.

4.1 User characteristics

The site's user pages allow members to include information about themselves. Whether users provided or not some type of information is potentially informative of the use they want to do of the website.

I constructed a dataset containing the information of whether users are providing the location, the personal website, a Linkedin profile, and a full name (i.e. an identifier with the form of "Name Surname"). In addition to these dummies I include the size of the biographical note that users can add to describe themselves, in terms of number of words, and the number of links provided in the biographical note. Table 2 reports summary statistics on these variables. It emerges that most users do not have a lot of information, but there is some heterogeneity.

	Share of users
Sample size	92,853
has full name	34.17 %
has website	16.67 %
has location	31.93 %
has Linkedin	1.54 %
has bio note	25.82 %
has links in bio	4.57 %

	net AboutMe	AboutMe	links AboutMe
Sample size	23,979	23,979	4,239
mean	21.92	31.12	2.84
std	33.27	47.66	3.93
min	1.00	1.00	1.00
25%	5.00	6.00	1.00
50%	10.00	14.00	1.00
75%	25.00	36.00	3.00
max	535.00	542.00	55.00

Table 2: Statistics of user characteristics. **(Left)** Share of users that have the given characteristic. **(Right)** Distribution of, respectively, number of words in the biographical note (net of *stopwords*), number of words in the biographical note (all), number of links in the biographical note, for users with a positive value. Sample include all users registered in the website at May 31^st, 2020

5 User Types

Users in online communities may differ substantially between each others on the motives that push them to participate. These heterogeneity induces different activity levels which can be observed in the data, and some patterns are quite standard across online communities. In general, a relatively small share of the registered users provide most of the content, inducing a very long tail in participation.

To give some statistics, in the community of English Language Learners there are 92853 registered users, but only 25607 (27.58%) users made at least one action. Table 5 reports the distribution of the total number of questions and answers published by each users (updated at the date of data retrieval). In average users post 2-3 questions and 4-5 answers, but the 75th quantile of the users published only 1 question and/or 1 answer. From figure 3 is instead possible to see that about 20% of active users produced 80% of questions and about 90% of answers (notice that may not be the same group of users doing both activities).

Platform designers need to take into account this diversity on the use of the website: while a relatively small share of participants provide most of the content, the success and longevity relies on the interactions of all different types of users. Megan Risdal, the lead product manager of Stack Overflow in 2019, writes[16]

> There are many types of people who use Stack Overflow. The main segments for Public Q&A are: Askers (including people who just look up answers to existing questions), Answerers,

[16] https://mrisdal.github.io/blog/posts/reflections-on-stack-overflow/

	Questions (num)	Answers (num)
number of users	25607	25607
mean	2.87	4.63
std	20.26	59.10
min	0.00	0.00
25%	0.00	0.00
50%	1.00	0.00
75%	1.00	1.00
max	949.00	4173.00

Table 3: Distribution of the number of posts published by the *active* users, i.e. the 27.58% of users that did at least one action. (Note that posting is not the only action considered, but the main one)

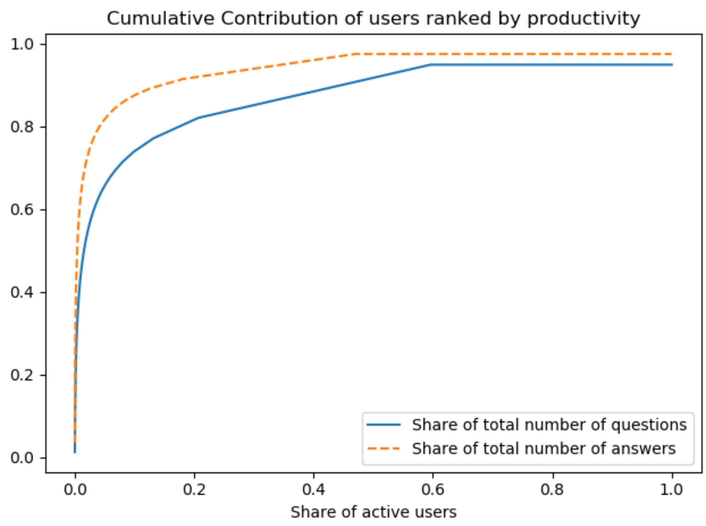

Figure 3: Cumulative share of questions/answers versus share of active users producing that content. Users are ranked by productivity, independently for questions and answers.

> Curators, and Moderators. [...] all of these segments (of real people!) interact with each other in feedback loops [...]. Any initiatives that only serve one segment put the community and product into imbalance.

To address this heterogeneity, I identify types of users using individual characteristics. The rationale is that the amount and kind of information that users provide may be used as discriminating factors for their motivations.

In a nutshell, the procedure I adopt is the following. First, I implement a Multiple Correspondence Analysis (MCA, similar to PCA but for categorical variables) to aggregate information and obtain new variables that keep most of the variation of the data. The objective is to 1) screen out variables that are not informative on the different types[17], and 2) transform the data so to obtain continuous variables instead of categorical. This is instrumental for the application of the second step: I implement the K-Means clustering algorithm to split the sample of individuals in the different types[18].
This process leads to the identification of three types.

More details on the procedure and on why I adopted this specific approach rather than alternatives are given in appendix A.3.
Table 4 provides summary statistics of individual characteristics by type. It is possible to notice that group 3 displays a lot of information: with respect to the other groups, they are more likely to any type of information, in particular a Linkedin profile, a biographical note with links, and a website. Type 2 users provide instead the least information, with most of them not displaying anything. Finally, type 1 are in between, tend to provide a small biographical note, the location, and a website, but not much else.

5.1 Heterogeneity in behavior

While types are only based on individual characteristics, in this section I present how behavior changes across groups.
In the panel, 3705 users are of type 1, 5414 of type 2 and 678 of type 3.

Badges.
Do types differ on the collection of badges? Badges are virtual medals rewarding the accomplishment a performance target. The accumulation of badges may suggest sensitivity to short term incentives. More challenging targets, i.e. silver and golden badges, point even more in this direction, as it is impossible

[17] The year of registration, for instance, seems to not be a relevant discriminating factor

[18] The K-Means clustering algorithm requires the number of clusters as input. To decide how many groups/types I should consider, I loop the K-Means process with the graphical representation of the individuals on the two first components of the MCA. I try several number of clusters and pick the one with the most intuitive separation of individuals in the graph.

user type	Num. users	Share of users that...					
		have full name	have website	have location	have linkedin	have bio	have links
1	23260	29.25%	47.18%	85.09%	0.00%	79.82%	2.88%
2	65134	35.31%	1.65%	9.04%	0.00%	1.73%	0.00%
3	4459	43.13%	76.74%	88.99%	32.05%	96.08%	80.06%

		net AboutMe	AboutMe	links AboutMe
user type	stat.			
1	25%	4.00	5.00	1.00
	50%	8.00	11.00	1.00
	75%	17.00	25.00	1.00
	count	18566.00	18566.00	669.00
	max	361.00	505.00	3.00
	mean	15.48	22.51	1.11
	min	1.00	1.00	1.00
	std	23.59	35.79	0.35
2	25%	9.00	15.00	
	50%	15.00	22.00	
	75%	25.00	37.00	
	count	1129.00	1129.00	
	max	208.00	397.00	
	mean	21.80	33.38	
	min	1.00	1.00	
	std	25.03	40.68	
3	25%	19.00	24.00	1.00
	50%	33.00	44.00	2.00
	75%	61.00	83.00	4.00
	count	4284.00	4284.00	3570.00
	max	535.00	542.00	55.00
	mean	49.90	67.83	3.16
	min	1.00	1.00	1.00
	std	51.56	71.11	4.20

Table 4: User characteristics by user type. The first table reports the share of user, for each type, to have the given information displayed. The second table reports the distribution of the size of the biographical note and of the links in it, conditional on having a biographical note/positive number of links.

to achieve them without tailoring behavior to the target. Figure 4 shows the average number of badges obtained by users in each group, where the vertical black bars are the standard errors of the means. It suggests that users more informative in their user pages are also the ones reacting more to short term incentives, type 3 obtaining more badges than type 1, and type 1 than type 2.

Time to reach the editing threshold.
The types show heterogeneity also on the probability, at each point in time, of having reached the delegation threshold.
I estimate the survival function, where the filing event is the achievement of the threshold number of points. Since in the data I use the value of the threshold changed, I estimate two survival functions, one for the users that registered before the change, and one for users that registered when the threshold was already in its final value. Figure 5 shows the plot of the survival functions for each type. Type 3 users obtain control the fastest

Share of production.
Table 5 reports instead the total and average production by type. The marginal contribution of type 3 users is the highest, followed by the one of type 1 users. The most relevant observation is on direct edits, where the 678 users of type 3 made nearly 60% of the total. Similar patterns can be identified in figure 6, where it is plotted the share of total content produced at each month by each type. It is possible to notice that type 3 users reduced their contribution in answering, while remaining the main editors of the website across time.

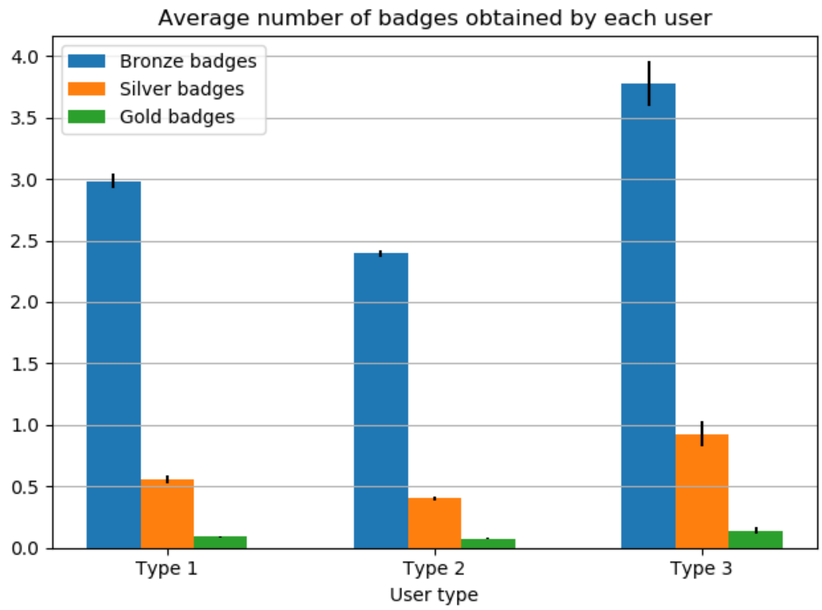

Figure 4: Average number of badges obtained by users in each type. Sample sizes: type 1 has 21192 users, type 2 22708, and type 3 4351.

Type	Num Users	num answers		num suggested edits		num direct edits	
		Total	Avg. per user	Total	Avg. per user	Total	Avg. per user
1	3705.0	63500.0	17.14	836.0	0.23	2272.0	0.61
2	5414.0	32511.0	6.00	309.0	0.06	465.0	0.09
3	678.0	18915.0	27.90	264.0	0.39	4022.0	5.93

Table 5: Total production by type, and average production per user.

Quality of editing.

Participation to elections.
Participation and win of an election may as well reveal information on the motives of participation. The participation in an election may signal in fact that the user has a specific commitment towards the community. To candidate in an election you need to have at least 300 reputation points, while to vote for candidates the requirements is of 150 points[19].
Figure 7 reports the number of candidates by type, and the number of winners. it is possible to notice that the total number of candidates mostly belong to type 1 and 3. Nevertheless, most of the users who have been finally elected belong to type 3.

5.2 Summary on types

Overall, the profile of each type emerges quite clearly form the descriptive evidence. The online community is in large part populated by anonymous users (type 2) that are not particularly active in production, and not much engaged. Little sensitivity to incentive induce small production, which in turns leads to longer average time to achieve the delegation threshold. Nevertheless, the size of this group is such that it still contributes for nearly 30% of the total production of answers.
On the contrary, type 3 users are few members but the most active. They provide a lot of information on their profile, suggesting important extrinsic motives. They produce the most, and provide the majority of the editing activity. Their very high activity may also justify the higher likelihood of winning elections.
Finally, Type 1 users are in between: they provide some information about themeselves, but no links or Linkedin profile, suggesting that they do not aim to signal outside of the platform. They contribute significantly, but are still achieving the delegation threshold in more time than users of type 3.

[19]For more details, see https://stackoverflow.blog/2010/12/02/stack-exchange-moderator-elections-begin/

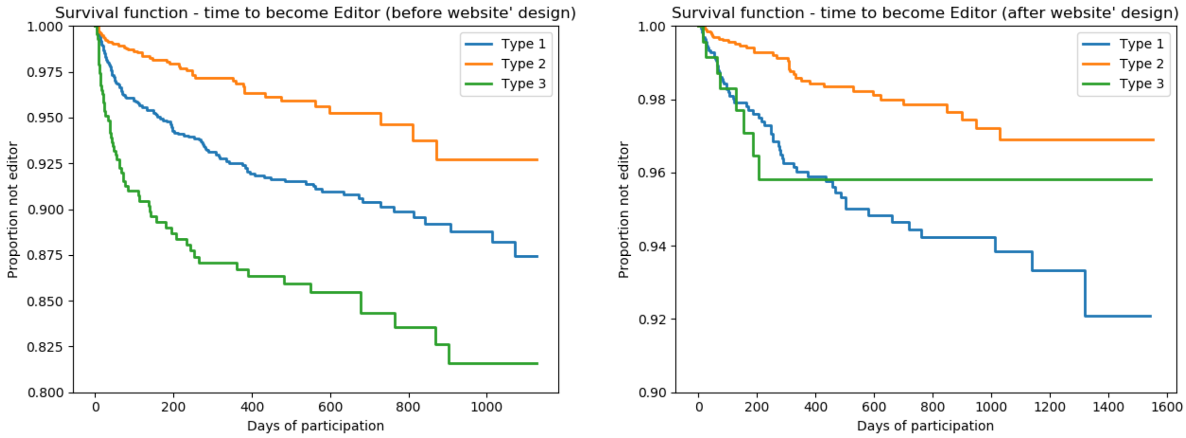

Figure 5: Survival function estimated on users who registered before the threshold change (left) and after the change (right). For the left graph, the data includes time series cut at the date when the reputation threshold changed.

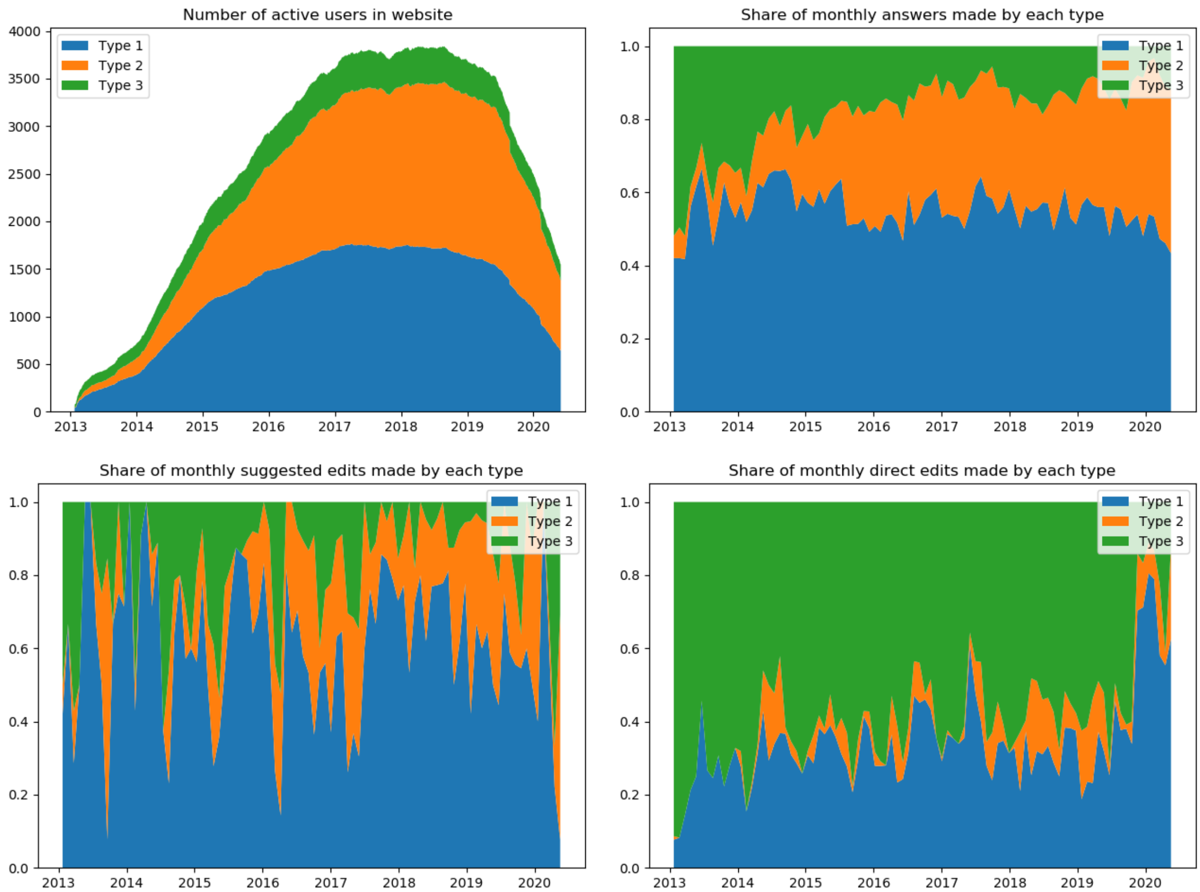

Figure 6: Total number of active users per type and across time (top left) and share of activity produced by each type.

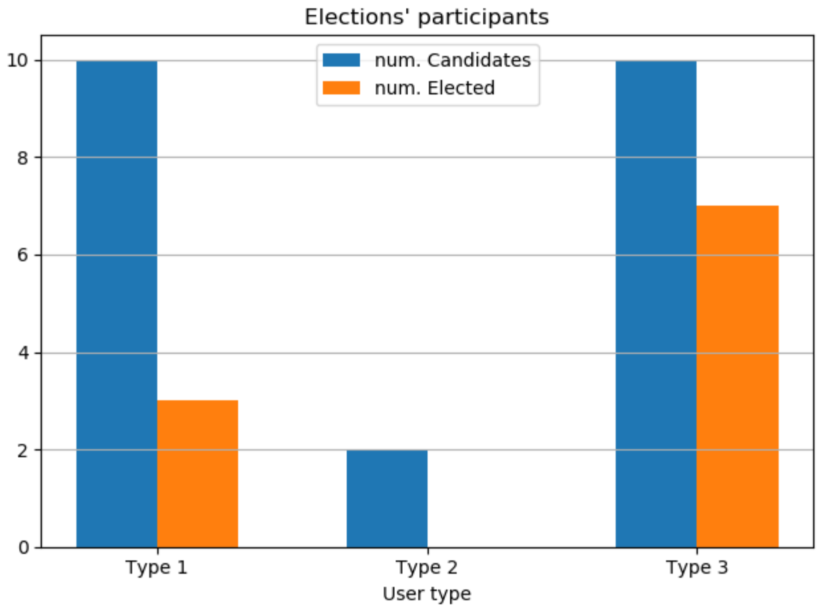

Figure 7: Number of candidates and, within those, the number of winners of the elections, by type.

6 Reduced Form Analysis

In this section, I will first provide some evidence that delegation affects users' behavior. Exploiting a change in the delegation rule, I separately investigate the behavioral response of users losing and gaining control rights. I then show why this is relevant for the success of the website, providing evidence that moderation has an important role for content quality.

6.1 Behavioral effect of delegation

To test preferences from observed behavior in a reduced form way, it is necessary to make assumptions on how preference affect choices. The assumption I will make in this section is the following:

> **Assumption**: The positive evaluation of control and autonomy over an action induces the choice of that action with higher frequency.

This assumption goes in line with evidence from the literature, as reported in the work by Sturm and Antonakis (2015) on the definition of power:

> [...] Generally, research has shown that power increases an action orientation and, thus, leads directly to the taking of action for those who possess it [...]

Testable implications are straightforward. Everything else equal, people that achieve more control and autonomy over an action should take the action more frequently. Vice-versa, a loss in control rights should go along with less activity.

6.1.1 loss of control rights

Due to a change in reputation requirements, users could loose the editing privilege. Stack Exchange websites in fact are rolled out in a beta version. When they move the the definitive version, the reputation thresholds to achieve the privileges increase[20], leading to an automatic adjustment of the allocated privileges.
For what concerns the privilege which allocates autonomy over editing, the threshold passed from 1000 to 2000 points, and everyone with an amount of points in between lost the privilege.

To test the hypothesis, I select users in the website that at the time of the graduation had less than 2000 points[21]. These users knew that, if they had 1000 points or may have reached 1000 points, but not 2000, they would have lost the editing privilege.

Figures 8 and 9 show the share of those people making a positive amount of edits and answers respectively. It is possible to see that users that loose the privilege, i.e. the orange bars, stop making edits, which confirms the hypothesis[22]. The same pattern is not observed in answering behavior. The share of *orange* users participating is not affected by the change in reputation requirements.

6.1.2 Achievement of control rights

To test the hypothesis I also look at behavior when users obtain more control. I exploit the discontinuity created by the allocation of control, and implement a staggered difference-in-difference regression[23]. I study behavior over two actions: editing, which is treated, and commenting, which is not affected by the privilege. Thse two action are somehow comparable, since both of them may be used to help improve the content, and both are not main ways to gain points.

[20]See section 4 for more details

[21]I consider the graduation date, rather than the date of the change in reputation thresholds, to account for potential anticipation effects: once the site graduate, users know that the reputation thresholds will be changed.

[22]This is not driven by the fact that everyone either reached the 2000 points, or never reached the 1000 points. Figure 30 in the appendix shows that every week had a positive number of users with an amount of point between 1000 and 2000.

[23]Sometimes called event study with two-way fixed effects.

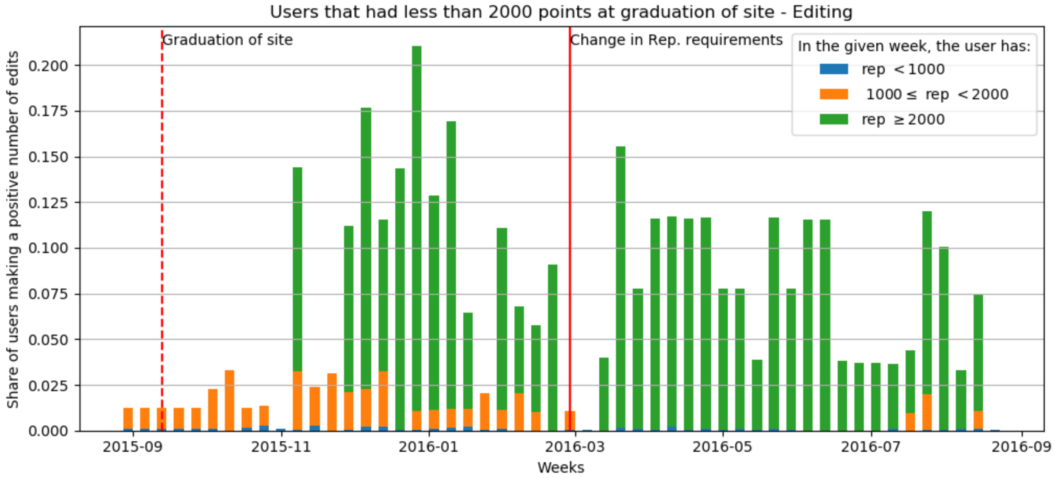

Figure 8: Number of users making a positive number of edits, out of the ones having less than 2000 points at the graduation week.

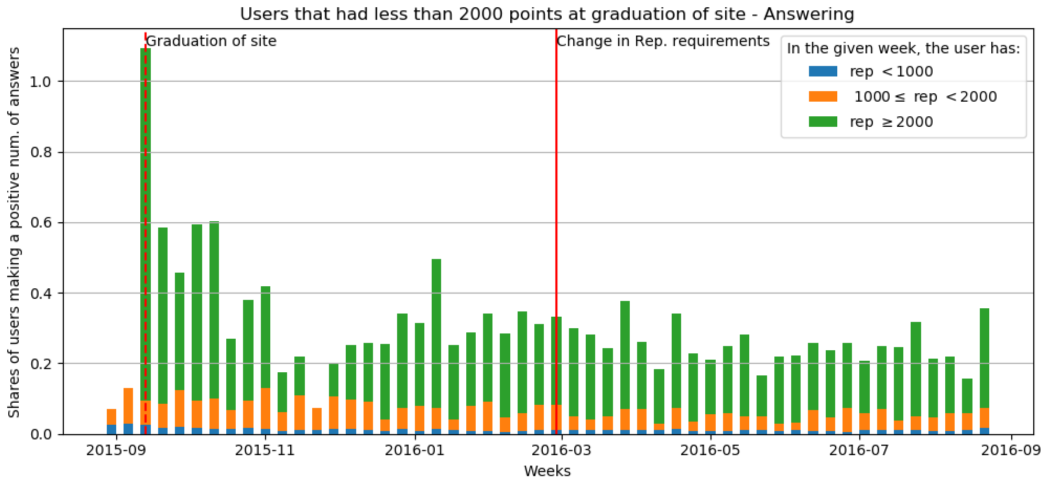

Figure 9: Number of users making a positive number of answers, out of the ones having less than 2000 points at the graduation week.

I estimate the following specification, where the running variable is the number of reputation points[24]

$$Y_{it} = \alpha_i + \gamma_t + \beta_{\bar{R}-r_{it}} + \delta_0 K_{it} + \delta_1 A_{it} + a_m + b_c + \varepsilon_{it}$$

where Y and K are either number of edits or number of comments, based on which is the outcome of interest, while A is number of answers made. α_i identifies the user fixed effect, γ_t the week fixed effect, and $\beta_{\bar{R}-r}$ identifies the fixed effect of being $r - \bar{R}$-distant from the threshold $\bar{R}$ in terms of reputation points, that is the fixed effect of having r reputation points. In the plot this will be the variable represented on the x-axis, where 0 identifies the 50-points reputation interval when the user was allocated control, 1 the next 50 points, and so on. Finally I include a dummy equal to 1 if the user is an elected moderator in time t, and a dummy equal to 1 if the user is a candidate in a moderators' election in time t (a_m and b_c) respectively. ε_{it} is an error term.

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

The results are shown in figure 10. It is possible to see that in the neighbors of the threshold, edits increase significantly, while the same clear pattern does not occur for the comments.

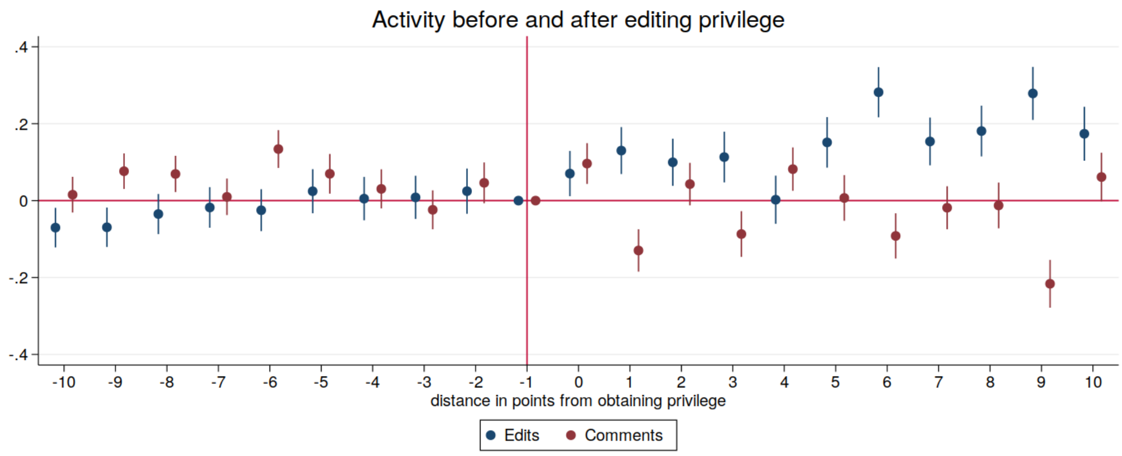

Figure 10: Estimates for the fixed effects of being in the nth reputation point interval above or below the threshold.

In section A.4 of the appendix I provide some robustness checks.

Heterogeneity.
Are these effects different across the different types of users? Figure 11 reports the estimates for the editing activity for the different types of users separately. It seems to suggest that the threshold effect, i.e. just on the neighborhood of the treatment, is driven by the Type 2 users, as for the others the effect vanishes. Type 3 users seem to have a long term effect. Figure 12 shows instead the behavioral response in commenting, where clear patterns do not emerge.

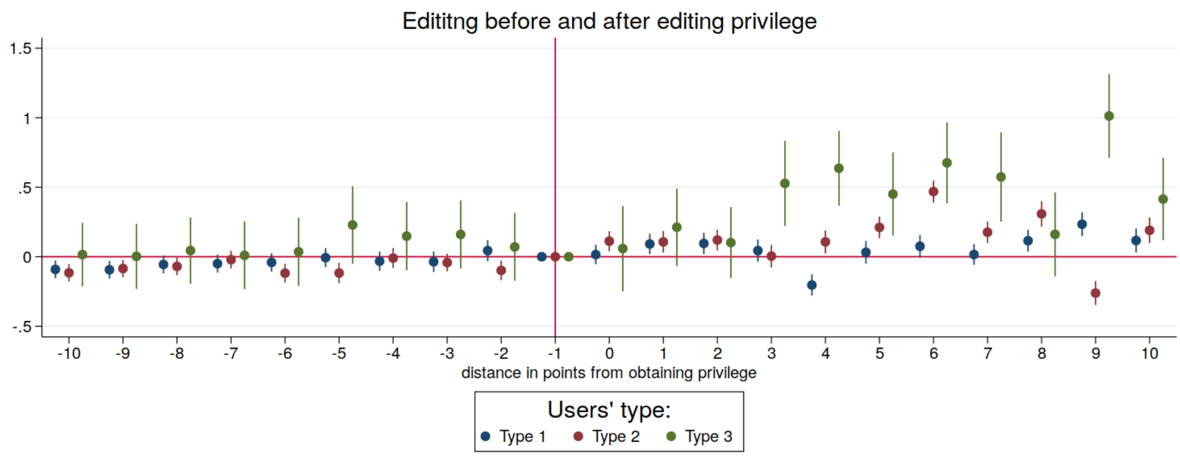

Figure 11: Estimates for the fixed effects of being in the nth reputation point interval above or below the threshold. Different estimates for each type of user.

6.2 Role of edits for website success

Incentivize editing has a relevant role for the website success, as moderation has a key role in content quality.

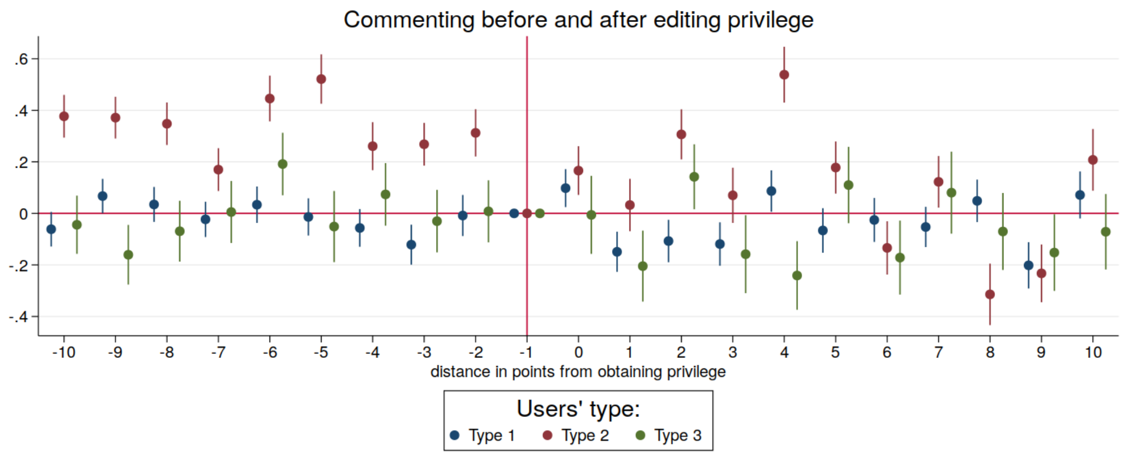

Figure 12: Estimates for the fixed effects of being in the nth reputation point interval above or below the threshold. Different estimates for each type of user.

I estimate the effect of additional edits on a post by running the following Poisson regression model:

$$\begin{aligned} \log(\lambda_i t) =& \alpha_j + \alpha_t + \sum_{k=1}^{6} \beta_k \mathbf{1}(NE_{jt} = k) \\ &\text{with} \quad \text{Num. Up-votes}_{jt} \sim \mathbb{P}(\lambda_i t) \end{aligned}$$

where j indexes published answers, and t the nth day after publication of post j. NE_{jt} refers to the number of times the post j has been edited in the t period after publication. In other words, NE_{jt} is equal to 0 in all periods before the answer receives the first edit, then it becomes 1, and remains equal to 1 for all periods up to the arrival of the second edit, and so on. α_j and α_t are post and period fixed effects.

Results are shown in table 6. When the post has received more edits, in average receive more up-votes.

	(1) numUpvotes	(2) numUpvotes
1 edit	0.573***	1.087***
	(16.11)	(50.47)
2 edits	1.131***	1.562***
	(13.00)	(34.97)
3 edits	1.525***	2.203***
	(6.52)	(24.56)
4 edits	2.520***	2.481***
	(7.43)	(14.29)
5 edits	1.917*	2.136***
	(2.21)	(5.12)
6 edits	-20.86***	-11.07
	(-20.96)	(-0.01)
constant	0.663***	
	(26.49)	
N	2757289	2757286
Errors clustered at author level	YES	NO
Post FE	NO	YES
Periods from creation FE	YES	YES

t statistics in parentheses

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

Table 6: Estimation for effect of edits on number of up-votes. Covariates are dummy variables equal to 1 if the given post at the given day has been edited # times. Specification 1 does not include posts that never received up-votes, not edits. Specifications 2, 3, and 4 exclude post that never received any up-vote. The estimation considers only the first 30 days after the publication of each given answer.

7 Dynamic Discrete Choice Model

The reduced form evidence is based on the assumption that certain behavior should follow from certain preferences. Anyway, a different approach, still based on revealed preferences, allows to relax this assumption. If different choices induce different outcomes, then a given choice reveals that an outcome is preferred from another. Based on this logic, we could measure if users value control by observing whether they undertake specific choices that will let them obtain control. In other words, when users start their participation, they have expectations on the benefit they will receive form their activity on the website, and they will behave such to maximize the expected utility.

The identification of the marginal utility of control could then be obtained by looking at what choices of contribution users make when they have an amount of points somehow close, but lower than the threshold. Certain actions will give to the user more points, in expectations, than others, inducing a faster arrival to the threshold. If participants value control and they care of obtaining the specific privilege that provides it, they will expect a higher discounted utility of their participation by choosing today actions that will allow to reach the threshold faster.

This identification is strictly based on the anticipation of future utility that an action will induce, compared to other actions, and for this reason is crucial to adopt a dynamic framework. At each point in time, users update their beliefs on the ability of a choice to make them reach the threshold, and take decisions accordingly.

7.1 Model setup

At each week[25], the user decides whether to participate in the online community and, if she does, she decides effort levels in two tasks, answering and editing. Effort is defined as a combination of quantity and quality of answers, and quantity of edits. An action choice in period t is then a vector:

$$\boldsymbol{\alpha_t} = \begin{bmatrix} \mathcal{NA}_t \\ \mathcal{QA}_t \\ \mathcal{NE}_t \end{bmatrix} \ni \mathcal{A}$$

where $\mathcal{NA}$ identifies the choice of quantity of answers, $\mathcal{QA}$ the choice of average quality of answers, and $\mathcal{NE}$ the choice of quantity of edits. $\mathcal{A}$ represents the choice set, including all possible combinations of effort levels in the two tasks.

Choices affect the utility in two ways: from one side the user pays the cost of effort, net of all the benefits that the actions provide in the given period. On the other, choices affect the transition and future realizations of states.
The net cost of effort is specific to the action made. the net cost of answering, for a user i in period t, is defined as:

$$CA_{it} \equiv \mathcal{QA}_{it} + \mathcal{NA}_{it}^{scarsity_{it}}$$

where $\mathcal{QA}$ is the average quality across the answers published by user i in period t, $\mathcal{NA}$ is the number of answers published in t, and *scarcity* is a variable that measure how few questions are still not answered, in the topic of expertise of user i. Appendix A.2 provides details on the construction of the variable scarcity. To note is that it is included in $[1, \infty]$, so that when there are many questions to answer, the cost tend to be linear, while with less questions available the cost becomes more and more convex in the quantity of answers.
The net cost of editing is instead the number of edits that the user decided to make:

$$CE_{it} \equiv \mathcal{NE}_{it}$$

The states that affect the utility are the accumulated reputation points (R), a dummy equal to 1 if the individual has control over editing and 0 otherwise ($Authority$) and a variable with the cumulative number of privileges obtained by the user ($cumT$), so to control for the possibility that users do not value authority per se, but rather value virtual vertical rewards (either perceiving them as virtual promotions, or as sequential targets in a game).

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

$$U_{it} = \beta_0 R_{it} + \beta_1 CA_{it} + \beta_2 CE_{it} + \beta_3 Authority_{it} + \beta_4 cumT_{it}$$

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

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

Timing of a period

As represented in figure 13, the timing is the following: the agent observes the values of the states realized at the end of the previous period, that is the total number of reputation points she has obtained, how many privileges she collected, and whether she has obtained control over editing or not. She then forms beliefs over the value of the states that may realize in the next periods, conditional to the possible choices she could make. She makes an effort decision over two tasks, and the flow payoff realizes. Finally, conditional on the choice made, the new value of the states realize.

[25] I will use the terms *period* and *week* interchangeably. 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 to human capital investment decisions. Examples of this literature are Arcidiacono et al. (2016), De Groote (2019).

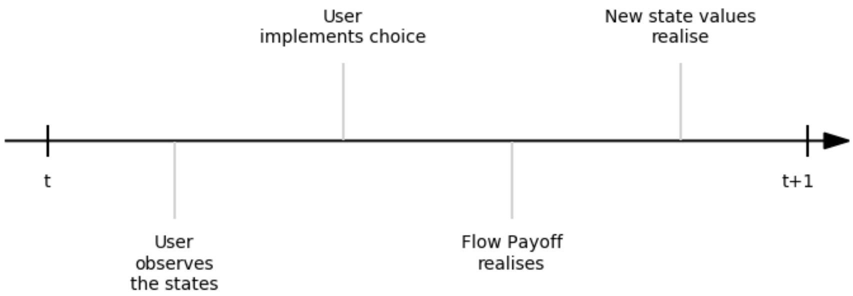

Figure 13: Timing of a period

7.1.1 Beliefs

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

Evolution of reputation points

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

Consider for simplicity the beliefs that the user form in the first period of participation t_0. He/she consider to choose a triplet $\{\mathcal{NA}, \mathcal{QA}, \mathcal{NE}\}$ of, respectively, number of answers, average quality of answers, and number of edits.

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

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

The expected values of these random variables, are given by:

$$\begin{aligned}
\lambda_{E,j,t_0} &= \exp\left(\beta_0 + \beta_1 \mathcal{QA}_{t_0} + \boldsymbol{EXP}_{t_0}\boldsymbol{\beta_2}\right) \\
\lambda_{U,j,t_0} &= \exp\left(\gamma_0 + \gamma_1 \mathcal{QA}_{t_0} + \gamma_2 \lambda_{E,j,t_0} + \boldsymbol{EXP}_{t_0}\boldsymbol{\gamma_3}\right) \\
\lambda_{D,j,t_0} &= \exp\left(\delta_0 + \delta_1 \mathcal{QA}_{t_0} + \delta_2 \lambda_{E,j,t_0} + \boldsymbol{EXP}_{t_0}\boldsymbol{\gamma_3}\right)
\end{aligned}$$

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

If the user, in its first period of participation, published $\mathcal{NA}_{t_0}$ answers, then he will expect to receive by the end of the period:

$$\Lambda_{U,t_0} = \mathcal{NA}_{t_0} \times \lambda_{U,j,t_0}$$
$$\Lambda_{D,t_0} = \mathcal{NA}_{t_0} \times \lambda_{D,j,t_0}$$

Which are, respectively, the total expected amount of up-votes and down-votes.

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

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

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

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

The answers produced in period t_0 may as well induce arrival of up-votes and down-votes in the next periods. This is model deterministically. Let Δt be the number of days passed from the publication day, such that if $t = t_0 + 1$, then $\Delta t = 1$. Then:

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

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

Given these assumptions, it results that a given effort decision made in one period, produces decreasing returns in the following ones. As a consequence, at a given point in time, the total amount of points that the user expects to receive depends on the aggregation of all these processes rising from each period's contributions. As a general rule, we then have that:

$$\Lambda_{U,t} = \Lambda_{U,t-1} \times \exp\left(\frac{-1}{\tau_U}\right) + \mathcal{NA}_t \times \lambda_{U,j,t}(\mathcal{QA}_t)$$
$$\Lambda_{D,t} = \Lambda_{D,t-1} \times \exp\left(\frac{-1}{\tau_D}\right) + \mathcal{NA}_t \times \lambda_{D,j,t}(\mathcal{QA}_t)$$

and

$$\mathbb{E}[\rho_t|\boldsymbol{\alpha}_t] = 10 \times \Lambda_{U,t} - 2 \times \Lambda_{D,t} + 2 \times \pi \times \mathcal{NE}_t.$$

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

$$\mathbb{E}[R_{t+1}|R_t, \boldsymbol{\alpha}_t] = R_t + \mathbb{E}[\rho_t|\boldsymbol{\alpha}_t]$$

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

As shown in the graphs in the appendix A.2, the availability of questions to answer increase monotonically over time. Users then expect a steady increase, given by an estimated parameter.

7.2 Application of the model

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

$\mathcal{NA}$	$\mathcal{QA}$	$\mathcal{NE}$
0.0	0.00	0.0
0.0	0.00	1.0
0.0	0.00	4.0
1.0	13.33	0.0
1.0	13.33	1.0
1.0	13.33	4.0
1.0	14.12	0.0
1.0	14.12	1.0
1.0	14.12	4.0
1.0	15.97	0.0
1.0	15.97	1.0
1.0	15.97	4.0
7.0	13.33	0.0
7.0	13.33	1.0
7.0	13.33	4.0
7.0	14.12	0.0
7.0	14.12	1.0
7.0	14.12	4.0
7.0	15.97	0.0
7.0	15.97	1.0
7.0	15.97	4.0

Table 7: Possible effort levels that users are allowed to choose in estimation.

levels of quality. Quantity of edits can take two possible levels. All options in the choice set are listed in the table 7.
The value of the possible levels are obtained by looking at the distribution of actions taken in the data by individuals at each week of participation. For what concerns the quantity of answers, I split the distribution at the 70^{th} quantile, corresponding to three answers, so to categorize effort between low (1 to 3 answers) and high (4 or more). I then select, as possible option for the user, the median values of these two categories, so either 1 or 7 answers. A similar process is made for quality and edits. The distribution of quality is split in three categories at the 33^{th} and 66^{th} quantiles. The median values are 13.33, 14.12, and 15.97. Finally, the distribution of number of edits is split at the 75^{th} quantile, leading to two categories: low effort, which includes 1 or 2 edits, and high effort, including 3 or more edits. The distribution of values within each category is plotted in figure 31 in the appendix.

State space. The state variables, affecting directly of indirectly the utility function, are let in their original values, i.e. I do not categorize and bin the state space. I preserve computational feasibility of the estimation by not computing ex-ante the transition probabilities[26].

CCPs. Conditional choice probabilities are computed before estimation via a static logit[27] Before estimation, the data is scaled so that each variable would be in the range $(0, 1)$. The scaling algorithm subtracts the minimum and divide by the difference between the maximum and the minimum.

7.3 Estimation

I split the estimation algorithms in two steps[28]. First, I compute, for each initial set of state values observed in the data , the expected value of the variables entering the utility function. The expected values are computed for all future periods necessary such that finite dependence holds, and under each possible choice that the user may take. This means that, with a sample of n observations, and 21 choices, I obtain 21 matrices $n \times m$, where m is the number of variables entering the flow utility.
In the second step, I compute and maximize the log-likelihood function. The derivation of the lihelihood

[26] Pre-computing transition probabilities, as, for instance, is done in Rust (1987), would require the computation of a $n \times n$ matrix, where n is the number of all possible combinations of state values. Including several variables which potentially can take many values, the size of this matrix quickly would not be computationally tractable

[27] Logistic regression in Scikit-learn (Pedregosa et al. 2011) with *saga* solver.

[28] Scott (2013) uses the same approach, but then does not employ maximum likelihood.

function is reported in appendix A.5. Since beliefs don't need to be computed during the maximization, the algorithm estimates the parameters in a significantly shorter time. I use the *BFGS* minimization algorithm, which numerically looks for the minimum by approximating the jacobian and hessian matrices of the likelihood function.

8 Results

8.1 First stage estimate of reduced form parameters

When users decide what action to take, they form beliefs on the arrival of points next period. For a given amount of answer and a given quality, they first predict the number of edits that they would receive in the publication date, and then the number of up-votes and down-votes that their content can receive next period.
For what concerns the arrival of edits on users' new content, the different specification estimates are reported in table 8 (the specification used in the structural model is the number 2). It is possible to notice that experience is negatively correlated with the arrival of edits. This means that indeed experience captures some skills that the quality variable is not able to measure: even controlling for quality, experienced users produce content that need to be corrected less.

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

t statistics in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

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

Once users have expectations of the number of edits they will receive, they predict how many up-votes and down-votes their content will receive in expectations. Estimates are reported in table 9 and 10 respectively. As expected, higher quality is correlated with more up-votes and less down-votes. Received edits have instead a positive coefficient in both cases. This could be explained because edits improve the quality, inducing more up-votes, but, at the same time, users may want to penalize content that was originally bad. Finally, more experienced users can expect more up-votes and less down-votes.

Once beliefs on the number of up-votes and down-votes arriving next period are formed, users expect how much this same content will produce points in the periods following the next. The returns in up-votes and down-votes are modeled as deterministically decreasing as shown in figures 14 and 15. The adopted model is the exponential, which was estimated via a non-linear fit. Table 11 reports the estimates.
In the figures, the red dots are the data values for λ_U and λ_D, and on the x-axis is reported Δt. The figures also compare the chosen model (exponential) with alternative ones, and present the decay both at daily and weekly level

Finally, the last parameter estimated in the first step is the rate of increase of availability of questions,

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

t statistics in parentheses

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

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

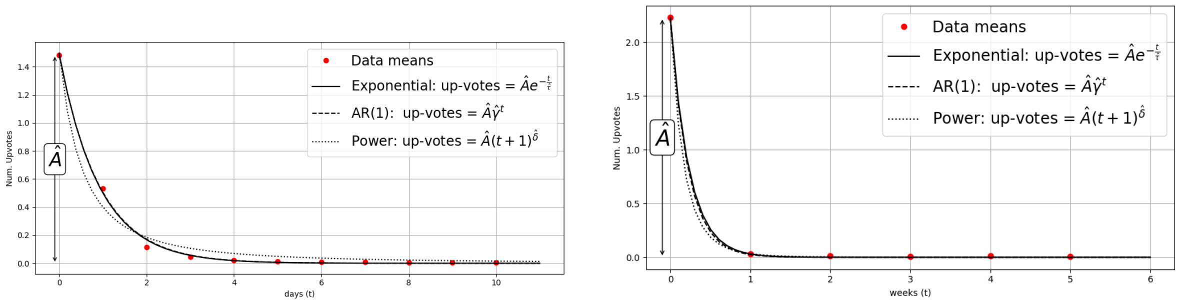

Figure 14: Return in up-votes from published content

in time. Since the community is increasing, and some questions remain unanswered, then users expect a steady increase in availability. Estimates are reported in table 12.

8.2 Flow payoff parameters

The flow payoff parameters obtained from the dynamic discrete choice model are presented in table 13. All type of users receive a positive marginal utility from collecting points. The marginal value of effort instead, that would be expected to be negative (we would be expect to be a marginal cost), takes instead different signs across types. The marginal cost of effort, in fact, is not well identified, and suffers from omitted variable bias. From not including in the specification all possible benefits that incentivise users to participate, the marginal cost of effort may be up-ward biased, since in the data the action may be taken even when it does not specifically increase the value of some other motivating device. A positive marginal utility of effort would be explained by the existence of omitted factors that drive the action. The action would be then taken even when, according to the adopted utility function, there is no specific reason to do so. A possible omitted factor in this case is the achievement of badges: as described in section 5, users of Type 3 gained a significant higher amount of badges than the other types, and they may edit answers to achieve those badges.

Table 14 reports estimates for a slight different specification of the flow utility, where interactions terms are included. In particular, the dummy variable capturing value of control is interacted with the cost of effort in answering and in editing. This different specification wants to test whether control may

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

t statistics in parentheses

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

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

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

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

period length	Number of additional available question next period
day	13.68
week	95.76

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

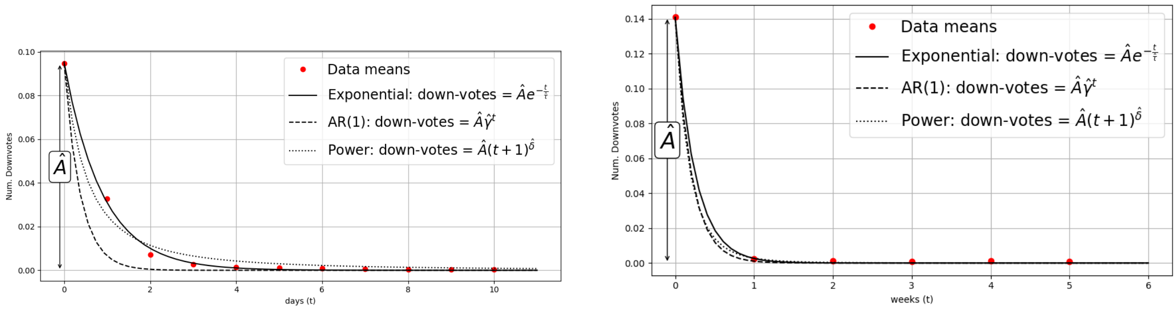

Figure 15: Return in down-votes from published content

Variables	(all sample)	(Type 1)	(Type 2)	(Type 3)
R	0.0074	0.0085	0.0051	0.0083
	(0.0001)	(0.0002)	(0.0004)	(0.0003)
CA	0.0004	0.0002	-0.064	0.0007
	(0.0002)	(0.0004)	(0.0166)	(0.0002)
CE	-0.6134	-3.8934	-9.2085	5.0916
	(0.1733)	(0.065)	(0.101)	(0.1476)
Tcum	-0.8409	-0.7986	-0.7548	-0.7899
	(0.0209)	(0.0289)	(0.0165)	(0.0517)
isEditor	1.2052	1.2239	1.2516	0.86
	(0.0593)	(0.0628)	(0.1649)	(0.1476)
N. users	9,783	3,700	5,407	676
Sample size	991,657	471,837	407,098	112,722

Table 13: Estimates for the flow payoff parameters considering the whole sample, or estimating separately for each type of user. Standard errors in parenthesis are computed from the numerically approximated hessian matrix.

not be valued only per se, but induce specific utility (or dis-utility) of given actions. This is specifically relevant when carrying the normative analysis, as more or less willingness to take action affects final amount of production. Estimates show that for all users the marginal value of effort increases, suggesting that the allocation of control induces the activation of other unobservables that drive more participation. This happens more for editing than for answering, and it is specifically marked for user type 2, that has in general a higher cost of participation.

Variables	(all sample)	(Type 1)	(Type 2)	(Type 3)
R	0.0069	0.0084	0.0067	0.0076
	(0.0002)	(0.0003)	(0.0004)	(0.0003)
CA	-0.0001	-0.0569	-0.2788	0.0006
	(0.0007)	(0.0219)	(0.0201)	(0.0003)
CE	-10.3311	-9.2302	-15.6076	-4.6256
	(0.0084)	(0.769)	(0.2748)	(0.035)
Tcum	-0.7745	-0.7028	-0.6145	-0.6637
	(0.0035)	(0.1609)	(0.0309)	(0.0271)
isEditor	1.3162	1.3106	1.3905	1.0767
	(0.0054)	(0.2689)	(0.0428)	(0.0661)
CA x isEditor	0.0609	0.1104	0.3687	0.0458
	(0.0043)	(0.0117)	(0.0214)	(0.0045)
CE x isEditor	12.2064	7.2181	11.3003	10.7403
	(0.0072)	(0.3908)	(0.2384)	(0.0751)
N. users	9,783	3,700	5,407	676
Sample size	991,657	471,837	407,098	112,722

Table 14: Estimates for the flow payoff parameters for the whole sample or by type of user. The specification includes interaction terms of the costs of actions with the control dummy. Standard errors in parenthesis, computed from the numerically approximated hessian matrix.

9 Counterfactual Analysis: Incentive effect of delegation

9.1 No anticipation of delegation

How total production would change if users would not be anticipating the allocation of control? In this counterfactual analysis users are still forward looking, but do not know that they will receive control with the achievement of a given amount of points.
Intuitively, if the delegation of control is valued and users adjust their participation to achieve it, the lack of anticipation completely erase its incentive effect.
Figure 16 reports the simulated amount of answers published in the platform, identifying the contribution of each type of users. The main result is that, under no anticipation of delegation, all the contribution made by users of type 2 would be missing. The reason of this effect is that users of type 2 are the ones with the highest cost of effort, but value the allocation of authority more than the others. Their participation is strongly incentivized by the expectation of receiving control over editing. Not expecting to receive any authority, they just loose all incentives to participate, and do not answer questions anymore. This effect lead to the fact that in average type 2 users will never reach the threshold, under the counterfactual scenario. Since the task of editing is mainly valued conditional on having full autonomy over it, type 2 users will also not produce any edit. This is graphically showed in figure 17. The trade-off from the platform perspective is anyway mainly based on the answering side, as anyway the contribution in edits of type 2 users is low in any case.
For the users of type 1 and 3 instead, even if they value authority, the incentive effect induced on their answering activity is limited, suggesting that are other factors the main drivers of their contribution. They still produce answers and reach the threshold. Once they obtain control, they increase their editing activity. In the simulation, none of the users produce edits before receiving control.

9.2 No delegation

If the company decides to not delegate at all any control, then it gets intuitively two outcomes. As for the no-anticipation counterfactual, users are not incentivized in production by the allocation of control. In particular type 2 users will be the most affected and will not answer questions. As shown in figure 18, the simulated lost in production is similar to the scenario with delegation but no anticipation. For what concerns edits instead, the obtaining of control is crucial to have individual participate, as otherwise their cost of effort is to high with respect to the benefits they receive. It follows that in the scenario without delegation, users will not produce edits, as shown in in figure 19.

On the choice of answers' quality, the simulated counterfactual show a slight decrease in quality, in particular for type 2 users. Figure 20 shows the proportion of answers published in each of the possible

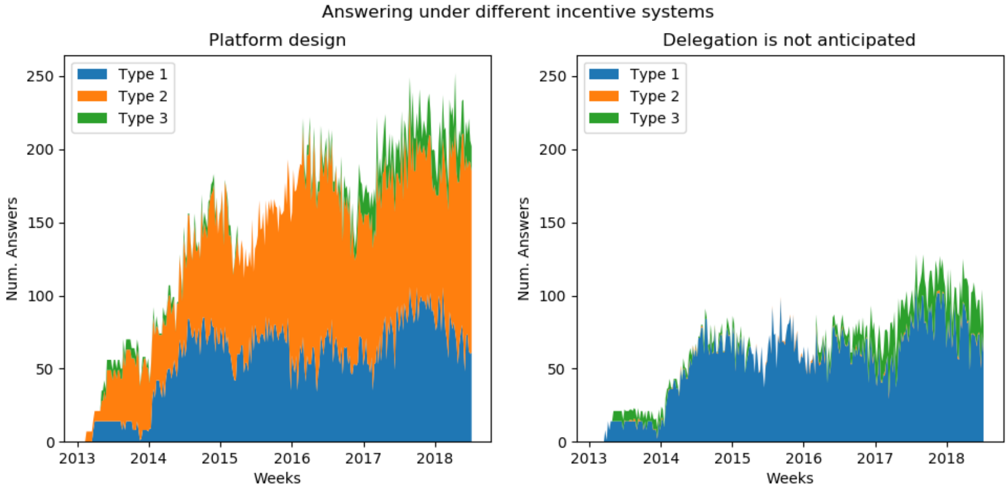

Figure 16: Simulated production under anticipated and not anticipated delegation. It is plotted the total number of Answer produced at each point in time, with different colors reporting production of a given type of users. The proportion of users for each type correspond to the distribution observed in the real data. Production made by 100 users with random starting date of participation, and 100 weeks of participation each. 38 users are of type 1, 55 of type 2, and 7 of type 3

quality levels, by each type in the different scenarios. It is possible to see that the share of answers of high quality shrinks particularly for type 2, but also type 3 users tend to partially substitute some high quality answers wih medium quality answers.

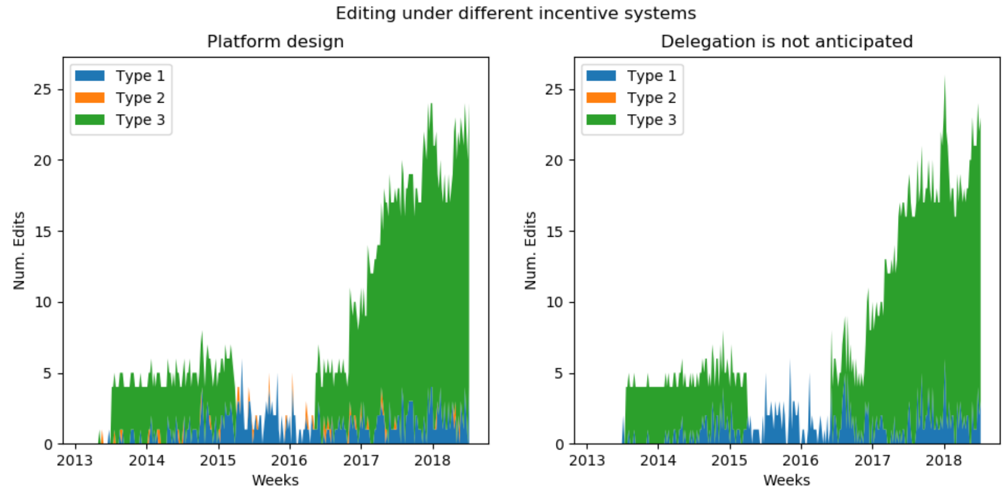

Figure 17: Simulated production under anticipated and not anticipated delegation. It is plotted the total number of Edits produced at each point in time, with different colors reporting production of a given type of users. The proportion of users for each type correspond to the distribution observed in the real data. Production made by 100 users with random starting date of participation, and 100 weeks of participation each. 38 users are of type 1, 55 of type 2, and 7 of type 3

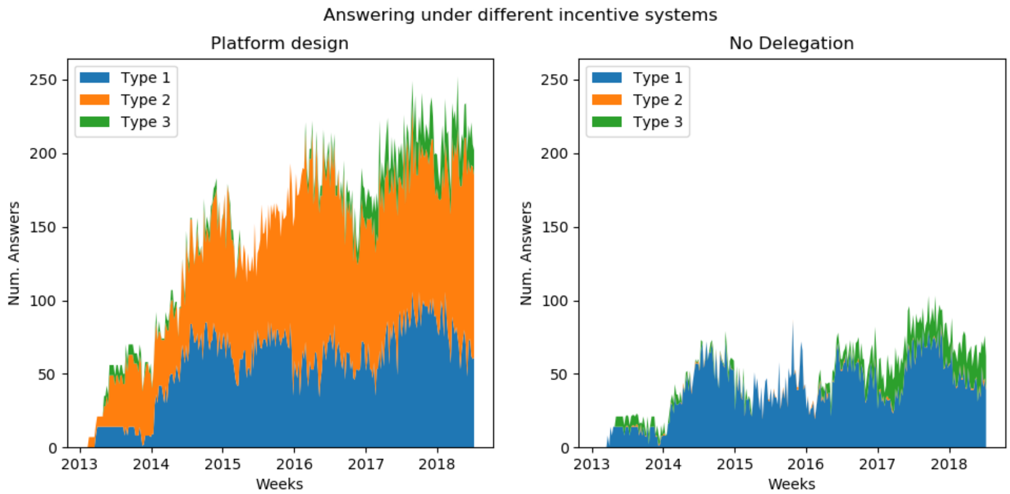

Figure 18: Simulated production under anticipated and not anticipated delegation. It is plotted the total number of Answer produced at each point in time, with different colors reporting production of a given type of users. The proportion of users for each type correspond to the distribution observed in the real data. Production made by 100 users with random starting date of participation, and 100 weeks of participation each. 38 users are of type 1, 55 of type 2, and 7 of type 3

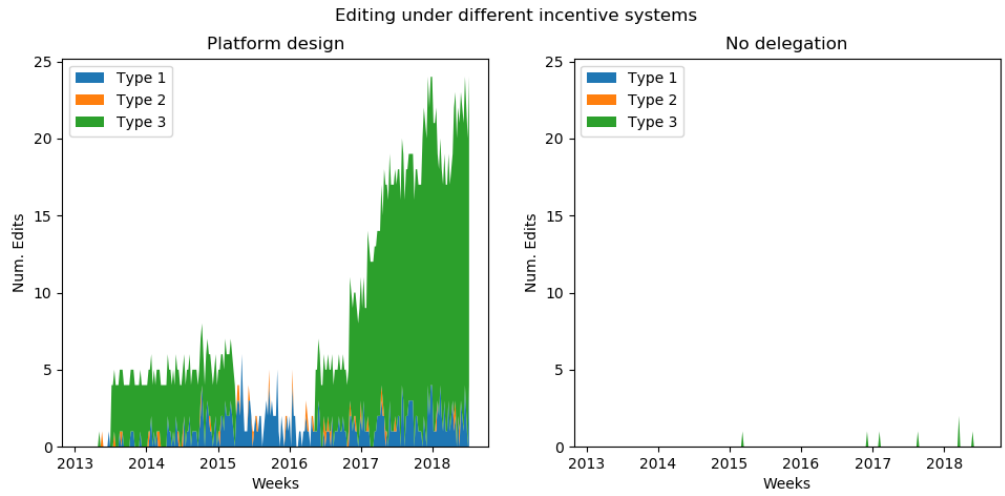

Figure 19: Simulated production under anticipated and not anticipated delegation. It is plotted the total number of Edits produced at each point in time, with different colors reporting production of a given type of users. The proportion of users for each type correspond to the distribution observed in the real data. Production made by 100 users with random starting date of participation, and 100 weeks of participation each. 38 users are of type 1, 55 of type 2, and 7 of type 3

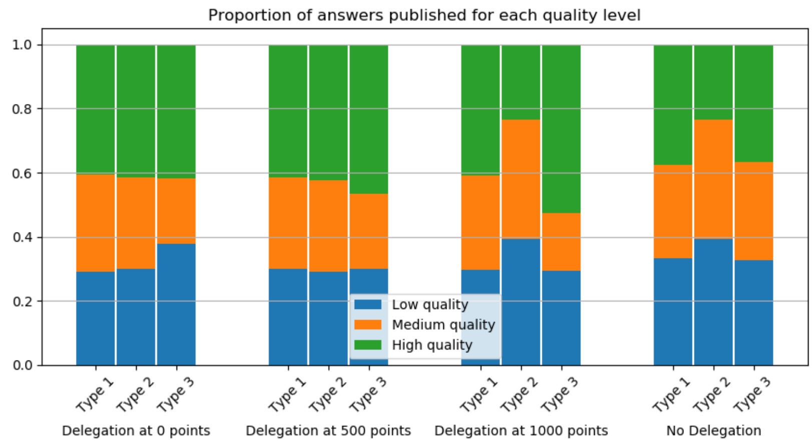

Figure 20: Share of answers published by each type, in each quality level. Separate set of bars identify different simulated scenarios.

10 Conclusion

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

I find that users increase their contribution when allocated more autonomy. Results show as well that individuals have a positive marginal utility from obtaining authority, but the long term effects on contribution is substantially heterogeneous across types of users. With counterfactual simulations, I show that the company, by not committing to a delegation system, would lose the participation of one type of users, inducing an important drop in total production.

References

Aghion, P., M. Dewatripont, and P. Rey (2004, March). Transferable control. *Journal of the European Economic Association 2*(1), 115–138. 4

Aghion, P. and J. Tirole (1997, February). Formal and real authority in organizations. *The Journal of Political Economy 105*(1), 1–29. 4

Arcidiacono, P., E. Aucejo, A. Maurel, and T. Ransom (2016). College attrition and the dynamics of information revelation. *Working Paper*. 21

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. 40, 41, 42

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. 4

Bandiera, O., M. C. Best, A. Q. Khan, and A. Prat (2020). The allocation of authority in organizations: A field experiment with bureaucrats. *Working paper*. 4

Bartling, B., E. Fehr, and H. Herz (2014, November). The instrinsic value of decision rights. *Econometrica 82*(6), 2005–2039. 4

Belenzon, S. and M. Schankerman (2015). Motivation and sorting of human capital in open innovation. *Strategic Management Journal 36*, 795–820. 5

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

Besley, T. and M. Ghatak (2008, May). Status incentives. *The American Economic Review: Papers & Proceedings 98*(2), 206–211. 4

Bester, H. (2009, November). Externalities, communication and the allocation of decision rights. *Economic Theory 41*(2), 269–296. 4

Bester, H. and D. Krähmer (2008, Autumn). Delegation and incentives. *RAND Journal of Economics 39*(3), 664–682. 4

Blanes I Vidal, J. and M. Möller (2007, Summer). When should leaders share information with their subordinates? *Journal of Economics & Management Strategy 16*(2), 251–283. 4

Chen, W., X. Wei, and K. Zhu (2017). Engaging voluntary contributions in online communities: A hidden markov model. *MIS Quarterly 42*(1), 83–100. 5

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). 5

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). 5

De Groote, O. (2019). A dynamic model of effort choice in high school. *Working Paper*. 21

Dessein, W. (2002, October). Authority and communication in organizations. *Review of Economic Studies 69*(4), 811–838. 4

Fairburn, J. A. and J. M. Malcomson (2001, January). Performance, promotion, and the peter principle. *The Review of Economic Studies 68*(1), 45–66. 4

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. 4

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. 4, 8

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. 4

Goes, P., C. Guo, and M. Lin (2016). Do incentive hierarchies induce user effort? evidence from an online knowledge exchange. *Information Systems Research*. 5

Greenacre, M. and J. Blasius (Eds.) (2006). *Multiple Correspondence Analysis and Related Methods*. Statistics in the Social and Behavioral Sciences Series. Chapman & Hall/CRC. 37

Hagberg, A. A., D. A. Schult, and P. J. Swart (2008, 08). Exploring network structure, dynamics, and function using networkx. In G. Varoquaux, T. Vaught, and J. Millman (Eds.), *Proceedings of the 7th Python in Science Conference (SciPy2008)*, Pasadena, CA USA, pp. 11–15. 47

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. 41

Hunter, J. D. (2007). Matplotlib: A 2d graphics environment. *Computing in Science & Engineering 9*(3), 90–95. 47

Jain, S., Y. Chen, and D. C. Parkes (2014). Designing incentives for online question-and-answer forums. *Games and Economic Behavior 86*. 5

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. 4

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. 5

Lê, S., J. Josse, and F. Husson (2008). Factominer: An r package for multivariate analysis. *Journal of Statistical Software, Articles 25*(1), 1–18. 47

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. 4

McKinney, W. (2010). Data structures for statistical computing in python. In S. van der Walt and J. Millman (Eds.), *Proceedings of the 9th Python in Science Conference*, pp. 56–61. 47

Nov, O. (2007, November). What motivates wikipedians? *Communications of the ACM 50*(11), 60–64. 4

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. 4

Pedregosa, F., G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research 12*, 2825–2830. 24, 47

Pikulina, E. S. and C. Tergiman (2020, May). Preferences for power. *Journal of Public Economics 185*(104173). 4

Rajan, R. G. and L. Zingales (1998, May). Power in a theory of the firm. *The Quarterly Journal of Economics 113*(2), 387–432. 4

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. 4, 5

Rust, J. (1987, September). Optimal replacement of gmc bus engines: an empirical model of harold zurcher. *Econometrica 55*(5), 999–1033. 24

Scott, P. T. (2013). Dynamic discrete choice estimation of agricultural land use. *Working Paper*. 24

Seabold, S. and J. Perktold (2010). Statsmodels: Econometric and statistical modeling with python. In *9th Python in Science Conference*. 47

Sturm, R. E. and J. Antonakis (2015, January). Interpersonal power: A review, critique, and research agenda. *Journal of Management 41*(1), 136–163. 3, 4, 16

Sun, L. and S. Abraham (2020). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Working Paper*. 17

Virtanen, P., R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. Jarrod Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. Carey, İ. Polat, Y. Feng, E. W. Moore, J. Vand erPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and S. . . Contributors (2020). SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods 17*, 261–272. 47

Xu, L., T. Nian, and L. Cabral (2020). What makes geeks tick? a study of stack overflow careers. *Management Science 66*(2), 587–604. 5

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. 5

Appendix A Details and Robustness

A.1 Construction of quality variable

The variable quality captures the variation of points received by an answer at its publication day explained by text characteristics.
Let $\boldsymbol{X}_j$ be a vector of text characteristics of an answer j, right after publication, so before any modification occurs. Let $\bar{t}_j$ be the publication date of answer j. I estimated the following linear model:

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

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

The vector of text characteristics includes:

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

Table 15 reports the estimates for the linear regression models used to predict the variable quality. The specification adopted corresponds to column (5).

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

t statistics in parentheses

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

Table 15: Regressions to predict the quality variables. Model finally used is in column (5).

A.2 Construction of scarcity variable

The construction of the variable scarcity follows several steps.

- Construct the variable *availability*, given by the cumulative number of questions appearing in the platform, that, each day, don't have yet an answer selected as best answer. This is equivalent for every users. The cumulative number of questions and the number of questions without an accepted answer are plotted in figure 21.
- Recover topics from question tags[29]. To do this, I first construct a graph of tags, where a link between two tags exists if the two tags appear at least once in the same question. The intensity of the links are given by the number of times that the two tags have appeared in a same question. I then identify topics using the Page rank algorithm[30], i.e. a topic will be those tags that are connected to the most other tags. I identify 6 topics, since the Page ranking value drops suddenly after those first 6 tags. The topics are: 'grammar', 'word-usage', 'meaning', 'sentence-construction', 'meaning-in-context', 'word-choice'. I then partition the graph around these 6 tags, using a Voronoi diagram. After this process I then have 6 topics, and a mapping from every tag to each of these topics. The word-clouds of tags related to each topic are plotted in figure 22.
- I allocate topics at the questions still to answer, recovered at the first bullet point: using the tags assigned to those questions, I obtain the share of each topic in each of the questions.
- I do a similar process for each user, on all questions he/she has answered, and recover the share of each topic in which he/she is expert about
- for each user i, I weight the available questions at each period t by his/her expertise, call this variable $avail_{it}$. Figure 23 shows the distribution of time of this variable, in average across users' lifetime in the website.

[29] Questioners can add tags when posting a question

[30] This is the Google search algorithm of the early times of the search engine

The variable scarcity is then defined as:

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

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

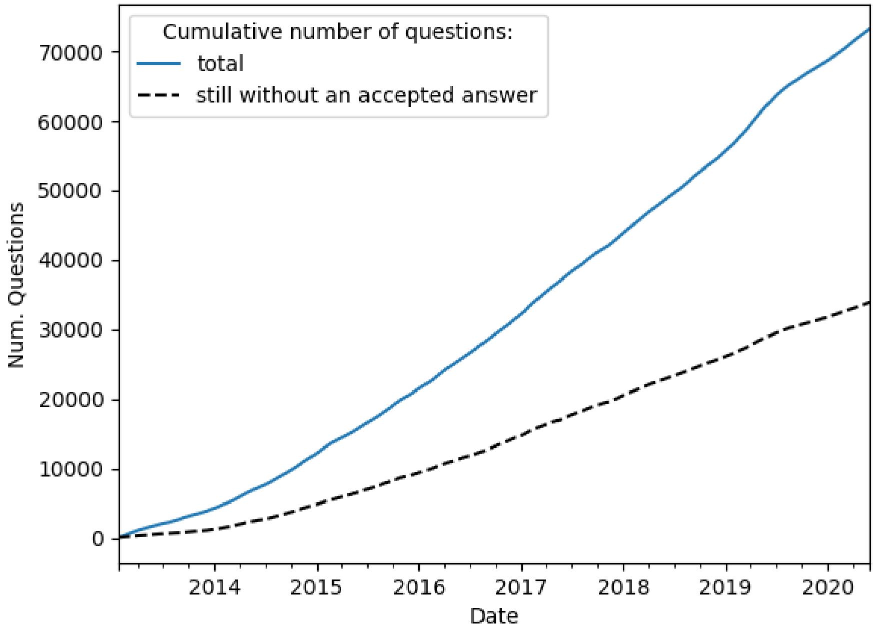

Figure 21: Cumulative number of questions in the platform, both total number and net of questions that have already selected an answer as *best answer*.

A.3 Construction of Types

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

Information aggregation. The most simple approach to reduce dimensionality would be to aggregate the variables via, for example, sum. Since the individual characteristics include dummy variables taking value 1 if the user decided to display some given information, as well as the length of the biographical description, summing over them gives a measure of the amount of information displayed. This approach turned to not be a good solution, as behavior is not linearly correlated with the amount of information. The aggregated variable is then not informative on the different types of users.
A common alternative is to perform the Principal Component Analysis (PCA). This approach transforms the data by creating orthogonal vectors, each containing the largest possible variance of the original variables. The first vector will be the most representative of the original variance, the second will be the most representative of the residual variance, and so on. PCA anyway relies on quantitative continuous variables, as it relies on the computation of the variance, and it is not suitable to dummy variables.

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

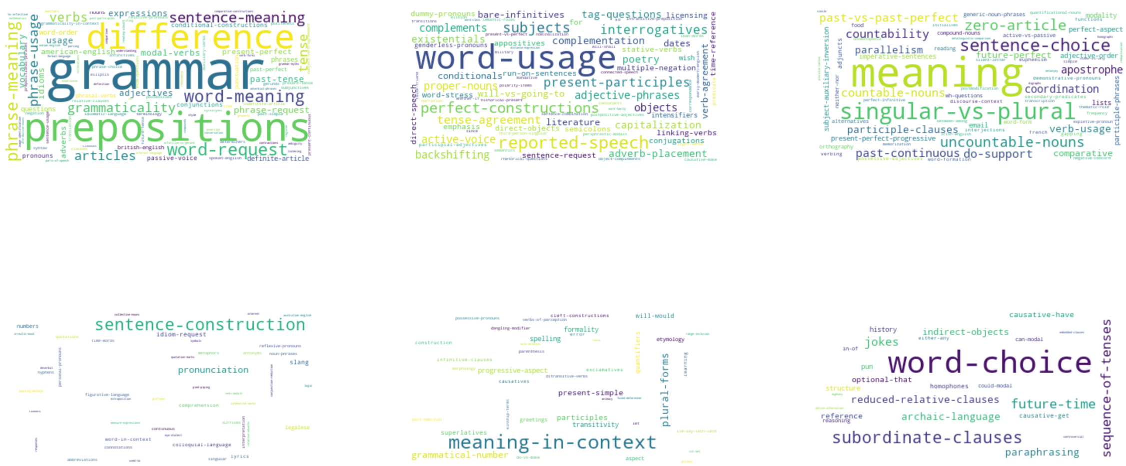

Figure 22: Word-clouds for each topic identified

original variables. Individual users can then be plotted in the reduced bi-dimensional space formed by each pair of dimensions. In the discussion that follows I will focus on the plane formed by the first and second dimensions. Note that, since this algorithm is applied to categorical variables, I bin in three groups the variables representing the length of the biographical note and the variable with the number of links appearing in the biographical note.
Figure 24 shows the variable representation in the first two dimensions space. First it is possible to notice, on the axes, that the first dimension contains about 17% of the information of the individual characteristics, while the second dimension about 8%. The location of the variables on the plain tells the extent to which that dimension include information from the given variables. It is possible to see that the length of the biographical note is the most important source of information for both dimensions, while the inclusion of location and website in the user page is only captured by the first dimension.
Figure 25 instead represents on the same dimensions the individuals, i.e. the sample of users. This graph may help to understand if individuals cluster in groups, based on the information of the first two factors. It is possible to observe that clear clusters are not emerging. Nonetheless, points are not displayed in an uniform cloud with respect to the axis. While some are grouping around the (0,0) point, meaning that they have characteristics close to the average of the sample, others appear on the positive side of the first dimension. Users appearing in the upper right quadrant are more likely to have a Linkedin profile, a website, and the location, compared to the average user, as well as longer biography with more links. Users in the bottom right quadrant are also more likely to have a website and the location, they tend to have a biography, but a short one.

Identification of groups. A typical clustering algorithm is the so called K-Means clustering. This algorithm requires the number k of groups that want to be identified, it picks k centroids (i.e. means of partitions of the observations) and updates the centroids so to minimize the within-cluster variance. This algorithm is also meant to work with continuous quantitative variables, so is not suitable to be directly applied on the original individual characteristics. I then apply the K-Means clustering procedure to the first 5 dimensions recovered after the application of the MCA procedure. These are continuous variable and still represent the information of the original data.
By choosing three clusters (i.e. $k = 3$), the resulting individual representation is shown in figure 26, with individuals colored based on the allocated cluster.

A.4 Reduced form - robustness checks

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

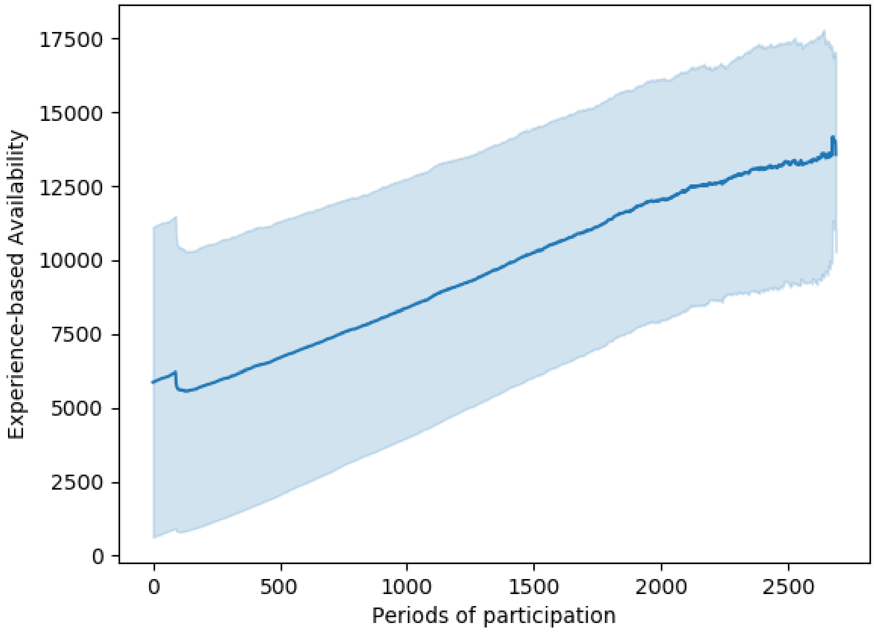

Figure 23: Average expertise-weighted availability of questions per period of participation on the website, across users. Shadow areas identify the standard deviation.

of section **??** around different thresholds.
In particular I consider the two privileges achieved just before and just after the allocation of authority. Figure 27 reports the estimates of the reputation-point intervals fixed effects, around the privilege "Established User". This privilege does not allocate any resource, and it is just a recognition. It is obtained with 750 points during the beta phase of the site, and with 1000 points during the final phase. It is possible to notice that right around the threshold it is not observed a significant increase in editing. Moving further from the threshold shows instead an increase, but pre-treatment effects seems to suggest the presence of a trend, rather than a causal effect of the treatment.
Finally, figure 28 shows again estimates for the same specification, but this time around the allocation of the "Creat Tag Synonyms" privilege. This privilege allows users to corrects tags. It is achieved either with 1250 points in the beta phase, or with 2500 points otherwise. Looking at the effects just in the neighborhood of the treatment, it is not really possible to identify a clear pattern.

A.5 Derivation of Likelihood function

Let $D \in \{1, 0\}$ be a binary variable that takes value equal to 1 when the user is given full ex-ante control over Edits. In addition, denote $\boldsymbol{d_t}$ a vector of dummy variables, $d_{\alpha t}$, for each possible choice $\boldsymbol{\alpha} \in \mathcal{A}$, such that $d_{\alpha t}$ is equal to 1 if in period t is selected choice $\boldsymbol{\alpha}$, and zero otherwise.
Choosing an action $\boldsymbol{\alpha^*}$ in period t, the one period flow utility of user i is then given by:

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

Where the vector $\boldsymbol{x_t}$ is described in section **??**.
The term $\varepsilon_{i\alpha^* t}$ is instead a choice specific utility term not measurable by the econometrician.

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

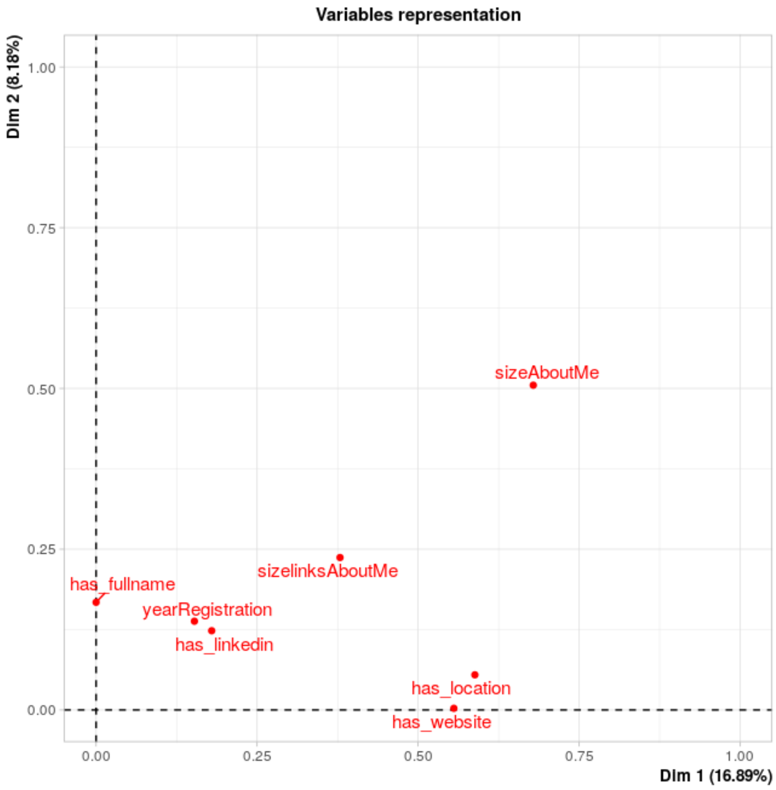

Figure 24: Variable representation on the first two dimensions plane, where the dimensions are obtained via the MCA of the individual characteristics.

A user selects a sequence of optimal decisions $\boldsymbol{d}^* \equiv \{\boldsymbol{d}_t^*\}_{t \leq T}$ that satisfies[31]:

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

where δ is a discount factor and, at each period t, the expectation is taken with respect to z_τ and ε_τ, for $\tau \geq t+1$.
In words, the agent, at each period, will choose whether to contribute in the platform and eventually what type of contribution to make, between producing content (answers), performing moderation task (edits), or both.

Identification and estimation

For the characterization of the problem I follow Arcidiacono and Miller (2011).
Define the ex-ante value function at period t as the discounted sum of the expected future payoff under optimal behavior, and before the shock ε_t is realized[32]. 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:

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

where the expectation is taken with respect to $\varepsilon_{\alpha t}$, and $f_{\alpha t}(z_{t+1}|z_t)$ is the probability that the vector of states will take a certain value in the next period, given the choice made. This transition probability

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

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

[32] 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.

Figure 25: Representations of users on the first two dimensions plane, where the dimensions are obtained via the MCA of the individual characteristics.

does not depend on all the history of past choices due to the assumptions made in the previous section. Define then the conditional value function $\nu_{\alpha t}(z_t)$ as the value function $V_t(z_t)$ for a given choice $\boldsymbol{\alpha}$ and net of the preference shock ε_t:

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

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

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

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

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

It follows that:

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

which can be rewritten as:

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

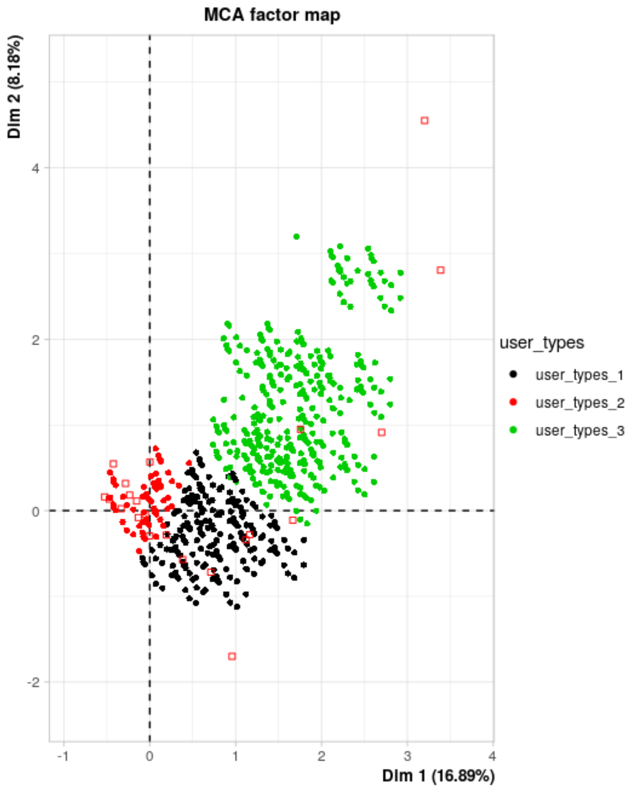

Figure 26: Representations of users on the first two dimensions plane, where the dimensions are obtained via the MCA of the individual characteristics. Colors refer to cluster groups identified with k-means clustering on the MCA dimensions.

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

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

To write the conditional value function as in 1 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 1 into the difference $\nu_{\alpha t}(z_t) - \nu_{\alpha' t}(z_t)$, all future periods after the sequence of choices will cancel out.

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

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

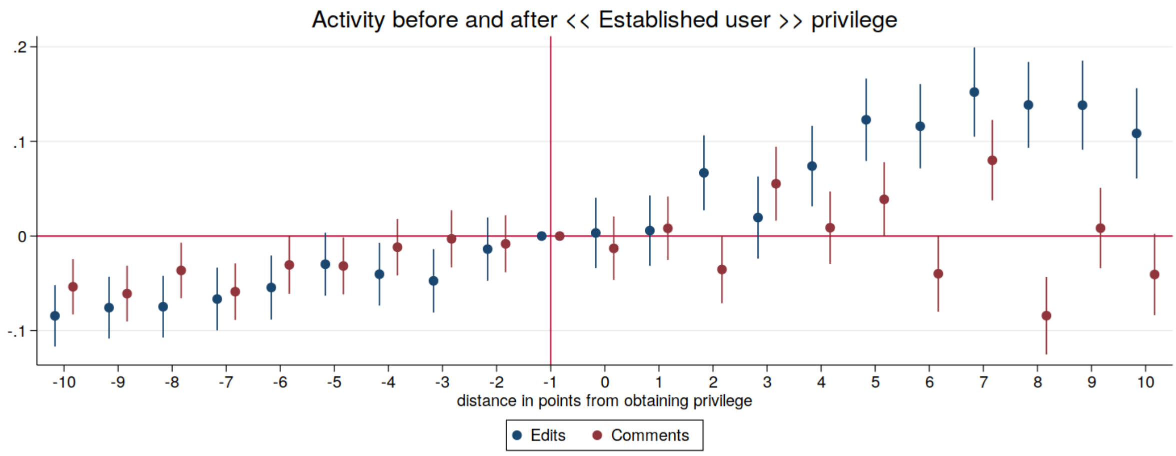

Figure 27: Estimates for reduced form effect around the *Establish User* privilege

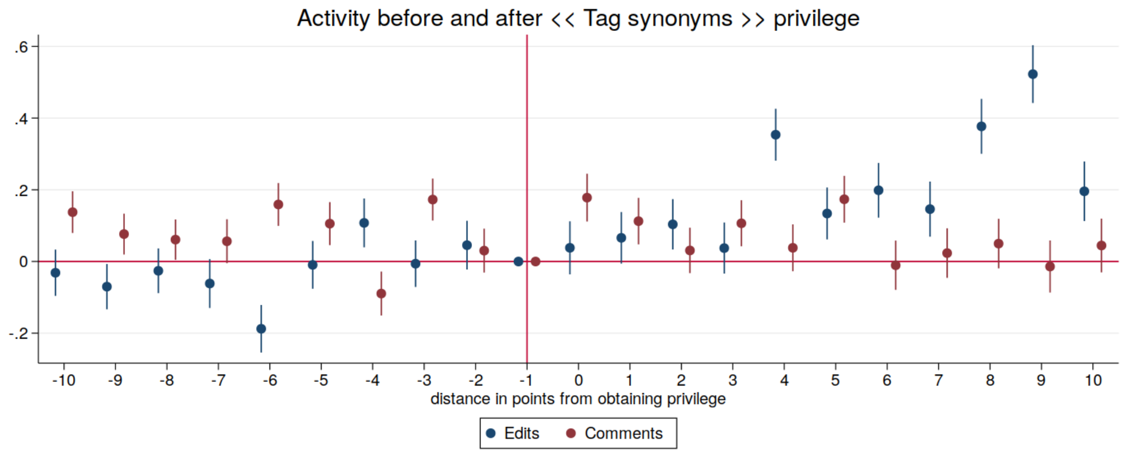

Figure 28: Estimates for reduced form effect around the *create tag synonyms* privilege

and the ex-ante value function as:

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

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

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

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

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

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

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

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

Appendix B Other figures

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

You gain reputation when:

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

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

You lose reputation when:

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

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

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

Figure 29: Rules to obtain or loose reputation in Stackexchange(https://stackoverflow.com/help/whats-reputation)

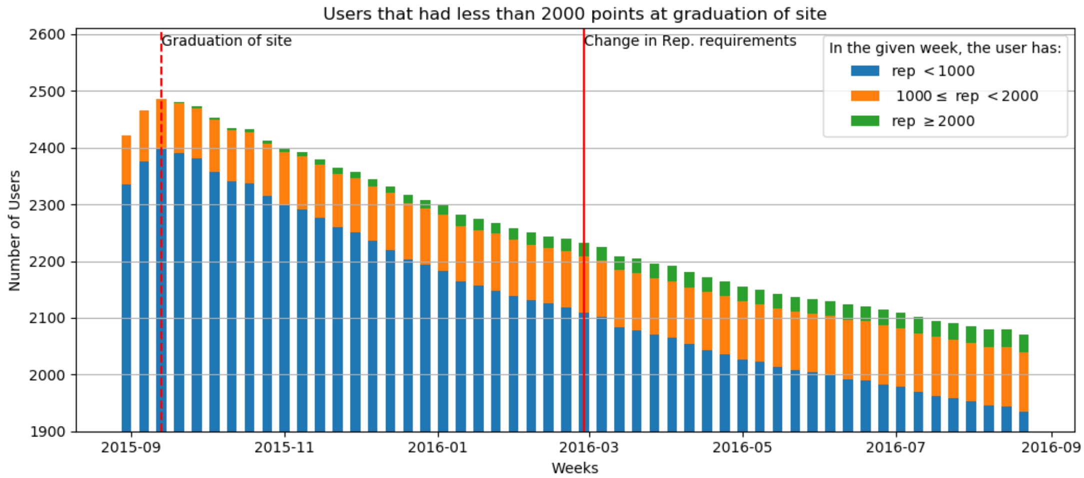

Figure 30: Number of users that have accumulated different amount of reputation points, conditional on having less than 2000 points at the graduation week. The decreasing value is due to exiting of the platform. It is possible to see that some users are reaching the level of 2000 points and they will not loose the privilege at the design date, some never reached the privilege, and others, the orange ones, loose it.

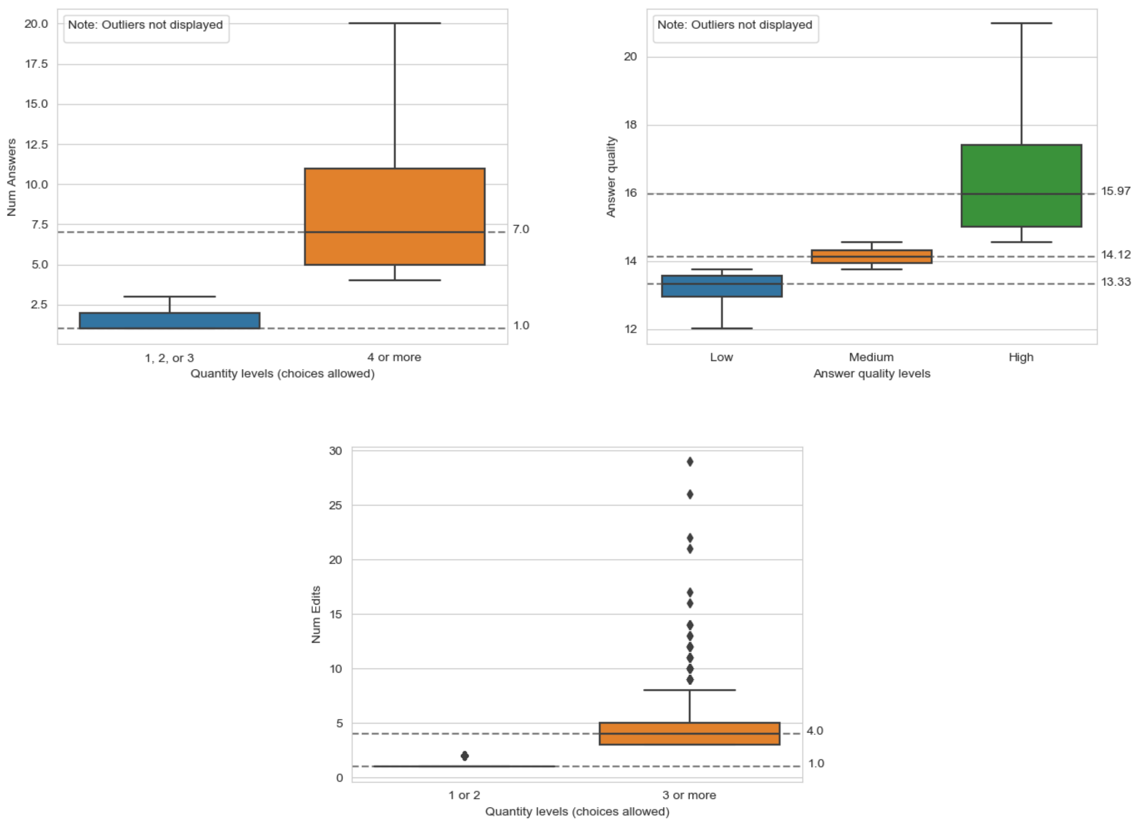

Figure 31: Categories of possible actions that users in the estimation are allowed to take, with the distribution of actual actions in each category. Values on the right vertical axis are the median value of each category, which make the set of options that users are allowed to choose.

Appendix C Credits for the software used

Pedregosa et al. (2011), Seabold and Perktold (2010), Hagberg et al. (2008), McKinney (2010), Lê et al. (2008), Virtanen et al. (2020), Hunter (2007)

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