

DISCUSSION PAPER SERIES

DP13575

**EXTERNALITIES IN KNOWLEDGE
PRODUCTION: EVIDENCE FROM A
RANDOMIZED FIELD EXPERIMENT**

Marit Hinnosaar, Toomas Hinnosaar, Michael
Kummer and Olga Slivko

INDUSTRIAL ORGANIZATION

EXTERNALITIES IN KNOWLEDGE PRODUCTION: EVIDENCE FROM A RANDOMIZED FIELD EXPERIMENT

Marit Hinnosaar, Toomas Hinnosaar, Michael Kummer and Olga Slivko

Discussion Paper DP13575

Published 08 March 2019

Submitted 06 March 2019

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **INDUSTRIAL ORGANIZATION**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Marit Hinnosaar, Toomas Hinnosaar, Michael Kummer and Olga Slivko

EXTERNALITIES IN KNOWLEDGE PRODUCTION: EVIDENCE FROM A RANDOMIZED FIELD EXPERIMENT

Abstract

Do contributions to online content platforms induce a feedback loop of ever more user-generated content or will they discourage future contributions? To assess this, we use a randomized field experiment which added content to some pages in Wikipedia while leaving similar pages unchanged. We find that adding content has a negligible impact on the subsequent long-run growth of content. Our results have implications for information seeding and incentivizing contributions, implying that additional content does not generate sizable externalities, neither by inspiring nor by discouraging future contributions.

JEL Classification: L17, L86, C93

Keywords: User-generated content, knowledge accumulation, Wikipedia

Marit Hinnosaar - marit.hinnosaar@gmail.com
Collegio Carlo Alberto and CEPR

Toomas Hinnosaar - toomas@hinnosaar.net
Collegio Carlo Alberto

Michael Kummer - m.kummer@uea.ac.uk
University of East Anglia

Olga Slivko - oslivko@gmail.com
Centre for European Economic Research (ZEW)

Externalities in Knowledge Production: Evidence from a Randomized Field Experiment*

Marit Hinnosaar[†] Toomas Hinnosaar[‡] Michael Kummer[§]
Olga Slivko[¶]

March 1, 2019

Abstract

Do contributions to online content platforms induce a feedback loop of ever more user-generated content or will they discourage future contributions? To assess this, we use a randomized field experiment which added content to some pages in Wikipedia while leaving similar pages unchanged. We find that adding content has a negligible impact on the subsequent long-run growth of content. Our results have implications for information seeding and incentivizing contributions, implying that additional content does not generate sizable externalities, neither by inspiring nor by discouraging future contributions.

Keywords: user-generated content, knowledge accumulation, Wikipedia

1 Introduction

Knowledge is a key input to many economic activities and a driver of economic growth (Romer, 1990; Grossman and Helpman, 1993; Jones, 1995). An increasing share of knowledge is created in the form of user-generated content: consumer feedback systems, discussion boards, Q&A sites, open-source software, social networks, and online information

*We are grateful to Chris Forman, Willa Friedman, Shane Greenstein, David Hugh-Jones, Giovanni Mastrobuoni, Ignacio Monzon, Juan S. Morales, Abhishek Nagaraj, Stephan Seiler, Ananya Sen, and Michael Zhang for valuable comments.

[†]Collegio Carlo Alberto and CEPR, marit.hinnosaar@gmail.com.

[‡]Collegio Carlo Alberto, toomas@hinnosaar.net.

[§]Georgia Institute of Technology & University of East Anglia, michael.kummer@econ.gatech.edu.

[¶]Centre for European Economic Research (ZEW), slivko@zew.de.

repositories, such as Wikipedia. Understanding the drivers of contributions to user-generated content has been an important question in economics and management for the past two decades (Lerner and Tirole, 2003).

In this paper, we investigate whether there are positive or negative externalities in user-generated content production. Understanding and quantifying such externalities has both policy and managerial implications. If content generation has positive externalities on future content generation, then information seeding¹ and paid contributions may have a high return on investment in terms of added stimulated growth (Aaltonen and Seiler, 2016). On the other hand, if content generation has negative externalities, then such policies may backfire and be not only ineffective but lead to worse eventual outcomes (Nagaraj, 2017).

Due to the reflection problem (Manski, 1993), externalities in content generation are difficult to identify. An externality occurs when a contribution by a user motivates other users to contribute (positive externality) or prevents further contributions (negative externality). Yet, the correlation in their contributions does not necessarily attest to an externality. A positive correlation may arise when users contribute because they were both exposed to the same external shock, such as a news article or a research finding. Similarly, a non-causal negative correlation may be caused by processes with periodic updates, such as elections or periodically updated statistics. To identify the causal effect, shocks to content growth and contributions must be independent over time. Without randomization provided by an experiment, the assumption is unlikely to be satisfied.

We estimate the causal impact of additional content on subsequent contributions using a randomized field experiment in Wikipedia. Randomization ensures that the addition of content is exogenous in terms of future content generation. As argued by the literature analyzing social interactions (including Manski (1993) and Aaltonen and Seiler (2016)), randomized experiments are the best way to cleanly identify causal relationships in such interactions.

The exogenous variation in our data is generated by a randomized field experiment, which was conducted in 2014.² The experiment added relevant content to randomly chosen Wikipedia pages while leaving similar pages unchanged. The treatment added about two paragraphs (approximately 2,000 characters) and one picture to each page in the treatment group. The pages were about mid-sized Spanish cities in different language editions of Wikipedia.

¹Nagaraj (2017) describes how such policies have been used by Wikipedia (seeding articles on more than 30,000 US cities from US Census Bureau data), OpenStreetMap (US Census maps), and Reddit (fake user accounts).

²A comprehensive description of the experiment is provided in Hinnosaar, Hinnosaar, Kummer, and Slivko (2017), who studied the impact of this treatment on real-world outcomes.

We use a dataset of Wikipedia editing histories which includes all versions of Wikipedia pages in the treatment and the control groups. Our dataset and the experimental setting allow us to analyze both short-term and long-term effects, up to four years after the experiment. Our main variable of interest is the outcome of content production—page length. To study the impact on the quality of content, we also analyze various other measures of editing activity, including the number of unique editors, number of edits, and the amount of content added and deleted.

Our main finding is that the additional content has a negligible impact on the subsequent long-run growth. The pages which were improved by adding about 2,000 characters of content, four years later are still longer by about the same amount. However, we do find some evidence that in the short-run the editing activity increased. In the first two years after the experiment, the treatment increased the number of Wikipedia users editing the treated pages and increased the number of edits. The increase in the number of users and edits was short-lived; there was no impact in the third and fourth year post-experiment. Moreover, the amount of content these users added was small and their edits were mostly limited to directly modifying the text added by the treatment. Our findings are robust to a set of alternative specifications, including alternative covariates and both cross-sectional and panel data frameworks. The experimental variation allows interpreting our findings as causal effects.

These findings have a clear policy and managerial implication—investments in information seeding and additional content contributions have a negligible cumulative effect on growth. Therefore information seeding and incentivizing contributions are a matter of direct cost-benefit analysis: they pay off if and only if the costs of creating the content are lower than the value of the new content. The additional costs or benefits via externalities that discourage or inspire future contributions are negligible.

Our paper contributes to the literature that studies externalities in user-generated content production. The closest to our work are Aaltonen and Seiler (2016) and Nagaraj (2017). Aaltonen and Seiler (2016) used detailed observational data from Wikipedia. Nagaraj (2017) used a natural experiment on OpenStreetMap, a Wikipedia-style digital map-making community, which started out with better seeding information in some regions compared to others for quasi-random reasons. The two papers arrive at contradicting conclusions, which warrant further investigation regarding this issue. Our paper is the first to study the question using a randomized field experiment, which allows causal identification of the underlying externalities.

More generally, the paper belongs to the literature that analyzes what drives contributions to user-generated content.³ The topics addressed in the prior related literature

³Other studies on Wikipedia have analyzed biases in Wikipedia’s content (Greenstein and Zhu, 2012,

include the role of personal gain (Shah, 2006), group size (Zhang and Zhu, 2011), networks (Fershtman and Gandal, 2011; Ransbotham, Kane, and Lurie, 2012), spillovers (Kummer, 2014), symbolic awards (Gallus, 2017), performance feedback (Huang, Burtch, Gu, Hong, Liang, Wang, Fu, and Yang, 2018), monetary rewards vs social motives (Sun, Dong, and McIntyre, 2017), contributor diversity (Ren, Chen, and Riedl, 2015), and economic conditions, such as unemployment (Kummer, Slivko, and Zhang, 2015) and migration (Slivko, 2018). More closely related to our work is Kane and Ransbotham (2016) who study the relationship between contributions and consumption of content. Using a large observational dataset from Wikipedia they document that more content increases contributions but the effect decreases over time. Our paper extends the literature by using variation from a randomized field experiment to measure the impact of additional content on future content generation. It confirms the finding by Kane and Ransbotham (2016) that the impact decreases over time, moreover, it finds that the long-term effect is negligible.

The structure of the paper is as follows. In the next section, we describe the experiment. Section 3 describes the data. Section 4 describes the empirical strategy and presents the results of the impact of treatment on the subsequent page length and on editing activity: the number of editors, the number of edits, and how much they edited. Section 5 discusses the connection between our findings and previous literature and limitations in our analysis. Section 6 concludes.

2 Experiment

The field experiment added content (text and photos) to randomly chosen Wikipedia pages. The sample consisted of 240 Wikipedia pages. Specifically, it consisted of the pages of 60 Spanish cities in the French, German, Italian, and Dutch editions of Wikipedia. The cities were all medium sized, excluding the largest like Madrid and Barcelona, and excluding also the smaller cities. The Wikipedia pages in these languages were relatively short—up to 24,000 characters in each of these four languages.

Each city and each language edition of Wikipedia was treated equally. For each city, its page was assigned to the treatment group in two randomly chosen languages. In each language edition of Wikipedia, 30 randomly chosen city pages were assigned to the treatment group. Specifically, to obtain balance in the treatment and control groups, the randomization was stratified.⁴ The 60 cities were divided into ten equal-sized groups. Within each group, each city was randomly assigned to one of six treatment arms. The

2017; Hinnosaar, 2018) and the impact of Wikipedia on market outcomes (Xu and Zhang, 2013; Hinnosaar, Hinnosaar, Kummer, and Slivko, 2017) and science (Thompson and Hanley, 2018).

⁴For further details of the randomization, see Hinnosaar, Hinnosaar, Kummer, and Slivko (2017).

six treatment arms were as follows: treat the city page in one of the six possible language pairs (French & German; French & Italian; French & Dutch; German & Italian; German & Dutch; Italian & Dutch). This resulted in a design where the number of pages which were treated equaled the number of those that remained in the control group.

The Wikipedia pages were treated mid-August, 2014. The treatment added about 2,000 characters of text and photos to each page in the treatment group. The added text and photos were mostly obtained from the corresponding Spanish and English language Wikipedia pages. Because all the pages were about Spanish cities, the Spanish Wikipedia typically contained more information than the other language versions. The English language version of the page, typically, was also more detailed than in the languages in the experiment. Hence, there was information available in Spanish and English pages that was missing from the other language editions of Wikipedia. The treatment translated that text and added it to the corresponding pages in the treatment group.

The treatment of the pages in Dutch Wikipedia was not successful. While in French, German, and Italian Wikipedia, the added text and photos survived well over time, all the additions to Dutch Wikipedia were deleted within 24 hours (by a single editor). Wikipedia allows anyone to edit. It also means that anyone can delete or undo the latest changes by reverting to a previous version of the page. This happened in the Dutch version of Wikipedia, where 24 hours after the treatment, all the pages looked as if they had never been treated. Therefore, we exclude Dutch pages from our analysis and restrict attention to the 180 pages in French, German, and Italian. Note that our main results do not change if Dutch pages are included in the analysis.

3 Data

Our main sources of data are the Wikipedia editing histories in the treated languages. An editing history contains the full text of each revision⁵ of each page starting from the creation of the page until the beginning of September 2018. Our sample consists of the 180 pages in the experiment, which are the pages of 60 cities in French, German, and Italian Wikipedia. In the following subsections, we describe the construction of the dataset and variables used in the analysis.

⁵A revision (or an edit) is a version of a Wikipedia article saved at a specific moment of time by a specific user. All revisions with the corresponding metadata, including full text, user, and timestamp, are preserved by Wikipedia and publicly available.

3.1 Page length

Our main outcome variable is the page length after the experiment. We measure page length in characters, including spaces and wiki markup commands.

Figure 1 presents average page length in the treatment and control groups. Until the experiment in August 2014, the average page length in the control and treatment groups was rather similar.⁶ The experiment added significant length to the pages in the treatment group. After the experiment, the difference has been relatively stable. An exception is a sharp increase in the mean of the treatment group in August 2016. This jump comes from the efforts of a single editor who worked hard to improve one page in French Wikipedia—the page of the city of Cordoba.⁷ Appendix A presents the same figure first, without French Cordoba (figure A.1a) and second, with the logarithm of page length (figure A.1b), both of which show no evidence of an increase in the treatment group average in 2016.




Figure 1: Average page length in the treatment and control groups

Notes: The number of observations is 90 in the control and 90 in the treatment groups. The experiment month (August 2014) is marked by the dashed vertical line.

Similar dynamics can be seen when looking at the changes separately by language

⁶The drop in both the treatment and control groups in early 2013 comes from technical changes in Wikipedia: Addbot removed about 2,000 characters from each page with an explanation similar to “Migrating 77 interwiki links, now provided by Wikidata”.

⁷By August 2016, the page of Cordoba in French Wikipedia was relatively mature, with length 19,426 characters (at the time 93% of the pages in our sample were shorter than that). During August 2016 this user increased the page length to 100,702 characters which is almost twice the length of the longest page at the time (57,076 characters). Our conclusions do not change if we exclude this page.

(figures 2a to 2c). As expected, the placebo test with the Dutch pages (figure 2d) shows that the assignment to treatment group had no impact. Page length is one possible output measure of knowledge production in Wikipedia. Similar dynamics as on figure 1 can also be seen on figure A.2 in appendix A which presents alternative measures of content: images and plain text (that is, html elements removed from the parsed text).

To make the treatment and the control groups comparable we subtracted the length of text added by the treatment from the length of pages in the treatment group. Table 1 presents summary statistics of page length after the experiment.

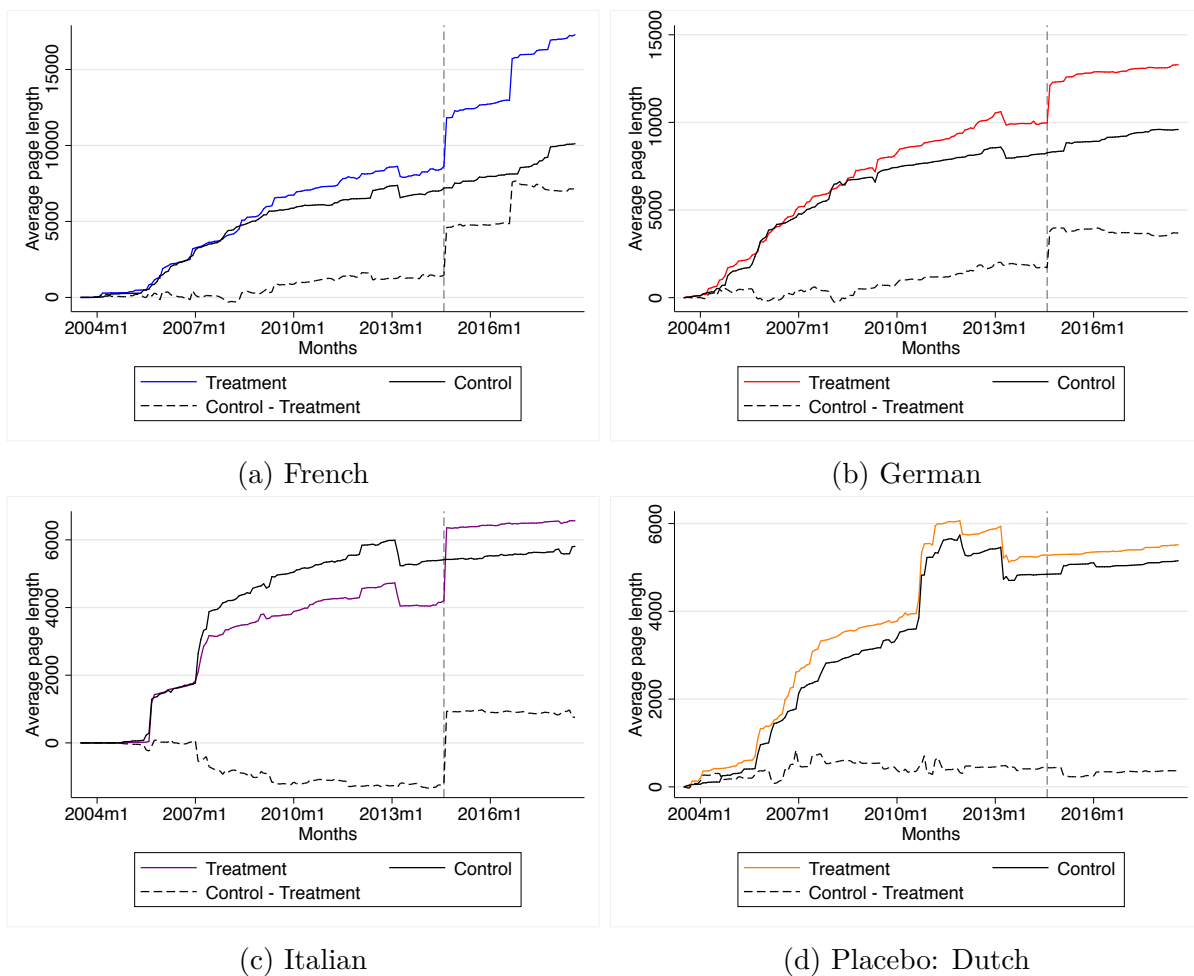


Figure 2: Average page length in the treatment and control groups, by language

Notes: On each figure, the number of observations is 30 in the control and 30 in the treatment groups. The experiment month (August 2014) is marked by the dashed vertical line.

3.2 Measures of editing activity

To construct the measures of editing activity, we start with 30,601 edits (revisions) from 180 Wikipedia pages. This includes all the edits except those generated as part of the

Table 1: Summary statistics of page length and editing activity post-treatment

Variable	Mean	SE	Min	10th	50th	90th	Max	Obs
Panel A: Page length minus the length of text added by treatment, post-treatment								
1st year	8012	6447	1260	1696	6588	16168	55378	180
2nd year	8640	9356	646	1663	6906	16302	98047	180
3rd year	9037	9709	627	1744	7358	17250	101428	180
4th year	9463	10141	611	1688	7554	18758	102688	180
Panel B: Average monthly number of unique users, post-treatment								
1st year	0.37	0.33	0.00	0.00	0.33	0.75	1.67	180
2nd year	0.32	0.35	0.00	0.00	0.17	0.71	2.58	180
3rd year	0.30	0.29	0.00	0.00	0.25	0.67	1.58	180
4th year	0.29	0.28	0.00	0.00	0.17	0.71	1.25	180
Panel C: Average monthly number of edits, post-treatment								
1st year	0.39	0.35	0.00	0.00	0.33	0.75	1.83	180
2nd year	0.34	0.41	0.00	0.00	0.21	0.75	3.67	180
3rd year	0.32	0.33	0.00	0.00	0.25	0.75	2.00	180
4th year	0.30	0.30	0.00	0.00	0.17	0.75	1.33	180
Panel D: Average monthly edit distance, post-treatment								
1st year	119.85	205.39	0.00	0.00	27.96	358.21	1520.83	180
2nd year	108.32	678.72	0.00	0.00	13.71	187.29	9027.83	180
3rd year	65.12	145.85	0.00	0.00	12.92	183.29	1273.92	180
4th year	112.12	500.08	0.00	0.00	11.12	167.58	5800.50	180

Notes: A unit of observation is a page (180 pages). Panel A reports summary statistics of page length (minus the length of text added by treatment) at the end of the 1st, 2nd, 3rd, and 4th year post-treatment. Summary statistics of post-treatment average monthly number of unique users are in panel B, average monthly number of edits in panel C, and average monthly edit distance in panel D. Page length and edit distance are measured in the number of characters.

treatment in the experiment. Following Aaltonen and Seiler (2016), we restrict the sample of edits in the following ways. First, we exclude edits by bots (about 30% of edits), these are non-human user accounts that generate automated edits. Specifically, we define bots as users whose username occurs in the list of bots (in the English, French, German or Dutch Wikipedias) or whose username includes “bot”. Second, we exclude reverts, which are edits that restore any previous version of the same page (about 7% of remaining edits). Third, we exclude vandalism (about 0.8% of remaining edits). We use the following criteria to classify an edit as vandalism: (a) an edit that only deletes text from the previous revision, and (b) the revision immediately after vandalism reverts the article back to a past revision. Then we are left with 19,586 productive edits generated by human users.

To analyze the impact of treatment on editing activity we construct three types of monthly measures that characterize how many people edited the pages, how many times they edited, and how much they edited. The first measure is the number of unique users

editing a page per month. We define a unique user by the username for registered users and by IP address for anonymous users. The second measure is the number of edits per month. To avoid double-counting of micro-edits,⁸ we first aggregate edits to day-user-page level, and then sum these up to month-page level. The third measure is edit distance, which we define as the number of characters an edit added plus the number of characters it deleted compared to the previous version of the page.⁹ We aggregate the edit distance measure to monthly level.

Panels B to D in table 1 present summary statistics of editing activity. Figure A.3 in appendix A describes the average editing activity in the treatment and control groups over time. Table A.1 in appendix A presents the comparison of pre-treatment page length and editing activity in the treatment group versus the control group. The table shows that there was no significant difference between the two groups before the treatment.

In addition to the aggregate measures of editing activity, we separate edits that directly modify the treatment text and those that modify other parts of the page. We classify edits into these two categories using a method similar to Hinnosaar, Hinnosaar, Kummer, and Slivko (2017). For each page in the treatment group, we use the diff algorithm between the revision before and after the treatment to determine treatment text—the exact text added by the treatment. For each revision post-treatment, using the diff algorithm between the treatment text and this revision, we check whether the revision deleted any part of the treatment text. If the revision didn’t delete anything from the treatment text, we classify the revision as one that edited other parts of the page.

4 Results

4.1 The impact on page length

Main empirical strategy. Because we are able to use experimental variation, we focus on cross-sectional estimation. Our empirical strategy compares page lengths after the experiment in the treatment and control groups. To make the pages comparable, we subtract the length of text added by the treatment from the length of pages in the treatment group. Hence, the estimates should be interpreted as the effect of treatment on page length after removing the mechanical increase created by the treatment. We estimate the

⁸Many Wikipedia editors save many revisions to the same page in a short period of time (because they save a revision after every small change while editing).

⁹Note that our edit distance measure differs from the Levenshtein edit distance by giving weight two to substitutions instead of weight one. For each edit, we calculate the edit distance using php FineDiff class at the granularity level of a character.

following regression:

$$Y_i = \beta_0 + \beta_1 TreatmentGroup_i + X_i + \varepsilon_i \quad (1)$$

where the outcome variable is the logarithm of page length (from which the length of the treatment text has been subtracted) of page i in the years post-experiment. The coefficient of interest is β_1 on $TreatmentGroup_i$, which is an indicator variable that takes value one if the page was assigned to the treatment group and zero if it was assigned to the control group. Covariates, X_i , include language and city fixed effects and logarithm of pre-treatment page length.

Main results. Table 2 presents the main results of the impact of treatment on page length up to four years after the experiment. Specifically, the outcome variables are the logarithm of page length (minus treatment text) at the end of the 1st year post-treatment (panel A), 2nd, 3rd, and the 4th year post-treatment (panels B to D). Columns 1–4, starting from no controls, successively add all the controls. The coefficients are stable across the specifications while adding covariates decreases standard errors.

The results in table 2 indicate that, once all the covariates have been included, pages in the treatment group compared to those in the control in the first years post-treatment are about 7% longer. By the fourth year, the treatment effect is not statistically significant, and the coefficient is smaller in magnitude. In all the years, the coefficient of the impact of treatment is less than 9% of the standard deviation (reported in table 2).

The results in table 2 suggest that by the end of the fourth year the impact of treatment on page length is negligible. After excluding the length of text added by the treatment, using the bounds of the 95% confidence interval implied by estimates in column 4 panel D, we are able to reject that by the end of the 4th year after the experiment, pages in the treatment group are more than 14.1% longer or more than 5.1% shorter than those in the control group.¹⁰ The estimation results are summarized graphically on figure 3.

Robustness and heterogeneity. Our analysis of robustness and heterogeneity is provided in online appendix A. We address robustness to inclusion of alternative covariates (table A.2) and heterogeneity across languages (table A.3) and across page characteristics (table A.4). In brief, we do not find any evidence that the treatment had any significant long term effect on page length.

First, we assess whether the results are sensitive to inclusion of alternative covariates. Table A.2 presents estimates from the same regressions as in table 2, but instead of city

¹⁰The upper bound: $0.045 + 1.96 \times 0.049 = 0.141$; and the lower bound: $0.045 - 1.96 \times 0.049 = -0.051$.

Table 2: The effect of treatment on page length

	Log. page length (minus treatment text)							
	Panel A: 1st year post-treatment				Panel B: 2nd year post-treatment			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
Treatment group	0.068	0.069	0.045	0.065*	0.071	0.075	0.047	0.071*
	(0.118)	(0.088)	(0.036)	(0.035)	(0.125)	(0.091)	(0.043)	(0.042)
Log. length before treatment			0.927***	0.851***			0.962***	0.853***
			(0.022)	(0.034)			(0.027)	(0.041)
Language FE	No	Yes	No	Yes	No	Yes	No	Yes
City FE	No	Yes	No	Yes	No	Yes	No	Yes
Mean dep. var.	8.699	8.699	8.699	8.699	8.723	8.723	8.723	8.723
SD dep. var.	0.792	0.792	0.792	0.792	0.836	0.836	0.836	0.836
Adj. R-squared	-0.004	0.510	0.909	0.922	-0.004	0.526	0.880	0.898
Observations	180	180	180	180	180	180	180	180

	Log. page length (minus treatment text)							
	Panel C: 3rd year post-treatment				Panel D: 4th year post-treatment			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
Treatment group	0.059	0.074	0.035	0.069	0.040	0.048	0.017	0.045
	(0.127)	(0.092)	(0.048)	(0.045)	(0.130)	(0.092)	(0.056)	(0.049)
Log. length before treatment			0.963***	0.857***			0.963***	0.830***
			(0.029)	(0.044)			(0.034)	(0.048)
Language FE	No	Yes	No	Yes	No	Yes	No	Yes
City FE	No	Yes	No	Yes	No	Yes	No	Yes
Mean dep. var.	8.762	8.762	8.762	8.762	8.795	8.795	8.795	8.795
SD dep. var.	0.847	0.847	0.847	0.847	0.868	0.868	0.868	0.868
Adj. R-squared	-0.004	0.524	0.858	0.889	-0.005	0.547	0.816	0.873
Observations	180	180	180	180	180	180	180	180

Notes: Each column presents estimates from a separate cross-section regression of 180 Wikipedia pages. The dependent variable is the logarithm of page length (minus the length of text added by the treatment) in the 1st year post-treatment (panel A), 2nd (panel B), 3rd (panel C), and the 4th year post-treatment (panel D). Standard errors are reported in parentheses. *** Indicates significance at 1 percent level, ** at 5 percent level, * at 10 percent level.

and language fixed effects, it includes stratification group dummies (see section 2). Results are identical to those without any fixed effects in table 2 (columns 1 and 3 in table 2).

Table A.3 re-estimates the regressions separately for each language. There is a small significant effect of treatment on page length only in French Wikipedia and only in the first two years post-treatment. There is no significant effect in other languages. Table A.4 re-estimates the regressions while dividing the sample by page age. While the estimates are insignificant, the magnitude of the coefficients in the first years post-treatment is larger in the case of younger pages.




Figure 3: The impact of treatment on page length

Notes: This figure presents point estimates (bars) and 95% confidence intervals (lines) of coefficients from four regressions in column 4 of table 2. The coefficients describe the impact of treatment on the logarithm of page length (from which the length of text added by the treatment has been subtracted) at the end of the 1st, 2nd, 3rd, and 4th year post-treatment. Unit of observation is a page (180 pages). For further details see table 2.

4.2 The impact on the number of users and edits

Main results. Table 3 presents the main results of the impact of treatment on the subsequent number of users and edits up to four years after the experiment. Each column presents estimates from similar cross-sectional regressions as in section 4.1, where the outcome variables measure editing activity in the 1st, 2nd, 3rd, and the 4th year after the experiment. The outcome variables are the average number of users (people editing the page) per month (panel A) and the average number of edits per month (panel B). In our preferred specification in table 3, covariates include city and language fixed effects. Additional covariates are included in the robustness analysis in online appendix A.

Panels A and B of table 3 show that the treatment increased the number of users and the number of edits during the first two years after the experiment. Specifically, the treatment increased the average monthly number of users editing the page by about 0.13 users (columns 2 and 4 in panel A) and increased the average monthly number of edits by 0.14 edits (columns 2 and 4 in panel B). These effects are 30–40% of a standard deviation.

However, these increases are only short-lived. In the third and fourth year, for both measures the effect of treatment is insignificant and the coefficients are small in magnitude. In the fourth year, the estimated increase is about 0.02 users per month, which is less

Table 3: The effect of treatment on subsequent number of users and edits

	1st year		2nd year		3rd year		4th year	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. Dependent variable: average number of users per month								
Treatment group	0.099**	0.122***	0.096*	0.119***	-0.006	0.014	-0.001	0.024
	(0.049)	(0.030)	(0.052)	(0.035)	(0.043)	(0.028)	(0.043)	(0.033)
Language FE	No	Yes	No	Yes	No	Yes	No	Yes
City FE	No	Yes	No	Yes	No	Yes	No	Yes
Mean dep. var.	0.368	0.368	0.317	0.317	0.295	0.295	0.292	0.292
SD dep. var.	0.329	0.329	0.353	0.353	0.287	0.287	0.285	0.285
Adj. R-squared	0.017	0.678	0.013	0.608	-0.006	0.618	-0.006	0.453
Observations	180	180	180	180	180	180	180	180
Panel B. Dependent variable: average number of edits per month								
Treatment group	0.109**	0.138***	0.119*	0.140***	-0.008	0.011	0.004	0.029
	(0.052)	(0.032)	(0.061)	(0.043)	(0.049)	(0.032)	(0.045)	(0.036)
Language FE	No	Yes	No	Yes	No	Yes	No	Yes
City FE	No	Yes	No	Yes	No	Yes	No	Yes
Mean dep. var.	0.389	0.389	0.335	0.335	0.319	0.319	0.305	0.305
SD dep. var.	0.351	0.351	0.412	0.412	0.325	0.325	0.302	0.302
Adj. R-squared	0.019	0.663	0.015	0.556	-0.005	0.623	-0.006	0.424
Observations	180	180	180	180	180	180	180	180
Panel C. Dependent variable: average number of edits per month excluding edits of the text added by treatment								
Treatment group	-0.003	0.022	0.058	0.072*	-0.056	-0.035	-0.050	-0.026
	(0.048)	(0.030)	(0.059)	(0.043)	(0.048)	(0.031)	(0.043)	(0.033)
Language FE	No	Yes	No	Yes	No	Yes	No	Yes
City FE	No	Yes	No	Yes	No	Yes	No	Yes
Mean dep. var.	0.333	0.333	0.305	0.305	0.295	0.295	0.278	0.278
SD dep. var.	0.318	0.318	0.398	0.398	0.319	0.319	0.287	0.287
Adj. R-squared	-0.006	0.641	-0.000	0.533	0.002	0.616	0.002	0.458
Observations	180	180	180	180	180	180	180	180

Notes: Each column presents estimates from a separate cross-section regression of 180 Wikipedia pages. Dependent variable is the average number of users or edits per month during the 1st year post-treatment (columns 1–2), 2nd (columns 3–4), 3rd (columns 5–6), and the 4th year post-treatment (columns 7–8). Dependent variables measure users (panel A), edits (panel B), and edits excluding those editing the text added by the treatment (panel C). Standard errors are reported in parentheses. *** Indicates significance at 1 percent level, ** at 5 percent level, * at 10 percent level.

than 9% of a standard deviation, and in the third year, the effect is even smaller (columns 6 and 8 in panels A and B). Figure 4 summarizes these results.

What do these editors and edits do in the first two years post-treatment if it has surprisingly little effect on page length (as we saw in section 4.1)? A natural explanation could be that the additional edits simply polish the text added by the treatment. To




Figure 4: The impact of treatment on the monthly average number of users and edits

Notes: This figure presents point estimates (bars) and 95% confidence intervals (lines) of coefficients from 12 regressions in columns 2, 4, 6, and 8 in panels A to C of table 3. The coefficients describe the impact of treatment on the average monthly number of users (panel A), edits (panel B), and edits excluding those of the text added by the treatment (panel C) in the 1st, 2nd, 3rd, and 4th year post-treatment. Unit of observation is a page (180 pages). For further details see table 3.

study this, we re-calculated the average number of edits per month while excluding edits that directly edited the text added by the treatment. The estimates using this outcome variable are presented in panel C of table 3 and figure 4c. The results show that when excluding the edits that directly affect the text added by the treatment then the treatment effect is much smaller. We conclude that a large share of the short-run increase in editing (seen in panels A and B of table 3) comes from editing the content added by the treatment.

Robustness and heterogeneity. In tables A.5 to A.7 in online appendix A we assess whether our findings reported in panels A and B in table 3 hold under different specifications. In brief, we find that the results are robust to including alternative covariates and estimating a panel data fixed effects model instead of cross-sectional regressions.

First, we assess whether the results are sensitive to the inclusion of alternative covariates. Table A.5 presents estimates from the same regressions as in table 3, but instead of city and language fixed effects, it includes stratification group dummies (see section 2). Results are identical to those without any fixed effects in table 3. Regressions in table A.6 include pre-treatment average number of users and edits. The results remain similar to those in table 3.

Table A.7 presents estimates from panel data fixed-effects regressions. The panel data fixed effects estimates are similar to those based on cross-sectional data. The treatment significantly increased the number of users and the number of edits in the first two years after the experiment but had no effect on the third and fourth year. The magnitude of the estimated effect in the first two years is only slightly larger compared to the cross-sectional framework: we estimate 0.14 users per month in the panel data regressions vs

0.12 in the cross-sectional regressions.

In tables A.8 to A.10 we re-estimate regressions in table 3 separately by each language. We find that there is heterogeneity across languages. In particular, the treatment has a positive significant effect on editing behavior only in French and German Wikipedia and there is no evidence of a positive effect in Italian Wikipedia.

To shed more light on where the short-run increase in editing activity comes from, we analyze the sample by splitting it in two ways. First, we distinguish edits by “old” users who have edited the same page prior to the experiment (i.e. users who might have a sense of “ownership” of the page) and “new” users who did not edit the page before the experiment. Table A.11 shows no evidence that the treatment has any effect on old users. Instead, we find that the short-run increase in the number of users comes from new users. It should be noted that the number of old users is small, which might make it difficult to detect the effect.

Second, we study the heterogeneous effects by the age of the Wikipedia pages. We hypothesized that in the early stages of page age, the effect of added content could be larger than in the later stages. We divided the sample by the median age of the page.¹¹ Estimates in table A.12 suggest that, indeed, in the case of younger pages the effect of treatment in the first year is larger.

4.3 The impact on edit distance

Main results. Table 4 presents the main results of the impact of treatment on the subsequent *edit distance*, which equals the total amount of text added plus the total amount of text deleted. By adding together the amount of text added and deleted, edit distance captures how polished the text is and in this way provides a measure of quality. As in the previous tables, each column presents estimates from a cross-sectional regression where the outcome variables measure editing activity in the 1st, 2nd, 3rd, and the 4th year after the experiment. The dependent variable in panel A is the average edit distance per month.

Estimates in table 4 panel A show that the treatment had no statistically significant effect on edit distance. Coefficients vary in sign and, with the exception of the second year, are small in magnitude, less than 15 characters per month (columns 2, 6, and 8). In the second year, the (statistically insignificant) coefficient estimate (in column 4) implies the treatment effect of 100 characters per month, which is still slightly less than 15% of a standard deviation.

Because the distribution of edit distance has a long tail, in panel B of table 4, we

¹¹Note that the sample is split to unequal groups because many pages were created on a same day.

Table 4: The effect of treatment on edit distance in the 4 years post-treatment

	1st year		2nd year		3rd year		4th year	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. Dependent variable: monthly average edit distance								
Treatment group	7.655 (30.699)	12.806 (31.301)	128.524 (101.004)	100.284 (105.570)	-15.897 (21.771)	-13.534 (21.692)	23.890 (74.735)	1.296 (79.265)
Language FE	No	Yes	No	Yes	No	Yes	No	Yes
City FE	No	Yes	No	Yes	No	Yes	No	Yes
Mean dep. var.	119.848	119.848	108.318	108.318	65.123	65.123	112.119	112.119
SD dep. var.	205.394	205.394	678.724	678.724	145.853	145.853	500.082	500.082
Adj. R-squared	-0.005	0.071	0.003	0.032	-0.003	0.115	-0.005	-0.005
Observations	180	180	180	180	180	180	180	180
Panel B. Dependent variable: monthly average capped edit distance								
Treatment group	13.644** (6.124)	17.282*** (5.609)	18.642* (9.776)	20.512** (8.654)	-6.174 (6.812)	-4.158 (5.294)	-2.823 (6.140)	0.485 (5.343)
Language FE	No	Yes	No	Yes	No	Yes	No	Yes
City FE	No	Yes	No	Yes	No	Yes	No	Yes
Mean dep. var.	39.173	39.173	35.008	35.008	32.183	32.183	28.758	28.758
SD dep. var.	41.530	41.530	66.061	66.061	45.675	45.675	41.095	41.095
Adj. R-squared	0.022	0.270	0.015	0.314	-0.001	0.463	-0.004	0.324
Observations	180	180	180	180	180	180	180	180

Notes: Each column presents estimates from a separate cross-section regression of 180 Wikipedia pages. Edit distance (panel A) equals the number of characters added plus the number of characters deleted. Capped edit distance (panel B) is calculated using individual edits which edit distance is capped at the 90th percentile. All dependent variables are averages across months and are measured in the number of characters. Standard errors are reported in parentheses. *** Indicates significance at 1 percent level, ** at 5 percent level, * at 10 percent level.

use an alternative edit distance measure calculated from individual capped edits. The individual edits are capped from above at the 90th percentile. The 90th percentile equals about 500 characters and is about 10 times larger than the median edit. In this way, the capped edit distance measure gives a smaller weight to long edits.

Estimates in panel B show that in the first two years after the experiment, the treatment increases the average capped edit distance. The magnitude of the effect is small—an increase of about 17 characters per month in the first year (column 2) and 21 characters per month in the second year (column 4). Provided that the average word length across languages in the experiment is about 10.4 characters, our treatment increased the edit distance by 1.5 words in the first and by two words in the second year.¹² The small magnitude of the short-term increase in edit distance is in line with the findings from panel C of table 3, which showed that most of the increase in editing comes from editing the

¹²Source: <http://www.ravi.io/language-word-lengths>.

content added by the treatment.

The treatment has no significant effect on capped edit distance in later years. In the later years, the coefficients are even smaller. In the fourth year, the estimated effect is about 0.5 characters per month (column 8), which is less than 1% of the standard deviation.

Robustness. Table 4 already illustrated that the estimates of the impact of treatment on edit distance are sensitive to the specific form of the outcome variable. In the following, we assess whether the results remain similar under different specifications. First, we re-estimate the cross-sectional regressions including alternative covariates. Second, instead of using the cross-sectional data, we use a panel and estimate a panel data fixed effects model. In brief, the results in all these robustness checks are close to those in table 4.

First, we assess whether the estimates are sensitive to including alternative covariates. Table A.13 presents estimates from the same regressions as in table 4, but instead of city and language fixed effects, it includes stratification group dummies (see section 2). Results are identical to those without any fixed effects in table 4 (columns 1, 3, 5, and 7 in table 4). In Table A.14 we re-estimate the regressions in table 4 including pre-treatment average edit distance measures as covariates. Note that the estimated coefficients on the added covariates have unstable signs and are not always statistically significant. But the estimates of the impact of treatment on edit distance remain similar to those in table 4.

Table A.15 presents estimates from panel data fixed-effects regressions. These results are similar to those based on cross-sectional data. According to the estimates, the treatment increased capped net edit distance by about 18 characters per month in the first year post-treatment and 23 characters per month in the second year. It had no statistically significant effect in the following years, and the coefficients are very small, suggesting an effect of fewer than 2 characters per month.

Taken together, the results in tables 4 and A.13 to A.15 show no significant effect of treatment on the raw edit distance measure and show a small positive impact on the capped edit distance only in the first two years after the treatment.

5 Discussion

Comparison with previous literature. Aaltonen and Seiler (2016) and Nagaraj (2017) studied the same question and reached two opposite conclusions. Aaltonen and Seiler (2016) used detailed observational data from Wikipedia and found that longer pages get more future edits. Their simulation suggests a large cumulative effect—due to the positive externality, Wikipedia has almost twice as much content than without. On the

other hand, Nagaraj (2017) used quasi-random variation from a natural experiment in a Wikipedia-style mapping service. He found that regions with better initial content received fewer contributions and had worse quality output in the long run (about ten percent higher error rate).

In contrast to these papers, we used variation from a randomized field experiment, which provides clean identification of the causal impact. Our results bridge the gap between their opposing conclusions. We find that the addition of content has a negligible impact on the subsequent long-run growth of content. However, we do find small and short-lived increases in page length and editing activity. To be precise, coefficients estimated by Aaltonen and Seiler (2016) implied that an additional 10,000 characters of text lead to 0.204 additional users per week, which corresponds to about 0.18 additional users per month for 2,000 characters added by the treatment. Our estimates for the first two years after the experiment are slightly smaller, about 0.12 additional users per month, and we find no significant long-run effects. Moreover, we also find that the short-run increase in editing activity is largely concentrated on the text added by the treatment. Taken together, our findings imply that in a setting where the variation comes from a randomized experiment the long-run externalities are negligible.

We would like to note that differences in the results could be explained not only by differences in the identification methods, but also by subtle differences in the research settings. In the following, we highlight three such differences. First, Wikipedia is a collaborative process, where it is not clear what the end product should look like. Each contribution can potentially signal to other editors the importance of a particular topic. For example, it is not clear which topics related to the Roman Empire (studied by Aaltonen and Seiler (2016)) are important enough to be covered in detail in Wikipedia. In this way, contributions become “votes” of importance. This effect is missing in the case of maps (studied by Nagaraj (2017)), where the final output should be more or less homogeneous in the level of detail. We studied city pages, which have a standardized structure and therefore the vote-of-importance effect should be smaller than in Aaltonen and Seiler (2016). Future research could further explore whether user-generated contributions also serve as signals or votes on the importance of a topic.

Second, in our setting, similar to Nagaraj (2017), the added content comes from an outside source and is not generated by the community itself.¹³ Instead, in the setting of Aaltonen and Seiler (2016), the content is created by the community. Content coming from an outside source might have a different impact than when it was generated by the existing users themselves.

Third, the key mechanism proposed by Nagaraj (2017) to explain the negative ex-

¹³Although, in our case, the added content is translated from the other language editions of Wikipedia.

ternality is the “ownership effect”, which plays a less prominent role in Wikipedia. Nagaraj (2017) suggested that contributors who added particular bridges or streets on the user-generated map may feel more responsible to keep these objects updated over time. Therefore, the treatment of adding more seeding information may backfire by not allowing the ownership of objects to arise naturally. Both Aaltonen and Seiler (2016) and our paper focus on textual content in Wikipedia, where ownership is less clear and we would thus expect the negative effect of adding content to be less prominent. Indeed, we find no evidence that old users (with a possible sense of ownership) change their editing behavior in response to the treatment, indicating that ownership plays a smaller role in Wikipedia editing. It would be interesting to see future research on the content ownership effect in a randomized experiment.

Implications. Many user-generated content platforms use managerial interventions that aim at motivating users to contribute new content. Examples include seeding the platform with initial content, compensating users for their contributions, or running campaigns to help to get the process started. Whether such policies should be used depends on whether the added content inspires an upward spiral of more user-generated content or whether it discourages future contributions. This choice is a critical managerial decision, because firm wikis, archives or Q&A forums all depend on sufficient provision of information. On the other hand, such interventions are costly and require committing resources that cannot be invested elsewhere. Hence, it is important to understand not only the direction but also the magnitude of any possible externality that added content might have on follow-on contributions.

Our findings have a clear policy and managerial implication. The additional content has a negligible cumulative effect on growth. Therefore information seeding and incentivizing contributions are a matter of direct cost-benefit analysis: they pay off if and only if the costs of creating the content are lower than the value of the new content to the users and eventually to the platform. The additional costs or benefits via externalities that discourage or inspire future contributions are negligible.

Limitations. Our analysis benefits from the clean identification due to the randomization, but it still faces limitations. The first concern is the generalizability of our results. As the above discussion already points out, the results from Wikipedia might not directly generalize to other user-generated content platforms. One relevant difference among the platforms is the magnitude of the contributor’s personal benefit. In Wikipedia, the personal benefit from contributing is likely to be smaller than in open-maps or open-source software. For example, a user of open-maps could directly benefit from correcting a mis-

take on a map, while a mistake in Wikipedia is unlikely to have any personal consequences.

In this paper, we measured knowledge production using page length, number of users, number of edits, and edit distance. The same method and data can be used to estimate the impact on other measures, for example, number of images, number of external or internal links, or subjective quality measures.

6 Conclusions

In this paper, we show that the addition of content has a negligible impact on the subsequent long-run growth of content. We identify the causal effect using exogenous variation from a randomized field experiment in Wikipedia. We find that the treatment, which added content to randomly chosen Wikipedia pages, had a negligible impact on the subsequent long-run growth. We do find some evidence of short-run increases in editing activity, in particular, increases in the number of edits and new editors in the first two years after the treatment. However, the amount of content these users added was small and most of their edits modified the content added by the treatment. Our results are robust to a large set of sensitivity checks, including alternative covariates and analysis in the cross-sectional and in the panel data fixed effects framework.

Our findings have a clear policy and managerial implication—information seeding and motivating content creation is not enough to generate a meaningful increase in future content generation. However, these policies are also not counterproductive as the discouragement effect on future contributions is also negligible. Therefore, it is simply a matter of direct cost-benefit analysis whether such policies pay off.


References

- AALTONEN, A., AND S. SEILER (2016): “Cumulative Growth in User-Generated Content Production: Evidence from Wikipedia,” *Management Science*, 62(7), 2054–2069.
- FERSHTMAN, C., AND N. GANDAL (2011): “Direct and indirect knowledge spillovers: the "social network" of open-source projects,” *The RAND Journal of Economics*, 42(1), 70–91.
- GALLUS, J. (2017): “Fostering Public Good Contributions with Symbolic Awards: A Large-Scale Natural Field Experiment at Wikipedia,” *Management Science*, 63(12), 3999–4015.

- GREENSTEIN, S., AND F. ZHU (2012): “Is Wikipedia Biased?,” *American Economic Review: Papers and Proceedings*, 102(3), 343–348.
- (2017): “Do Experts or Crowd-based Models Produce More Bias? Evidence from Encyclopedia Britannica and Wikipedia,” *MIS Quarterly*, forthcoming.
- GROSSMAN, G. M., AND E. HELPMAN (1993): *Innovation and growth in the global economy*. MIT press.
- HINNOSAAR, M. (2018): “Gender Inequality in New Media: Evidence from Wikipedia,” *manuscript*.
- HINNOSAAR, M., T. HINNOSAAR, M. E. KUMMER, AND O. SLIVKO (2017): “Wikipedia Matters,” *manuscript*.
- HUANG, N., G. BURTCH, B. GU, Y. HONG, C. LIANG, K. WANG, D. FU, AND B. YANG (2018): “Motivating User-Generated Content with Performance Feedback: Evidence from Randomized Field Experiments,” *Management Science*, forthcoming.
- JONES, C. I. (1995): “R & D-based models of economic growth,” *Journal of Political Economy*, 103(4), 759–784.
- KANE, G. C., AND S. RANSBOTHAM (2016): “Content as Community Regulator: The Recursive Relationship Between Consumption and Contribution in Open Collaboration Communities,” *Organization Science*, 27(5), 1258–1274.
- KUMMER, M. (2014): “Spillovers in networks of user generated content: Pseudo-experimental evidence on Wikipedia,” *manuscript*.
- KUMMER, M. E., O. SLIVKO, AND X. M. ZHANG (2015): “Economic downturn and volunteering: Do economic crises affect content generation on Wikipedia?,” *manuscript*.
- LERNER, J., AND J. TIROLE (2003): “Some Simple Economics of Open Source,” *The Journal of Industrial Economics*, 50(2), 197–234.
- MANSKI, C. F. (1993): “Identification of Endogenous Social Effects: The Reflection Problem,” *The Review of Economic Studies*, 60(3), 531–542.
- NAGARAJ, A. (2017): “Information Seeding and Knowledge Production in Online Communities: Evidence from OpenStreetMap,” *manuscript*.
- RANSBOTHAM, S., G. C. KANE, AND N. H. LURIE (2012): “Network Characteristics and the Value of Collaborative User-Generated Content,” *Marketing Science*, 31(3), 387–405.

- REN, Y., J. CHEN, AND J. RIEDL (2015): “The Impact and Evolution of Group Diversity in Online Open Collaboration,” *Management Science*, 62(6), 1668–1686.
- ROMER, P. M. (1990): “Endogenous technological change,” *Journal of Political Economy*, 98(5, Part 2), S71–S102.
- SHAH, S. K. (2006): “Motivation, Governance, and the Viability of Hybrid Forms in Open Source Software Development,” *Management Science*, 52(7), 1000–1014.
- SLIVKO, O. (2018): “‘Brain Gain’ on Wikipedia: Immigrants Return Knowledge Home,” *manuscript*.
- SUN, Y., X. DONG, AND S. MCINTYRE (2017): “Motivation of User-Generated Content: Social Connectedness Moderates the Effects of Monetary Rewards,” *Marketing Science*, 36(3), 329–337.
- THOMPSON, N., AND D. HANLEY (2018): “Science is shaped by wikipedia: Evidence from a randomized control trial,” *manuscript*.
- XU, S. X., AND X. ZHANG (2013): “Impact of Wikipedia on Market Information Environment: Evidence on Management Disclosure and Investor Reaction,” *MIS Quarterly*, 37(4), 1043–1068.
- ZHANG, X., AND F. ZHU (2011): “Group Size and Incentives to Contribute: A Natural Experiment at Chinese Wikipedia,” *The American Economic Review*, 101(4), 1601–1615.

A Online Appendix: Additional figures and tables




(a) Page length without Cordoba in French

(b) Logarithm of page length

Online Appendix Figure A.1: Robustness: Page length

Notes: The number of observations used to calculate the average is 90 in the control group and 89 (Figure A.1a) or 90 (Figure A.1b) in the treatment group. The experiment month (August 2014) is marked by dashed vertical line.




(a) Images


(b) Plain text

Online Appendix Figure A.2: Other output measures


Notes: The number of observations is 90 in the control and 90 in the treatment groups. The experiment month (August 2014) is marked by dashed vertical line. *Plain text* is obtained by removing html elements from the parsed text.



(a) Number of users



(b) Number of edits



(c) Edit distance

Online Appendix Figure A.3: Average input measures in the treatment and control groups per month

Notes: The number of observations is 90 in the control and 90 in the treatment groups. The experiment month (August 2014) is marked by dashed vertical line.

Online Appendix Table A.1: Comparison of pre-treatment characteristics in the treatment group versus the control group.

	Control group mean (1)	Treatment group mean (2)	t-test p-value (3)	Wilcoxon test p-value (4)	Obs. (5)
Log. length before treatment	8.586	8.611	0.842	0.655	180
Aver. # of users before treatment	0.378	0.333	0.321	0.285	180
Aver. # of edits before treatment	0.396	0.351	0.341	0.364	180
Aver. edit dist. before treatment	79.013	76.786	0.882	0.738	180
Aver. capped edit dist. before treatment	42.511	38.238	0.470	0.583	180

Notes: Column 1 and 2 present the means of pre-treatment values of variables, separately for the control and the treatment group. Column 3 presents the p-value of the t-test for whether the difference between the control and treatment groups is significantly different from zero. Column 4 presents the p-value of the corresponding Wilcoxon rank-sum test. Column 5 presents the number of observations used in each test.

Online Appendix Table A.2: Robustness: The effect of treatment on page length, alternative controls. Dependent variable: Logarithm of page length (minus treatment text).

	1st year		2nd year		3rd year		4th year	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treatment group	0.068 (0.115)	0.045 (0.036)	0.071 (0.122)	0.047 (0.044)	0.059 (0.123)	0.035 (0.048)	0.040 (0.126)	0.017 (0.056)
Log. length before treatment		0.936*** (0.024)		0.968*** (0.029)		0.969*** (0.032)		0.959*** (0.037)
Group FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Mean dep. var.	8.699	8.699	8.723	8.723	8.762	8.762	8.795	8.795
SD dep. var.	0.792	0.792	0.836	0.836	0.847	0.847	0.868	0.868
Adj. R-squared	0.046	0.908	0.048	0.876	0.048	0.855	0.057	0.811
Observations	180	180	180	180	180	180	180	180

Notes: Each column presents estimates from a separate cross-section regression of 180 Wikipedia pages. The dependent variable is the logarithm of page length (minus the length of text added by the treatment) in the 1st year post-treatment (columns 1-2), 2nd (columns 3-4), 3rd (columns 5-6), and the 4th year post-treatment (columns 7-8). All regressions include stratification group dummies, *GroupFE*. Standard errors are reported in parentheses. *** Indicates significance at 1 percent level, ** at 5 percent level, * at 10 percent level.

Online Appendix Table A.3: The effect of treatment on page length, by language. Dependent variable: Logarithm of page length (minus treatment text).

	1st year		2nd year		3rd year		4th year	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. Sample: pages in French Wikipedia								
Treatment group	0.256 (0.201)	0.144* (0.080)	0.328 (0.216)	0.210** (0.096)	0.285 (0.219)	0.168 (0.105)	0.223 (0.224)	0.109 (0.125)
Log. length before treatment		0.807*** (0.046)		0.850*** (0.055)		0.842*** (0.060)		0.821*** (0.071)
Mean dep. var.	8.865	8.865	8.925	8.925	9.003	9.003	9.072	9.072
SD dep. var.	0.782	0.782	0.847	0.847	0.852	0.852	0.869	0.869
Adj. R-squared	0.011	0.843	0.021	0.808	0.012	0.772	-0.000	0.693
Observations	60	60	60	60	60	60	60	60
Panel B. Sample: pages in German Wikipedia								
Treatment group	0.170 (0.160)	0.011 (0.047)	0.142 (0.160)	-0.015 (0.050)	0.145 (0.161)	-0.013 (0.051)	0.147 (0.157)	-0.007 (0.053)
Log. length before treatment		0.942*** (0.038)		0.936*** (0.040)		0.939*** (0.041)		0.913*** (0.043)
Mean dep. var.	9.002	9.002	9.021	9.021	9.049	9.049	9.075	9.075
SD dep. var.	0.621	0.621	0.618	0.618	0.621	0.621	0.609	0.609
Adj. R-squared	0.002	0.914	-0.004	0.904	-0.003	0.900	-0.002	0.887
Observations	60	60	60	60	60	60	60	60
Panel C. Sample: pages in Italian Wikipedia								
Treatment group	-0.223 (0.193)	0.008 (0.009)	-0.258 (0.204)	-0.016 (0.034)	-0.253 (0.205)	-0.011 (0.038)	-0.249 (0.213)	0.000 (0.047)
Log. length before treatment		0.989*** (0.006)		1.032*** (0.023)		1.034*** (0.025)		1.064*** (0.031)
Mean dep. var.	8.231	8.231	8.222	8.222	8.233	8.233	8.238	8.238
SD dep. var.	0.751	0.751	0.795	0.795	0.798	0.798	0.827	0.827
Adj. R-squared	0.006	0.998	0.010	0.973	0.009	0.967	0.006	0.952
Observations	60	60	60	60	60	60	60	60

Notes: Each column presents estimates from a separate cross-section regression of 60 Wikipedia pages. The dependent variable is the logarithm of page length (minus the length of text added by the treatment) in the 1st year post-treatment (columns 1-2), 2nd (columns 3-4), 3rd (columns 5-6), and the 4th year post-treatment (columns 7-8). Sample is restricted to pages in French Wikipedia (panel A), pages in German Wikipedia (panel B) or pages in Italian Wikipedia (panel C). Standard errors are reported in parentheses. *** Indicates significance at 1 percent level, ** at 5 percent level, * at 10 percent level.

Online Appendix Table A.4: The effect of treatment on page length, by page age. Dependent variable: Logarithm of page length (minus treatment text).

	1st year		2nd year		3rd year		4th year	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. Sample: pages with age below median								
Treatment group	0.239 (0.188)	0.109 (0.085)	0.244 (0.189)	0.114 (0.084)	0.237 (0.186)	0.114 (0.091)	0.142 (0.190)	0.026 (0.110)
Log. length before treatment		0.825*** (0.054)		0.831*** (0.054)		0.796*** (0.058)		0.760*** (0.070)
Language FE	No	Yes	No	Yes	No	Yes	No	Yes
Mean dep. var.	8.638	8.638	8.660	8.660	8.710	8.710	8.761	8.761
SD dep. var.	0.732	0.732	0.735	0.735	0.724	0.724	0.734	0.734
Adj. R-squared	0.010	0.800	0.011	0.807	0.010	0.765	-0.008	0.666
Observations	60	60	60	60	60	60	60	60
Panel B. Sample: pages with age above median								
Treatment group	-0.018 (0.151)	0.029 (0.020)	-0.016 (0.162)	0.030 (0.039)	-0.030 (0.166)	0.015 (0.041)	-0.010 (0.170)	0.034 (0.045)
Log. length before treatment		0.978*** (0.016)		1.009*** (0.030)		1.010*** (0.031)		1.020*** (0.035)
Language FE	No	Yes	No	Yes	No	Yes	No	Yes
Mean dep. var.	8.730	8.730	8.754	8.754	8.788	8.788	8.812	8.812
SD dep. var.	0.822	0.822	0.883	0.883	0.904	0.904	0.930	0.930
Adj. R-squared	-0.008	0.982	-0.008	0.941	-0.008	0.940	-0.008	0.931
Observations	120	120	120	120	120	120	120	120

Notes: Each column presents estimates from a separate cross-section regression of Wikipedia pages. The dependent variable is the logarithm of page length (minus the length of text added by the treatment) in the 1st year post-treatment (columns 1-2), 2nd (columns 3-4), 3rd (columns 5-6), and the 4th year post-treatment (columns 7-8). Sample is restricted to pages with age below median (panel A) or above median (panel B). The sample is split to unequal groups because many pages were created on a same day. Standard errors are reported in parentheses. *** Indicates significance at 1 percent level, ** at 5 percent level, * at 10 percent level.

Online Appendix Table A.5: Robustness: The effect of treatment on subsequent number of users and edits, alternative controls

	1st year (1)	2nd year (2)	3rd year (3)	4th year (4)
Panel A. Dependent variable: average number of users per month				
Treatment group	0.099** (0.047)	0.096* (0.051)	-0.006 (0.042)	-0.001 (0.041)
Group FE	Yes	Yes	Yes	Yes
Mean dep. var.	0.368	0.317	0.295	0.292
SD dep. var.	0.329	0.353	0.287	0.285
Adj. R-squared	0.067	0.058	0.024	0.070
Observations	180	180	180	180
Panel B. Dependent variable: average number of edits per month				
Treatment group	0.109** (0.051)	0.119* (0.060)	-0.008 (0.048)	0.004 (0.043)
Group FE	Yes	Yes	Yes	Yes
Mean dep. var.	0.389	0.335	0.319	0.305
SD dep. var.	0.351	0.412	0.325	0.302
Adj. R-squared	0.067	0.041	0.022	0.072
Observations	180	180	180	180

Notes: Notes: Each column presents estimates from a separate cross-section regression of 180 Wikipedia pages. Dependent variable is the average number of users (panel A) or edits (panel B) per month during the 1st year post-treatment (column 1), 2nd (column 2), 3rd (column 3), and the 4th year post-treatment (column 4). All regressions include stratification group dummies, *Group FE*. Standard errors are reported in parentheses. *** Indicates significance at 1 percent level, ** at 5 percent level, * at 10 percent level.

Online Appendix Table A.6: Robustness: The effect of treatment on subsequent number of users and edits, past measures as controls

	1st year		2nd year		3rd year		4th year	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. Dependent variable: average number of users per month								
Treatment group	0.136*** (0.032)	0.136*** (0.027)	0.134*** (0.036)	0.134*** (0.033)	0.026 (0.030)	0.021 (0.028)	0.029 (0.030)	0.037 (0.031)
Aver. # of users before treatment	0.813*** (0.053)	0.477*** (0.095)	0.841*** (0.060)	0.509*** (0.115)	0.693*** (0.049)	0.233** (0.097)	0.667*** (0.050)	0.451*** (0.111)
Language FE	No	Yes	No	Yes	No	Yes	No	Yes
City FE	No	Yes	No	Yes	No	Yes	No	Yes
Mean dep. var.	0.368	0.368	0.317	0.317	0.295	0.295	0.292	0.292
SD dep. var.	0.329	0.329	0.353	0.353	0.287	0.287	0.285	0.285
Adj. R-squared	0.573	0.733	0.529	0.661	0.526	0.632	0.493	0.517
Observations	180	180	180	180	180	180	180	180
Panel B. Dependent variable: average number of edits per month								
Treatment group	0.145*** (0.036)	0.150*** (0.030)	0.158*** (0.044)	0.154*** (0.041)	0.025 (0.034)	0.018 (0.031)	0.033 (0.034)	0.041 (0.035)
Aver. # of edits before treatment	0.795*** (0.057)	0.417*** (0.098)	0.887*** (0.070)	0.514*** (0.133)	0.748*** (0.053)	0.237** (0.100)	0.646*** (0.053)	0.420*** (0.112)
Language FE	No	Yes	No	Yes	No	Yes	No	Yes
City FE	No	Yes	No	Yes	No	Yes	No	Yes
Mean dep. var.	0.389	0.389	0.335	0.335	0.319	0.319	0.305	0.305
SD dep. var.	0.351	0.351	0.412	0.412	0.325	0.325	0.302	0.302
Adj. R-squared	0.530	0.706	0.478	0.603	0.522	0.638	0.449	0.482
Observations	180	180	180	180	180	180	180	180

Notes: Each column presents estimates from a separate cross-section regression of 180 Wikipedia pages. Dependent variable is the average number of users (panel A) or edits (panel B) per month during the 1st year post-treatment (columns 1-2), 2nd (columns 3-4), 3rd (columns 5-6), and the 4th year post-treatment (columns 7-8). All regressions include as a covariate the average number of users or edits pre-treatment. Standard errors are reported in parentheses *** Indicates significance at the 1 percent level, ** at 5 percent level, * at 10 percent level.

Online Appendix Table A.7: Robustness: The effect of treatment on subsequent number of users and edits, panel data FE regression

	# users		# edits	
	(1)	(2)	(3)	(4)
Treatment group, post-treatment 1st year	0.144*** (0.033)	0.144*** (0.033)	0.154*** (0.037)	0.154*** (0.037)
Treatment group, post-treatment 2nd year	0.141*** (0.037)	0.141*** (0.036)	0.164*** (0.045)	0.164*** (0.044)
Treatment group, post-treatment 3rd year	0.039 (0.032)	0.039 (0.032)	0.037 (0.036)	0.037 (0.035)
Treatment group, post-treatment 4th year	0.044 (0.034)	0.044 (0.033)	0.049 (0.037)	0.049 (0.037)
Month FE	Yes	Yes	Yes	Yes
Month-language FE	No	Yes	No	Yes
Page FE	Yes	Yes	Yes	Yes
Mean dep. var.	0.338	0.338	0.356	0.356
SD dep. var.	0.706	0.706	0.773	0.773
Observations	18360	18360	18360	18360

Notes: A unit of observation is a page-month pair. Dependent variable is the number of users (columns 1-2) or edits (columns 3-4) per month. All regressions include page fixed effects and either month fixed effects or month-language fixed effects. *Treatment group, post-treatment 1st year* is an indicator variable that takes value one during the first year post-treatment if the page belongs to the treatment group and zero otherwise; and similarly for other years. The sample is a balanced sample that starts from the time period when all the 180 pages existed. Standard errors, reported in parentheses, are clustered by page (180 pages). *** Indicates significance at the 1 percent level, ** at 5 percent level, * at 10 percent level.

Online Appendix Table A.8: The effect of treatment on subsequent number of users and edits, pages in French Wikipedia

	1st year (1)	2nd year (2)	3rd year (3)	4th year (4)
Panel A. Dependent variable: average number of users per month				
Treatment group	0.139 (0.091)	0.183 (0.113)	-0.017 (0.082)	-0.003 (0.080)
Mean dep. var.	0.392	0.372	0.325	0.326
SD dep. var.	0.357	0.443	0.315	0.306
Adj. R-squared	0.022	0.027	-0.017	-0.017
Observations	60	60	60	60
Panel B. Dependent variable: average number of edits per month				
Treatment group	0.147 (0.093)	0.244* (0.141)	-0.028 (0.096)	-0.003 (0.081)
Mean dep. var.	0.429	0.408	0.353	0.332
SD dep. var.	0.367	0.557	0.371	0.312
Adj. R-squared	0.024	0.033	-0.016	-0.017
Observations	60	60	60	60

Notes: Each column presents estimates from a separate cross-section regression of 60 Wikipedia pages in French Wikipedia. Dependent variable is the average number of users (panel A) or edits (panel B) per month during the 1st year post-treatment (column 1), 2nd (column 2), 3rd (column 3), and the 4th year post-treatment (column 4). Standard errors are reported in parentheses. *** Indicates significance at the 1 percent level, ** at 5 percent level, * at 10 percent level.

Online Appendix Table A.9: The effect of treatment on subsequent number of users and edits, pages in German Wikipedia

	1st year (1)	2nd year (2)	3rd year (3)	4th year (4)
Panel A. Dependent variable: average number of users per month				
Treatment group	0.219*** (0.082)	0.131 (0.080)	0.097 (0.068)	0.064 (0.072)
Mean dep. var.	0.510	0.357	0.360	0.354
SD dep. var.	0.332	0.314	0.264	0.280
Adj. R-squared	0.096	0.028	0.018	-0.004
Observations	60	60	60	60
Panel B. Dependent variable: average number of edits per month				
Treatment group	0.239** (0.092)	0.139 (0.084)	0.108 (0.077)	0.083 (0.080)
Mean dep. var.	0.533	0.372	0.396	0.375
SD dep. var.	0.373	0.329	0.299	0.308
Adj. R-squared	0.089	0.029	0.017	0.002
Observations	60	60	60	60

Notes: Each column presents estimates from a separate cross-section regression of 60 Wikipedia pages in German Wikipedia. Dependent variable is the average number of users (panel A) or edits (panel B) per month during the 1st year post-treatment (column 1), 2nd (column 2), 3rd (column 3), and the 4th year post-treatment (column 4). Standard errors are reported in parentheses. *** Indicates significance at the 1 percent level, ** at 5 percent level, * at 10 percent level.

Online Appendix Table A.10: The effect of treatment on subsequent number of users and edits, pages in Italian Wikipedia

	1st year (1)	2nd year (2)	3rd year (3)	4th year (4)
Panel A. Dependent variable: average number of users per month				
Treatment group	-0.061 (0.053)	-0.025 (0.069)	-0.097 (0.066)	-0.064 (0.063)
Mean dep. var.	0.203	0.221	0.201	0.196
SD dep. var.	0.205	0.265	0.258	0.243
Adj. R-squared	0.006	-0.015	0.020	0.001
Observations	60	60	60	60
Panel B. Dependent variable: average number of edits per month				
Treatment group	-0.058 (0.054)	-0.028 (0.072)	-0.106 (0.070)	-0.069 (0.068)
Mean dep. var.	0.204	0.225	0.208	0.207
SD dep. var.	0.209	0.277	0.273	0.264
Adj. R-squared	0.003	-0.015	0.021	0.001
Observations	60	60	60	60

Notes: Each column presents estimates from a separate cross-section regression of 60 Wikipedia pages in Italian Wikipedia. Dependent variable is the average number of users (panel A) or edits (panel B) per month during the 1st year post-treatment (column 1), 2nd (column 2), 3rd (column 3), and the 4th year post-treatment (column 4). Standard errors are reported in parentheses. *** Indicates significance at the 1 percent level, ** at 5 percent level, * at 10 percent level.

Online Appendix Table A.11: The effect of treatment on subsequent number of edits by old versus new users

	1st year		2nd year		3rd year		4th year	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. Dependent variable: Average number of edits by old users per month								
Treatment group	-0.004 (0.016)	0.006 (0.015)	-0.010 (0.009)	-0.011 (0.009)	-0.006 (0.012)	-0.007 (0.012)	0.004 (0.015)	0.008 (0.015)
Language FE	No	Yes	No	Yes	No	Yes	No	Yes
City FE	No	Yes	No	Yes	No	Yes	No	Yes
Mean dep. var.	0.070	0.070	0.037	0.037	0.041	0.041	0.041	0.041
SD dep. var.	0.107	0.107	0.061	0.061	0.080	0.080	0.103	0.103
Adj. R-squared	-0.005	0.201	0.001	0.042	-0.004	0.044	-0.005	0.147
Observations	180	180	180	180	180	180	180	180
Panel B. Dependent variable: Average number of edits by new users per month								
Treatment group	0.113** (0.044)	0.133*** (0.029)	0.129** (0.057)	0.151*** (0.040)	-0.002 (0.045)	0.018 (0.027)	-0.000 (0.038)	0.022 (0.031)
Language FE	No	Yes	No	Yes	No	Yes	No	Yes
City FE	No	Yes	No	Yes	No	Yes	No	Yes
Mean dep. var.	0.319	0.319	0.299	0.299	0.278	0.278	0.264	0.264
SD dep. var.	0.301	0.301	0.387	0.387	0.298	0.298	0.255	0.255
Adj. R-squared	0.030	0.637	0.022	0.568	-0.006	0.662	-0.006	0.416
Observations	180	180	180	180	180	180	180	180

Notes: Each column presents estimates from a separate cross-section regression of 180 Wikipedia pages. Dependent variable is the average number of edits by *old users* (panel A) or *new users* (panel B) per month during the 1st year post-treatment (columns 1–2), 2nd (columns 3–4), 3rd (columns 5–6), and the 4th year post-treatment (columns 7–8). *Old users* are defined as those who had edited the page before treatment, and *New users* are those who had not. Standard errors are reported in parentheses. *** Indicates significance at 1 percent level, ** at 5 percent level, * at 10 percent level.

Online Appendix Table A.12: The effect of treatment on subsequent number of users per month, by page age

	1st year		2nd year		3rd year		4th year	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. Sample: pages with age below median								
Treatment group	0.172*** (0.054)	0.176*** (0.055)	0.103 (0.070)	0.134* (0.076)	0.042 (0.053)	0.034 (0.055)	-0.033 (0.054)	-0.032 (0.054)
Language FE	No	Yes	No	Yes	No	Yes	No	Yes
Mean dep. var.	0.303	0.303	0.232	0.232	0.246	0.246	0.217	0.217
SD dep. var.	0.224	0.224	0.275	0.275	0.206	0.206	0.207	0.207
Adj. R-squared	0.136	0.134	0.019	0.033	-0.007	-0.001	-0.011	-0.020
Observations	60	60	60	60	60	60	60	60
Panel B. Sample: pages with age above median								
Treatment group	0.062 (0.067)	0.072 (0.060)	0.093 (0.069)	0.106 (0.075)	-0.029 (0.058)	-0.035 (0.060)	0.015 (0.057)	0.019 (0.055)
Language FE	No	Yes	No	Yes	No	Yes	No	Yes
Mean dep. var.	0.401	0.401	0.359	0.359	0.320	0.320	0.330	0.330
SD dep. var.	0.367	0.367	0.381	0.381	0.318	0.318	0.310	0.310
Adj. R-squared	-0.001	0.306	0.007	0.178	-0.006	0.188	-0.008	0.187
Observations	120	120	120	120	120	120	120	120

Notes: Each column presents estimates from a separate cross-section regression of Wikipedia pages. Dependent variable is the average number of users per month during the 1st year post-treatment (columns 1–2), 2nd (columns 3–4), 3rd (columns 5–6), and the 4th year post-treatment (columns 7–8). Sample is restricted to pages with age below median (panel A) or above median (panel B). The sample is split to unequal groups because many pages were created on a same day. Standard errors are reported in parentheses. *** Indicates significance at 1 percent level, ** at 5 percent level, * at 10 percent level.

Online Appendix Table A.13: Robustness: the effect of treatment on subsequent edit distance, alternative controls

	1st year (1)	2nd year (2)	3rd year (3)	4th year (4)
Panel A. Dependent variable: monthly average edit distance				
Treatment group	7.655 (30.583)	128.524 (101.061)	-15.897 (21.408)	23.890 (73.634)
Group FE	Yes	Yes	Yes	Yes
Mean dep. var.	119.848	108.318	65.123	112.119
SD dep. var.	205.394	678.724	145.853	500.082
Adj. R-squared	0.002	0.002	0.031	0.024
Observations	180	180	180	180
Panel A. Dependent variable: monthly average capped edit distance				
Treatment group	13.644** (6.116)	18.642* (9.786)	-6.174 (6.699)	-2.823 (6.004)
Group FE	Yes	Yes	Yes	Yes
Mean dep. var.	39.173	35.008	32.183	28.758
SD dep. var.	41.530	66.061	45.675	41.095
Adj. R-squared	0.024	0.012	0.032	0.039
Observations	180	180	180	180

Notes: Each column presents estimates from a separate cross-section regression of 180 Wikipedia pages. Edit distance (panel A) equals the number of characters added plus the number of characters deleted. Capped edit distance (panel B) is calculated using individual edits which edit distance is capped at the 90th percentile. All dependent variables are averages across months and are measured in the number of characters. All regressions include stratification group dummies, *Group FE*. Standard errors are reported in parentheses. *** Indicates significance at 1 percent level, ** at 5 percent level, * at 10 percent level.

Online Appendix Table A.14: Robustness: the effect of treatment on subsequent edit distance, pre-treatment average edit distance as a control

	1st year		2nd year		3rd year		4th year	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. Dependent variable: monthly average edit distance								
Treatment group	7.929 (30.731)	14.054 (31.036)	130.293 (100.588)	103.817 (105.085)	-14.643 (20.124)	-14.895 (21.084)	25.234 (74.399)	-0.361 (79.345)
Aver. edit dist. before treatment	0.123 (0.154)	-0.369* (0.210)	0.794 (0.503)	-1.043 (0.712)	0.563*** (0.101)	0.402*** (0.143)	0.603 (0.372)	0.489 (0.538)
Language FE	No	Yes	No	Yes	No	Yes	No	Yes
City FE	No	Yes	No	Yes	No	Yes	No	Yes
Mean dep. var.	119.848	119.848	108.318	108.318	65.123	65.123	112.119	112.119
SD dep. var.	205.394	205.394	678.724	678.724	145.853	145.853	500.082	500.082
Adj. R-squared	-0.007	0.087	0.012	0.042	0.143	0.165	0.004	-0.006
Observations	180	180	180	180	180	180	180	180
Panel A. Dependent variable: monthly average capped edit distance								
Treatment group	15.554*** (5.549)	17.414*** (5.629)	22.063** (8.594)	21.242** (8.527)	-3.341 (5.599)	-3.601 (5.158)	-0.533 (5.284)	1.048 (5.204)
Aver. capped edit dist. before treatment	0.447*** (0.070)	0.072 (0.120)	0.801*** (0.109)	0.394** (0.182)	0.663*** (0.071)	0.300*** (0.110)	0.536*** (0.067)	0.304*** (0.111)
Language FE	No	Yes	No	Yes	No	Yes	No	Yes
City FE	No	Yes	No	Yes	No	Yes	No	Yes
Mean dep. var.	39.173	39.173	35.008	35.008	32.183	32.183	28.758	28.758
SD dep. var.	41.530	41.530	66.061	66.061	45.675	45.675	41.095	41.095
Adj. R-squared	0.199	0.266	0.241	0.335	0.326	0.491	0.258	0.360
Observations	180	180	180	180	180	180	180	180

Notes: Each column presents estimates from a separate cross-section regression of 180 Wikipedia pages. Edit distance (panel A) equals the number of characters added plus the number of characters deleted. Capped edit distance (panel B) is calculated using individual edits which edit distance is capped at the 90th percentile. All dependent variables are averages across months and are measured in the number of characters. All regressions include as a covariate pre-treatment average edit distance or capped edit distance. Standard errors are reported in parentheses. *** Indicates significance at 1 percent level, ** at 5 percent level, * at 10 percent level.

Online Appendix Table A.15: Robustness: The effect of treatment on subsequent edit distance measures, panel data FE regression

	Edit distance		Capped edit distance	
	(1)	(2)	(3)	(4)
Treatment group, post-treatment 1st year	9.882 (33.347)	9.882 (33.005)	17.917*** (6.419)	17.917*** (6.413)
Treatment group, post-treatment 2nd year	130.751 (100.353)	130.751 (100.515)	22.914*** (8.638)	22.914*** (8.580)
Treatment group, post-treatment 3rd year	-13.670 (21.110)	-13.670 (21.168)	-1.901 (5.920)	-1.901 (5.918)
Treatment group, post-treatment 4th year	26.117 (74.428)	26.117 (74.629)	1.449 (5.932)	1.449 (5.809)
Month FE	Yes	Yes	Yes	Yes
Month-language FE	No	Yes	No	Yes
Page FE	Yes	Yes	Yes	Yes
Mean dep. var.	88.936	88.936	37.271	37.271
SD dep. var.	1086.255	1086.255	134.352	134.352
Observations	18360	18360	18360	18360

Notes: A unit of observation is a page-month pair. Dependent variable is monthly edit distance (columns 1-2) or capped edit distance (columns 3-4). Edit distance equals the number of characters added plus the number of characters deleted. Capped edit distance is calculated using individual edits which edit distance is capped at the 90th percentile. All dependent variables are averages across months and are measured in the number of characters. All regressions include page fixed effects and either month fixed effects or month-language fixed effects. *Treatment group, post-treatment 1st year* is an indicator variable that takes value one during the first year post-treatment if the page belongs to the treatment group and zero otherwise; and similarly for other years. The sample is a balanced sample that starts from the time period when all the 180 pages existed. Standard errors, reported in parentheses, are clustered by page (180 pages). *** Indicates significance at the 1 percent level, ** at 5 percent level, * at 10 percent level.