

Economic Downturn and Volunteering: Does a Crisis Affect Content Generation on Wikipedia?

Michael Kummer, Olga Slivko, Michael Zhang *

October 18, 2015

Abstract

In this paper, we address the impact of surging unemployment on online public good provision. Specifically, we ask how drastically increased unemployment affects voluntary contributions of content to the online encyclopedia Wikipedia. We put together a monthly country-level data set, which combines country specific economic outcomes with data on contributions to the online encyclopedia. As a source of exogenous variation in the economic state we use the various economic crises in European countries following the financial crisis in the US, which started in September 2008. We find that economic downturn is associated with more viewership, which channels higher participation of volunteers in Wikipedia expressed in editing activity and content growth. We provide evidence for increased information search online or online learning as a potential channel of the change in public goods provision, which is a potentially important side effect of economic downturn.

Keywords: public goods, unemployment, online platforms, user generated content.

JEL Classification Numbers: D29, D80, H41, J60, L17.

*Address: L7, 1 68161 Mannheim, telephone: +49(0)621123-358, e-mail: *kummer@zew.de*, *slivko@zew.de*, *zhang@ust.hk*.

1 Introduction

The importance of knowledge inputs for economic growth is hard to overestimate. Yet, little is known about how knowledge generation is related to the economic state of an economy, to economic downturn or to unemployment. In this paper, we address the impact of surging unemployment on online public good provision. Specifically, we ask how drastically increased unemployment rates in the aftermath of the European economic crisis in 2008, affected voluntary contributions of content to the online encyclopedia Wikipedia.

Unemployment may affect the willingness for contributing to online public goods in contradicting ways. On the one hand, unemployed people might have more spare time. Hence, as they spend more time online while searching for information, they might exert some effort and contribute to online knowledge. On the other hand, changes in contributing behavior might be due to employed people who spend more time online in order to improve their chances to keep their job. Alternatively, employed people may face larger pressure in their paid job. Therefore, if their contributions to online public goods were the main source, the net effect on the public goods provisions during economic crisis could be negative. A decrease in the public goods provision would threaten to further aggravate the long-term effects of increased unemployment. This effect would be especially problematic if the public good is an important input to the knowledge in the economy.

We put together a monthly country-level dataset which combines country specific economic outcomes with data on contributions to the online encyclopedia. As source of exogenous variation in the economic state we use the financial crisis and the ensuing global recession, which broke out with the collapse of Lehman Brothers in September 2008 (Verick and Islam, 2010). While (un)employment was relatively robust in some countries, the recession massively affected others and led to surging unemployment rates. We examine the effect of this event on the provision of effort and knowledge on the public good Wikipedia. We apply two frameworks: Our main specification is based on a difference-in-difference approach, but we also run specification tests using fixed-effect OLS regression analysis with monthly data on unemployment in Europe.

We find that increased unemployment is associated with higher participation of volunteers in Wikipedia and with an increased rate of content generation. More editors join the community and also the number of highly active users increases, suggesting that existing editors increase their

activity. We also find that the number of edits per article increases, and slightly weaker evidence for an increased overall content growth.

One of the most problematic consequences of the recent economic crises in Europe was their strong impact on unemployment. Our study sheds light on potentially productive voluntary online time use and its relationship to unemployment. For contributing knowledge to Wikipedia, individuals do not need any special skills. Contributions might vary from checking typos and rephrasing the existing text to adding pieces of information related to professional interest or hobbies. Yet, through their contribution to the online public good, the unemployed might acquire valuable writing and computer skills. At the same time, our paper addresses the change in public goods provision as a potential side effect of unemployment.

The rest of the paper is structured as follows. Section 2 presents previous studies related to my research question. Section 3 describes the dataset and Section 4 discusses the empirical approach. Section 5 conducts the empirical analysis on the country level, while Section 6 reports the analyzes of the German district level data. Section 7 discusses the obtained results, limitations and avenues for further research before Section 8 concludes.

2 Related Research

The present study contributes to the economic literature on public goods provision by focusing on the relationship between a potential rise in the unemployment and contributions to the largest online encyclopedia, Wikipedia. This encyclopedia is produced collaboratively and is accessible to anyone with an Internet connection. Wikipedia can be regarded as a modern public good by definition, since it is non-excludable and non-rival (Hess and Ostrom (2003)). There is evidence that Wikipedia is becoming a standard reference source. The popularity of Wikipedia (6th most visited website) is a clear indication that many people are interested in its content.

The private incentives for voluntary public goods provision have been addressed in existing theoretical studies through the prism of the interplay between the effect of the audience size and the free-riding incentives (Andreoni (1988), Andreoni (1989), Andreoni (1990), Andreoni (2007)). In Andreoni (2007) the provision of public goods is shown to be congestible, e.g. an increase in the number of recipients increases the total giving but at a lower rate. This finding was further

supported by empirical evidence for open-source software and online peer productive communities. Comino et al. (2007) find that the size of the “community of developers” in open-source projects increases the chances of progress but this effect decreases as the community gets larger. They conclude this to be a signal of possible coordination problems. Zhang and Zhu (2010) use exogenous variation in the recipient group size on Wikipedia (a block by the Chinese government) to show the importance of the recipient group size for individual incentives for knowledge provision. Besides, in online contexts like ours, social interactions or social image are suggested to be incentives for public good provision (Kandel and Lazear (1992), Algan et al. (2013)) as well as the heterogeneity of skills or other characteristics of the public good providers (Wang and Zudenkova (2014)). Furthermore, since the late 1980s researchers have increasingly contrasted theoretical models with experimental studies in the lab. The main insights of this extensive literature have been surveyed by Vesterlund (2006).

We adopt an alternative perspective on the public goods provision asking whether individuals having more time would allocate a part of it for charity. In this paper we use the recent economic crisis in Europe as an exogenous shock to available spare time. We analyze whether people who get unemployed or (had to) reduce their working time, would dedicate themselves to online knowledge provision. Several studies have looked at how unemployed allocate their time considering a range of potential time uses (Knabe et al. (2010), Krueger and Mueller (2012), Aguiar et al. (2012), Aguiar et al. (2013)). Aguiar et al. (2012) provide an extensive review of the literature on time use and life-cycle behaviour of households. Although unemployed people have more time to spend on leisure, they are less satisfied with life and during specific activities (Knabe et al., 2010). Krueger and Mueller (2012) find unemployed not only to be more sad compared to when they are employed, but they also spend 35 percent of the change in time spent working less in leisure at the time of reemployment. Here, leisure includes computer and Internet use. Aguiar et al. (2012) use the American Time Use Survey (ATUS) to analyze trends in time allocation. They state that since the 1960s, individuals spend more time on leisure. This category includes personal use of computer by definition as well as other activities such as watching television or engaging in sports. By analyzing time diary data from four different countries, Burda and Hamermesh (2010) come to the similar result that only a small share of the additional time of getting unemployed is used for home production, also indicating that unemployed spend more time on other activities such as computer

use. The ATUS analysis of Aguiar et al. (2013) focusses on the period of the global recession in the late 2000s decade in particular. Their results confirm the previous ones. They finds that more than 50 percent of the additional time is spent on leisure activities, yet two-thirds are absorbed by watching TV and sleeping. More interestingly, roughly two percent of the foregone market hours are allocated to civic and religious engagement.

The time use studies usually include email and other computer uses among leisure activities. Online leisure time is found to be a substitute for work since most of the time spent online is spent on social networks, online games, email and portals (Wallsten (2013)). Moreover, the young people are found to spend more time online. These findings are complemented by Goldfarb and Prince (2008) who show that, conditional on having Internet access, poorer people spend more time online than wealthy people as they have a lower opportunity cost of time. At times of economic crisis both these groups of people can be threatened by increased unemployment rates or decreased salaries. Taken together, these three facts lead us to expect that people who suddenly have more time spend at least a part of it online.

The literature suggests at least some examples of correlation between contribution to public goods and unemployment. As for the civic public good, Uslaner (2002) finds indications that unemployment is negatively correlated with both religious as well as secular volunteering. However, this study with data from the U.S. and Canada uses only cross sectional data and thus gives no information about effects of rising unemployment over time. Uslaner's results are consistent with the findings of Freeman (1997), who finds that volunteers are predominantly people "with higher potential earnings or greater demands on their time: the employed, married persons, those with larger families, persons in the 35-54 peak earnings ages, the more highly educated, professionals and managers". Moreover, among men, working more hours is even positively correlated with participation in volunteering. This is also in line with Taniguchi (2006), who studies the effect gender differences and employment on volunteering using the National Survey of Midlife Development in the United States (MIDUS) 1995-1996. His results suggest that unemployment has a negative effect on men's volunteering, which is not the case for women. Moreover, working part-time and working full time makes no difference in men's efforts in volunteering.

What do these results imply for our study? On the one hand, the observed shift in time allocation towards more computer use and increased civic engagement might lead to increased

provision of public goods. This increase could be caused by two effects: Firstly, people who haven't been active might begin searching for information on the Internet and discover the encyclopedia. Consequently they might also become interested in volunteering. Secondly, previous contributors might dispose of more time when getting unemployed and reallocate it partially to contributing to the public information good. On the other hand, it might be also the case that individuals who provide knowledge on Wikipedia do this independently of their employment status. In this case, the exogenous increase in unemployment will not result in any changes in the dynamic of contributions.

3 Data

To study the consequences of the various European crises on online knowledge generation we proceed on two levels. First we analyze a sample of European countries, which were affected by one of the crises in the year 2008 and compare them to relatively unaffected countries. Second, we analyze German district level data in greater detail and focus on the relationship of increased unemployment to (German) Wikipedia contributions from this district.

3.1 Wikipedias on the Country/Language Level

Our analysis uses Wikipedia monthly statistics provided by the Wikimedia Foundation. These statistics include the number of Wikipedians, the number of articles in Wikipedia, database sizes, number of words, and readership statistics for all language versions of Wikipedia. To study the relationship of country level unemployment on an entire Wikipedia, we need to focus on countries which have an (ideally) unique language. For example, some of the most heavily affected countries, such as Spain and Portugal, had to be excluded since their languages are spoken not exclusively in these countries, but also in Latin America and elsewhere. Therefore, measurement of the effect of unemployment on the activity on Wikipedia in those countries would be distorted by contributions from e.g. Latin America.

Table 7 shows the final set of Wikipedia language versions used in this paper. The share of language speakers who live in the corresponding country of origin varies from 50 percent to 99 percent (see column 1). As a substitute for the Spanish Wikipedia, we add the Catalan version,

which is also actively promoted by the Catalan population. We excluded another Spanish region, the Basque Country, because of the elevated activity of automated scripts, “bots”, in the Basque Wikipedia. According to the Wikimedia Foundation, 75% of all edits and 50% of all new articles in the Basque Wikipedia are made by bots. Bots are active in other Wikipedia editions as well, but not on such a high level¹. The final sample consists of 22 Wikipedia language editions. In addition to the largest European Wikipedias, we included the small Wikipedias from Iceland and Ireland, which are both countries that were heavily affected by the European economic crisis. We also add the Japanese and Korean Wikipedias to have benchmark countries with rather stable economies during the last decade.

Table 1 gives an overview over the countries in the sample. It also clarifies which countries we consider *affected* by the crisis and which we considered unaffected. Countries were considered to be affected by the crisis, if they experienced a significant decrease in hours worked, an increase in unemployment, had extensive coverage in the media, and as stated in EU reports. We devised separate classifications based only on individual criteria, which all gave similar results. The onset of the crisis was defined to be the beginning of the media coverage. Focusing on hours worked rather than unemployment is advantageous, because unemployment rates are not available monthly for some countries. For Catalonia, we have used monthly data for the registered unemployed population and quarterly unemployment rates, which are combined for interpolation to get monthly unemployment. For Russia, we have only quarterly unemployment data, so we linearly interpolated the missing values for Russia as well.²

To measure contributions to Wikipedia we focus on five variables for which we have monthly statistics: (1) aggregate views per month, (2) the number of active Wikipedians with a modest number of monthly edits ranging from 5 to 100, (3) the number of active Wikipedians with more than 100 monthly edits, (4) edits per article, and (5) the content growth of a corresponding language edition of Wikipedia in terms of words. Having several measures for contributions allows to analyze several possible effects of unemployment. If some unemployed people become editors in order to assume a new role as a substitute for the work in the labor force or in order to acquire new skills (text editing, studies in some area), the number of Wikipedians would increase with economic crises.

¹<http://stats.wikimedia.org/EN/BotActivityMatrixCreates.htm>

²For Catalonia: <http://www.idescat.cat/economia/inec?tc=5&id=0607&dt=201405> and for Russia http://www.gks.ru/wps/wcm/connect/rosstat_main/rosstat/en/main/

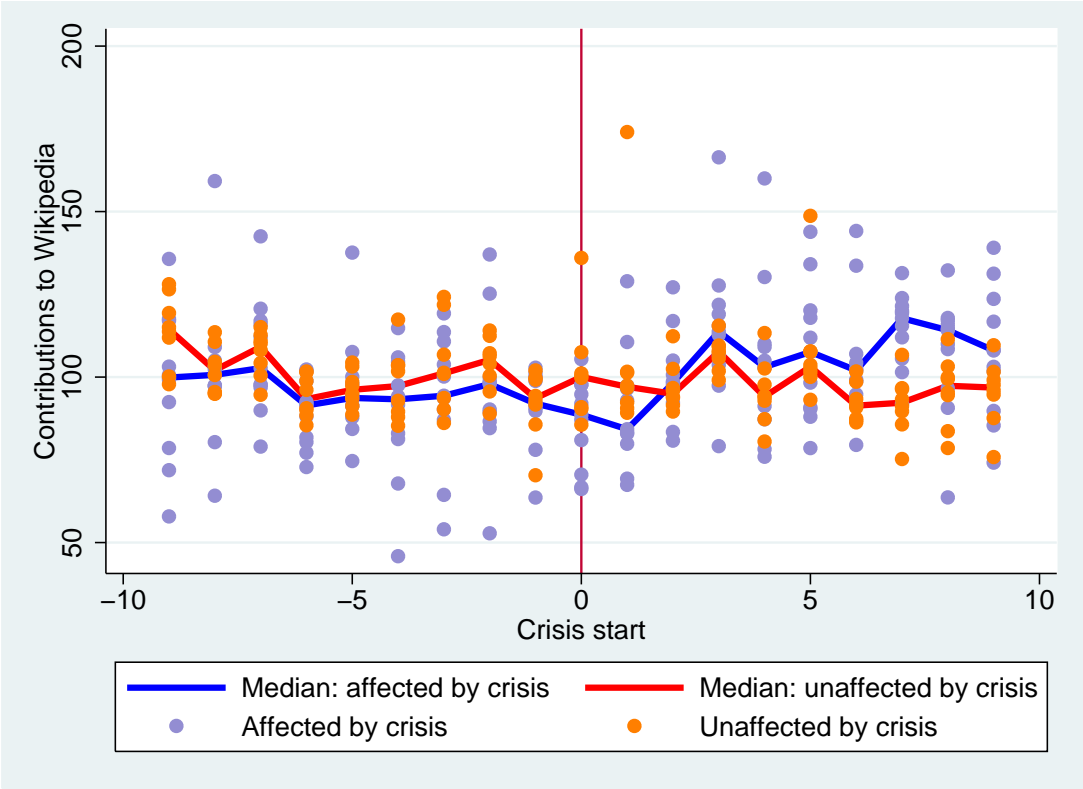
Table 1: Crisis Indicators: unemployment rates and the difference between them (%)

	Affected by crisis	Crisis start	Unemp.rate,%	Change in Unempl.,%
Bulgarian	yes	Oct 2008	6	3
Catalan	yes	Sept 2008	11	10
Czech	yes	Oct 2008	5	4
Danish	no	.	4	4
Dutch	no	.	3	1
Finnish	no	.	7	6
German	no	.	8	2
Greek	yes	June 2009	9	5
Hungarian	yes	March 2009	9	4
Icelandic	yes	Oct 2008	5	10
Italian	yes	May 2009	8	4
Japanese	no	.	4	2
Norwegian	no	.	3	1
Polish	no	.	8	2
Romanian	yes	Oct 2008	6	2
Russian	yes	Oct 2008	7	4
Slovene	yes	Oct 2008	5	2
Swedish	no	.	7	5
Turkish	yes	Oct 2008	11	6

NOTES: This Table shows, how countries unemployment rates were affected during the crisis. Affected countries were identified either by a sharp increase in unemployment or a decrease in the hours worked in the economy.

Further, if already existing contributors lose their job, they might spend more time online, which could lead to an increase in the number of active Wikipedians. The size of the database, the amount of words, and mean edits per article could result from both of these effects, an increasing number of casual Editors or an increasing number of active Wikipedians. Note that the data only contain activities of registered users, so anonymous users' edits and article creations are disregarded in this section of the paper, but they will play a prominent role when using data on German districts.

Figure 1: Monthly Development of Words contributed



NOTES: The figure shows monthly content growth measured in words added. The median across the 22 Wikipedias in our sample is shown as line. The other Wikipedias are shown as scatters. Groups of affected and unaffected countries are shown separately, and the time spans 9 months before and after the crisis.

The data for the control variable for using Internet is provided by the World Bank and states the share of people (per 100 persons) who access Internet at home, via any device type and connection. To obtain monthly values, we linearly extrapolated the yearly data. This is an important control variable here, because language versions of Wikipedia differ not only in their size of articles and editors but they are also at different stages of development. This fact might be due to country-

specific factors (Rask (2008)) or due to technological factors like Internet penetration.

Figure 1 gives a descriptive account of one of the key outcomes in Wikipedia: monthly growth measured in words added. The median across the 22 Wikipedias in our sample is shown as line and displayed together with the other Wikipedias, which are shown as scatter plot. The groups of affected and unaffected countries are shown separately 9 months before and after the crisis. The graph shows that the growth of the Wikipedias was developing on a very similar dynamic with a very light decrease in monthly contributions. Before the crisis affected (to be) countries grew slightly slower than the unaffected countries, whereas after the crisis content growth in the affected countries was slightly faster than in unaffected ones. More outcomes are presented in the Appendix in figure 3. The patterns are similar for views, edits/article and active Wikipedians, but not for casual editors. For this variable we see a difference in the trends, that must be accounted for in the regression analysis.

3.2 German District Level Data

For the analysis at the level of German districts, we combine economic indicators on the German district level with data on contributions to German Wikipedia.

Using a large data set which is available online and contains the revision history of all articles from German Wikipedia, we aggregate individual contributions to German Wikipedia to compute total contributions by districts. For that, we mapped the IP-addresses associated to edits to the corresponding German districts. However, due to the specifics of data storage on Wikipedia, publicly available data contain IP-addresses of edits only for those contributors who skipped the log-in procedure before saving their contributions, i.e. only for anonymous contributions. Therefore, our measures of contributions to Wikipedia at the district level account only for anonymous activity, which could be considered occasional and relatively small in terms of content generated. In terms of overall editing activity on German Wikipedia, anonymous edits represent 15% of all edits during the period of our analysis, years 2008-2009.

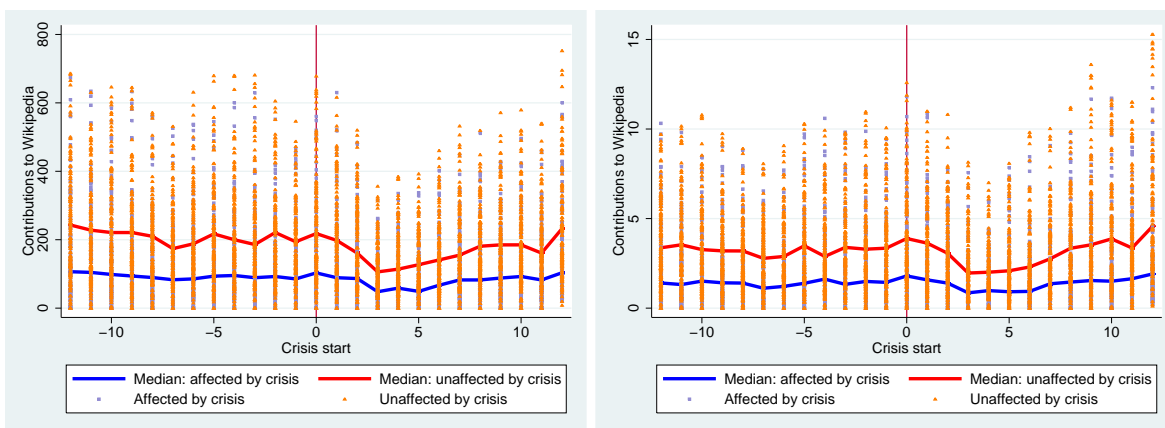
For German districts, we choose January 2009 as the moment when the crisis becomes significant for the German economy. In this month, the German government announced the necessity to combat the crisis and suggested a new policy measure, the extension of the pre-existing “Kurzarbeit” program (temporary part-time). As a result, German companies were subsidized if they decided to

keep their employees during the crisis. This was achieved by reducing the employees' working time while largely maintaining their monthly remuneration (subsidy and part-time wage combined).

The districts, relatively affected and unaffected by the economic crisis in Germany, are defined based on changes in their unemployment rate after the crisis. To have sufficient variation between affected and unaffected districts, we define the 33% of districts with the largest changes in unemployment as affected by the crisis. The 33% of districts with the lowest and even, in some instance, negative changes are defined as unaffected and used as control group for our estimation. The shares of affected districts per German State (Land) are displayed in Table 9. Surprisingly, the highest shares of affected districts can be observed in traditionally economically strong industrial German states, such as Bavaria or Baden-Wuerttemberg. Also weaker states like Thuringia had a large share of affected districts.

Figure 2 shows the evolution of contributions in the districts over a 24-month interval, before and after the economic shock affected the economy. In the absence of shocks, contributions in both groups of districts follow the same trend. After the shock, contributions in the unaffected districts drop while in the affected districts the reduction is less severe.

Figure 2: Development of Main Outcomes on the District Level



NOTES: The figure shows the median values of the dependent variable for the groups of affected and unaffected districts 6 months before and after the crisis. The left panel shows the development of the *number* of edits), while the right panel shows the *total length* of edits over time.

4 Empirical analysis

Our goal is to examine whether the sudden increase in available time that results from the loss of jobs triggers contributions to Wikipedia. The relationship between unemployment and contributions to Wikipedia is analyzed in two frameworks.

In our main specification (cf. section 4.1) we rely on a difference-in-difference approach. The economic crisis is used as a source of exogenous variation to available time in the economic system, and we compare content generation in affected countries to content generation in unaffected countries. In a later section we repeat this strategy comparing German districts where the crisis was felt stronger than in other districts.

To test our specification, we estimate whether the unemployment rate is related to the online knowledge generation also *before* the onset of the crisis. This is done using simple OLS-Regressions which we describe in the second subsection.

4.1 Difference in Difference

In our main specification we use a difference-in-differences (DiD) approach. The first difference compares content generation before and after the shock, and the second difference compares content generation in affected countries to content generation in unaffected countries, which did not experience large variations in unemployment. Using this strategy potentially allows us to measure the impact of additional spare time on contributions to Wikipedia over a given time interval, while controlling for all other possible sources of influence. The central assumption we need to make for the DiD, is that the changes in the readership and contribution activity are indeed due to the crisis and not due to some other underlying factors which correlate with the timing of the crisis.

The difference-in-difference regression is:

$$Contributions_{it} = \beta After_t + \gamma (After_t \times Affected_i) + \mu_i + \nu_t + \epsilon_{it} \quad (1)$$

$After_t$ and $Affected_i$ are dummy variables. $Affected_i$ separates the countries that were affected by the economic crises from the unaffected ones. $After_t$ equals one if the time period is after t_0 . As the variable $Affected_i$ does not vary over time, it drops out in the fixed-effects specification. The coefficient of interest is γ for the crossterm of these two dummies, which measures the difference-

in-differences.

The dependent variable $Contributions_{it}$ measures contributions to Wikipedia as captured by several variables. These are readership (Wikipedia article views in column (1)) and four indicators of contributions: in column (2) the number of active Wikipedians with about 5 to 100 monthly edits and in column (3) with more than 100 edits. Furthermore we analyze (4) edits per article and (5) the total number of words in the database. We run the DiD regression for each of these variables on the country fixed effect and the two dummy variables in the regression.

The validity of the DiD approach relies on the definition of an exogenous shock as well as the groups of treated by the shock and the control group. The groups of countries affected and unaffected by the economic crises were defined according to news in the press, reports by the European Commission or the OECD, and also the information on English Wikipedia regarding whether a country experienced economic crisis. According to these criteria, the list of countries which were affected by the crises includes Bulgaria, Catalonia, Czech Republic, Greece, Hungary, Italy, Romania, Russia, Slovakia, Turkey, Ukraine. The countries with stronger economy, including Denmark, Netherlands, Finland, Germany, Japan, Norway, Poland and Sweden were relatively not affected by the economic crises or took measures to prevent drastic increases of unemployment.

The timing of the shock, specifically the onset of the crisis for affected countries is defined as the month when they were hit by the crisis. We can retrieve this moment in time in two ways. First, when the European economic crisis after September 2008 hit a given country, this event was widely discussed in the media. Therefore, we combined information on the media with that on English Wikipedia regarding whether a country experienced economic crisis. Alternatively, we select the onset of the crisis based on changes in the monthly numbers of hours worked in the economy. The time of (empty) treatment in unaffected countries is set to September, 2008 which is the time when the US and the European economic crises were widely discussed in the media. In a robustness check, we set this date one year later. This does not influence our findings.

4.2 OLS Regression

In the specification test we check that there is no positive correlation between unemployment and contributions already before the onset of the crisis. For these models we rely on fixed effects OLS-regressions, which, like in the DiD, analyze readership (Wikipedia article views in column

(1)) and four indicators of contributions to Wikipedia (columns (2)-(5)). We regress each of these variables on the unemployment rates and hours worked in the countries, corresponding to Wikipedia language editions. The regression equation then is given by:

$$Contributions_{it} = \beta Unemployment\ Rate_{it} + \gamma Controls_{it} + \mu_i + \nu_t + \epsilon_{it},$$

where i stands for the Wikipedia language edition, t is the month and γ is a scalar of parameters, each corresponding to a control variable. Year and month dummies, as well as fixed effects are included to rule out time trend effects and individual unobserved heterogeneity. Internet penetration is included to control for the population's access to Internet which varies strongly among European countries.

Note, that the scope of the OLS regression is limited, as it can only indicate the presence of correlation between additional spare time due to unemployment and contributions to Wikipedia. Yet, it is important to check that this relationship did not exist before the crisis, since crisis is believed to hit countries with a weaker economy more likely.

5 Results

In what follows we discuss our baseline results. For these tables the onset of the economic crises is defined based on a sharp decrease in hours worked and major media events. For the estimation, we take a twelve month interval (Table 2) and a 24-month interval (Table 3). Twelve months cover six months before and six months after the onset of the crisis and the 24-month interval covers twelve months before and after. In the countries affected by the economic crises the readership of Wikipedia in the corresponding language (Views) increases within six months after the crisis by 14.5 per cent (Table 2). Similarly, we observe here a 11 per cent increase in the number of editors who contribute regularly but with only 5-100 monthly contributions. While there is a spike in readership of Wikipedia after more people get unemployed, some of these readers (a small share) become contributors of knowledge to Wikipedia. It is worth mentioning that any person with Internet access can become a contributor to Wikipedia as any kind of contributions is possible, from correcting typos and rephrasing sentences to adding a line to the topic related closely to the person, which can be a town of residence as well as a local tradition.

Table 2: DID Regression for the period of 6 months before and 6 months after the crisis

	(1)	(2)	(3)	(4)	(5)
	Views	Active 5-100e.	Active \geq 100e.	Edits p.article	Words growth
After crisis	-11.05 (9.390)	-8.134 (4.801)	0.206 (3.592)	-3.399* (1.686)	2.090 (6.944)
Affected after crisis	14.49** (6.448)	10.87* (5.280)	2.059 (4.074)	3.665** (1.582)	5.275 (8.870)
Time trend	1.191 (0.850)	0.790* (0.393)	0.490 (0.382)	1.348*** (0.166)	0.453 (0.580)
Constant	109.8*** (2.415)	107.1*** (4.420)	102.0*** (1.904)	94.61*** (0.773)	104.8*** (4.865)
Month Dummies	yes	yes	yes	yes	yes
Quarter Dummies	yes	yes	yes	yes	yes
Year Dummies	yes	yes	yes	yes	yes
Observations	247	247	247	247	247

NOTES: The table contains different measures of contributions to Wikipedia in each column: (1) views of Wikipedia, (2) the number of active Wikipedians (with at least 5 edits), (3) the number of very active Wikipedians (with more than 100 edits), (4) the average number of edits per article, (5) the new words added. All indicators of contributions to Wikipedia are normalized such that the mean value of the variable across all periods is considered 100%. The rest of the monthly values are computed as a percentage of this value. The variable of interest, which represents the treatment effect, *affected_after*, is an interaction term between dummies for the countries that are affected by the crisis with the time dummy indicating the period after the crisis. All specifications include month and year dummies, and a common time trend. Standard errors, clustered by countries, are in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

The results for the 24-month time interval (in Table 3) show that this initial increase in readership and the number of editors is sustainable and converts into a sustainable increase in content generation. The relative increase in readership stays at 13 % and the increase in regular editors with modest amount of activity becomes stronger, with 14.7 per cent. However, over the horizon of a full year, we also see an increase in very active editors (≥ 100 edits; +11.4%). We also find an increase in both edits/article (+6.1%) and overall number of words added to the encyclopedia (+13.4%). These results are in line with a mechanism whereby the contributions to a public good increase as a result of more time spent while reading information online. At least some people, after initially consuming the content, transform themselves from consumers into contributors of online knowledge.

5.1 Specification Tests

We ran several tests to check the validity of our specification. Most importantly we verify that the unemployment rate is not positively correlated with contributions already *before* the crisis. This is important, because the crisis is likely going to hit weaker economies harder. Hence, if

Table 3: DID Regression for the period of 12 months before and 12 months after the crisis

	(1)	(2)	(3)	(4)	(5)
	Views	Active 5-100e.	Active \leq 100e.	Edits p.article	Words growth
After crisis	-7.942 (5.705)	-6.634 (4.036)	-3.525 (2.870)	-3.458* (1.822)	-5.758 (5.127)
Affected after crisis	13.15* (7.006)	14.70*** (4.571)	11.42** (4.386)	6.181** (2.688)	13.43** (5.706)
Time trend	0.207 (0.424)	0.159 (0.186)	0.130 (0.174)	1.129*** (0.0872)	0.287 (0.422)
Constant	108.3*** (3.972)	88.25*** (3.326)	96.81*** (2.956)	84.70*** (1.652)	101.0*** (4.728)
Month Dummies	yes	yes	yes	yes	yes
Quarter Dummies	yes	yes	yes	yes	yes
Year Dummies	yes	yes	yes	yes	yes
Observations	429	475	475	475	475

NOTES: The table contains different measures of contributions to Wikipedia in each column: (1) views of Wikipedia, (2) the number of active Wikipedians (with at least 5 edits), (3) the number of very active Wikipedians (with more than 100 edits), (4) the average number of edits per article, (5) the new words added. All indicators of contributions to Wikipedia are normalized such that the mean value of the variable across all periods is considered 100%. The rest of the monthly values are computed as a percentage of this value. The variable of interest, which represents the treatment effect, *affected_after*, is an interaction term between dummies for the countries that are affected by the crisis with the time dummy indicating the period after the crisis. All specifications include month and year dummies, and a common time trend. Standard errors, clustered by countries, are in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

contributions to Wikipedia were correlated to unemployment before the crisis, then we could not exploit economic crises to study how an increase in unemployment. We would simply capture this preexisting correlation and erroneously attribute it to the crisis.

Hence we run an OLS regression of contributions on unemployment 12 months prior to the crisis. The results are shown in Table 8, which contains different measures of contributions to Wikipedia in each column: (1) views of Wikipedia, (2) the number of active Wikipedians (with at least 5 edits), (3) the number of very active Wikipedians (with more than 100 edits), (4) the average number of edits per article, (5) the new words added. All indicators of contributions to Wikipedia are normalized such that the mean value of the variable across all periods is considered 100%. The other values are normalized as a percentage of this value. The independent variable of interest, *Unemployment*, is the normalized monthly unemployment rate. All specifications include month, quarter and year dummies, and a common time trend.

The coefficient of the variable of interest, the unemployment rate, is not statistically significantly different from zero for most outcomes. It is positive (and not significant) only for active editors with

more than 100 monthly edits. For all other dependent variables the point estimates are negative and only for views this relationship is statistically significant. We consider this to be no evidence of correlation between unemployment and contributions to Wikipedia before the shock

Next we check whether the channel through additional views can indeed be responsible for the subsequent increase in the number of editors and content growth. Table 4 shows the OLS-Results when regressing views and activity on Wikipedia over the 24-month period, twelve months before and after the onset of crisis. The table shows the relationship between views of Wikipedia and different measures of contributions to Wikipedia in each column: (1) the number of active Wikipedians (with at least 5 edits), (2) the number of very active Wikipedians (with more than 100 edits), (3) the average number of edits per article, (4) the new words added. The independent variable of interest, *Views*, is the normalized with respect to the monthly average, and standard errors are clustered by countries.

Table 4: OLS Regression for the period of 12 months before and 12 months after the crisis

	(1)	(2)	(3)	(4)
	Active 5-100e.	Active more 100e.	Edits p.article	Words growth
Views	0.305** (0.136)	0.257** (0.0909)	0.0307 (0.0308)	0.237*** (0.0612)
Time trend	0.0185 (0.0638)	0.275*** (0.0566)	0.926*** (0.0580)	0.392 (0.295)
Constant	66.73*** (13.74)	69.99*** (9.140)	89.75*** (3.050)	76.47*** (7.422)
Month Dummies	yes	yes	yes	yes
Year Dummies	yes	yes	yes	yes
Observations	429	429	429	429

NOTES: The table shows the relationship between views of Wikipedia and different measures of contributions to Wikipedia in each column: (1) the number of active Wikipedians (with at least 5 edits), (2) the number of very active Wikipedians (with more than 100 edits), (3) the average number of edits per article, (4) the new words added. All indicators of contributions to Wikipedia are normalized such that the mean value of the variable across all periods is considered 100%. The rest of the monthly values are computed as a percentage of this value. The independent variable of interest, *Views*, is the normalized with respect to the monthly average. All specifications include month, quarter and year dummies, and a common time trend. Standard errors, clustered by countries in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

The results in Table 4 confirm that views are a crucial predictor for all edit related outcomes except the number of edits/article. One percent more views is associated with more active editors (+0.3% active and +0.25% very active editors) and more content growth as measured in words. Note, though that these OLS results cannot be interpreted causally, though. Exploring an exoge-

nous variation in the amount of spare time in the economy allows us to get a more reliable causal effect. Another general problem in the regressions is the small number of countries that can be used in the regression. As a result we only have a small number of panel observations, which are also very heterogeneous. In order to remedy this problem, we study data from the German district level in the next section.

Next, we performed a set of robustness checks to ensure that the moment of the economic crisis that we define for the unaffected countries does not affect the results, while the moment of the crisis for the affected countries should matter. Setting up the crisis moment as September, 2009 indeed does not affect our baseline results. Then, defining a placebo treatment, for example, setting the economic crisis for all countries at September, 2007 vanishes our baseline results, suggesting that the found effects are indeed due to the change in spare time of individuals due to unemployment.

Finally we checked the robustness of our OLS approach by using the figures of unemployment among young people (15-24 years old) as an explanatory variable. One would expect, young people are more likely to use Internet and, consequently, to contribute to online public goods than the elder generations. The results suggest that indeed the magnitude and significance of the unemployment effect is indeed larger for youth unemployment. These regressions are very similar to the results for German district level data and hence, for reasons of space, these results are only shown in the next section.

6 Empirical Analysis of the Regional Level

In addition to the analysis at the country-level, we examine our results in a framework where treated and untreated subjects are more similar than countries. To achieve a similar institutional context we conduct an analysis at the level of German administrative districts, called “Kreise.”

We perform the same econometric estimation as in equation (1) where contribution of district i in month t is measured by the total number of edits or the sum of bytes generated by anonymous contributors from the IP-addresses belonging to district i . Again, month, quarter and year dummies, as well as fixed effects are included to rule out time trends and district unobserved heterogeneity.

Germany is treated as a country less affected by the crisis in our country-level analysis, therefore, we expect that the general trend of contributions will be negative. However, in districts with higher unemployment the negative trend could be mitigated due to more time spent online.

6.1 Results at the regional level

The estimation results for German districts are shown in Table 5. They support our hypothesis that, while contributions to German Wikipedia fall after the crisis, in districts with higher unemployment additional activity mitigates the negative overall trend.

Table 5: DID Regression for German Districts (6 months before and after the crisis)

	(1)	(2)
	# Edits (norm.)	Contribution (KB) (norm.)
After	-32.62*** (2.954)	-16.53*** (4.286)
Affected after crisis	6.188** (2.515)	5.327 (3.913)
Constant	104.4*** (1.600)	110.2*** (2.886)
Year dummies	Yes	Yes
Quarter dummies	Yes	Yes
Month dummies	Yes	Yes
Observations	3341	3341

NOTES: The table contains different measures of contributions to Wikipedia in each column: (1) the number of revisions, (2) the total contribution length in kb. The variable of interest, which represents the treatment effect, *Affected after crisis*, is an interaction term between dummies for the districts that are affected by the crisis with the time dummy indicating the period after the crisis. All specifications include month and year dummies. Standard errors, clustered by districts, in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

The table shows the results for our main measures of contributions to Wikipedia in each respective column: (1) the number of revisions, (2) the total contribution length in kb. Since these data were self collected and are based only on anonymous contributions of not (yet) registered users, we can only collect data on these two measures. As before, the coefficient of interest, which measures the treatment effect, belongs to the crossterm *Affected after crisis*. It is the interaction term between dummies for the districts that are affected by the crisis with the time dummy indicating the period after the crisis.

The results for a 12-month interval suggest that contributions fall after the crisis start in January 2009 in Germany by 32.6 per cent (the number of edits) and by 16.5 per cent (the amount of bytes generated). However, in districts affected by a rise in unemployment this trend is mitigated by additional activity on Wikipedia, with the difference in differences before and after the crisis being 6.2 per cent of edits. The DiD coefficient for edit length is of similar size but insignificant.

In addition to our baseline specification at the district level, we perform a robustness check to see whether the effect we find is indeed channeled by unemployed who are spending more time online. We redefine our affected districts based on youth unemployment, the rate of unemployed individuals between the ages of 18 and 23. We expect them to get involved in online volunteering more easily. The results in Table 6 support our finding, however the effects are very similar to the previous results. While contributions to Wikipedia decrease overall, this decrease is 6.1 per cent lower in districts with higher youth unemployment.

Our data set at the district level contains the time of each edit. We used it to calculate total amounts of edits made during the working time, from Monday to Friday in the interval from 9 a.m. to 6 p.m., and in the remaining leisure hours including the weekend. Table 12 presents the results of this analysis. The coefficient is a little bit larger for both outcomes and remains significant only for the number of edits. An analogous DiD regression for contributions during working hours does not show an effect. Overall we find evidence that our finding on the positive effect of unemployment on contributions to Wikipedia is driven by edits made during leisure time, a finding which clearly offers many opportunities for interesting further research.

Table 6: DID Regression with robust estimator of variance (before and after 6 months)

	(1)	(2)
	# Edits (norm.)	Contribution (KB) (norm.)
After	-34.12*** (2.867)	-22.45*** (4.277)
Affected after crisis	6.116** (2.553)	5.936 (3.991)
Constant	102.9*** (1.508)	109.0*** (2.969)
Year dummies	Yes	Yes
Quarter dummies	Yes	Yes
Month dummies	Yes	Yes
Observations	3315	3315

NOTES: The table contains different measures of contributions to Wikipedia in each column: 1) the number of revisions, (2) the total contribution length in kb. The variable of interest, which represents the treatment effect, *Affected after crisis*, is an interaction term between dummies for the districts that are affected by the crisis with the time dummy indicating the period after the crisis. All specifications include month, quarter and year dummies. Standard errors, clustered by districts, in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

7 Discussion, Limitations and Further Research

We analyzed the relationship of unemployment and public goods provision online. In a nutshell, we find that increased unemployment is associated with higher participation of volunteers in Wikipedia and an increased rate of content generation. Articles are read more frequently and the number of highly active users increases, suggesting that also existing editors increase their activity. We also find robust evidence that the number of edits per article increases, and slightly weaker support for an increased overall content growth. We find the overall effect to be rather positive than negative, which is reassuring news if the encyclopedia functions as an important knowledge base for the economy. A negative overall effect would be an alarming side effect of the crisis after all. However, our findings open up a large array of further questions.

On the one hand they suggest that higher unemployment may be associated with greater volunteering activity and productive time usage. Yet we cannot fully answer how this mechanism works exactly. Particularly, it seems that new editors begin to acquire new capabilities and devote their time to producing public goods. While we observe overall content growth, we could not find robust

evidence for an increase in the number of *new* articles per day (not shown). This suggests that the increased participation is focused on adding to the existing knowledge, rather than providing new one. Doing so requires less experience than providing new one or than creating new articles, which may be interpreted as a sign of learning by the new contributors.

The question whether unemployment can result in increased provision of public (online) goods and private learning is crucial, given that we observe accelerating labor substitution due to digitization. Especially if a part of the liberated capacity can result in increased knowledge documentation and generation, this may be a positive surprise. On the other hand, if additional time due to unemployment is predominantly wasted it might point to the need of a more active management of these resources.

While we are able to test our hypotheses from several angles and to show the robustness of our findings some limitations cannot easily be overcome. Most importantly, we exploit the economic crisis as source of exogenous variation in the economic state and the unemployment rates. This is based on the following assumptions for identification: First we require the testable assumptions, that the countries we use in the comparison would have sufficiently development over time absent any shock. This requires that the Wikipedias are sufficiently similar and that the countries are somewhat homogeneous with respect to other economic and societal developments in the period of observation. Second, we assume that the increased editing activity that we observe is due to the crisis and not due to other factors, which might again vary across countries. Despite our focus on European countries, this assumption may potentially be quite strong when looking at only a monthly interval.

To mitigate this, we verified whether we can see the effects that were found on the country level on regional district-level data on Germany. Arguably the institutional, the macroeconomic and political setup is likely to be more homogeneous when looking at German districts rather than European Countries. However, two remarks concerning this analysis. The regional analysis allowed only for a restricted set of available dependent variables, since we can only determine average properties of the regional edits, and we have to focus on the number of edits and the average length of edits in bytes. This is because using IP-addresses of anonymous contributions and hence counting edits/article or the number of experienced editors becomes meaningless. Second, the use of IP-addresses implies that we can only look at a restricted set of all contributions. These come most

likely from new or occasional users, because experienced users typically edit under their user name. This is both a limitation and a blessing at a same time, since the restricted group of inexperienced editors is very interesting in the context of our study.

A final limitation of the German Data comes from the fact, that the unemployment rates in Germany increased much less than in Europe. In addition to the generally higher robustness of the economy, a reduced working hour scheme, called “Kurzarbeit” was heavily applied. Further research could aim at augmenting similar data set with information on how many firms in the region used “Kurzarbeit” during the crisis. This could allow to disentangle the effects of more disposable time vs. the effect of increased unemployment.

In addition to our assumptions we have to point out, that even if these assumptions for identification are satisfied, we can only provide indicative evidence on what drives this content generation. Is it the unemployed or the employed who contribute? This question cannot be answered at all for the country level data. For the German district level data we could provide an indication, that additional edits are not generated during working hours, but during leisure time. But does this reflect that the working editors increase activity or does it reflect a shift in behavior of editors, who previously contributed from their work and prefer not to take this risk any longer?

Another fruitful avenue for further research could investigate what is actually written. This question has to remain unanswered at the current stage of research. Maybe people simply write *about* the crisis? This seems unlikely, given the overall growth that we observe. However, smaller or larger fraction of the additional readership and content generation in the affected countries might be a direct increase in demand for economic information or the consequence of updating the encyclopedia with current events. Alternatively increased editing activity might be dedicated to improving the overall quality of articles or individual user might contribute to their favorite topic of interest, which they also find enjoyable to write about.

To answer who makes the edits we would need user and editor level data, and to see what they write we would have to analyze articles on their content level. Further research could analyze the nature of contributions and which type of articles are edited. Also to what extent district specific articles are being improved or whether articles related to affected professions This is beyond the scope of this paper and, especially on the get edited would be very computationally intensive, but might lead to interesting additional insights from further research. On the country level this is

almost unthinkable though, since the data available are too highly aggregated.

More fine grained data, on the user level ideally, would not only allow us to look at what information is being searched and which edits are made, it would also allow contrasting Wikipedia editing activity with other ways how newly unemployed use their additional time.

8 Conclusion

In this paper, we analyze the relationship between the economic downturn and the provision of online public good, i.e. online knowledge. In the times of digitization driven labor substitution it is a crucial question whether unemployed people invest their additional time in the provision of public (online) goods. We exploit the European economic crisis, beginning in 2008, as source of exogenous variation in the economic state that affected unemployment rates. In doing so, we focus on the contribution to online public goods as a side effect of an economic crisis.

Specifically, we use a difference in differences strategy around the time of the European economic crises to examine the dynamics of content generation on the most important online public good, Wikipedia. For affected countries, we find a (relative) increase in readership and contributions in the aftermath of the crisis. This increase appears to be more strongly correlated with youth unemployment, and the relationship becomes important only after the crisis.

We verify whether we can see similar effects on district-level data from Germany. By looking at German districts we can conduct a more fine grained analysis, with more units of observations in a more homogeneous institutional context. For this analysis, we located the origin of anonymous contributions to the German Wikipedia (by district) and matched them with regional information on unemployment. We find similar results as on the country level, and show that the effect is driven by additional contributions during leisure time rather than during working hours.

Taken together, these findings point to more public goods giving as a side effect of the crisis, but they cannot unequivocally be attributed to those who lose their employment.

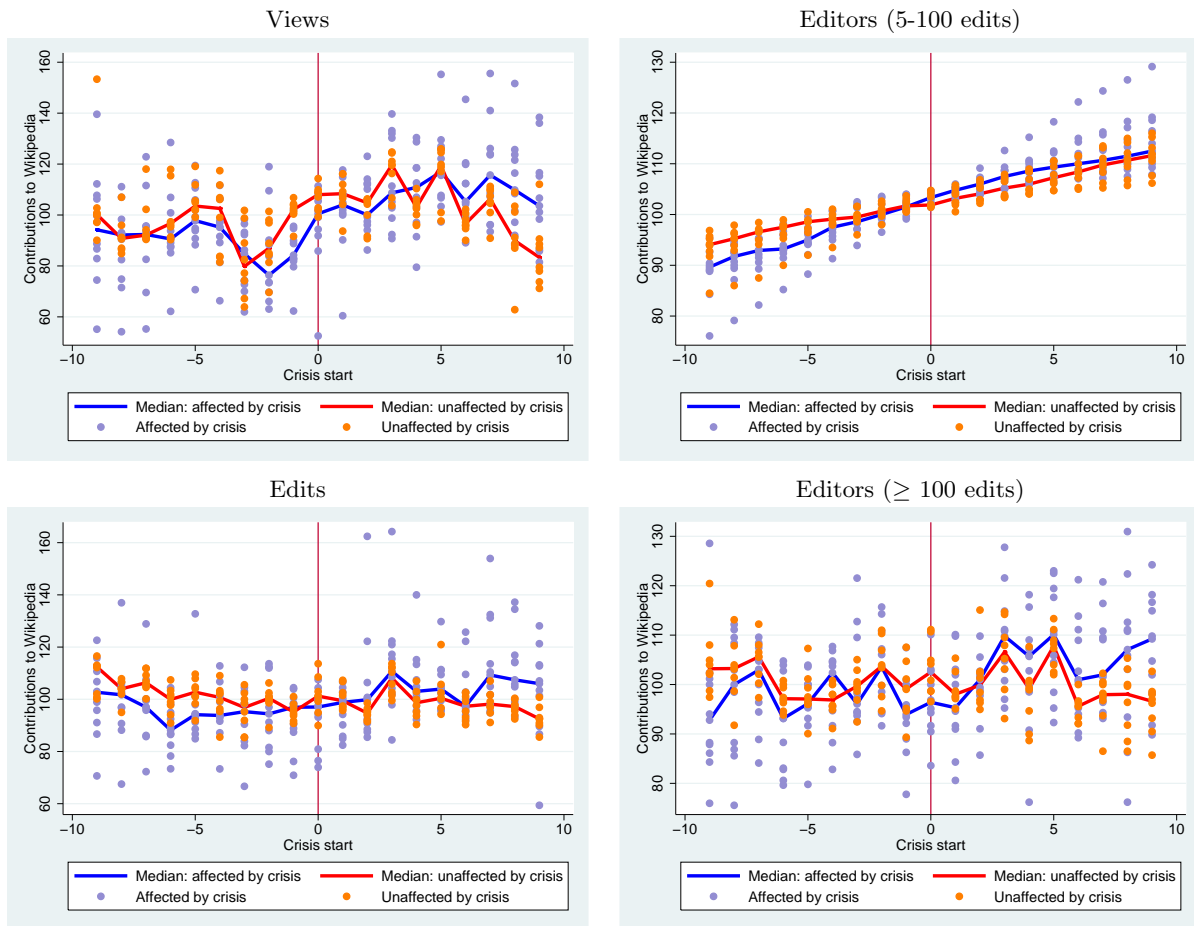
References

- Aguiar, Mark, Erik Hurst, and Loukas Karabarbounis**, “Recent developments in the economics of time use,” *Annu. Rev. Econ.*, 2012, 4 (1), 373–397.
- , —, and —, “Time use during the great recession,” *The American Economic Review*, 2013, 103 (5), 1664–1696.
- Algan, Yann, Yochai Benkler, Mayo Fuster Morell, and Jérôme Hergueux**, “Cooperation in a Peer Production Economy Experimental Evidence from Wikipedia,” in “Workshop on Information Systems and Economics, Milan, Italy” 2013, pp. 1–31.
- Andreoni, James**, “Privately provided public goods in a large economy: the limits of altruism,” *Journal of Public Economics*, 1988, 35 (1), 57–73.
- , “Giving with impure altruism: applications to charity and Ricardian equivalence,” *The Journal of Political Economy*, 1989, pp. 1447–1458.
- , “Impure altruism and donations to public goods: a theory of warm-glow giving,” *The economic journal*, 1990, pp. 464–477.
- , “Giving gifts to groups: How altruism depends on the number of recipients,” *Journal of Public Economics*, 2007, 91 (9), 1731–1749.
- Burda, Michael C and Daniel S Hamermesh**, “Unemployment, market work and household production,” *Economics Letters*, 2010, 107 (2), 131–133.
- Comino, Stefano, Fabio M Manenti, and Maria Laura Parisi**, “From planning to mature: On the success of open source projects,” *Research Policy*, 2007, 36 (10), 1575–1586.
- Freeman, Kassie**, “Increasing African Americans’ participation in higher education: African American high-school students’ perspectives,” *Journal of Higher Education*, 1997, pp. 523–550.
- Goldfarb, Avi and Jeff Prince**, “Internet adoption and usage patterns are different: Implications for the digital divide,” *Information Economics and Policy*, 2008, 20 (1), 2–15.
- Hess, Charlotte and Elinor Ostrom**, “Ideas, artifacts, and facilities: information as a common-pool resource,” *Law and contemporary problems*, 2003, pp. 111–145.
- Kandel, Eugene and Edward P Lazear**, “Peer pressure and partnerships,” *Journal of political Economy*, 1992, pp. 801–817.
- Knabe, Andreas, Steffen Rätzel, Ronnie Schöb, and Joachim Weimann**, “Dissatisfied with Life but Having a Good Day: Time-use and Well-being of the Unemployed*,” *The Economic Journal*, 2010, 120 (547), 867–889.
- Krueger, Alan B and Andreas I Mueller**, “Time use, emotional well-being, and unemployment: Evidence from longitudinal data,” *The American Economic Review*, 2012, 102 (3), 594–599.
- Rask, Morten**, “The reach and richness of Wikipedia: Is Wikinomics only for rich countries?,” *First Monday*, 2008, 13 (6).
- Taniguchi, Hiromi**, “Mens and Womens Volunteering: Gender Differences in the Effects of Employment and Family Characteristics,” *Nonprofit and Voluntary Sector Quarterly*, 2006, 35 (1), 83–101.
- Uslaner, Eric M**, “Religion and civic engagement in Canada and the United States,” *Journal for the Scientific Study of Religion*, 2002, 41 (2), 239–254.
- Verick, Sher and Iyanatul Islam**, “The great recession of 2008-2009: causes, consequences and policy responses,” Technical Report, Discussion paper series//Forschungsinstitut zur Zukunft der Arbeit 2010.
- Vesterlund, Lise**, “Why do people give,” *The nonprofit sector: A research handbook*, 2006, 2, 168–190.
- Wallsten, Scott**, “What are we not doing when we’re online,” Technical Report, National Bureau of Economic Research 2013.
- Wang, Chengsi and Galina Zudenkova**, “Non-Monotonic Group-Size Effect in Repeated Provision of Public Goods,” Available at SSRN 2378866, 2014.
- Zhang, Xiaoquan and Feng Zhu**, “Group size and incentives to contribute: A natural experiment at Chinese Wikipedia,” *American Economic Review*, Forthcoming, 2010, pp. 07–22.

A Appendix

A.1 Country Level Analysis

Figure 3: Comparison of both groups in one plot.



NOTES: This figure illustrates a comparison of the trends in the key outcomes before and after the respective European economic crises. The dependent variables are the views/language, and edits/article on the left and active Wikipedians (5 to 100 and more than 100 on the right.) Median values of the dependent variable for the groups of affected and unaffected countries 9 months before and after the crisis

Table 7: Wikipedia key variables within the period of 12 months before and 12 months after the crisis

	Language speakers (m)	In main country, %	Views per speaker	Wikipedians, %	Active 5-100 edits, %	Active > 100 edits, %
Bulgarian	8.16	86.05	2	0.02	10.7	3.5
Catalan	4.08	.	3	0.06	13.3	3.9
Czech	10.62	97.93	4	0.04	12.3	2.8
Danish	5.52	97.42	3	0.06	9.9	2.3
Dutch	21.94	71.54	6	0.06	9.6	2.2
Finnish	5.39	94.58	10	0.13	9.7	2.3
German	78.25	89.21	12	0.10	9.0	1.4
Greek	13.43	79.65	1	0.02	6.0	1.8
Hungarian	12.61	78.06	2	0.04	12.8	3.3
Icelandic	0.24	94.32	9	0.16	11.8	5.4
Italian	63.66	90.64	5	0.04	10.2	2.2
Japanese	122.06	99.13	8	0.03	12.7	1.5
Norwegian	4.74	97.85	6	0.13	9.7	2.1
Polish	38.66	94.66	8	0.04	10.8	2.3
Romanian	23.78	83.67	1	0.01	12.1	2.9
Russian	167.33	81.87	1	0.01	17.0	3.4
Slovene	2.09	91.60	4	0.07	12.6	2.8
Swedish	9.20	96.12	7	0.09	11.0	2.4
Turkish	70.81	93.92	1	0.01	10.6	1.9
Total	34.87	89.90	5	0.06	11.2	2.7

Columns (3)-(6) are means of the interval 12 months before to 12 months after crisis

Sources: *stats.wikimedia.org*

A.1.1 Specification Tests and Alternative Specifications

Table 8: OLS Regression for the period of 12 months before the crisis

	(1)	(2)	(3)	(4)	(5)
	Views	Active 5-100e	Active \geq 100e	Edits p.article	Words growth
Unemployment rate (normalized)	-0.208*** (0.0687)	-0.0910 (0.167)	-0.0249 (0.0877)	0.0103 (0.0221)	-0.178 (0.204)
Time trend	0.130 (0.339)	0.300 (0.422)	0.345 (0.289)	1.420*** (0.117)	0.0133 (0.484)
Constant	124.9*** (10.33)	98.44*** (13.07)	97.82*** (8.370)	83.44*** (2.204)	113.6*** (15.82)
Month Dummies	yes	yes	yes	yes	yes
Year Dummies	yes	yes	yes	yes	yes
Observations	182	228	228	228	228

NOTES: The table contains different measures of contributions to Wikipedia in each column: (1) views of Wikipedia, (2) the number of active Wikipedians (with at least 5 edits), (3) the number of very active Wikipedians (with more than 100 edits), (4) the average number of edits per article, (5) the new words added. All indicators of contributions to Wikipedia are normalized such that the mean value of the variable across all periods is considered 100%. The remainder of the monthly values are computed as a percentage of this value. The independent variable of interest, *Unemployment*, is the normalized monthly unemployment rate. All specifications include month, quarter and year dummies, and a common time trend. Standard errors, clustered by countries are in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

A.2 German districts

Figure 4: German administrative units on the district level



Table 9: Crisis Indicators: unemployment rates and the difference between them (%)

	Share of affected districts	Unemp.rate,%	Change in Unempl.,%
Baden-Wuerttemberg	0.63	4.69	1.13
Bavaria	0.69	4.46	1.25
Brandenburg	0.58	13.21	1.12
Bremen	0.50	13.23	0.40
Hamburg	.	8.33	0.83
Hessen	0.17	6.71	0.64
Lower Saxony	0.10	7.85	0.68
Mecklenburg-Western Pomerania	0.77	14.15	1.65
North Rhine-Westphalia	0.29	8.50	0.76
Rhineland-Palatinate	0.40	6.18	0.99
Saarland	0.00	6.62	0.86
Saxony	0.85	12.78	1.86
Saxony-Anhalt	0.80	13.72	1.34
Schleswig-Holstein	0.18	8.10	0.75
Thuringia	0.76	11.35	1.86

Table 10: Relationship between unemployment and activity on Wikipedia before the crisis

	(1)	(2)
	# Edits (norm.)	Contribution (KB) (norm.)
Unemployment	-5.680 (4.318)	-3.503 (5.979)
Year dummies	Yes	Yes
Quarter dummies	Yes	Yes
Month dummies	Yes	Yes
Observations	1541	1541

NOTES: The table contains OLS-fixed effects regressions for different measures of contributions to Wikipedia in each column: (1) the number of revisions, (2) the total contribution length in kb. The independent variable of interest is *Unemployment rate* for each district. The time range is 6 months before the onset of the economic crisis. All specifications include month and year dummies. Standard errors, clustered by districts in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

- *Before the shock*, there is *no* evidence on correlation between economic situation and contributions to Wikipedia

Table 11: Relationship between unemployment and activity on Wikipedia after the crisis

	(1)	(2)
	# Edits (norm.)	Contribution (KB) (norm.)
Unemployment	4.059* (2.254)	5.645 (3.928)
Year dummies	Yes	Yes
Quarter dummies	Yes	Yes
Month dummies	Yes	Yes
Observations	1554	1554

NOTES: The table contains OLS-fixed effects regressions for different measures of contributions to Wikipedia in each column: (1) the number of revisions, (2) the total contribution length in kb. The independent variable of interest is *Unemployment rate* for each district. The time range is 6 months after the onset of the economic crisis. All specifications include month and year dummies. Standard errors, clustered by districts in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

- *After* the shock, there *is* correlation between economic situation and contributions to Wikipedia

Table 12: DID for contributions made in leisure time of the day (6pm-9am and weekends)

	(1)	(2)
	# Edits (norm.)	Contribution (KB) (norm.)
After	-35.08*** (3.281)	-19.45*** (5.068)
Affected after crisis	7.939*** (2.707)	6.378 (4.381)
Constant	109.4*** (2.111)	105.8*** (3.048)
Year dummies	Yes	Yes
Quarter dummies	Yes	Yes
Month dummies	Yes	Yes
Observations	3377	3377

NOTES: The table contains different measures of contributions to Wikipedia in each column: 1) the number of revisions, (2) the total contribution length in kb. The variable of interest, which represents the treatment effect, *Affected after crisis*, is an interaction term between dummies for the districts that are affected by the crisis with the time dummy indicating the period after the crisis. The time range is 6 months before and after the crisis. All specifications include month, quarter and year dummies. Standard errors, clustered by districts, in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

- The increase in editing activity does not occur during working hours but in the evenings and on weekends.