

Economic Downturn and Volunteering: Do Economic Crises Affect Content Generation on Wikipedia?*

Michael Kummer, Olga Slivko, Michael Zhang[†]

June 5, 2016

Abstract

In this paper, we address the impact of the economic downturn and a consecutive increase in unemployment on online public good provision. Specifically, we ask how the European economic crisis affected voluntary contributions of content to the online encyclopedia Wikipedia. We study this question on the country-level and on the level of German districts.

First, we put together a unique monthly country-level data set, which combines country specific economic outcomes with data on contributions to Wikipedia. Our source of exogenous variation is the effect of the financial crisis in the US in September 2008, which affected the European countries with varying intensity. We find that the economic downturn is associated with more viewership, which channels higher participation of volunteers in Wikipedia expressed in editing activity and content growth. In a second step, we compare Wikipedia activity in affected and unaffected German districts, where we can track anonymous and registered edits and their timing. With higher unemployment first the number of anonymous edits increases and second the number of registered editors.

Keywords: online platform, Wikipedia, public goods, unemployment, user generated content.

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

* We benefitted from discussions with Sinan Aral, Irene Bertschek, Vivek Ghosal, Shane Greenstein, Marit Hinnoaar, Toomas Hinnoaar, Markus Mobius, Marianne Saam, Patrick Schulte, Frank Verboven Steffen Viète, Joel Waldfogel, Michael Ward. We thank participants in several presentations at ZEW and in the following conferences: 13th ICT (Mannheim), SEEK Workshop on the Digital Economy (Moncalieri), 8th ICT (ParisTech), EARIE 2015 (Munich) and the 5th SEEK Conference (Mannheim). The authors are grateful to the Wikimedia Foundation and Frédéric Schütz for access to the Wikipedia data. We are grateful to Manfred Knobloch for initial support with data processing. Ruetger Egolf, Justus Steins, and Alaina Totten provided outstanding research assistance. Financial support from SEEK 2014 is gratefully acknowledged.

[†]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 innovation and 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 rising unemployment on online knowledge generation and public goods provision. Specifically, we ask how increased unemployment rates after the European economic crisis in 2008, affected voluntary content generation to the online encyclopedia Wikipedia.

The economic downturn may affect the willingness for contributing to Wikipedia and other online public goods in contradicting ways. On the one hand, unemployed have more spare time and so spend more time online while searching for information. Similarly, employed may engage in online learning to strengthen their cv. While browsing online, individuals can learn about Wikipedia, view the articles and, potentially, some of them could start to exert some effort and begin to contribute to online knowledge. On the other hand, both employed and unemployed people may face larger economic pressure and so they could choose to dedicate their time to profit-making activities. In this case, the net effect on the public knowledge repository during economic crisis could be negative. A decrease of knowledge provision would threaten to reduce the economy's innovative potential or further aggravate the long-term effects of increased unemployment. This effect would be especially problematic if the public knowledge good is an important input to the knowledge in the economy.

To estimate the effect of interest we analyze both the European country-level and the German district-level data. The dataset for European country-level analysis 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 . While economic conditions were 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-differences approach, but we also run specification tests using fixed-effect OLS regression analysis with monthly data on unemployment in Europe.

In a second step, we analyze a different data set base on regional information on the German district level. We put together data on 402 German districts and are able to distinguish between anonymous contributions, and contributions of registered editors. Moreover we can analyze whether contributions are made during work hours or in the leisure time. We find that the economic downturn leads to higher participation of volunteers and an increase in content generated on Wikipedia. More editors participate in the production 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.

In ongoing research we aim at contrasting the effect in the contributions to Wikipedia with changes in the overall browsing patterns that followed after increases in the European unemployment rates. Moreover, we are collecting additional data to pin down regional and industry specific effects in knowledge creation, by gathering additional information on Wikipedia contributions in Germany and exploiting the category matching.

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 interests or hobbies. Yet, through their contribution to the knowledge resource, the unemployed might acquire valuable writing and computer skills. The contributions could also serve as a useful input for knowledge and innovation, by helping to increase the knowledge stock in the economy.

The remainder of the paper is structured as follows. Section 2 presents previous studies related to our 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 analysis 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

We contribute to the economic literature on public goods provision by providing an additional evidence on the private motivation for contributing to an online public good, and show how online volunteering changes with economic downturn.

The previous theoretical and empirical studies analyzed the private incentives for voluntary public goods provision from the perspective of the interplay between the free-riding incentives and social effects (Andreoni (1988), Andreoni (1989), Andreoni (1990), Andreoni (2007)). For example, 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 public goods provision but at a decreasing rate. This finding received empirical support in the context of online public goods, such as open-source software and online peer productive communities (Kandel and Lazear (1992), Comino et al. (2007), Zhang and Zhu (2011), Algan et al. (2013)). 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. Zhang and Zhu (2011) show the importance of the recipient group size for individual incentives for knowledge

provision using exogenous variation in the recipient group size on Wikipedia (a block by the Chinese government).¹

Our study explores the incentive to contribute to an online public good from an individual perspective, in addition to already studied social effects. We analyze whether due to the economic downturn individuals change their time allocation. Specifically, we use the recent economic crisis in Europe in the aftermath of the financial crisis in the US in 2008 as an exogenous shock to the time spent online. As unemployment rises during the economic crisis, people who become unemployed or (have to) reduce their working time, could be increasing their online time. This additional time could be spent browsing on the Internet, searching for information, and, to some extent, contributing content to online platforms.

Previous studies have looked at how the unemployed allocate their time, considering a range of potential beneficial and wasteful 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 be spent on leisure, they are less satisfied with life and specific activities (Knabe et al., 2010). Krueger and Mueller (2012) find that previously unemployed sharply decrease time devoted to leisure activities at the time of reemployment (by 35 per cent of the time now allocated to working). In their paper, 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 find that more than 50 per cent 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.

Our paper focuses specifically on time spent for providing online 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.

Online leisure time was shown to be a substitute for work since most of the time spent online is

¹In addition, 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).

spent on social networks, online games, email and portals (Wallsten (2013)). Moreover, young people 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 experience a sudden increase in time available, will, at least, partially reallocate it on online activities.

Economic downturn may also lead to a decrease in online content generation. For the civic public good, unemployment is shown to be negatively correlated with both religious as well as secular volunteering (Freeman (1997), Uslander (2002)). However, Uslander (2002) uses cross-sectional data from the U.S. and Canada and thus gives no information about effects of rising unemployment over time. Freeman (1997) also 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, we could observe changes in content generation during the economic downturn due to an increase in browsing or learning behavior by those individuals who still hold their jobs. In this case, the effect can go in both directions. Employed individuals can decrease contributions as their opportunity cost of time is higher in the crisis. However, if online volunteering would follow the same pattern as civic volunteering, it could experience growth. Our contribution is to shed light on these questions by analyzing how the economic crisis in Europe affected online knowledge generation, and specifically, contributions to Wikipedia.

3 Data

To study how the recent European crises affected online knowledge generation we will use data on both the country and the district levels. In the first part of the paper 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. In the second part we analyze German district level data. This is useful, since it allows us to repeat the analysis of the relationship of interest in a more detailed level. Using data from over 400 German districts, we show that the relationship of interest is also found using more units of observation from a more homogenous institutional environment.

3.1 Wikipedias on the Country/Language Level

The country level analysis is based on 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 the United Kingdom, Spain and Portugal, had to be excluded since their languages are spoken not exclusively in these countries, but all over the world. 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 (or the United States/Australia and other countries with many speakers of English).

Table 6 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 level². 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,

²<http://stats.wikimedia.org/EN/BotActivityMatrixCreates.htm>

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	no	.	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
Slovakian	no	.	10	5
Slovene	no	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.

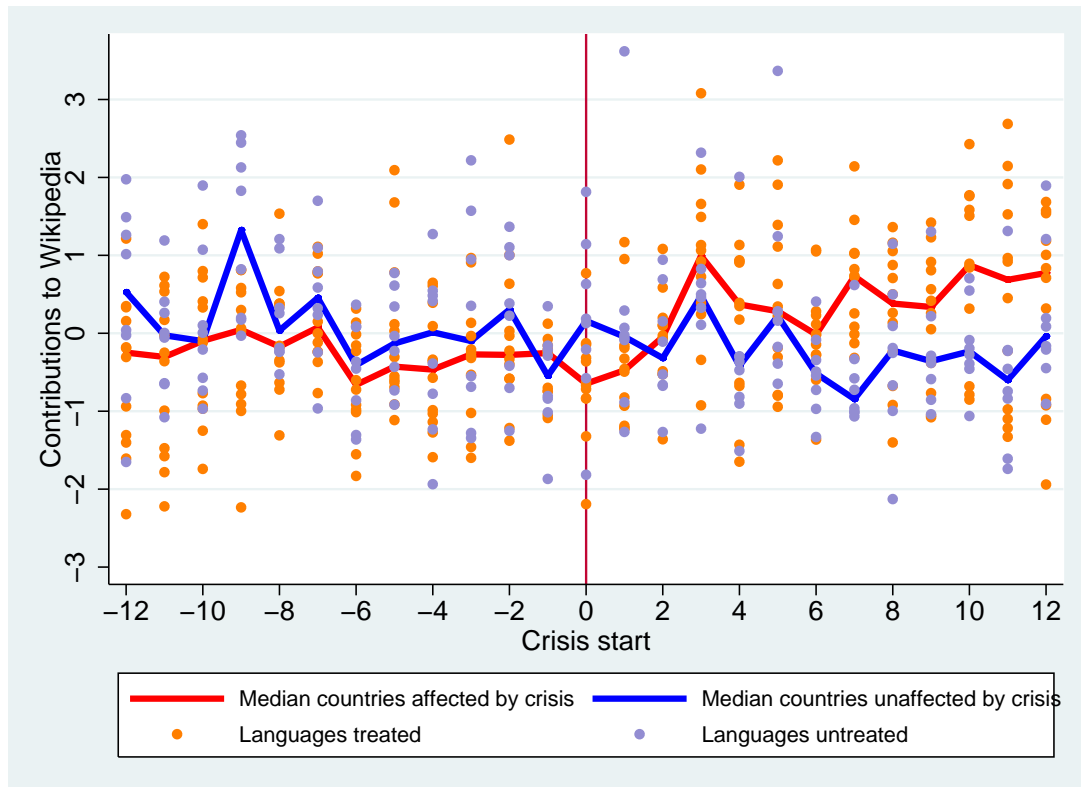
had extensive coverage in the media, and as stated in EU reports. We devised separate classifications based only on individual criteria, which all gave similar results. The onset of the crisis was defined to be the beginning of the media coverage. Focusing on hours worked rather than unemployment is 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.³

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 us 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. 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

³For 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/

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 12 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 of groups of affected and unaffected language editions of Wikipedia in our sample are displayed separately 12 months before and after the crisis together with the other Wikipedias, which are shown as scatter plot. According to the graph, the growth of the Wikipedias was developing on a very similar dynamic with a very light decrease in monthly contributions. Before the crisis, countries that would be affected grew slightly slower than the unaffected countries, whereas

after the crisis content growth in the affected countries was slightly faster than in unaffected ones. The patterns are similar for views, edits per 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.

One of the main concerns about the country level data above is the fact that the countries are quite heterogenous both culturally and economically and we can only use 19 units of observation. We address these concern by repeating the analysis on the German district level, and the data for this analysis are described in the subsequent subsection.

3.2 German District Level Data

To test if our results continue to hold when the administrative setting is more homogeneous, we repeated the analysis with data on the more than 400 German districts. This also gives us a richer data set with more units of observation.

Despite the fact that the German economy was relatively robust to the economic crisis, there was considerable variation in how different districts were affected. We exploit this variation to repeat the analysis above. 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. While we do not believe that anonymous edits are representative for all editing activity, we deem it highly relevant editing activity for our research question. Unregistered edits are typically made by very occasional or unexperienced editors and thus capture contributions of relatively new editors.

Figure 2: Development of Main Outcomes on the District Level

Number of Edits Added Text (Length)

NOTES: The figure shows the median values of the number of edits (left) and average added length per edit (right) for the groups of affected and unaffected districts 12 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.

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 8. Surprisingly, the highest shares of affected districts can be observed in traditionally economically strong industrial German states, such as Bavaria or Baden-Wuerttemberg. Weaker states like Thuringia also had a large share of affected districts.

Figure 2 shows the evolution of contributions in the districts over a 24-month interval, with 1 year before and 1 year 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 the reduction is less severe in the affected districts.

4 Empirical analysis

Our goal is to examine whether a 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-differences 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 approach for the comparison of German districts where the crisis was felt stronger than in other districts.

To test our specification, we check whether the relationship between the unemployment rate and online knowledge generation already existed *before* the onset of the crisis. This is done using simple OLS-Regressions which we describe in the second subsection.

4.1 Difference-in-Differences

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 a 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 in 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. It is important to check that this relationship did not exist before the crisis, since crises are believed to have a greater likelihood to hit countries with a weaker economy.

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 the start of economic crisis covered by major media sources. For the estimation, we take a 24-month interval (Table 12). Twenty-four months cover twelve months before and twelve months after the onset of the crisis. The results 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, (6) the number of hyperlinks set between articles within Wikipedia site and (7) references from Wikipedia articles to external sources. All indicators of contributions to Wikipedia are normalized with respect to their mean and standard deviation values such that the coefficients represent the changes in the dependent variables in standard deviations. All specifications include month and year dummies, and a common time trend.

Table 12 shows the results when we run the difference in differences analysis period by period. The table has 8 dependent variables: (i) Article Views, (ii) Active editors with 5-100edits (iii) Active editors

Table 2: DID Regression for the period of 12 months before and 12 months after the crisis

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Views	Active 5-100e.	Active more 100e.	Edits p.article	Words growth	Int. links	Ext. links
After treatment	-0.0741 (0.240)	-0.397 (0.366)	0.0558 (0.260)	0.00286 (0.0514)	-0.243 (0.266)	0.0127 (0.0188)	0.0147 (0.0385)
Treated countries after T	0.344 (0.273)	1.089*** (0.350)	0.789** (0.332)	0.0313 (0.0226)	0.699** (0.255)	-0.0247*** (0.00856)	-0.0143 (0.0156)
Time trend	-0.346*** (0.115)	0.198 (0.185)	0.296** (0.134)	0.139*** (0.0108)	0.188 (0.139)	0.137*** (0.00350)	0.125*** (0.00479)
Constant	14.19*** (4.075)	-8.016 (6.705)	-12.14** (5.112)	-2.155*** (0.431)	-7.017 (5.127)	-1.788*** (0.127)	-1.375*** (0.212)
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	451	500	500	500	500	500	500

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, (6) hyperlinks to Wikipedia articles, (7) links to external sources. 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, *Treated countries after T*, 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: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

with more than 100 edits, (iv) edits per article, (v) growth of the total data in Wikipedia, (vi) new articles per day, (vii) internal links and (viii) external links. For each of these dependent variables we show the cross-term of treatment and a given period, i.e. the difference in differences coefficient for every period. The reference period stretches from period -11 to -5. The four first coefficients measure the "effect" before treatment actually begins: "Trtd -4m" measures the effect of affected economies, 4 months (3, 2 and 1 month) before the peak of the crisis. Similarly "Trtd 0m" stands for the peak month and then 1, 2, 3,... months after the crisis. When looking at the results, it is clear that there is no large and systematic effect article in views, but we do see an effect in the number of active editors (both, 5-100 edits and most pronounced editors with more than 100 edits). Furthermore, we measure more Edits/article, but this increased dynamic is already notable before unemployment increases and fades after 5 or 6 months. After that we observe an effect in newly created articles. This two effects are flanked by a slightly less robust but quite consistent increase in content growth after the peak increase in unemployment.

Table 12 shows the results when we run the difference in differences analysis period by period. The table has 8 dependent variables: (i) Article Views, (ii) Active editors with 5-100edits (iii) Active editors with more than 100 edits, (iv) edits per article, (v) growth of the total data in Wikipedia, (vi) new articles per day, (vii) internal links and (viii) external links. For each of these dependent variables we show the crossterm of treatment and a given period, i.e. the difference in differences coefficient for every period. The reference period stretches from period -11 to -5. The four first coefficients measure the "effect" before treatment actually begins: "Trtd -4m" measures the effect of affected economies, 4 months (3, 2 and 1 month) before the peak of the crisis. Similarly "Trtd 0m" stands for the peak month and then 1, 2, 3,... months after the crisis. When looking at the results, it is clear that there is no large and systematic effect article in views, but we do see an effect in the number of active editors (both, 5-100 edits and most pronounced editors with more than 100 edits). Furthermore, we measure more Edits/article, but this increased dynamic is already notable before unemployment increases and fades after 5 or 6 months. After that we observe an effect in newly created articles. This two effects are flanked by a slightly less robust but quite consistent increase in content growth after the peak increase in unemployment.

The results for the 24-month time interval (in Table 12) show that this initial increase in readership and the number of editors is sustainable and converts into a sustainable increase in content generation. An increase in regular editors with modest amount of activity becomes stronger, 1.1 standard deviations. However, over the horizon of a full year, we also see an increase in very active editors (≥ 100 edits; 0.8 std), words added to the encyclopedia (0.7 std). 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 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 7, which contains the regression coefficients of the independent variable of interest, *Unemployment*, on different measures of contributions to Wikipedia in each column. All specifications include month 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 and edits per article. For all other dependent variables the point estimates are negative and insignificant. 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 3 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 normalized with respect to the monthly average, and standard errors are clustered by countries.

The results in Table 3 confirm that views are a crucial predictor for all edit related outcomes except the number of edits/article. An increase in views by one standard deviation is associated with more active editors (0.3 std for active and 0.2 std for very active editors) and more content growth as measured in words (0.2 std). These OLS results cannot be interpreted causally. Exploring an exogenous 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 October, 2009 indeed does not affect

Table 3: OLS Regression for the effect of views during the period of 12 months before and 12 months after the crisis

	(1)	(2)	(3)	(4)	(5)	(6)
	Active 5-100e.	Active more 100e.	Edits p.article	Words growth	Int. links	Ext. links
Views	0.266** (0.110)	0.202* (0.105)	-0.0103 (0.0116)	0.174** (0.0772)	0.0114** (0.00489)	-0.0117 (0.0129)
Time trend	-0.0131 (0.235)	0.223 (0.215)	0.122*** (0.0141)	0.253 (0.179)	0.139*** (0.00373)	0.132*** (0.00555)
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes
Month dummies	Yes	Yes	Yes	Yes	Yes	Yes
Observations	451	451	451	451	451	451

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, (5) hyperlinks to Wikipedia articles, (6) links to external sources. 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 monthly youth unemployment rate. All specifications include month and year dummies, and a common time trend. Standard errors, clustered by countries in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

our baseline results. By defining a placebo treatment, such as setting the economic crisis 18 months later, our baseline results vanish. This suggests that our results are, indeed, due to the change in spare time of individuals due to unemployment.

Our most important “robustness check” repeats the entire analysis on German district level data. Running the analysis with German districts, allows us to address the concern that European countries are very heterogeneous. Too large heterogeneity casts doubt on whether these countries can serve as counterfactuals for each other. Moreover the country level analysis can only be performed for 22 units of observation and hence constrains our flexibility in testing analysis on subsets of the data. Both of these concerns can be addressed by analyzing the German district level. Since this analysis is very comprehensive, we present it in the next section.

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 the magnitude and significance of the unemployment effect is 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 also conducted the analysis at the level of German administrative districts, called “Kreise.” Doing so allows us to address the concern that European countries may be too heterogeneous to serve as counterfactuals for each other. Running the analysis with German districts, we scrutinize the robustness of our results in a framework where treated and untreated subjects are more similar than countries. Moreover, we can use many more units of observation and achieve a very similar institutional context.

We perform the same econometric estimation as in equation (1) where contribution of district i in month t is measured by the amount (the total number of edits or the sum of bytes) of anonymous contributors from the IP-addresses belonging to district i as well as activity indicators and the number of registered contributors. Importantly, for the registered contributors we are able to match their contributions to German districts only if they have voluntarily revealed their geographical location on their own user pages containing their wiki-profiles. These registered users are responsible for about 18 per cent of total contributions in the observed time period, and together with 17.5 per cent of anonymous contributions, our sample covers about 35 per cent of contributions to German Wikipedia. As in previous regressions, month and year dummies, as well as fixed effects are included to rule out time trends and district unobserved heterogeneity.

Overall our findings suggest that total contributions to German Wikipedia fall after the crisis, while in districts with higher unemployment the negative overall trend is slower. This is logical, since Germany was less affected by the crisis, a negative general trend of contributions could be expected. However, it appears that in districts with higher unemployment, the negative trend was mitigated, which might be due to more time spent online.

6.1 Results at the regional level

The main results for German districts are shown in Table 4.

The table shows the results for our main measures of contributions to Wikipedia in each respective column: (1) the number of anonymous edits, (2) the total anonymous contribution length in kb, the same indicators for registered users in (3) and (4), and (5) the number of registered contributors who reveal their geographical location on user pages. As before, the coefficient of interest, which measures the treatment effect, belongs to the cross-term *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 6-month interval suggest that anonymous edits fall after the crisis starts in January 2009 in Germany by 1 standard deviation as well as do contributions of registered editors. However, in districts affected by a rise in unemployment this negative trend is mitigated by additional activity on Wikipedia, with the difference in differences before and after the crisis about 0.3 standard deviations in anonymous content and 0.5 standard deviations in the number of registered active editors in the districts.

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 15 and 25. We expect them to get involved in online volunteering more easily. The results in Table 5 indicate that younger unemployed are not those responsible for our results. As we have unemployment rate for those youngsters who are included into labour force, this might indicate that contributions to Wikipedia are due to activity of individuals with higher education. Remarkably, anonymous contributions rise in the districts with higher youth unemployment, but not the number of registered users.

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 11 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

Table 4: DID Regression for German Districts (6 months before and after the crisis)

	(1)	(2)	(3)	(4)	(5)
	# An.edits (norm.)	An.contribution, KB(norm.)	# Reg.edits (norm.)	Reg.contribution, KB(norm.)	# Reg users
After treatment	-1.082*** (0.0755)	-0.669*** (0.0906)	-0.285* (0.145)	-0.287** (0.131)	-1.128*** (0.335)
Treated districts after T	0.267*** (0.0672)	0.172** (0.0744)	0.0517 (0.130)	0.0862 (0.116)	0.524** (0.207)
Constant	0.125*** (0.0440)	-0.0677 (0.0520)	-0.0601 (0.0654)	-0.140** (0.0586)	7.858*** (0.119)
Month dummies	Yes	Yes	Yes	Yes	Yes
Year dummies	Yes	Yes	Yes	Yes	Yes
Observations	3367	3367	2223	2210	1878

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: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 5: DID Regression for German Districts (6 months before and after the crisis) where affected districts are determined by changes in youth unemployment

	(1)	(2)	(3)	(4)	(5)
	# An.edits (norm.)	An.contribution, KB(norm.)	# Reg.edits (norm.)	Reg.contribution, KB(norm.)	# Reg users
After treatment	-1.047*** (0.0791)	-0.511*** (0.0970)	-0.207 (0.139)	-0.270** (0.124)	-0.919** (0.357)
Treated districts after T	0.146** (0.0699)	-0.0109 (0.0755)	-0.110 (0.125)	-0.00810 (0.111)	0.253 (0.241)
Constant	0.0989** (0.0450)	-0.131*** (0.0480)	-0.00927 (0.0697)	-0.109* (0.0611)	6.997*** (0.108)
Year dummies	Yes	Yes	Yes	Yes	Yes
Month dummies	Yes	Yes	Yes	Yes	Yes
Observations	3341	3341	2210	2197	1845

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: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

driven by edits made during leisure time, a finding which clearly offers many opportunities for interesting further research.

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. With higher unemployment, articles are read more frequently and the number of highly active users increases, suggesting that existing editors also increase their activity. Moreover, we find robust evidence that content grows at a higher pace. The overall effect seems 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.

Potentially, higher unemployment may be associated with greater volunteering activity and productive time usage. Yet we cannot fully answer how this mechanism works. Particularly, it seems that new editors begin to acquire new capabilities and devote their time to producing public goods. As the number of *new* articles created per day between articles on Wikipedia also increase (not shown in the baseline results), the increased participation is focused on adding to the existing knowledge as well as providing new topics. However, a decrease in the amount of set hyperlinks between Wikipedia articles and no effect on the amount of external links might indicate that the increase in content generated observed is due to the activity of unexperienced contributors or minor edits. Unfortunately, less links in the content might also suggest a decrease in content quality, which is then a negative consequence of more activity.

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, several limitations cannot easily be overcome. For example, 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 used in the comparison would have sufficient development over time assuming no 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 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 and political setup is likely to be more homogeneous when looking at German districts rather than at European countries. However, two remarks concerning this analysis apply. First, the regional analysis is based on the IP-addresses of anonymous contributions, which allowed only for a restricted set of available dependent variables. Specifically, we can only determine two average properties of the edits in any given region, i.e. we have to focus on the *number of edits* and the *average length of edits* in bytes. This is because computing statistics like the *number of active editors*, or *edits per article* becomes meaningless when edits of registered authors are neglected. 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.

Nevertheless, ongoing research aims to expand the scope of this analysis, by gathering additional location on active users. Thus we aim at pinning down regional and industry specific effects in knowledge creation, by gathering additional information on Wikipedia contributions in Germany and exploiting the category matching.

A final limitation of the German Data comes from the fact, that the unemployment rates in Germany increased much less than in other European countries. 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 disentanglement of 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 users 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 get edited would be very interesting. This is beyond the scope of this paper and, especially on the article level, this analysis is 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 on how newly unemployed use their additional time.

8 Conclusion

In this paper, we analyze the relationship between the economic downturn and the provision and documentation of knowledge online public good, i.e. online knowledge. In the times of digitization driven labor substitution it is a crucial question whether unemployed individuals could invest their additional time in the provision of (online) public 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 the world's most important online knowledge repository 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 worldwide online public good, Wikipedia. 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.

Overall we find a positive effect of the European economic crisis on contributions to the Wikipedias of countries which were affected. This effect does not lead to overall accelerated contribution, but softens the deceleration in contributions that we observe across most Wikipedias after 2009. On the level of 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. On the level of German Districts, we observe an increase in (a) anonymous edits and (b) in the number of registered users who reveal their origin. Moreover, we see that only anonymous edits are driven by youth unemployment. Finally, the effect is generally driven by additional contributions during leisure time rather than during working hours.

Taken together, our findings point to (relatively) more knowledge provision as a side effect of the crisis. However, the effect cannot unequivocally be attributed to those who lose their employment. In ongoing research we aim at shedding light at the precise channel of the effect and analyze who provides the additional contributions. We are gathering additional category specific data to pin down sector specific effects. Moreover we are currently expanding the scope of our research strategy to browser data to understand how general browsing patterns changed after the crisis. Thus we hope to shed light on how unemployed people use their time online.

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.
- 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.
- Zhang, Xiaoquan Michael and Feng Zhu**, “Group size and incentives to contribute: A natural experiment at Chinese Wikipedia,” *American Economic Review*, 2011, 101 (4), 1601–1615.

A Appendix

A.1 Country Level Analysis

Table 6: 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	11	0.11	8.5	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
Slovakian	5.19	91.56	2	0.03	11.5	3.6
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	33.39	89.99	5	0.06	11.1	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 7: OLS Regression for the period of 12 months before the crisis

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Views	Active 5-100e.	Active more 100e.	Edits p.article	Words growth	Int. links	Ext. links
Unemployment rate (norm.)	0.0257 (0.149)	-0.0652 (0.167)	-0.0855 (0.132)	0.0156 (0.0283)	-0.112 (0.183)	-0.00464 (0.0131)	0.0277 (0.0316)
Time trend	-0.448*** (0.156)	0.521** (0.232)	0.460*** (0.148)	0.153*** (0.0169)	0.188 (0.143)	0.139*** (0.00537)	0.111*** (0.00973)
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	191	240	240	240	240	240	240

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, (6) hyperlinks to Wikipedia articles, (7) links to external sources. 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, *Unemployment*, is the normalized monthly unemployment rate. All specifications include month and year dummies, and a common time trend. Standard errors, clustered by countries are in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

A.2 German districts

Figure 3: German administrative units on the district level



Table 8: Crisis Indicators: unemployment rates and the difference between them (%)

	Share of affected districts	Unemp.rate,%	Change in Unempl.,%
1	0.72	4.67	1.18
2	0.69	4.49	1.25
3	0.00	13.80	0.76
4	0.56	12.89	1.10
5	1.00	10.64	1.36
6	.	8.33	0.83
7	0.17	6.79	0.61
8	0.12	7.84	0.68
9	0.77	14.15	1.65
10	0.35	7.97	0.81
11	0.44	5.93	0.98
12	0.00	6.62	0.86
13	0.84	12.78	1.87
14	0.80	13.72	1.34
15	0.18	7.98	0.71
16	0.80	10.88	1.94
17	0.32	9.17	0.95

Table 9: Relationship between unemployment and activity on Wikipedia during 6 months before the crisis

	(1)	(2)	(3)	(4)	(5)
	# An.edits (norm.)	An.contribution, KB(norm.)	# Reg.edits (norm.)	Reg.contribution, KB(norm.)	# Reg users
Unemployment	-0.101 (0.0847)	-0.0288 (0.0976)	-0.0307 (0.107)	0.0888 (0.101)	0.0312 (0.162)
Year dummies	Yes	Yes	Yes	Yes	Yes
Month dummies	Yes	Yes	Yes	Yes	Yes
Observations	1553	1553	1027	1021	879

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 independent variable of interest is *Unemployment rate* for each district. 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 10: Relationship between unemployment and activity on Wikipedia during 6 months after the crisis

	(1)	(2)	(3)	(4)	(5)
	# An.edits (norm.)	An.contribution, KB(norm.)	# Reg.edits (norm.)	Reg.contribution, KB(norm.)	# Reg users
Unemployment	0.0746 (0.0626)	0.0198 (0.0614)	-0.117* (0.0609)	-0.0687 (0.0712)	0.0396 (0.134)
Year dummies	Yes	Yes	Yes	Yes	Yes
Month dummies	Yes	Yes	Yes	Yes	Yes
Observations	1566	1566	1038	1032	867

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 independent variable of interest is *Unemployment rate* for each district. 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 11: DID Regression for German Districts for contributions made in leisure time of the day (6pm-9am and weekends)

	Working time				Leisure time			
	(1) A. Edits	(2) A. Bytes	(3) R. Edits	(4) R. Bytes	(5) A. Edits	(6) A. Bytes	(7) R. Edits	(8) R. Bytes
After treatment	-0.881*** (0.0817)	-0.507*** (0.0857)	-0.214 (0.149)	-0.265* (0.145)	-1.060*** (0.0765)	-0.553*** (0.0931)	-0.203 (0.140)	-0.151 (0.122)
Treated districts after T	0.188*** (0.0662)	0.144** (0.0719)	0.0743 (0.128)	0.121 (0.110)	0.301*** (0.0654)	0.145** (0.0696)	0.0282 (0.128)	0.0834 (0.112)
Constant	0.0756 (0.0489)	0.00412 (0.0564)	-0.106 (0.0685)	-0.116* (0.0666)	0.151*** (0.0495)	-0.0652 (0.0541)	-0.0660 (0.0616)	-0.152*** (0.0547)
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3367	3367	2158	2132	3367	3367	2197	2197

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 increase in editing activity does occurs both during working hours and in the evenings or weekends, but the leisure time activity seems to affect our results most.

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

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Views	Act 5-100e	Act \geq 100e	Edits/art	Growth	New art/d	Int. links	Ext. links
Trtd -4m	0.295 (0.208)	0.452* (0.219)	0.493* (0.254)	0.0955* (0.0541)	-0.170 (0.285)	-0.0768 (0.215)	-0.000783 (0.0225)	-0.0939*** (0.0273)
Trtd -3m	0.170 (0.213)	0.661* (0.323)	0.422 (0.334)	0.118* (0.0575)	-0.145 (0.233)	-0.228 (0.249)	-0.0107 (0.0280)	-0.0982** (0.0348)
Trtd -2m	-0.0691 (0.300)	0.782** (0.343)	0.533 (0.319)	0.107* (0.0574)	0.0374 (0.340)	0.174 (0.297)	-0.0196 (0.0321)	-0.0775* (0.0406)
Trtd -1m	-0.481 (0.300)	0.599* (0.331)	0.137 (0.310)	0.0964* (0.0477)	0.274 (0.180)	-0.0839 (0.217)	-0.0316 (0.0363)	-0.0669 (0.0458)
Trtd 0m	-0.168 (0.279)	0.706* (0.405)	0.506 (0.398)	0.126** (0.0466)	-0.140 (0.306)	-0.00819 (0.255)	-0.0317 (0.0352)	-0.0813* (0.0427)
Trtd +1m	0.214 (0.324)	1.034*** (0.330)	0.823** (0.339)	0.138*** (0.0446)	0.251 (0.328)	0.191 (0.247)	-0.0472 (0.0322)	-0.0893** (0.0417)
Trtd +2m	0.490 (0.297)	1.126*** (0.326)	0.872** (0.350)	0.118** (0.0435)	0.468* (0.261)	0.416 (0.306)	-0.0509 (0.0304)	-0.0839** (0.0388)
Trtd +3m	0.234 (0.383)	1.258** (0.471)	1.003** (0.461)	0.142** (0.0534)	1.018*** (0.354)	0.190 (0.400)	-0.0388 (0.0273)	-0.0729* (0.0365)
Trtd +4m	0.510 (0.379)	1.252*** (0.389)	1.137** (0.447)	0.127** (0.0489)	0.686* (0.367)	0.321 (0.328)	-0.0325 (0.0232)	-0.0650* (0.0328)
Trtd +5m	0.731** (0.307)	1.130** (0.431)	1.208** (0.464)	0.0862** (0.0315)	0.567 (0.444)	1.006* (0.517)	-0.0260 (0.0196)	-0.0490 (0.0287)
Trtd +6m	0.336 (0.360)	1.095** (0.406)	1.142** (0.403)	0.0542 (0.0360)	0.839*** (0.254)	1.024*** (0.327)	-0.0242 (0.0158)	-0.00761 (0.0216)
Trtd +7m	0.255 (0.394)	1.418*** (0.428)	1.010** (0.439)	0.0184 (0.0367)	1.234*** (0.344)	1.007** (0.427)	-0.00839 (0.0150)	0.0232 (0.0191)
Trtd +8m	0.407 (0.444)	1.897*** (0.465)	1.143* (0.559)	-0.00323 (0.0365)	0.767** (0.356)	0.719** (0.335)	-0.00966 (0.0133)	0.0251 (0.0162)
Trtd +9m	0.711 (0.430)	1.765*** (0.470)	1.041* (0.517)	-0.0419 (0.0459)	0.601 (0.397)	0.621 (0.381)	-0.0147 (0.0178)	0.0287 (0.0177)
Trtd +10m	0.278 (0.413)	1.505*** (0.467)	1.180** (0.503)	-0.0686* (0.0338)	1.128** (0.490)	1.390*** (0.464)	-0.00720 (0.0174)	0.0461** (0.0213)
Trtd +11m	0.184 (0.436)	1.671*** (0.522)	0.979* (0.530)	-0.114** (0.0440)	1.442*** (0.480)	1.337*** (0.446)	0.00293 (0.0161)	0.0541** (0.0212)
Trtd +12m	0.571 (0.417)	1.747*** (0.550)	1.628*** (0.487)	-0.137** (0.0579)	1.050** (0.489)	1.178** (0.499)	0.0174 (0.0226)	0.0644** (0.0290)
Time trend	0.0449*** (0.0141)	-0.0383* (0.0221)	-0.0160 (0.0201)	0.135*** (0.000927)	-0.0199* (0.0112)	-0.0591*** (0.0116)	0.136*** (0.000378)	0.134*** (0.000741)
Constant	-0.878*** (0.253)	-0.237 (0.297)	-0.435* (0.225)	-1.789*** (0.0160)	-0.166 (0.174)	0.615*** (0.160)	-1.763*** (0.00969)	-1.734*** (0.0139)
Months	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	451	500	500	500	500	500	500	500

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, ie. growth of the entire Wikipedia, (6) new articles per day, (7) internal links and (8) external links. 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, $Trtd + Xm$, is an interaction term between dummies for the countries that are affected by the crisis with the time dummy indicating the period X months after the crisis. All specifications include month and year dummies, and a common time trend. Standard errors, clustered by countries, are in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.