Unemployment and Digital Public Goods Contribution

Michael Kummer
University of East Anglia, Georgia Institute of Technology & Centre for European Economic Research (ZEW), ,

Olga Slivko
Leibniz ZEW - Centre for European Economic Research, olga.slivko@zew.de,

Xiaoquan Zhang
CUHK Business School, ,

Economic crises have a harmful effect on employment. However, while the resulting loss of jobs has been shown to have many negative consequences for the affected individuals, it may also push them into new activities, such as provision of service to their communities. In this paper, we show how individuals engage in *socially useful* activities after an increase in unemployment. Specifically we document increased online content generation at Wikipedia, the world's largest user generated knowledge repository. Leveraging German district-level and European country-level unemployment data we analyze the relationship between the economic crisis in 2008-2010 and contributions to Wikipedia. We find increased socially valuable activity in the form of knowledge acquisition and contributions to Wikipedia. For German districts, we observe an increase in the rate of content generation on Wikipedia in more severely affected districts. These effects are even stronger at the European country level. Our findings suggest that public goods provision increases as a positive side effect of economic crises.

1   Introduction

Economic crises, like the one that followed the bankruptcy of Lehman Brothers in 2008, have often had very negative effects on the economy such as severely reduced income and massive layoffs. The impact of such economic crises is likely to be negative on organizations and communities that rely on donations and volunteers. But how do individuals change their charitable behavior when they are confronted with increased levels of unemployment? Specifically, how do individuals adjust their volunteering time in the time of economic downturns?

Prior studies summarized in List (2011) examine the relation between economic crises and monetary donations. The literature generally found that during economic downturns, money donations did not decrease as significantly as macroeconomic declines, but increased substantially following economic upturns. As a result, during the past half century, money donations not only outpaced, but doubled, the growth of Standard & Poor's 500. While such evidence is available for monetary donations, less is known for contributions of personal time and effort, volunteering, a second important form of giving to a charitable cause.

This paper examines the relation between economic crises and volunteering time. This question is of critical importance to organizations and platforms that coordinate volunteers to jointly promote charitable causes. Such organizations need to form adequate expectations about how much volunteering time they can expect in the future to efficiently plan the scale and scope of their commitments and activities.

Understanding the relationship between economic crises and volunteering also matters from a societal perspective. First, many volunteering organizations provide valuable public goods that range from disaster relief to publicly accessible knowledge, which matter to the wider economy. Second, individual working hours have been decreasing for decades in many societies, and several recent innovations and economic developments suggest that the trend to ever decreasing needs for human labor could continue. Hence, understanding the relationship between unemployment and volunteering time sheds light on the extent to which reduced demand for labor can free up resources that can become an input for the provision of the greater public good.

To study the relationship of interest, we analyze how economic crises and unemployment affect one particular type of volunteering: contributions to online public goods. Similar to traditional forms of charitable giving, economic crises severely limit resources for online public goods provision. However, while unemployment statistics and charitable money donations are easily captured and measured, statistics on volunteering time for the provision of public goods are often hard to obtain. To tackle this challenge, we study contributions to Wikipedia, one of the world's largest digital public good. We collect a comprehensive dataset on the economic crisis and empirically examine the impact of unemployment on Wikipedia contributions. We find that economic downturns are associated with increased contributions to the online public good Wikipedia.

We identify the effect of interest by examining the changes in the pattern of digital public goods provision during the period of significant social and economic crisis in Europe that followed the economic crisis in 2008-2010. Because this crisis was largely unexpected, it works as a natural experiment: participants in the economy took various actions to react to increased unemployment and we use a difference-in-differences (DID) framework to tease out the resulting effects.

Charitable contributions are a significant part of the economy, even without considering the value of volunteering time. Money donations in the United States now exceed about 2% of the gross domestic product ($314 billion in real 2008 dollars) (List 2011). The Corporation for National and Community Service estimates that 63 million Americans (about 25% of the population) volunteered a total of 8 billion hours of service in 2016.[1] These estimates do not include contributions to online public goods and therefore severely underestimate the time and value of overall volunteering.

In this paper we aim at filling this gap and at estimating the extent of online volunteering. While offline volunteering usually requires physically going somewhere and registering, and might even require some initial training, in the digital sphere much smaller contributions can be made any time of the day, and can already have a visible effect and make a difference. Moreover, it is easy to give online volunteering a try, perhaps even anonymously, and to turn away again if the activity turns out

---

[1]https://independentsector.org/resource/the-value-of-volunteer-time/.

to be less rewarding than expected. As the fifth-most-visited website, Wikipedia receives numerous views and volunteers from all over the world contribute their time and knowledge. The platform, thus, represents a perfect setting for studying online volunteering. Wikipedia is non-excludable and non-rival and can thus be regarded as a digital public good (Hess and Ostrom 2003). It is both collaboratively produced and universally accessible to anyone with Internet access, and it was highly active both before, during and after the economic crisis. Moreover, the wiki software records every edit in the article history, which is why Wikipedia provides an ideal arena to analyze the relationship between economic pressure from unemployment and volunteering behavior. We stress that similar patterns could apply to other digital content platforms such as Github, Stack Overflow, or Android mobile applications. Moreover, in the coming years, the rapid development of artificial intelligence (AI) will call for a rise of online volunteering platforms, such as Zooniverse, in which crowds participate in the creation of data sets subsequently used for training AI algorithms for research in many areas. Therefore, the potential value of the outcome of online volunteering and its societal impact is expected to grow drastically in the next years.

The effect of an economic crisis on online volunteering is not straightforward *ex ante*. On the one hand, the observed shift in time allocation towards more computer use and increased civic engagement might lead to the *increased* provision of public goods, thus more contributions to Wikipedia. Previous contributors might be able to allocate more time after they are displaced from their jobs and contribute their time to the public information good. Others, who were not aware of Wikipedia might begin searching for information on the internet and discover Wikipedia. Consequently, they might become interested in volunteering. Even individuals who still have jobs might do more online search for useful information and end up contributing to Wikipedia. On the other hand, the crisis may lead to *reduced* contributions to Wikipedia, because contributing time to public goods is clearly not an obvious reaction when people's jobs are threatened. Employed and unemployed individuals may feel threatened by social decline and might find it difficult to contribute to Wikipedia during a period of large-scale unemployment as the opportunity cost of their time is higher. Our contribution is to shed light on these questions by empirically analyzing how unemployment affects online volunteering time, and, specifically, contributions to Wikipedia.

We find that unemployment in German districts leads to higher participation by volunteers and increased content generation on Wikipedia. Both the number of participating editors and the number of highly active users increase. The number of edits to Wikipedia articles increases, and there is slightly weaker evidence for increased growth of overall content. We replicate this study with an EU country-level dataset. The results for European countries are consistent with those at the district-level in Germany. We find even stronger effects on content generation because of the higher variation in unemployment in Europe during the crisis.

While it is hard to draw final conclusions on the identity of the new editors, several of our findings are in line with a mechanism by which the crisis motivated a share of the population from regions with higher unemployment to begin to edit and, over time, existing participating editors also increased their activity. Because Wikipedia functions as an important knowledge base for the economy, our results document a new and somewhat valuable side effect of the economic crisis, that has been previously overlooked by policy makers.

The remainder of the paper is structured as follows. The next section reviews the literature. Section 3 describes the dataset, and section 4 discusses the empirical approach. Sections 5 and 6 report results on contributions to Wikipedia from German districts and European countries, respectively. Section 7 discusses the findings and limitations of our study, with suggestions for further research, and section 8 concludes.

2   Related Research

Our paper contributes to three major streams of the literature and in what follows we discuss these contributions one by one.

Motives to Contribute to Public Goods: Our paper adds to the literature that discusses the role of intrinsic and extrinsic motives for contribution to public goods highlighting an additional source of public goods contributions, the unemployment. Existing theoretical and empirical studies analyze the private incentives for voluntary public goods provision from the perspective of the interplay between free-riding incentives and social effects (e.g. Andreoni 2007).[2] Other research on public goods

---

[2]In Andreoni (2007), the provision of public goods is shown to be subject to congestion. That is, an increase in the number of recipients increases the total provision of public goods but at a decreasing rate.

contribution establishes the link between wealth/income to donations to the public good. Such works commonly assume that only those people who are doing well contribute to public goods and that transfers from the rich to the poor are organized through this mechanism (cf. List (2011) for a recent systematization and overview).

The more recent literature on public goods provision has established that other forces, such as social pressure, guilt, or sympathy, may play important roles in the decision to contribute to public goods (Andreoni 1988, 1989, 1990, 2007). This finding receives 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, 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.[3] Chen et al. (2010) propose social comparisons as the motivating mechanism for contributions of movie ratings to the online community MovieLens.

The Effects of Unemployment on Volunteering: The second strand of literature our paper relates to tackles the effect of unemployment on individual activities and on volunteering, in particular.

Previous research suggests that unemployment may lead to a decrease in volunteering in general. For civic public goods, unemployment has been shown to be negatively correlated with both religious and secular volunteering (Uslaner 2002, Freeman 1997), but the data in these studies did not feature within variation over time.[4] Lim and Laurence (2015) further distinguish between formal volunteering and informal helping behavior and show that they both decline during the recession of 2008 - 2009

---

[3]In 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).

[4]Uslaner (2002) uses cross-sectional data from the US and Canada, but cannot use variation from changes in unemployment over time. The outcomes of survey-based studies suggest that volunteers are relatively wealthier, predominantly male and and economically more active (Freeman 1997), and that unemployment is negatively correlated with volunteering by men, but not by women (Taniguchi 2006).

in the UK, while their individual-level analysis suggests that particularly unemployed individuals are more likely to contribute to both formal and informal volunteering activities. Similarly, looking at the aggregate effects, we also observe a general decrease in content contributions for Germany during the recession. However, unemployment moderates this common negative trend, and we show that in regions more severely hit by unemployment, the decrease in contributions is lower due to the influx of new contributors who did not belong to the community (anonymous contributors) at the time.

More generally, the literature discusses how the unemployed reallocate their time, considering a wide range of potential uses of time (Aguiar and Hurst 2007, Knabe et al. 2010, Krueger and Mueller 2012, Aguiar et al. 2012, 2013). Although unemployed people have more time to spend on leisure, they are less satisfied with life and specific activities (Knabe et al. 2010).[5] While only 2% of the foregone market hours are allocated to civic and religious engagement (Aguiar et al. 2013), the studies suggest that most of the additional online leisure time was spent on social networks, online games, email and portals, and young people spent more time online (Wallsten 2013).[6]

---

[5]This finding is in line with Pissarides (1992) who noted that recently unemployed face a threat of a permanent loss of skills and subsequent social decline as time proceeds. Krueger and Mueller (2012) find that the previously unemployed sharply decrease their hours devoted to leisure activities at the time of reemployment (by 35% of the time now allocated to working). In their paper, leisure includes computer and Internet use. The American Time Use Survey (ATUS) analysis of Aguiar et al. (2013) focuses on the period of the global recession in the late 2000s, and confirms earlier findings. They find that more than 50% of the additional time is spent on leisure activities, yet two-thirds are absorbed by watching TV and sleeping. This is in line with Aguiar et al. (2012) analysis of ATUS data and the changing trends in time allocation. Since the 1960s, individuals spend more time on leisure, which included personal computing (but also includes watching television or engaging in sports). Burda and Hamermesh (2010) analyze time diary data and conclude that only a small share of the additional time is used for home production. Unemployed spent more time on other activities such as computer use.

[6] These findings are complemented by Goldfarb and Prince (2008), who show that, conditional on having internet access, poorer people spend more time online, as their opportunity cost for time is lower than that of wealthier people. At times of economic crisis, both younger and poorer people can be threatened by increased unemployment rates or decreased salaries.

While "online" and "offline volunteering" share many features it is plausible, that many of our findings apply to offline volunteering. However, even though "online volunteering" is an integral part of general volunteering, they are not exactly the same. A key difference between online and offline volunteering is the flexibility in terms of the minimum necessary time spent on contributions and effort required to contribute. While contributing to online public good certainly has a lower hurdle, we also note that increased contributions may be a by-product of increased consumption of online information due to recession, so that volunteering becomes a side-effect of self-interested online activities. In that sense, our results suggest that the negative impact of economic crises on charitable activities from the previous literature (Uslaner 2002, Freeman 1997) might not carry over to information goods. Therefore, provided that unemployment increases time spent online, we could expect even a stronger effect on volunteering "online", as compared to volunteering "offline".

In our analysis, different from the previous literature, we focus on observational data from online contributions to the largest online encyclopedia, Wikipedia. We exploit variation in unemployment from the recent economic crisis in Europe in the aftermath of the financial crisis in 2008. This crisis generated an exogenous shock to the time spent online because people who were laid off during the economic crisis had to reduce their working hours. Our data, hence, provide an extraordinary opportunity to shed light on the relationship between economic pressure from unemployment and volunteering behavior as an instance of public goods provision by less successful economic agents.

Unemployment might drive pro-social behaviour if people feel threatened by unemployment and contribute in order to learn or maintain their skills, and maybe also to increase their self-esteem. This would be consistent with the research about pro-social behavior, which has documented a positive impact of regular volunteering on subjective well-being and happiness (Post 2005, Borgonovi 2008, Cattan et al. 2011, Binder and Freytag 2013, Dunn et al. 2008, Aknin et al. 2013, Liang et al. 2001).

We record increased online volunteering during a period of high unemployment on observational data that combine contributions of content to Wikipedia with data from a shock to unemployment. Moreover, other considerations such as prestige, respect, or guilt are relatively less important in this

setting, because the livelihood and well-beings of the volunteers are jeopardized. The recorded increase in contributions to Wikipedia during the economic crisis highlights a novel channel through which private motivations foster contributions to an online public good. Note that a similar mechanism could be at play on other digital content platforms such as Github, Stack Overflow, or Android mobile applications.

Summarizing our contributions, our study (1) uses a novel dataset of German and European Wikipedia contributions to study unemployment's effect, (2) contributes to the literature on public goods by highlighting a new channel that fosters content contributions to an important digital public good, and (3) identifies a socially valuable activity that results from higher unemployment.

3 Data

Our main analysis of the relationship between unemployment and contributions of content is based on German data that we collected at the district (NUTS 3) level. Despite the fact that the German economy held up relatively well during the economic crisis, different districts were affected differently by the rise in unemployment as well as by the "Kurzarbeit" (temporarily reduced working hours) program. In our analysis, we combine economic indicators of unemployment and reduced working hours at the district level with data on online activity and contributions to the German Wikipedia in the districts.

3.1 Unemployment and Reduced Working Hours in Germany in 2009-2010

In January 2009 the German government announced the need to combat the crisis. The German unemployment rate started rising early in 2009; in addition, many companies applied the extended Kurzarbeit program.[7] As a result, the government proposed to address the crisis by massively expanding the existing Kurzarbeit program. According to the rules of the program, employers experiencing

[7]In the last decade, German economy was constantly growing. Therefore, the unemployment rate had a general trend of decreasing. Rather than focusing on the absolute unemployment rate, we focus on changes in the unemployment rate and examine how the changes affect economic agents' behaviors. In years 2009-2010 the decreasing trend of unemployment rate experienced a shock of an increase and then continued decrease after the crisis. This period is very interesting as it shows that in the generally growing German economy, there was a financial crisis that changed the trend of unemployment.

a negative demand shock could activate reduced working hours for their employees. They would keep paying employees according to their hours worked, and the government would pay the workers about 60% of the foregone income. In January 2009, this program was extended from 6-12 to 18 and later 24 months, and its scope was broadened to cover a much larger number of industries (Walz et al. 2012).[8] Overall, the seasonally adjusted unemployment rate rose by 1%, and 300,000 people participated in Kurzarbeit. We thus define January 2009 as the onset of the great economic crisis and the moment when the crisis becomes significant for the German economy.

Figure 1 Unemployment and reduced working hours in Germany in 2009-2010

NOTES: Combined, the trends show that the rise in unemployment between 2009 and 2010 in Germany corresponded to a massive application of the Kurzarbeit program in the same period. Source: *Bundesagentur fuer Arbeit*.

[8]The Kurzarbeit program existed before the official announcement of the financial crisis. In general, the period of application is six months. However, under exceptional economic conditions, the program can be extended. As a measure of combating the economic crisis, German government varied this extension period. Thus, from January to June 2007 employers could use this program to retain their important employees for up to 15 months if the company faces a temporal reduction in demand. Then, this period was reduced to 12 months, but as Germany officially entered into recession, in January 2009, this period was extended to 18 months. Six months later, this period was extended to 24 months for employees who joined the program during the first six months of the recession. As a result, the workers of industries experiencing the negative shock could apply reduced working hours up to two years starting in January 2009.

Figure 2   Unemployment and Youth Unemployment in Germany in 2009-2010
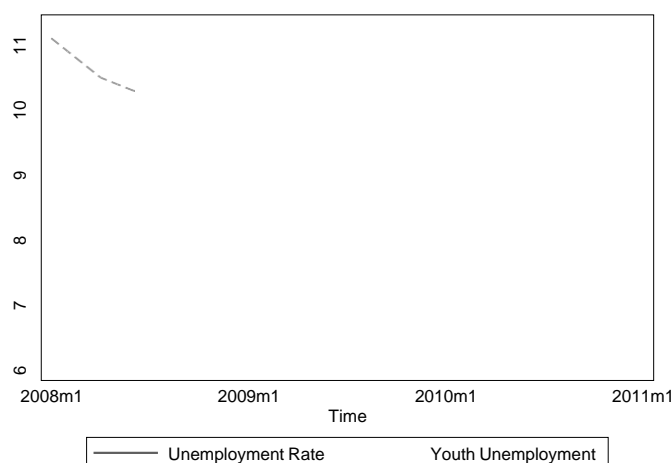
NOTES: The figure shows the rise in unemployment between 2009 and 2010 in Germany, and separates youth unemployment from all other workers. Like in all other European countries, young workers were more affected. Source: *Bundesagentur fuer Arbeit*.

We obtained monthly data on the number of unemployed, unemployment rates, and participants in the Kurzarbeit programs at the district level and generated a district-level dataset. Across all 16 German federal states, we observe 402 districts ('Kreise') shown in Figure 3. Our main estimation data are based on the 402 German Kreise, which are similar to to midsize US counties.[9] Table 1 summarizes the monthly panel data at the Kreise level. We observe 402 districts (Kreise) in the six months before and after the shock. The average unemployment rate is 7.5%, and unemployment increased on average by 1%. The smallest change in unemployment was a decrease by 1.1%, and the maximum increase was 3.4%. In the same table we also show our variables that allow us to analyze content generation on Wikipedia.[10]

Whether the districts are affected or unaffected by the economic crisis is defined based on changes in their unemployment rate after the crisis. To have sufficient variation between affected and unaffected districts, we rank the districts in terms of the change of unemployment rate and define the top 30%

[9]Georgia, with roughly 10 million residents, has more than 150 counties, while California, with more than 40 million residents, has only 58 counties. Germany, with 80 million citizens, has 402 Kreise so the average Kreise has approximately 200,000 residents.

[10]Specifically we observe the number of edits and the amount of content (in kilobytes) and we distinguish edits from registered users and anonymous IP addresses.

Table 1   Summary Statistics: Main Variables in the Regression Dataset

District Level Economic Variables	Mean	Sd	Min	P10	P50	P90	Max
GDP/Capita (1000 EUR)	29	12	13	18	25	44	94
Inhabitants (1000s)	200	232	0	70	144	349	3443
Total Hours Worked per Year (mln)	142	185	26	50	96	247	2397
# Unemployed (1000s)	8.3	14	.71	1.9	5.3	16	245
Unemployment Rate (%)	7.5	3.7	1.4	3.5	6.7	13	20
Unemployment Rate Men	7.8	3.9	1.1	3.5	6.9	14	20
Unemployment Rate Women	8	3.8	2	3.7	7.1	13	22
Youth Unemployment (age ≤ 25,%)	7.3	3.7	.67	3.1	6.6	13	22
# Businesses using Kurzarbeit	78	106	1	5	45	183	1371
# Individuals on Kurzarbeit Scheme	1761	2862	1	34	693	4658	25431

Treatment Indicators	Mean	Sd	Min	P10	P50	P90	Max
Change in Unemployment rate (%)	1.1	.68	-1.1	.24	1	1.9	3.4
Dummy: Affected	.5	.5	0	0	.5	1	1
Change in # of Workers on Kurzarbeit	3484	3903	-238	586	2079	8377	24963
Dummy: Kurzarbeit Shock	.5	.5	0	0	.5	1	1

Content Generation	Mean	Sd	Min	P10	P50	P90	Max
# Registered Users	3.7	13	0	0	1	8	218
# Reverted Revisions	.021	.16	0	0	0	0	3
# Edits from District	441	1061	1	37	159	986	18347
# Anonymous Edits	201	437	1	32	115	362	8061
# Registerd Edits	239	758	0	0	3	596	11598
Total Added Content (KB)	127	306	.004	5.3	40	288	4845
Total Added Content – Anonymous (KB)	61	152	.004	4.3	27	122	3138
Total Added Content – Registered (KB)	66	215	0	0	.22	164	4524

Local Content Generation	Mean	Sd	Min	P10	P50	P90	Max
# Businesses Using Kurzarbeit	77	98	1	5	45	185	927
# Part Time Employees	27	48	.01	.6	10	73	706
# Local Edits (about district)	641	1082	19	104	278	1445	13548
# Local Edits – Anonymous	108	193	0	15	49	239	2226
# Local Edits – Registered	522	887	16	81	223	1184	11162
Total Added Local Content (KB)	9.8	20	.14	.9	3.5	22	268
Total Local Content – Anonymous (KB)	.0059	.015	-.23	.00018	.0025	.015	.27
Total Local Content – Registered (KB)	.031	.07	-1.2	.00097	.011	.084	.8

NOTES: The unit of observations is Kreis i in month t. The time variable is normalized and runs from -6 months to 6 months after the shock. No. of Observations in the top and middle panel = 5226; No. of Kreise = 402; No. of States = 16. The lowest panel shows the distribution of the variables that measure local content generation (*about* districts; No. of Obs = 3120).

Figure 3   German administrative units at the district level
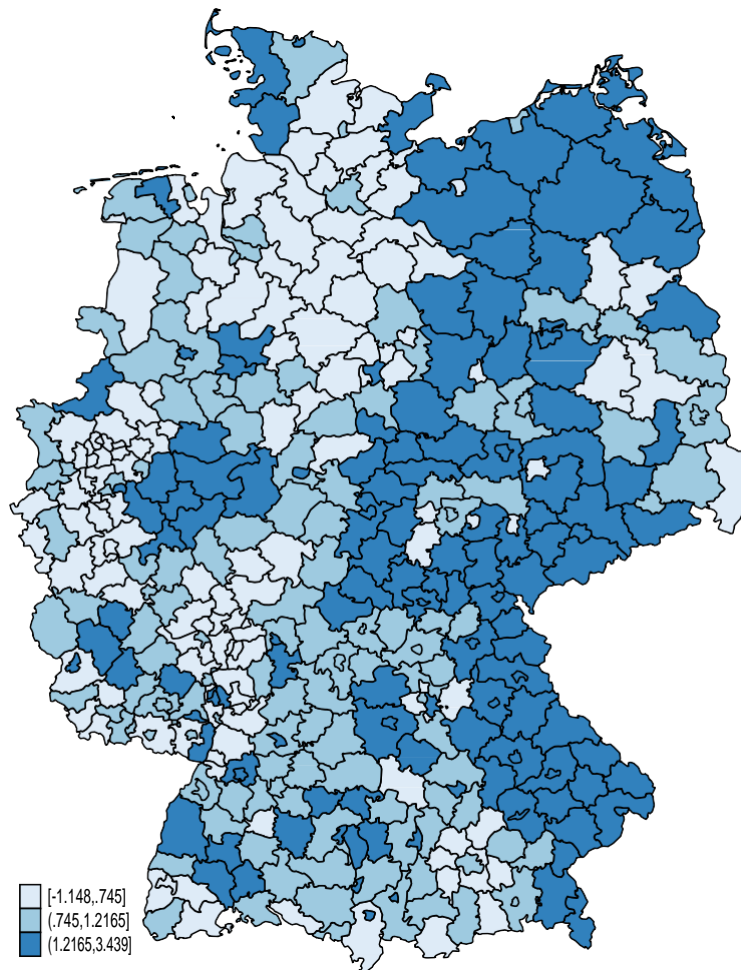
NOTES: The Figure shows the German "NUTS 3" units, which is the level of granularity at which we observe Wikipedia Edits. The darker districts experienced a relatively stronger increase in the unemployment rate than the lighter colored ones (when evaluating the six months after the onset of the crisis and comparing them to the six months before).

as affected by the crisis. The 30% of districts with the lowest, sometimes even negative, changes are defined as unaffected and used as the control group for our estimation. Note that the magnitude of the treatment was considerable, because unemployment increased by approximately 1.1% more *on average* in the affected districts.[11]

Appendix Table 11 shows a comparison of the two groups before the crisis began. This comparison highlights the big difference in the changes in unemployment rate. Departing from initially similar unemployment rates, the affected group experienced an average increase of 1.86% in the unemployment

---

[11]The standard error is less than **0.05**; results are available upon request.

rate while the unaffected group experienced an average increase of approximately 0.33% only. While similar on most parameters, the affected districts have generally smaller population and slightly lower income per capita. The difference in the number of inhabitants is also reflected in the number of aggregate edits. Even though per capita differences are much smaller, theses differences stress the importance of using a difference-in-differences strategy and of analyzing the similarity of the trends in district-level editing activities before the crisis.

In Table 2 we aggregate the unemployment rates at the state level to show how unemployment varied across the 16 states. The table also shows the shares of affected districts per German state. Our definition of crisis based on change in the unemployment rate implicitly controls for the baseline of economic status of the states. As a result, the highest shares of affected districts can be observed in traditionally economically strong industrial German states, such as Bavaria or Baden-Wuerttemberg. Weaker states such as Thuringia also had a large share of affected districts.

3.2   Contributions to Wikipedia

In our analysis, we focus on three different measures of contributions to Wikipedia: (1) additions of content (2) deletions of content and (3) editing activity. To quantify additions and deletions, we measure the bytes that were added or deleted to the platform in a given month. It is useful to quantify additions and deletions separately to shed light on the nature of the effort that individuals exert. To quantify editing activity we counted the number of edits. Moreover, when constructing our three measures of content generation on Wikipedia, we were able to distinguish contributions from anonymous and registered users and calculated the content generation of both groups separately. In what follows, we describe how we matched anonymous contributions to districts via the recorded IP addresses, and how we used registered contributors that reveal their location to match them to their corresponding districts.

Anonymous Contributions to German Wikipedia at the District Level: Based on a large dataset that contains the revision history of all articles of German Wikipedia, we aggregate individual

Table 2 **Unemployment Indicators in 16 German States**

	Workers	Unemployment Rate			Districts
	(1)	(2)	(3)	(4)	(5)
	# Unemployed (1000s)	All Workers (%)	Youth (age ≤ 25) (%)	Change (%)	Share Affected
Baden-Wuerttemberg	254.1	4.69	4.13	1.15	0.62
Bavaria	291.2	4.46	4.08	1.18	0.73
Berlin	232.6	13.80	15.36	1.03	0.00
Brandenburg	169.6	13.21	13.27	0.87	0.55
Bremen	37.1	13.23	11.83	0.61	0.50
Hamburg	75.2	8.33	7.72	0.79	0.00
Hessen	209.2	6.71	6.78	0.78	0.18
Lower Saxony	304.3	7.85	7.91	0.68	0.11
Mecklenburg-Western Pomerania	120.1	13.74	12.41	1.05	0.83
North Rhine-Westphalia	774.7	8.47	8.11	0.90	0.26
Rhineland-Palatinate	120.4	6.18	6.71	1.01	0.42
Saarland	37.5	6.62	6.47	1.01	0.00
Saxony	277.2	12.80	12.59	1.61	0.82
Saxony-Anhalt	169.7	13.72	13.32	1.40	0.80
Schleswig-Holstein	108.4	8.10	8.62	0.72	0.10
Thuringia	135.3	11.35	10.53	1.76	0.76

NOTES: This table shows mean values of unemployment indicators, the number of unemployed and the rates, as well as difference in the unemployment rate before and after the shock and the share of districts affected by the shock for each German state.

monthly contributions and compute total contributions by districts. For this aggregation, we map the IP addresses associated with edits to the corresponding German districts.[12]

In terms of overall editing activities on German Wikipedia, anonymous edits represent about 16% of all edits during our period of analysis (2008-2010). While we do not suppose that anonymous edits are representative for all editing activities, we deem it highly relevant for our research question, because anonymous edits are typically made by occasional or unexperienced editors. Thus, our measures of anonymous contributions to Wikipedia at the district level account either for contributions by newcomers or for occasional and relatively small contributions in terms of the content generated.

Registered Contributions at the German District Level: In addition to anonymous editing activities, we collect information on the location of registered Wikipedia contributors, whenever they reveal it publicly on their user talk/profile page. We thus match almost 25% of the registered edits by users who edited Wikipedia under their username, to a district of origin. Given that these registered edits were made by editors with a well-developed user talk/profile page, we consider them representative of edits by very active Wikipedia users, thus covering the other side of the spectrum.

The middle panel of Table 1 shows both anonymous and registered monthly edits from a district to German Wikipedia, together with the number of registered users and the number of reverted edits. Because registered users are very active, they generate a much larger amount of content than anonymous users do. We also show the total number of edits that we could match to each district via one of the two approaches. Together, we could thus associate about 35% of the activities on the German-language Wikipedia to a district of origin. By analyzing a significant part, but not all of the edits by German contributors we introduce an additional identification assumption to our analysis. Specifically, we assume that registered Wikipedia editors who did not add a regional identifier to their profile adjust their editing behavior in similar ways as editors who did add such an identifier. Even though we deem this assumption plausible, we cannot directly test it because we cannot match users without identifier to any regions.

[12]Contributions associated with IP addresses are made only by contributors who skipped the log-in procedure – that is, only by "anonymous" contributors.

In our German district-level analysis, we do not examine the contributions made by German users to Wikipedia versions of other languages such as Austrian, Swiss and English. This data limitation is an issue only if users in affected German districts change their cross-country contribution behavior in a systematically different way than unaffected users, which we believe is unlikely. Even if there is a systematic difference, our country-level analysis would address this issue.[13]

We verify that the large uncovered share of edits does not distort our results indirectly, by analyzing edits about districts in a robustness check (see next paragraph and section 5.1.6). In this analysis of district-specific content, we can match the edits of all users who edit in categories that can be associated with local interests in a district.[14]

Content about Districts To examine the full activity on Wikipedia at the level of districts, we look deeper into the topical content of contributions and identify content that is specific to a particular district. To that end, we take advantage of the "category tree," which is Wikipedia's extensive categorization tool. To attribute articles on Wikipedia to German districts, we used the fact that contributors assign each article on Wikipedia to different categories. To match the articles to the districts, we use Wiki-topics (categories) and Wikipedia's hierarchical category tree. To do so we extract the ID of all pages that belong to a district's local interest category and recompute all our measures of monthly district-specific activities on Wikipedia based only on the local interest pages. We are thus able to track contributions that refer to our specific districts by the category of each article.

Table 3 shows an example of the categories that contain the name of the district Verden (in Lower Saxony), in northwest Germany. These categories include the district's geography, famous people, sports, buildings, churches, and water sources. In our data, we find 82 categories per district on

---

[13]It is possible that some measurement errors may still exist in the outcome variables. However, as long as these errors are random and are uncorrelated with other variables in the model, we would end up with inflated standard errors and have a bias against finding significant results.

[14]Moreover, we analyze data from country-level models that also consider all contributions by registered users, and the results are even stronger. We report these results in section 6.

average. After we identified the district-specific categories, we used the set of Wikipedia articles in these categories to analyze district-specific content generation. This approach allows us to consider all contributions to these district-specific articles, including all contributions by authors who could not be matched to a region.

Table 3   Local Categories Describing Wikipedia Content about the German District Verden

ID	District	Full Article Name
3361	Verden	Ehemalige_Gemeinde_(Landkreis_Verden)
3361	Verden	Gewässer_im_Landkreis_Verden
3361	Verden	Unternehmen_(Landkreis_Verden)
3361	Verden	Kirchengebäude_im_Landkreis_Verden
3361	Verden	Verkehr_(Landkreis_Verden)
3361	Verden	Ort_im_Landkreis_Verden
3361	Verden	Geographie_(Landkreis_Verden)
3361	Verden	Person_(Verden_(Aller))
3361	Verden	Bauwerk_im_Landkreis_Verden
3361	Verden	Verden_(Aller)
3361	Verden	Achim_(Landkreis_Verden)
3361	Verden	Benutzer:aus_dem_Landkreis_Verden
3361	Verden	Domherr_(Verden)
3361	Verden	Blender_(Landkreis_Verden)
3361	Verden	Sport_(Landkreis_Verden)
3361	Verden	Bischof_von_Verden
3361	Verden	Landkreis_Verden
3361	Verden	Person_(Landkreis_Verden)
3361	Verden	Fürstbischof_(Verden)
3361	Verden	Landkreis_Verden_nach_Gemeinde
3361	Verden	Benutzer:aus_Verden

NOTES: The table shows an example of how we identified subcategories that consisted of articles with content *about* a district (here: "Verden" district).

The resulting measures of contributions of district-specific knowledge are shown in the bottom panel of Table 1. Again, we show both anonymous and registered monthly edits *about* a district on German-language Wikipedia. Because these measures include the edits of all users, the ratio of anonymous to registered edits is about 1:6. The share of contributions by registered users to local content is smaller than contributions to all contents in the middle panel of Table 1. The variable *Total contributions* simply measures the sum of both types of edits about a district.

European Country Level Analysis In Section 6, we present an analysis of data at the level of European countries. The analysis supports the findings based on German districts. Because these data originated from a different data source, we defer the discussion of the European data to that section.
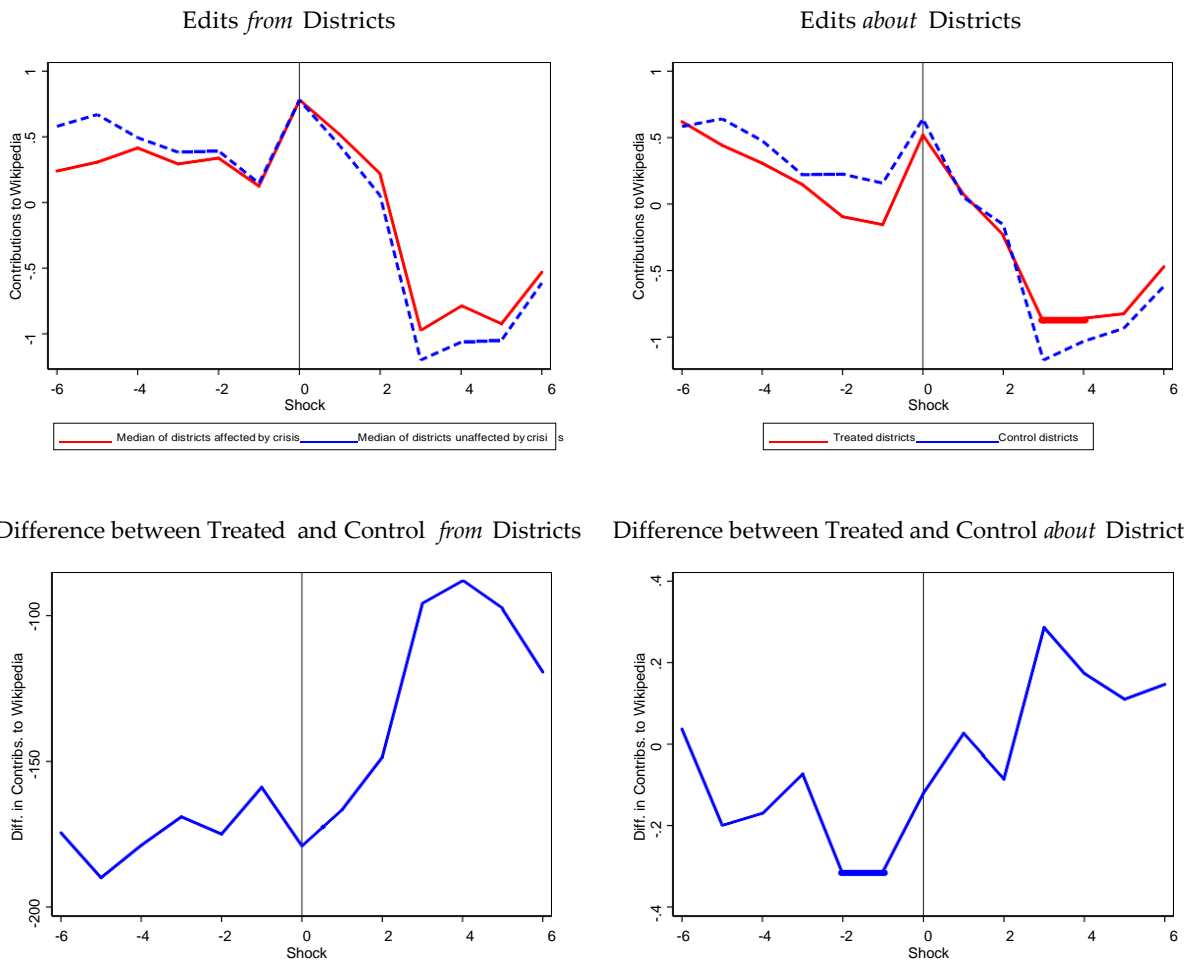
3.3 Descriptive Evidence

Before describing our empirical approach, we present a descriptive visualization of our data on the German districts (Kreise) in Figure 4. The left panel shows anonymous edits *from* the districts, while the right panel shows anonymous edits *about* these districts over time. We show the district-specific anonymous editing behavior six months before and after the onset of the crisis (highlighted by the thin red line at $t = 0$). The two figures in the upper row show the median values of the normalized number of edits, while the lower row shows the absolute difference between affected and unaffected districts.

Contributions *from* and *about* affected districts (i.e., treatment group) increased relaive to those from and about unaffected districts (i.e., control group). A comparison of the normalized trends reveals that both anonymous contributions (from districts) and registered contributions (about districts) experienced a relative increase in the districts that were more affected by the crisis.

4 Empirical Approach

We analyze the relationship between the economic crisis and the voluntary contributions of online knowledge to Wikipedia in two frameworks. First, we check whether the variation in the unemployment rates can explain the variation in content generation. We use fixed-effects ordinary least square (OLS) regressions and analyze whether unemployment is a key driver of increased online knowledge contributions.

Second, in our main specification, we rely on a difference-in-differences approach. The shock to unemployment is used as a source of exogenous variation in "disposable time" in the German economic system, and we compare content generation in German districts where the crisis was felt more strongly to the other, less affected, districts. In Section 6, we adopt this approach to compare severely affected European countries with countries that experienced only a moderate increase in unemployment. Compared to the country-level analysis, the analysis of German districts in this section allows us to focus on a framework with many units of observations (almost 400) in a homogeneous institutional context. The advantage of country-level analysis lies in the considerably larger variation in European unemployment rates.

Figure 4   Development of the Main Outcomes for Content at the District Level

Edits *from* Districts

Edits *about* Districts

Difference between Treated and Control *from* Districts

Difference between Treated and Control *about* Districts

NOTES: The figure shows the median values of the normalized number of edits *from* users in the district (left) and median values of the normalized number of edits *to the district's* "local interest" content (right). The upper-panel graphs show the median of the normalized number of edits for the affected and unaffected districts separately. The figure is based on monthly data six months before and after the crisis. The lower-panel graphs show the absolute difference in the number of edits by the two user groups.

4.1   Difference-in-Differences Framework

Our main specification uses a difference-in-differences (DID) approach to the data from German administrative districts (Kreise). The first difference compares content generation on Wikipedia before and after the shock, and the second difference compares content generation in affected by the rise in unemployment German districts to content generation in relatively unaffected districts. As outcomes of interest we analyze the number of edits and the size of contributions in bytes by anonymous and registered users. This identification strategy allows us to measure the impact of the economic crisis 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 in the DID framework is that the changes in the readership and contribution activities are indeed due to the crisis, rather than to some unobservable factors that correlate with the timing of the crisis. Moreover, treated and untreated districts have to share their pretrend dynamics, which we consider plausible, given the visual evidence.

The unit of observation in our data is a district in Germany with all corresponding statistics (e.g., unemployment rates) and aggregated contributions to Wikipedia observed every month before and after the official beginning of the economic crisis in Germany, which was announced in January 2009. The estimated equation is as follows:

$$Contributions_{it} = a_i + \beta\,(AfterT_t \times Treated_i) + X_{it}\,\gamma + v_t + s_{it}, \tag{1}$$

$AfterT_t$ and $Treated_i$ are dummy variables. $Treated_i$ separates districts that were affected by the economic crises from those unaffected. $AfterT_t$ equals 1 if the time period is after the shock t_0. As the variable $Treated_i$ does not vary over time, it drops out in the fixed-effects specification. Similarly, in the analysis for German regions the crisis moment is determined by switching from December, 2008 to January, 2009, therefore, the effect of $AfterT_t$ is captured by one of the time dummies, v_t. The coefficient of interest is the coefficient for the interaction term between these two dummies, β, which measures the DID after the shock to unemployment. In all regressions, as before, we include month dummies as well as district fixed effects of German districts to rule out district-specific unobserved heterogeneity and common time trends. Like before, the vector γ is a vector of parameters, to allow including additional (time-varying) control variables in X_{it}, i.e. district characteristics.

The contribution of district i in month t are measured in several ways. As was explained in subsection 3.2, we distinguished contributions from anonymous or registered users, and we computed (1) added bytes, (2) deleted bytes, and (3) the number of edits for both groups.

4.2 Fixed-Effects Panel Regressions

In a robustness check, we use an alternative estimation approach in which we rely on our panel data and use fixed-effects OLS regressions. We regress the same dependent variables as before on the unemployment rates in the districts. The regression equation is:

$$Contributions_{it} = a_i + \beta \; Unemployment \; Rate_{it} + X_{it} \; \gamma + v_t + s_{it}, \quad\quad (2)$$

where i refers to the districts, t is the time period (year and month). The coefficient of interest is β, which measures the relationship between changes in unemployment and online content contributions in the districts. District fixed effects, a_i, as well as month dummies, v_t, are included to rule out individual unobserved heterogeneity and time-trend effects. The vector γ is a vector of parameters, each corresponding to one of the control variables in X_{it}, i.e. district characteristics. Note that X_{it} is empty in the main specification, because the use of a fixed effects design precludes using time-invariant controls and, hence, the number of variables in X_{it} is quite limited (GDP/capita, and the shares of employees that work in manufacturing and in the service sector). However, in Appendix Tables 13 (and 14) we include the available time-varying controls.

5 Results for German Districts

In this section we present the results of our analysis at the district level. The subsequent section (Section 6) repeats the analogous analysis at the country level.

5.1 Within Districts

In this subsection we present the results for all content generated by users *within* German districts. Overall total contributions to German-language Wikipedia fall after the crisis, but our findings suggest that in districts with higher unemployment the negative overall trend is slower. In the next subsection we present the results for content *about* districts. The results of this section focus on content contributed to Wikipedia, added or deleted, that has not been reverted. In Wikipedia, a revert means that the content contributed recently by a user can be removed if the user fails to satisfying the guidelines (e.g., vandalism, editing wars). The analysis of the content reverts is included in the Appendix (Tables 15 and 16) and is also briefly discussed in this section. It provides us an insight on how valuable the added content is or whether the deletion of the content was appropriate. For example, an increase in reverts after the crisis could indicate that a substantial share of content added due to the crisis is non-valuable and represents vandalism, or that an inflow of new users to the platform triggered editing wars in the community.

5.1.1 Difference-in-Differences Analysis Our main results on the relationship between the unemployment rate and contributions to Wikipedia are presented in Table 5. These results are based on a DID specification for German districts. The observed outcomes include content generated anonymously (Col. 1 - 3) and by registered users (Col. 4 - 6). For both groups of users, we differentiate between the amount of content (in kilobytes), added (Col. 1 and 4) and deleted (Col. 2 and 5), as well as the editing activity, measured by the number of edits (Col. 3 and 6). All dependent variables are transformed into logarithms, and we control for monthly dummies to capture the country-wide temporal dynamics. The coefficient of interest, which measures the treatment effect, is the cross-term *Treated districts after T*. It is the interaction term between dummies for districts that are affected by the crisis with the time dummy indicating the period after the crisis.

The results for a 12-month interval (6 months before and 6 months after the shock) suggest several interesting patterns. In affected districts that experienced an increase in unemployment we can see an increase in anonymous content contributions, but not in contributions by registered users, who are actual members of the community. Our DID estimates suggest that this effect is about 6% in edits and 10% in kilobytes. Interestingly, the contributions of valuable content that remained on article pages are driven by content addition, not by deletion.[15]

The results suggest an increase in participation of the anonymous members of the Wikipedia, and, moreover, the main source of this increase are anonymous and, likely, new users to the community. Inevitably, this provoked an increase in content contributions which do not satisfy the guidelines, and also potential vandalism or editing wars. Tables 15 and 16 (in the Appendix) present the effect of the shock to unemployment on the number of vandal edits and on the share of malicious content contributions in total content contributions, correspondingly. As we can see the number of vandal edits in the affected districts slightly increases in absolute numbers (Table 15). However, the share

[15]Two of the dependent variables, namely, the number of edits performed by anonymous and registered users, could be alternatively estimated using a count data model. Due to distribution properties (the distributions of variables are overdispersed) we chose the negative binomial model for an alternative estimation. Table 26 in the Appendix presents the results, which are consistent with the linear estimations in the baseline specification.

of malicious content is not affected (Table 16). The results suggest that a slight increase in malicious activity by anonymous users during the crisis did not bring significant quality deterioration in relative terms.

The picture becomes slightly more nuanced after we decompose the treatment effect after the shock into effects by months after treatment (see Table 4). In this table, we can see that the positive effect on the number of anonymous edits is the strongest over the first four months after the shock. As for the size, in the month of the shock the amount of contributed kbytes rises by about 9%, and in the second month by almost 10%. The effect on anonymous edits fades in the fifth month after the shock. Content additions, as opposed to content deletions, show a positive dynamics in the districts affected by the unemployment rise, and the effect on edits gets stronger as the effect of the crisis strengthens.

Overall, we find that total contributions to the German-language Wikipedia fall after the crisis, but additional activity by new users mitigates the negative overall trend in districts with increased unemployment.

As mentioned in the Data section, in our DID regressions we compare the top 30% of districts which are affected strongly by the unemployment shock with the bottom 30%, to ensure the districts in treatment and control groups are sufficiently different. The choice of the cutoff point is determined by the tradeoff between the need to compare groups of districts with substantially different change in unemployment, and at the same time to use as many observations as possible. We perform an analysis of how this choice may affect the results of our study and show in Table 17 (in Appendix) that our results are robust to the choice of the cutoff point as soon as the districts that are just in the median of distribution are eliminated. As a dependent variable, we choose our main result from Col. (3) in Table 5. This result is replicated in Col. (5) of Table 17, and Col.(4) here shows that we could still eliminate fewer districts in the middle, using in the regression top 40% with the highest increase in unemployment and bottom 40% with the lowest increase in unemployment and the effects of interest would remain. As expected, Col. (6) and (7) show that we could also further eliminate the districts in the middle of distribution and get even stronger effects.

Table 4   Difference-in-Differences Regressions for German Districts.

	Anonymous Content			Registered Content		
	(1)	(2)	(3)	(4)	(5)	(6)
	Log (Added (KB))	Log (Deleted (KB))	Log (# Edits)	Log (Added (KB))	Log (Deleted (KB))	Log (# Edits)
Treated in Month -3	0.025	0.153	0.054	-0.429*	-0.190	-0.124
	(0.100)	(0.196)	(0.037)	(0.255)	(0.209)	(0.103)
Treated in Month -2	-0.041	0.128	0.048	-0.093	0.133	-0.021
	(0.102)	(0.203)	(0.035)	(0.263)	(0.220)	(0.113)
Treated in Month -1	-0.064	0.433**	0.049	-0.029	0.195	0.021
	(0.102)	(0.182)	(0.040)	(0.269)	(0.238)	(0.118)
Treated in Month 0	0.070	0.100	0.091**	-0.168	0.029	-0.031
	(0.102)	(0.190)	(0.038)	(0.333)	(0.282)	(0.145)
Treated in Month 1	0.072	-0.195	0.096**	-0.369	-0.228	-0.107
	(0.098)	(0.201)	(0.038)	(0.297)	(0.256)	(0.140)
Treated in Month 2	0.258**	0.111	0.104***	-0.285	-0.051	-0.127
	(0.105)	(0.212)	(0.038)	(0.314)	(0.290)	(0.149)
Treated in Month 3	0.084	-0.025	0.100*	-0.049	-0.076	-0.069
	(0.115)	(0.271)	(0.052)	(0.293)	(0.274)	(0.139)
Treated in Month 4	0.131	-0.063	0.128**	-0.044	0.044	-0.024
	(0.107)	(0.240)	(0.050)	(0.301)	(0.296)	(0.145)
Treated in Month 5	-0.007	0.102	0.029	-0.011	0.199	0.013
	(0.109)	(0.250)	(0.050)	(0.295)	(0.271)	(0.138)
Treated in Month 6	0.010	-0.266	0.052	0.052	0.074	-0.101
	(0.105)	(0.228)	(0.053)	(0.315)	(0.285)	(0.150)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3120	3120	3120	3120	3120	3120
Number of Districts	240	240	240	240	240	240
R^2	0.10	0.17	0.44	0.01	0.01	0.01

NOTES: The table shows the per-period results of our main difference-in-difference analysis that contrasts affected and unaffected German districts around the onset of the European Financial Crisis in 2009. The columns contain different measures of contribution activity to Wikipedia: (1) the total length of anonymously added content (in KB), (2) the total length of anonymously deleted content (in KB), (3) the number of edits (revisions) by anonymous editors, (4) the total length of content added by registered users (in KB) (5) the total length of content deleted by registered users (in KB), (6) the number of edits (revisions) by registered editors. The variables of interest that decompose the treatment effect are *Treated in month 1* to *Treated in month 6*. These variables are interactions between a dummy for districts that are affected by the crisis with time dummies that correspond to the respective period after the crisis. All specifications include time period (month-year) dummies. Time Range: 6 Months Before and After the Crisis Start. Standard errors, clustered by districts, are in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 5   Difference-in-Differences Regressions for German Districts.

	Anonymous Content			Registered Content		
	(1)	(2)	(3)	(4)	(5)	(6)
	Log (Added (KB))	Log (Deleted (KB))	Log (# Edits)	Log (Added (KB))	Log (Deleted (KB))	Log (# Edits)
Treated Districts After T.	0.110**	-0.145	0.060**	-0.012	-0.036	-0.041
	(0.054)	(0.113)	(0.027)	(0.197)	(0.187)	(0.101)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3120	3120	3120	3120	3120	3120
Number of Districts	240	240	240	240	240	240
R^2	0.06	0.16	0.44	0.00	0.01	0.00

NOTES: The table shows the results of our main difference-in-difference analysis that contrasts affected and unaffected German districts around the onset of the European Financial Crisis in January 2009. The columns contain different measures of contribution activity to Wikipedia: (1) the total length of anonymously added content (in KB), the total length of anonymously deleted content (in KB), (3) the number of edits (revisions) by anonymous editors, (4) the total length of content added by registered users (in KB) (5) the total length of content deleted by registered users (in KB), (6) the number of edits (revisions) by registered editors. The variable of interest, which captures the treatment effect, *Treated districts after T* (reatment), 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. Observations range from 6 months before to 6 months after the onset of the crisis. All specifications include time period (month-year) dummies. Standard errors, clustered by districts, are in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	Anonymous Content			Registered Content		
	(1)	(2)	(3)	(4)	(5)	(6)
	Log (Added (KB))	Log (Deleted (KB))	Log (# Edits)	Log (Added (KB))	Log (Deleted (KB))	Log (# Edits)
Unemployment Rate (%)	0.032	0.093*	0.027**	-0.014	-0.027	-0.023
	(0.023)	(0.054)	(0.011)	(0.062)	(0.059)	(0.030)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	5226	5226	5226	5226	5226	5226
Number of Districts	402	402	402	402	402	402
R^2	0.05	0.17	0.44	0.00	0.01	0.00

NOTES: The table shows the results of a fixed-effects analysis that directly relates Wikipedia contributions to the unemployment rate of German districts around the onset of the European Financial Crisis in 2009. The columns contain different measures of contribution activity to Wikipedia: (1) the total length of anonymously added content (in KB), (2) the total length of anonymously deleted content (in KB), (3) the number of edits (revisions) by anonymous editors, (4) the total length of content added by registered users (in KB) (5) the total length of content deleted by registered users (in KB), (6) the number of edits (revisions) by registered editors. The independent variable of interest is *Unemployment rate* in period t for each district. Fixed Effects Regressions ranging from 6 Months Before to 6 Months After the Crisis Start. All specifications include time period (month-year) dummies. Standard errors, clustered by districts, are in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

5.1.2 Robustness check: Alternative Estimation Strategies Our first set of robustness checks investigates how sensitive our results are to our estimation strategy. We first use a linear fixed-effects panel regression framework. Table 6 shows the results for regressing several measures of content generation on the unemployment rate and district fixed effects.[16] As before, each column presents estimation results for our main measures of editing activity and contributions to Wikipedia: content generated anonymously (Col. 1 - 3) and by registered users (Col. 4 - 6). As before, we present the results for the anonymously added and deleted amount of content (in kilobytes) and the number of edits by anonymous users, as well as the same measures for content contributed by registered users. The results suggest that an increase in the regional unemployment rate is strongly related to the variation in anonymous contributions to Wikipedia. A 1% increase in unemployment rate is associated with a 3% increase in anonymous edits. Moreover, it is related to content deletion, though the significance of the coefficient is marginal. Based on average editing activity on German Wikipedia in 2009 and 2010, that corresponds to almost 4,340 additional edits and 55 megabytes over the six-month period after the shock we observe. Due to limitations in the availability of control variables, which are collected for German districts on the yearly basis, Table 6 does not include any control variables. Table 13 (in Appendix) demonstrates that the inclusion of the available control variables does not affect the results.

In addition to the OLS Regressions, we verified robustness to using the fuzzy DID approach (De Chaisemartin and D'Haultfoeuille 2017). These results are shown in Table 12 in Appendix B.

5.1.3 Robustness check: Controlling for District-Specific Pre-treatment Trends We now perform a check in which we control for any district-specific patterns in the contributions to Wikipedia before the crisis. Such district-specific patterns, if systemic, could interfere with the assumptions of our natural experiment setting. To implement this check, we compute each district's individual pre-treatment trend for each dependent variable before the shock. We then extrapolate

---

[16]The fixed effect essentially covers all available control variables, since macroeconomic indicators such as the population structure or internet penetration do not vary month to month.

each district's pre-treatment trend to the period after the shock and include this new variable in the regression equations. We show the corresponding results in Table 7. By construction, it should be significant and it allows to control out any trends in the dependent variables that would have been there in the absence of treatment. Despite the inclusion of the pretreatment trends in the dependent variable in our regressions, our results on our main dependent variables of interest remain unchanged when district-specific pre-treatment trends in the contributions to Wikipedia are controlled for.

5.1.4 Robustness check: Source of Change in Contributions Additionally, we investigate the channels behind the changes in contributions in affected districts. First, we explore the time stamps of edits in our sample. Based on the time of contribution, we aggregate edits made during working hours, from Monday to Friday from 8 a.m. to 6 p.m. (local time in Germany), and in the remaining leisure hours of the day, including the weekend.

Appendix Table 18 shows that the pattern for anonymous editing activity the effect of treatment remains similar during both working and leisure hours. The pattern corresponds exactly to our main results during leisure time and during working hours the increase is found for the number of edits, but not for total content generation.[17] Taken together these results suggest that the additional content could be generated by non-Wikipedians who skip the registration procedure (anonymous users) and, with the crisis, spend more time online in both working and leisure hours, and as well as by committed Wikipedians who contribute more during working hours.

We also explore how district-specific characteristics could moderate the treatment effects of interest. Tables 19 and 20 (in Appendix) show that the districts with a larger share of manufacturing in the total output (above the median) are not driving the results. The rise in unemployment touched districts with specialization in manufacturing as well as in services, and in all affected regions Internet behaviour, including contributions of online knowledge, changed. On the contrary, in the districts where the higher shares of households have access to a high-speed Internet (16 Mbit per second and

[17]Registered editors increased their activity only during working time (results not included, but available upon request).

more) contributions to Wikipedia increased stronger than in other affected by the crisis districts. This finding is consistent with our interpretation that the increased readership of Wikipedia channeled knowledge contributions.

5.1.5 Robustness Check: Additional Shock in Part Time Employment (Kurzarbeit)

To shed more light on the mechanism, we exploit the fact that Germany addressed the crisis with a special part time labour program (Kurzarbeit), which was aimed at preventing layoffs.[18] Because of this program, many workers were effectively working considerably shorter hours for many months, without having to look for a new job, and without bearing the high cost of being unemployed. This second type of labor market adjustment, allows us to deepen our analysis in two ways. First, we can analyze whether "normal" unemployment and Kurzarbeit have different effects. Second, districts which experienced both shocks were arguably most strongly affected by the economic crisis.

We analyze the interplay of unemployment and Kurzarbeit by dividing our districts in four groups: (1) completely unaffected districts, (2) districts that experienced an increase in the number of part-time employees, (3) districts that had an increase in unemployment and (4) districts that experienced both shocks. In Table 8 we analyze the four groups of districts separately. The table shows the three groups of districts with at least one shock separately ("no shock" serves as the baseline). The results show that the effects are driven by districts that were affected by both shocks. One shock alone did not result in an increase in contributions to Wikipedia.

This result corroborates our main finding, by highlighting that the result is driven by more strongly affected districts which experienced both shocks. Since neither of the shocks individually has any significant effect, we cannot discern whether anonymous editors increase their activity because they have more spare time, or because they experience the pressure of unemployment. In Appendix Table 22, we show that these results do not change when controlling for economic activity and the industrial structure of a district.

---

[18]Rather than laying off their workers, firms could massively reduce working hours and the state would compensate the workers for a part of their income loss.

Table 7   Robustness: Difference-in-Differences with Pretrends for German Districts (Range: 6 Months Before and 6 Months After the Crisis Start)

	Anonymous Content			Registered Content		
	(1)	(2)	(3)	(4)	(5)	(6)
	Log (Added (KB))	Log (Deleted (KB))	Log (# Edits)	Log (Added (KB))	Log (Deleted (KB))	Log (# Edits)
Treated Districts After T.	0.097*	-0.217*	0.046*	0.017	-0.062	-0.036
	(0.058)	(0.115)	(0.027)	(0.189)	(0.184)	(0.094)
Pretrend	0.098***	0.121***	0.123***	0.215***	0.184***	0.247***
	(0.034)	(0.027)	(0.031)	(0.049)	(0.049)	(0.046)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3120	3120	3120	3120	3120	3120
Number of Districts	240	240	240	240	240	240
R^2	0.11	0.17	0.44	0.06	0.04	0.08

NOTES: The table shows the results of our main difference-in-difference analysis that contrasts affected and unaffected German districts around the onset of the European Financial Crisis in 2009. The columns contain different measures of contribution activity to Wikipedia: (1) the total length of anonymously added content (in KB), (2) the total length of anonymously deleted content (in KB), (3) the number of edits (revisions) by anonymous editors, (4) the total length of content added by registered users (in KB) (5) the total length of content deleted by registered users (in KB), (6) the number of edits (revisions) by registered editors. The variable of interest, which captures the treatment effect, *Treated districts after T* (reatment), 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 time period (month-year) dummies. Standard errors, clustered by districts, are in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 8   Contrasting Unemployment with Part Time Labor (Kurzarbeit) for German Districts.

	Anonymous Content			Registered Content		
	(1)	(2)	(3)	(4)	(5)	(6)
	Log (Added (KB))	Log (Deleted (KB))	Log(# Edits)	Log (Added (KB))	Log (Deleted (KB))	Log(# Edits)
Treated (KA) X After	-0.048	-0.002	0.002	0.294	0.013	0.102
	(0.065)	(0.139)	(0.030)	(0.289)	(0.287)	(0.163)
Treated (UR) X After	-0.019	-0.108	0.047	0.158	0.055	0.063
	(0.088)	(0.181)	(0.044)	(0.287)	(0.270)	(0.141)
Treated (KA+UR) X After	0.158**	-0.171	0.073**	0.229	0.120	0.046
	(0.069)	(0.153)	(0.037)	(0.242)	(0.240)	(0.129)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3120	3120	3120	3120	3120	3120
Number of Districts	240	240	240	240	240	240
R^2	0.06	0.16	0.44	0.01	0.01	0.00

NOTES: The table shows the results of our main difference-in-difference around the *end* of the European Financial Crisis in 2009. The analysis contrasts affected and unaffected German districts,6 months before and after crisis start and the columns contain different measures of contribution activity to Wikipedia: (1) content added by *'anonymous edits'* (length in KB) (2) content deleted by anoymous users (in KB) (3) number of edits (revisions) by anonymous editors (4) content added by registered editors from the district (length in KB) (5) content deleted by 'registered edits' (in KB) (6) number of edits (revisions) by registered editors. The variables of interest, which capture the treatment effect, are *Treated (KA) X After*, *Treated (UR) X After* and *Treated (KA + UR) X After*. *Treated (KA) X After* is an interaction term between the dummy for the districts that are affected only by an increase in part-time labour arrangements (Kurzarbeit) with the time dummy indicating the period after the crisis. Similarly, *Treated (UR) X After* is an interaction term between the dummy for the districts that are affected only by unemployment and the dummy for the time period after the crisis. Finally, *Treated (KA + UR) X After* captures the treatment effect of both treatments: it is an interaction term between a dummy for districts that were affected by both an increase in unemployment and the use of part-time labour arrangements. All specifications include time period (month-year) dummies. Standard errors, clustered by districts, in parentheses:
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Article submitted to *Information Systems Research*; manuscript no.   Revision of ISR-2018-163.R1

Author: *Article Short Title*

5.1.6 Robustness Check: Local Knowledge Contributions to German Wikipedia In

this robustness check we focused on edits *about* a district rather than from editors *from* a district. Specifically, we identify pages about institutions, points of interest and local infrastructures such as bicycle paths, which are clearly specific to each district. Other examples of such pages are monuments, hiking trails, rail tracks, and regional soccer clubs. To identify such pages and match them to the districts we use Wiki-topics (categories) and Wikipedia's hierarchical category tree. We identified all articles that mention the names of any German districts in their titles and recompute all variables based on the resulting set of "local interest pages." We then evaluate if we can observe similar patterns of more activities on these pages as the job market worsened at the onset of the crisis, in Jan 2009. As before, we use changes in the part-time work rates and unemployment rate to distinguish *affected and unaffected* districts.[19]

The results of this estimation are shown in Appendix Table 23 and 24. These tables show the difference-in-differences approach from equation 1, and maintain the structure of our baseline Table 5. Hence, the observed outcomes in Columns 1-6 are the same as in the first six columns of the previous tables: Activity of anonymous users (Cols. 1-3) and activity of registered users (Cols. 4-6). Note however, that we capture the edits of *all* registered users now, since the edits are already matched to each district via the set of (local interest) articles under study. As before, we differentiate between additions (Cols. 1 & 4), deletions (Cols. 2 & 5). The # Edits, (Cols. 3 & 6) and the variable of interest is the cross term.

The findings corroborate the analysis in the previous subsection. A Wikipedia-wide decrease in content generation is mitigated in affected districts with a spike in unemployment and part time labor. As in the main specification, the relative increase of content is driven by an increase in anonymous editing. Our DID-estimates suggest that anonymous edits rose by approximately 6% and the additional amount of content (in Kbytes) rose by approximately 10%. However, confirming the findings in the previous robustness check, these patterns emerge only for the rigorous definition of a

[19]The magnitude of the treatment was considerable, because unemployment rate is approximately 1.1% higher *on average* in the affected districts ($se < 0.05$).

double spike in both the unemployment rate and part time labor (Kurzarbeit). Unlike in the main specification approach, this effect translates to a significant overall effect on all edits (not shown), but, as before, the results on anonymous editing continue to hold when we control for district-specific pre-treatment trends (in Appendix Table 25).

6 Analysis for European Countries

To examine the robustness and the potential to generalize our results from the German dataset, we perform an additional analysis for a sample of European countries. We focus on countries that were affected by the economic crises during the period 2008-2009 and compare them to countries relatively less affected by the economic crises or those where successful policies were implemented to prevent the deepening of the crises. This allows us to examine the relationship of interest in a setting where the variation in the intensity of shocks to unemployment is much higher than across German districts. Another advantage of this analysis is that we can consider total edits to Wikipedia and, hence, we have more and better measures of Wikipedia content, such as views, new words, or active users with various frequency of contributions. We can even observe average edits per article, and the average number of hyperlinks set between articles within Wikipedia and to external web-sites. Different from the previous analysis, we can also leverage the well-established differential timing of different countries entering crisis.

Before presenting the empirical specification and the results (in subsections 6.2 and 6.3), we briefly discuss our data which covers twenty European countries and Wikipedia editions in languages spoken predominantly in those countries.

6.1 Country-level Data

To analyze content generation on the European country level, we combine data on European countries' labour markets with aggregate contributions to various versions of Wikipedia of the corresponding countries. Contributions to Wikipedia come from Wikipedia's monthly statistics for different language editions of Wikipedia 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 between country level unemployment and Wikipedia, we need to focus on countries that have a 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, our efforts to measure 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).

Appendix Table 27 shows the Wikipedia language versions that we could use in this paper. The share of language speakers who live in the corresponding country of origin varies from 71% to 99% (see column 1). To substitute for the Spanish Wikipedia, we add the Catalan version, which is also actively promoted by the Catalan population. We exclude another Spanish region, the Basque Country, because of the elevated activity of automated scripts, or "bots", in the Basque Wikipedia. According to the Wikimedia Foundation, 75% of all edits and 50% of all new articles in the Basque Wikipedia were created by bots in 2009. Bots are active in other Wikipedia editions as well, but not at such a high level.[20] We also exclude Ireland because people in Ireland mostly speak English and only 45% of the speakers of Irish live within the country. The final sample consists of 20 Wikipedia language editions. In addition to the largest European versions of Wikipedia, we include the small ones such as the one for Iceland, which is a country that was heavily affected by the European economic crisis. Beyond European languages, we also include Japanese version of Wikipedia, because Japan is the only country for which we were able to find monthly total working hours in the economy and monthly unemployment rates.[21]

From Wikipedia's language statistics we retrieve seven relevant indicators of user activity: (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

---

[20]http://stats.wikipedia.org/EN/BotActivityMatrixCreates.htm

[21]Japan is not a European country. During the economic crisis, Japan was not severely affected. We include it as a control. Including or excluding Japan from our dataset does not qualitatively change our results.

edits, (4) average edits per article, and (5) the content growth of a corresponding language edition of Wikipedia in terms of words, (6) the number of hyperlinks between the articles in Wikipedia, and (7) the number of references from Wikipedia to external websites. Having several measures for contributions allows us to analyze different aspects of the increased contributions to Wikipedia to both the quantity and quality of the content of the online encyclopedia. For example, the growth in the number of words would indicate more content on Wikipedia, while more edits per article could mean that with increased participation, articles on average got more attention for further improvement.

Appendix Table 28 gives an overview over the countries in the sample. It also clarifies which countries we consider affected or unaffected by the crisis. We consulted the European Central Bank reports about the crisis 2008-2009 specifically to find information about whether a country was affected by drastically increased unemployment or reductions in hours worked, and also when the crisis started. Countries were considered to be affected by the crisis, if they experienced a significant decrease in hours worked, or an increase in unemployment. For the beginning of the shock we looked at the months of 2008 or 2009 mentioned in the reports and also at the country level statistics on hours worked. In the data we would see when there is a significant decrease in hours worked and we would take the second month of a sustained decrease in hours worked as crisis onset month. In our country-level analysis, we choose the crisis start as October 2008 when the crisis officially started for the affected countries. This time point is generally considered the point in time when most counties that suffered from the crisis were officially "in recession." This allows us to compare between affected countries (e.g., Greece and Bulgaria, etc.) and those less affected ones (e.g., Germany and Norway, etc.) who did not enter crisis in that month. This country-level crisis start differs from that of the regional-level analysis of Germany (January, 2009), for whom we have more precise data available.[22] This choice of the timing of the shock for the countries from the control group does not qualitatively

[22]The public debt crisis did not begin in Germany, but in countries that had weaker public finances (e.g. Ireland, Portugal or Eastern European Countries). These countries experienced either a drastic shock in the exchange rate (e.g. Poland, Romania), or a sudden spike in interest rates to refinance their public debt (e.g. Portugal and Greece).

change the final results. Furthermore, we exclude the month in which we estimate the crisis begin

from estimations to make sure our classification procedure does not drive our estimated effects.

Figure 5   Monthly Development of Words Contributed
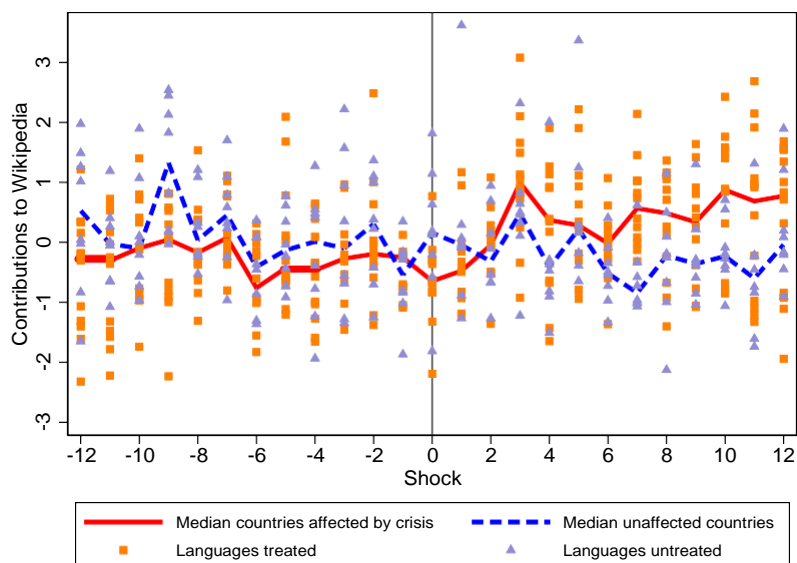
NOTES: The figure shows monthly content growth measured in words added. The median values for affected and unaffected countries across the 20 language versions of Wikipedia in our sample are shown as the two lines. The time spans 12 months before and after the crisis.

Figure 5 gives a descriptive account of one of the key outcomes in Wikipedia: monthly growth

measured in words added. We separately show the medians of the affected and unaffected language

editions of Wikipedia in our sample 12 months before and after the crisis. We show the medians

together with scatter plot of different language versions of Wikipedia. The graph illustrates that,

Some countries experienced both. These events unfolded in fall 2008, then affected each country at a different time. Therefore, we set the crisis start as October 2008 when the crisis officially started. This allows us to compare the affected countries (e.g., Greece and Bulgaria, etc.) with unaffected countries (e.g., Germany and Norway, etc.) Unlike crisis-affected countries in Europe, Germany was not directly affected. The demand for German exports dwindled as a consequence of the changed exchange rates and fiscal consolidation efforts of the affected countries. Hence, German firms and the government only reacted after January 2009 by expanding the program for temporary part time labor, or by laying off staff. This difference is important, because the crisis was not felt and unemployment was not an issue before January 2009 in Germany.

before the crisis, countries that would be affected grew slower than the unaffected countries, whereas after the crisis content growth in the affected countries was faster than in unaffected ones. The patterns are similar for views, edits per article and active Wikipedians, but not for occasional 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. While we cannot easily deal with this issue at the country level, the figure above with the sharp behavioral change gives some confidence that economic crisis played a role in changing people's incentives to contribute to Wikipedia.

6.2 Empirical Approach

At the country level, as with district level in Germany, we estimate the difference-in-differences model. The regression equation is given by:

$$Contributions_{it} = a_i + \beta \left(AfterT_t \times Treated_i \right) + v_t + s_{it} \tag{3}$$

The unit of observation is country i (and its corresponding Wikipedia language edition) in month t. The dependent variable $Contributions_{it}$ measures contributions to Wikipedia (such as page views, the number of editors who contribute minor and major changes, hyperlinks, etc.). We use logarithm transformation on all our measures of contributions to Wikipedia.[23] $AfterT_t$ and $Treated_i$ are dummy variables. $Treated_i$ distinguishes between countries that were stronger or weaker affected by the economic crises. $AfterT_t$ equals one after the month of the shock t_0. As the variable $Treated_i$ does not vary over time, it drops out in the fixed-effects specification. The coefficient of interest is β for the interaction term of these two dummies, which measures the treatment effect of interest. The country-fixed effects a_i, and time-fixed effects v_t are also included in the regressions. Since we have only twenty units of observation at the country level, we use a 24-month interval, which covers 12 months before and 12 months after the onset of the crisis.

[23] Our findings do not change if we normalize with respect to their mean and standard deviation values such that the coefficients represent the changes in the dependent variables in standard deviations.

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. However, we omitted the month when the crisis began in all specifications, because the crisis broke out gradually in most countries rather than on the first day of the month. By leaving out the month of the onset, we ensure that the treatment period is clearly after the onset of the crisis and the pre-treatment period is clearly before the onset. Moreover, omitting that month also gives credit to the notion that an economic crisis typically takes some time to pick up its full momentum.

6.3 Results

The results of the baseline DID estimation are shown in Table 9. Each column shows the results for one of our seven dependent variables measuring activity on Wikipedia: (1) Article Views, (2) Active editors with 5-100 edits (3) Active editors with more than 100 edits, (4) edits per article, (5) growth of total data in Wikipedia, (6) internal links, and (7) external links. The coefficient of interest *Treated countries after T* suggests that, after the shock, there is a 14% increase in the number of active users with few monthly contributions and a 13% increase in active users who heavily edit Wikipedia, contributing more than 100 edits per month. Contributions of new words to language editions of Wikipedia grow by 13%. Articles received on average 6% more edits and 14% more links to external sources of information on the web.

Because the country level analysis allows us to analyze more measures of activity on Wikipedia than at the district-level, we can attempt to take the analysis one level further. Specifically, we can generate additional insight into the mechanisms that enable increased content provision to Wikipedia. In Table 10, we highlight the role of viewership as a key mediating factor by analyzing the relationship between viewership and content growth. The table shows the results when using a fixed-effects panel analysis in which we regress activity on Wikipedia on views over a 24-month period, 12 months before and 12 months after the onset of crisis. Again, in each column of Table 10, we show our six different measures of contributions to Wikipedia.[24]

---

[24](1) number of active Wikipedians (with at least 5 edits), (2) number of very active Wikipedians (with more than 100 edits), (3) average number of edits per article, (4) new words added, (5) number of internal hyperlinks between articles on Wikipedia, and (6) number of references from Wikipedia articles to an external website.

Table 9   Country Level Difference-in-Differences Regressions for the Period of 12 Months Before and 12 Months After the Crisis.

	(1) Ln (Views)	(2) Ln (Active Editors (5-100e))	(3) Ln (Very Act. Editors (≤ 100 e))	(4) Ln(Avg.Edits per Article)	(5) Ln (New Words)	(6) Ln (# Wikilinks)	(7) Ln (# Ext. Links)
After treatment	0.163***	-0.003	-0.009	0.095***	0.039	0.249***	0.309***
	(0.030)	(0.032)	(0.038)	(0.014)	(0.047)	(0.015)	(0.027)
Treated countries after T	0.146**	0.142***	0.127***	0.056**	0.133**	0.057	0.141**
	(0.068)	(0.044)	(0.042)	(0.025)	(0.058)	(0.037)	(0.053)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	432	480	480	480	480	480	480
Number of Countries	20	20	20	20	20	20	20
R^2	0.39	0.20	0.29	0.87	0.11	0.93	0.92

NOTES: The table shows the effect of increased unemployment on contributions to the corresponding country's Wikipedia in Europe for the following indicators: (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) hyperlinks to external sources. All measures of contributions to Wikipedia are in logs, and the month of the estimated crisis onset was omitted from the regressions. 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 time (month - year) fixed effects. Standard errors, clustered by countries, are in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 10

Table 10 **Country Level Content Contributions: Fixed Effects Regressions for the Effect of Views during the Period of 12 Months Before and 12 Months**

After the Crisis.

	(1) Ln (Active Editors (5-100e))	(2) Ln (Very Act. Editors (≤ 100 e))	(3) Ln(Avg.Edits per Article)	(4) Ln (New Words)	(5) Ln (# Wikilinks)	(6) Ln (# Ext. Links)
Log (Views)	0.306***	0.184**	0.020	0.203***	0.150***	0.158***
	(0.103)	(0.086)	(0.026)	(0.058)	(0.049)	(0.050)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	432	432	432	432	432	432
Number of Countries	20	20	20	20	20	20
R^2	0.25	0.22	0.85	0.11	0.94	0.92

NOTES: The table shows the relationship between a country's Wikipedia views and different measures of content contribution to Wikipedia. Fixed Effects Regressions for the Effect of Views during the Period of 12 Months Before and 12 Months After the Crisis The independent variable of interest is *Log Views*. In each column we show a different measure of contribution activity: (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 in logs. All specifications include time (month - year) fixed effects, and exclude the period of treatment. Standard errors, clustered by countries in parentheses: *** $p < \mathbf{0.01}$, ** $p < \mathbf{0.05}$, * $p < \mathbf{0.1}$.

The results in Table 10 confirm that views are a crucial predictor for edit-related outcomes except the number of edits per article. An increase in views by 1% is associated with more active editors (0.31% and 0.18%) and more words in Wikipedia articles (0.2%). Moreover, views are positively related to our measures of content quality, the number of internal links set between Wikipedia articles and external links to information sources. These findings show the role of views as a key mediating factor for additional content generation.

To illustrate the magnitude of the effects we found, we consider the example of Italy, which was one of the countries that was most severely affected by the rise in unemployment. The average unemployment rate in Italy over the observed period was 7.6%, and it increased by 1.2%, which is equivalent to approximately 300,000 additional unemployed people.[25] After the shock, the number of editors with few edits grows by 14%, and, as suggested by our results in Table 9, we would observe 0.14 * 2,440 = 342 additional editors.

6.4 Specification Tests

We ran several tests to check the validity of our specification. First, we decompose the interaction term of interest $AfterT_t \times Treated_i$ into $\sum \mathbb{I}(Year\_Month = t) \times Treated_i$, which allows us to run the DID analysis period by period. The results in Table 30 show the cross-terms between an indicator for treatment and an indicator with a value of 1 if the observation is in a given period. We use monthly intervals for the three months before and four months after treatment and two-month intervals thereafter. The reference period ranges from 4 to 12 months prior to treatment. All coefficients compare the difference between treated and control observations in a given month to the same difference in the reference period. The first three coefficients measure the effect before treatment actually begins and are expected to be nonsignificant. The subsequent coefficients show how the difference between the treated and the controls develops as the treatment takes effect after the shock.

The results suggest that, after the shock in the 12 subsequent months, the viewership of Wikipedia in the treated countries increased 8-20% more rapidly than in the unaffected countries, although the

[25]Based on a workforce of 24.5 million. See http://data.worldbank.org/indicator/SL.TLF.TOTL.IN?locations=IT/.

effect becomes significant with a lag. Similarly, the growth of Wikipedia content, as measured in additional words, grew substantially more rapidly, but was clearly stronger only in the later months after the shock, 12-22%. The number of active users with few monthly edits increased more rapidly in treated countries starting in month 2, and the DID varies approximately 14-23% for affected authors from those who are unaffected. For these measures, the coefficients referring to the months before the shock indicate that the treated observations were not different from the controls before the shock, suggesting proper identification. Very active authors with more than 100 edits contribute more before the onset of the crisis. These results also confirm positive effects for both active users who contribute few and many edits and suggest a pattern in which users tried editing when the crisis began and became active editors a few months later, creating effects that become larger in later months after the shock. In contrast, we find a systematically increasing difference over time for edits per article and links to outside references, and thus we cannot reject that they are simply on different trends. Moreover, the coefficients for the differences before the crisis suggest that the treated countries' Wikipedia editions have fewer users, especially for active users with few monthly edits, than the control ones.

In the Appendix, we introduce a few additional specification tests to examine the robustness of the results. First, we control for country-specific trends in various measures of Wikipedia content development before the shock (Table 29). Even after accounting for country-specific trends, our results hold on the number of new active users with small monthly contributions, for monthly views, and for contributed text ("word growth").

In the next specification test, we verify that the unemployment rate was not positively correlated with contributions *before* the crisis. This is important, because the crisis is likely to hit weaker economies harder. If contributions to Wikipedia were correlated with unemployment before the crisis, then we could not exploit the European economic crisis to study how an increase in unemployment affects contributions to Wikipedia. We would simply capture the preexisting correlation and erroneously attribute it to the crisis. Hence, in Table 31, we show an OLS fixed-effects panel regression

of contributions on unemployment 12 months prior to the crisis. The table lists the regression coefficients of the independent variable of interest, *Unemployment*, on our measures of contributions to Wikipedia. All specifications include time dummies. The coefficient of the variable of interest, the unemployment rate, is not significantly different from zero. We consider this to offer evidence of no correlation between unemployment and contributions to Wikipedia before the shock.

Finally, we check the robustness of our OLS approach by using the rate of unemployment among young people (15-24 years old) as an explanatory variable. As expected, young people are more likely to use the internet and, consequently, to contribute to online public goods than older generations. The results (see Table 32) suggest that both magnitude and significance of the unemployment effect are larger for the young population.

Our country-level analysis adds value to our baseline results from German districts, because it allows us to exploit greater variation in the changes in economic hardship and proves that the results we find hold not only for Germany but can be viewed as generally valid across many countries.

7   Discussion, Limitations, and Further Research

Human history includes many instances of social and economic advancement. Each such social and economic restructuring brings new methods of production and consumption, but one unfortunate consequence of the progress is that workers with older skills are displaced from their jobs. Our results support that human beings respond to structural economic shifts by reallocating their time to peer production-related activities.

In this paper, we analyze the relationship between unemployment and the provision of online public goods during the economic crisis in Europe in 2008-2009. We find that higher unemployment is associated with higher participation by volunteers in Wikipedia and an increased rate of content generation. Our findings are based on a comparison across German districts and across European countries. We exploit that some districts/countries were affected by relatively large increases in unemployment while others were not, and we show that, Wikipedia articles were read more frequently in areas where unemployment increased. The increase in readership was followed by more edits of

anonymous or casual editors ("beginners"), and subsequently the number of highly active users grew. Over time, we observe increased content growth.

Our main analysis is based on a comparison of German districts and a country-level analysis of European economies. At the German district level, districts with higher increases in unemployment had relatively more contributions than less-affected districts, when the overall downward trend in total contributions to German-language Wikipedia during the crisis are controlled for. Our estimates suggest that we observe almost 4,340 additional edits and 55 megabytes over the six-month period after the shock. At the European level, contribution growth was more rapid where the crisis hit harder. We stress several aspects of our findings: First, the effects are consistently found for edits *from* German districts (Table 5), for edits *about* German districts (Table 24), and at the European level (Table 9). Second, the pattern of contributions in Germany and at the European level aligns, because we find that a downward trend in generally less-affected Germany was partially mitigated in districts with higher unemployment. Finally, the effects are sizable and the estimated increases in editing activities typically range from 5% to 20%. For example, our country-level analysis suggests that the number of casual editors (with 5 to 100 edits/month) grew by 9.5% to 14%, as suggested by our results in Tables 9 and 29. For the Italian-language Wikipedia, this would mean about 300 additional editors, with 5 to 100 edits every month. Thus, the overall effect suggests that the threat of unemployment is associated with increased online contribution of public goods. Because Wikipedia functions as an important knowledge base for the economy, our results document a new and somewhat valuable side effect of the economic crisis.

Our findings open up further questions. Potentially, higher unemployment may be associated with greater volunteering and productive time usage. Research about pro-social behavior has documented a positive impact of regular volunteering on subjective well-being, and happiness, which even increased over time if regular volunteering is sustained (Borgonovi 2008, Binder and Freytag 2013).[26] Moreover,

---

[26]Related research showed that prosocial giving and donations were found to have a strong effect on happiness, and they appear to be deeply ingrained human behaviors (Dunn et al. 2008, Aknin et al. 2013). Separate research found happiness to be a strong predictor of health (Sabatini 2014), and there is a strong and more general relationship between intergenerational transfers and aging in social species (Kaplan and Robson 2002, Lee 2003)

giving and volunteering can have particularly strong positive effects in older and unhappy people, for example via esteem enhancement (Post 2005, Cattan et al. 2011, Binder and Freytag 2013, Liang et al. 2001).

The relationship of unemployment and online volunteering for Wikipedia highlights a positive energy that might fruitfully be channeled into Wikipedia and similar projects in times of economic crisis. These results are in line with previous findings about the positive effects of volunteering. As such, online volunteering could complement public employment schemes, which are frequently considered useful for inclusion and social skills, but are also expensive and of doubtful efficiency (Fervers 2018). Contributing to Wikipedia might be a way that might enhance the cognitive skills and computer skills of those who decide to edit. Likewise, by growing a project that helps millions of others, such contributors might improve their well-being and self-esteem.

Note, however, that the positive effects of prosocial actions have been shown to depend on the volunteer's autonomous motivation for helping (Weinstein and Ryan 2010). Moreover, the entry threshold for prospective contributors might be high. These findings should be taken into account when designing any interventions that try to leverage the potential of online volunteering. A way to do so would be a coordinated effort of local authorities and the Wikipedia community to integrate prospective new users who are temporarily unemployed, but would volunteer to contribute local knowledge to Wikipedia (or contribute their time in similar constructive ways). Moreover, we cannot fully answer how this mechanism works. In particular, it seems that new editors begin to acquire new capabilities and devote their time to contributing to online public goods. As more new articles are created in Wikipedia every day, the increased participation is focused on adding to existing knowledge as well as introducing new topics.

While it is hard to draw definitive conclusions on the identity of the new editors, several of our findings are in line with a mechanism by which the crisis motivated new users to begin to edit and, over time, existing participating editors also increased their activity. First, the effect is strongest in districts that are most severely affected, by both unemployment shock and the adoption of the Kurzarbeit

program. Second, the increase is clearly driven by anonymous users, but a large share of the edits seem to be more valuable than other anonymous edits.[27] While the number of reverted edits increases, the relative share of reverted edits does not significantly increase. Moreover, there are no robust effects on content deletions. Third, editing by anonymous users precede more registrations of new users at the European level, and we provide some evidence that the content generation comes from edits that are made during leisure time in Germany. Meanwhile, we found much weaker effects in the number of hyperlinks between Wikipedia articles and to external sources. When pretrends are controlled for, the coefficient for Wikipedia links even turns negative, and the coefficient for external links shows nonsignificant effects. These findings might indicate that the observed increase in content generated is due to the activity of inexperienced contributors or minor edits. Unfortunately, fewer links in the content might also suggest a decrease in content quality, which is then a negative consequence of more activity.[28] In the German data, this pattern is mirrored by a small increase in unproductive edits that had to be reverted, but not in the share of these edits in all edits. This indicates that there was no negative effect on the quality of content production due to unemployment. Taken together, these results suggest that a share of the population from regions that faced unemployment first increased their use of online knowledge repository, Wikipedia, and afterward began to contribute to the public knowledge good.

Although 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 use the economic crisis as a source of exogenous variation in the economic state and the unemployment rates. This strategy is based on the following identification assumptions. First, contributions to Wikipedia should not be correlated with the likelihood that the countries would be affected by the crisis. A specification test

[27]While we also find a slightliy increased frequency of reverting (which suggests a lack of experience), previous research on Wikipedia by Anthony et al. (2009), suggests that a large share of anonymous editors are occasional contributors with shorter but very reliable contributions, especially when they are domain experts.

[28]Wikipedia insists that all contributions be based on verifiable knowledge. Hence, the number of external links is an indication that the contributions are backed by reliable sources.

in the Appendix (Table A5) provides first supporting evidence for this assumption. Second, using districts and countries as controls requires that the various Wikipedia editions be sufficiently similar and that the districts/countries be somewhat homogeneous with respect to economic and social developments in the period of observation. The qualitatively similar findings in both the country-level and the district-level analyses offer some confidence. Clearly, the institutional, macroeconomic, and political setup is more homogeneous for German districts than for European countries. Also, assuming similar Wikipedia editions is no problem for German districts, because the Wikipedia under study is the *same*.

However, two concerns remain about this analysis, because the regional analysis is based on the IP-addresses of anonymous contributions. First, using IP addresses of anonymous contributors allowed for only a restricted set of available dependent variables. Specifically, we can only determine two measures of the efforts in any given region, that is we have to focus on the *number of edits* and the *length of edits* in kilobytes. This is because computing statistics such as the *number of active editors*, or *edits per article* becomes meaningless when we can observe only a part of the edits (from anonymous editors and editors with self-reported location). Second, the use of IP addresses implies that we can look only at a specific set of contributions. These contributions most likely come from new or occasional users, because experienced users typically edit under their user name. In our research, we mitigate this problem by examining edits by registered users who reveal their (self-reported) location.

In addition, even if our identifying assumptions are satisfied, we can only provide indicative evidence on whether the additional content generation is driven by unemployed or employed users. Although this question cannot even be addressed for the country level, for the German district-level data, we could provide an indication that additional edits are more likely during leisure time.[29] Still, it remains unclear as to whether the employed users increase their activity during leisure hours or whether the unemployed prefer to contribute during the hours that we classified as leisure times.

[29]In line with models from labor economics, household production could be shifted to working hours, and additional activities of the recently unemployed would thus be observed during leisure times.

Another fruitful avenue for further research could investigate what is actually written. This question has to remain unanswered at the current stage of our research. Perhaps people are simply writing *about* the crisis. Such a pattern seems unlikely, given the overall growth that we observe. However, a smaller or larger proportion 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 discover who makes the edits, we would need user- and editor-level data, and, to see what they write, we would have to analyze articles for their content. Further research could analyze the nature of contributions and the type of articles that are edited. Also, the extent to which district-specific articles are improved or whether articles related to affected professions are edited would be very interesting. These questions are beyond the scope of this paper and, especially at the article level, this analysis is computationally intensive but might lead to interesting additional insights in further research. At the country level, such an analysis is almost unthinkable though, because the data available are too highly aggregated. More fine-grained data, ideally at the user level, would allow us to look at which information is searched and which edits are made. Beyond that, we could contrast Wikipedia editing activities with other ways in which the newly unemployed use their additional time.

8 Conclusion

In this paper, we study how individuals reallocate their time to the provision of online public goods when faced with increased rates of unemployment. We observe a moderate increase in socially valuable volunteering in the form of contributions to Wikipedia. The findings are consistent at both the European and the German district-level. The patterns are suggestive of a creative and constructive potential that is freed up as a positive side-effect of unemployment, that might carry over to job displacement in general.

The question whether unemployment can result in an increased provision of public (online) goods and private learning is crucial, given that we observe accelerating labor substitution due to digitization. If a part of the liberated capacity results in increased knowledge documentation and generation, this may be a positive surprise. Even though it remains to be tested how easily the effects we found carry over to digitization-induced job displacement, we highlight a constructive pattern. Our findings show that individuals reacted to increased unemployment during the European economic crisis by reallocating their time to production-related activities, and contributing more to Wikipedia.

The results obtained in our study have important practical and policy implications. During the recession that started in late 2008, UK newspapers found increased volunteering. For example, BBC News reported that volunteering agencies such as Community Service Volunteers (CSV) and Youth-Net saw increased inquiries and applications.[30] But at the same time, the Guardian questions if high youth unemployment may pose a serious challenge for charities to recruit and retain volunteers.[31] Similarly, in the USA, while the Washington Post argued that volunteering had increased despite the recession, the New York Times found decreased volunteering during the same time.[32] These newspaper findings were mostly based on data on offline and formal volunteering and missed important information about online and informal volunteering. To reconcile the mixed findings, Lim and Laurence (2015), using a large survey dataset, suggest that overall formal and informal volunteering both declined in the UK since 2008. They found that the decline is more salient in communities that suffer from social and economic disadvantages and this decline cannot be explained by individual hardship following financial insecurity. Our paper differs from the prior work in two important ways: first, our dataset is based on observed rather than self-reported contributions. Second, we focus on online

[30]http://news.bbc.co.uk/2/hi/uk_news/8008428.stm, accessed November 2018.

[31]https://www.theguardian.com/voluntary-sector-network/2012/jun/07/unemployment-charities-volunteering, accessed November 2018.

[32]see http://www.washingtonpost.com/wp-dyn/content/article/2010/06/15/AR2010061501449.html?noredirect=on and https://www.nytimes.com/2009/08/27/us/27volunteer.html, respectively. Accessed November 2018.

and informal contributions. These two differences allow us to offer additional and complementary evidence. While in the offline world, people's contributions to volunteering work may increase their individual employability, volunteering online through contributions to Wikipedia may not have direct individual benefits, but it creates a positive externality of generating useful knowledge.

The valuable knowledge contributions that occur due to worsened economic conditions represent a beneficial side effect that has been overlooked by policymakers. Our results suggest that measures aimed at encouraging knowledge contributions could have a strong effect on volunteering. Policymakers could enhance these beneficial effects by encouraging the active groups in the society to contribute online knowledge more systematically. While financial aids may jump-start economic recovery, government support on contributions to online public goods will provide valuable channels for knowledge exchange and help encourage skill upgrades in the society, turning a crisis into an opportunity. To this end, it is important to better measure online and informal contributions (in addition to those reported by traditional volunteering agencies such as YouthNet and the Red Cross) so that the increased contributions to knowledge can be better captured in official statistics.

Acknowledgments

References

Aguiar, Mark, Erik Hurst. 2007. Measuring trends in leisure: the allocation of time over five decades. *The Quarterly Journal of Economics* **122**(3) 969–1006.

Aguiar, Mark, Erik Hurst, Loukas Karabarbounis. 2012. Recent developments in the economics of time use. *Annu. Rev. Econ.* **4**(1) 373–397.

Aguiar, Mark, Erik Hurst, Loukas Karabarbounis. 2013. Time use during the great recession. *The American Economic Review* **103**(5) 1664–1696.

Aknin, Lara B, Christopher P Barrington-Leigh, Elizabeth W Dunn, John F Helliwell, Justine Burns, Robert Biswas-Diener, Imelda Kemeza, Paul Nyende, Claire E Ashton-James, Michael I Norton. 2013. Prosocial spending and well-being: Cross-cultural evidence for a psychological universal. *Journal of Personality and Social Psychology* **104**(4) 635.

Algan, Yann, Yochai Benkler, Mayo Fuster Morell, Jérôme Hergueux. 2013. Cooperation in a peer production economy experimental evidence from wikipedia. *Workshop on Information Systems and Economics, Milan, Italy*. 1–31.

Andreoni, James. 1988. Privately provided public goods in a large economy: the limits of altruism. *Journal of Public Economics* **35**(1) 57–73.

Andreoni, James. 1989. Giving with impure altruism: applications to charity and ricardian equivalence. *The Journal of Political Economy* 1447–1458.

Andreoni, James. 1990. Impure altruism and donations to public goods: a theory of warm-glow giving. *The economic journal* 464–477.

Andreoni, James. 2007. Giving gifts to groups: How altruism depends on the number of recipients. *Journal of Public Economics* **91**(9) 1731–1749.

Anthony, Denise, Sean W Smith, Timothy Williamson. 2009. Reputation and reliability in collective goods: The case of the online encyclopedia wikipedia. *Rationality and Society* **21**(3) 283–306.

Binder, Martin, Andreas Freytag. 2013. Volunteering, subjective well-being and public policy. *Journal of Economic Psychology* **34** 97–119.

Borgonovi, Francesca. 2008. Doing well by doing good. the relationship between formal volunteering and self-reported health and happiness. *Social science & medicine* **66**(11) 2321–2334.

Burda, Michael C, Daniel S Hamermesh. 2010. Unemployment, market work and household production. *Economics Letters* **107**(2) 131–133.

Cattan, Mima, Eddy Hogg, Irene Hardill. 2011. Improving quality of life in ageing populations: What can volunteering do? *Maturitas* **70**(4) 328 – 332. doi:https://doi.org/10.1016/j.maturitas.2011.08.010. URL http://www.sciencedirect.com/science/article/pii/S0378512211003197.

Chen, Yan, F Maxwell Harper, Joseph Konstan, Sherry Xin Li. 2010. Social comparisons and contributions to online communities: A field experiment on movielens. *American Economic Review* **100**(4) 1358–98.

Comino, Stefano, Fabio M Manenti, Maria Laura Parisi. 2007. From planning to mature: On the success of open source projects. *Research Policy* **36**(10) 1575–1586.

De Chaisemartin, Clément, Xavier D'Haultfoeuille. 2017. Fuzzy differences-in-differences. *The Review of Economic Studies* **85**(2) 999–1028.

Dunn, Elizabeth, Lara Aknin, Michael I Norton. 2008. Spending money on others promotes happiness. *Science (New York, N.Y.)* **319** 1687–8. doi:10.1126/science.1150952.

Fervers, Lukas. 2018. Can public employment schemes break the negative spiral of long-term unemployment, social exclusion and loss of skills? evidence from germany. *Journal of Economic Psychology* **67**(C) 18–33.

Freeman, Richard B. 1997. Working for nothing: The supply of volunteer labor. *Journal of Labor Economics* **15**(1, Part 2) S140–S166.

Goldfarb, Avi, Jeff Prince. 2008. Internet adoption and usage patterns are different: Implications for the digital divide. *Information Economics and Policy* **20**(1) 2–15.

Hess, Charlotte, Elinor Ostrom. 2003. Ideas, artifacts, and facilities: information as a common-pool resource. *Law and contemporary problems* 111–145.

Kandel, Eugene, Edward P Lazear. 1992. Peer pressure and partnerships. *Journal of political Economy* 801–817.

Kaplan, Hillard S., Arthur J. Robson. 2002. The emergence of humans: The coevolution of intelligence and longevity with intergenerational transfers. *Proceedings of the National Academy of Sciences* **99**(15) 10221–10226. doi:10.1073/pnas.152502899. URL https://www.pnas.org/content/99/15/10221.

Knabe, Andreas, Steffen Rätzel, Ronnie Schöb, Joachim Weimann. 2010. Dissatisfied with life but having a good day: Time-use and well-being of the unemployed*. *The Economic Journal* **120**(547) 867–889.

Krueger, Alan B, Andreas I Mueller. 2012. Time use, emotional well-being, and unemployment: Evidence from longitudinal data. *The American Economic Review* **102**(3) 594–599.

Lee, Ronald D. 2003. Rethinking the evolutionary theory of aging: Transfers, not births, shape senescence in social species. *Proceedings of the National Academy of Sciences* **100**(16) 9637–9642. doi:10.1073/pnas.1530303100. URL https://www.pnas.org/content/100/16/9637.

Liang, Jersey, Neal M Krause, Joan M Bennett. 2001. Social exchange and well-being: is giving better than receiving? *Psychology and aging* **16**(3) 511.

Lim, Chaeyoon, James Laurence. 2015. Doing good when times are bad: volunteering behaviour in economic hard times. *The British journal of sociology* **66**(2) 319–344.

List, John A. 2011. The market for charitable giving. *Journal of Economic Perspectives* **25**(2) 157–80.

Pissarides, Christopher A. 1992. Loss of skill during unemployment and the persistence of employment shocks. *The Quarterly Journal of Economics* **107**(4) 1371–1391.

Post, Stephen G. 2005. Altruism, happiness, and health: It's good to be good. *International journal of behavioral medicine* **12**(2) 66–77.

Sabatini, Fabio. 2014. The relationship between happiness and health: evidence from italy. *Social Science & Medicine* **114** 178–187.

Taniguchi, Hiromi. 2006. Men's and women's volunteering: Gender differences in the effects of employment and family characteristics. *Nonprofit and Voluntary Sector Quarterly* **35**(1) 83–101.

Uslaner, Eric M. 2002. Religion and civic engagement in canada and the united states. *Journal for the Scientific Study of Religion* **41**(2) 239–254.

Vesterlund, Lise. 2006. Why do people give. *The nonprofit sector: A research handbook* **2** 168–190.

Wallsten, Scott. 2013. What are we not doing when we're online. Tech. rep., National Bureau of Economic Research.

Walz, G, BJ Buiskool, M Collewet, J de Koning, Country expert France, O Calavrezo. 2012. Short-time working arrangements during the crisis and lessons to learn. *Zoetermeer: Panteia* .

Weinstein, Netta, Richard M Ryan. 2010. When helping helps: autonomous motivation for prosocial behavior and its influence on well-being for the helper and recipient. *Journal of personality and social psychology* **98**(2) 222.

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

A Comparison of the Treatment and Control Group

Table 11   Summary Statistics on Articles by Treatment Status, before Treatment.

	Treated	Control	All
GDP/Capita (1000 EUR)	27.91 (11.17)	31.41 (15.46)	29.66 (13.59)
Inhabitants (1000s)	143.4 (97.74)	238.2 (173.1)	190.7 (148.2)
# Unemployed (1000s)	6.205 (5.531)	9.409 (9.411)	7.804 (7.877)
Unemployment Rate (%)	7.426 (4.031)	7.519 (3.364)	7.473 (3.712)
Unemployment Rate Men	7.668 (4.333)	8.249 (3.845)	7.958 (4.106)
Unemployment Rate Women	8.927 (4.559)	8.548 (3.638)	8.738 (4.128)
Youth Unemployment (age ≤ 25,%)	7.089 (4.274)	7.265 (3.671)	7.177 (3.984)
# Businesses using Kurzarbeit	17.35 (23.81)	17.13 (17.64)	17.24 (20.96)
# Part time employees	276.7 (488.1)	225.7 (377.5)	251.2 (437.0)
Change in Unemployment rate (%)	1.855 (0.474)	0.331 (0.325)	1.095 (0.864)
Dummy: Affected	1 (0)	0 (0)	0.501 (0.500)
# Registered Users	1.666 (4.992)	5.302 (11.58)	3.480 (9.089)
# Edits from District	219.5 (381.9)	614.1 (818.0)	416.4 (667.5)
# Anonymous Edits	126.9 (129.2)	303.5 (325.7)	215.0 (262.8)
# Registerd Edits	92.63 (317.0)	310.6 (679.0)	201.4 (540.4)
Total Added Content (KB)	82.47 (198.5)	226.3 (329.1)	154.2 (280.9)
Total Added Content – Anonymous (KB)	48.23 (130.1)	110.6 (152.0)	79.35 (144.8)
Total Added Content – Registered (KB)	34.24 (134.3)	115.7 (264.6)	74.87 (213.5)
Observations	1425		

NOTES: The table shows the means of the main variables by treatment status in the months before treatment was applied. The unit of observations is district (Kreis) i in month t. Column 1 shows treated districts and Column 2 the control group. The time variable is normalized and runs from -6 to -1.; no. of obs. = 1425; no. of districts in estimation = 240. Standard errors in parenthesis.

B   Robustness Checks for the Main Results

Table 12   Fuzzy Differences-in-Differences Regressions for German Districts.

	LATE	Std. Error	[95% Conf. Interval]	
Dependent Variable: Log (Added (KB))				
Wald-DID	0.088	0.084	-0.054	0.237
Wald-TC	0.100	0.087	-0.044	0.270
Dependent Variable: Log (# Edits)				
Wald-DID	0.052	0.034	0.007	0.120
Wald-TC	0.068	0.035	0.016	0.146

NOTES: The table displays the results of Wald-DID and time-corrected Wald ratio. Sample size 3,120 observations. Standard errors are clustered at the district level.

Table 13   Fixed Effects Regressions for the Relationship Between Unemployment and Activity on Wikipedia.

	Anonymous Content			Registered Content		
	(1)	(2)	(3)	(4)	(5)	(6)
	Log (Added (KB))	Log (Deleted (KB))	Log (# Edits)	Log (Added (KB))	Log (Deleted (KB))	Log (# Edits)
Unemployment Rate (%)	0.034	0.094*	0.027**	-0.015	-0.028	-0.023
	(0.023)	(0.054)	(0.011)	(0.061)	(0.059)	(0.030)
GDP/Capita	0.003	-0.015	-0.002	-0.042	-0.075**	-0.030
	(0.011)	(0.026)	(0.006)	(0.041)	(0.037)	(0.021)
Share Empl. Manuf.	-0.150	-0.131	-0.060	-0.140	-0.091	-0.104
	(0.172)	(0.369)	(0.090)	(0.724)	(0.485)	(0.277)
Share Empl. Serv.	-0.121	-0.195	-0.036	-0.116	-0.151	-0.121
	(0.171)	(0.358)	(0.089)	(0.713)	(0.466)	(0.270)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	5226	5226	5226	5226	5226	5226
Number of Districts	402	402	402	402	402	402
R^2	0.05	0.17	0.44	0.01	0.01	0.01

NOTES: The table shows the results of a fixed-effects analysis that directly relates Wikipedia contributions to the unemployment rate of German districts around the onset of the European Financial Crisis in 2009. Observations range from 6 Months Before to 6 Months After the Crisis Start. The columns contain different measures of contribution activity to Wikipedia: (1) the total length of anonymously added content (in KB), (2) the total length of anonymously deleted content (in KB), (3) the number of edits (revisions) by anonymous editors, (4) the total length of content added by registered users (in KB) (5) the total length of content deleted by registered users (in KB), (6) the number of edits (revisions) by registered editors. The independent variable of interest is *Unemployment rate* in period t for each district. All specifications include time period (month-year) dummies. Standard errors, clustered by districts, are in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 14   Difference-in-Differences Regressions for German Districts (Range: 6 Months Before and After the Crisis Start).

	Anonymous Content			Registered Content		
	(1)	(2)	(3)	(4)	(5)	(6)
	Log (Added (KB))	Log (Deleted (KB))	Log (# Edits)	Log (Added (KB))	Log (Deleted (KB))	Log (# Edits)
Treated Districts After T.	0.109*	-0.139	0.059**	-0.020	-0.038	-0.044
	(0.056)	(0.112)	(0.028)	(0.193)	(0.186)	(0.101)
GDP/Capita	0.005	-0.020	0.002	-0.025	-0.093*	-0.040
	(0.013)	(0.033)	(0.008)	(0.057)	(0.051)	(0.029)
Share Empl. Manuf.	-0.037	-0.201	-0.025	0.030	-0.469	-0.142
	(0.204)	(0.459)	(0.115)	(0.992)	(0.646)	(0.387)
Share Empl. Serv.	-0.007	-0.270	-0.011	0.093	-0.505	-0.144
	(0.200)	(0.436)	(0.114)	(0.989)	(0.634)	(0.384)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3120	3120	3120	3120	3120	3120
Number of Districts	240	240	240	240	240	240
R^2	0.06	0.16	0.44	0.01	0.01	0.01

NOTES: The table shows the results of our main difference-in-difference analysis that contrasts affected and unaffected German districts around the onset of the European Financial Crisis in January 2009. The columns contain different measures of contribution activity to Wikipedia: (1) the total length of anonymously added content (in KB), (2) the total length of anonymously deleted content (in KB), (3) the number of edits (revisions) by anonymous editors, (4) the total length of content added by registered users (in KB) (5) the total length of content deleted by registered users (in KB), (6) the number of edits (revisions) by registered editors. The variable of interest, which captures the treatment effect, *Treated districts after T* (reatment), 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 time period (month-year) dummies. Standard errors, clustered by districts, are in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

B.1   Analysis of Reverts

Table 15   **Difference-in-Differences Regressions for German Districts (Range: 6 Months Before and After the Crisis Start): Analysis of Reverted Content.**

	Anonymous Users		Registered Users	
	(1)	(2)	(3)	(4)
	Log (Deleted (KB))	Log (# Edits)	Log (Deleted (KB))	Log (# Edits)
Treated Districts After T.	0.188	0.054*	0.111	0.004
	(0.261)	(0.030)	(0.112)	(0.017)
Time FE	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes
Observations	3120	3120	3120	3120
Number of Districts	240	240	240	240
R^2	0.08	0.08	0.00	0.00

NOTES: The table shows the results of our main difference-in-difference analysis that contrasts affected and unaffected German districts around the onset of the European Financial Crisis in January 2009. The columns contain different measures of malicious content deletions on Wikipedia: (1) the total length of anonymously deleted content (in KB), (2) the number of edits (revisions) by anonymous editors deleting Wikipedia articles, (3) the total length of content deleted by registered users (in KB), (4) the number of edits (revisions) by registered editors deleting Wikipedia articles. The variable of interest, which captures the treatment effect, *Treated districts after T* (reatment), 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 time period (month-year) dummies. Standard errors, clustered by districts, are in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 16   Difference-in-Differences Regressions for German Districts (Range: 6 Months Before and After the Crisis Start): Analysis of Reverted Content

	Anonymous Users		Registered Users	
	(1)	(2)	(3)	(4)
	Share (Deleted (KB))	Share (# Deleting Edits)	Share (Deleted (KB))	Share (# Deleting Edits)
Treated Districts After T.	-0.005	0.000	0.010	0.003
	(0.011)	(0.001)	(0.008)	(0.002)
Time FE	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes
Observations	3120	3120	1676	1685
Number of Districts	240	240	152	153
R^2	0.03	0.01	0.01	0.01

NOTES: The table shows the results of our main difference-in-difference analysis that contrasts affected and unaffected German districts around the onset of the European Financial Crisis in January 2009. The columns contain different measures of malicious content deletions on Wikipedia: (1) the total length of anonymously added content (in KB), the total length of anonymously deleted content (in KB), (3) the number of edits (revisions) by anonymous editors, (4) the total length of content added by registered users (in KB). The variable of interest, which captures the treatment effect, *Treated districts after T* (reatment), 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 time period (month-year) dummies. Standard errors, clustered by districts, are in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

B.2 Sensitivity Analysis

Table 17   **Difference-in-Differences Regressions for German Districts (Range: 6 Months Before and After the Crisis Start): Sensitivity Analysis.**

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dropped	+/-0%	+/-5%	+/-10%	+/-15%	+/-20%	+/-25%	+/-30%
Treated districts after T	0.024	0.028	0.038*	0.046*	0.060**	0.055*	0.072**
	(0.020)	(0.021)	(0.023)	(0.024)	(0.027)	(0.030)	(0.035)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	5213	4693	4173	3653	3120	2600	2080
Number of Districts	401	361	321	281	240	200	160
R^2	0.44	0.46	0.45	0.44	0.44	0.43	0.42

NOTES: The table presents the results of our main difference-in-difference analysis varying the cutoff percentiles for the German districts affected by the crisis vs those that are less affected. The columns contain the results for different cutoff points. The dependent variable of the sensitivity analysis is the number of edits made by anonymous editors in German districts. The variable of interest, which captures the treatment effect, *Treated districts after T* (reatment), 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 time period (month-year) dummies. Standard errors, clustered by districts, in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

B.3 Heterogeneity of Treatment Effects

Table 18 Difference-in-Differences Regressions for German Districts for Content Contributed in Working Time and Leisure Time of The Week.

	Working Time - Anonymous Content			Leisure Time - Anonymous Content		
	(1)	(2)	(3)	(4)	(5)	(6)
	Log (Added (KB))	Log (Deleted (KB))	Log (# Edits)	Log (Added (KB))	Log (Deleted (KB))	Log (# Edits)
Treated Districts After T.	0.096	-0.023	0.057*	0.125*	-0.070	0.071**
	(0.074)	(0.136)	(0.031)	(0.068)	(0.146)	(0.031)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3120	3120	3120	3120	3120	3120
Number of Districts	240	240	240	240	240	240
R^2	0.04	0.17	0.37	0.05	0.11	0.33

NOTES: The table shows the results of our main difference-in-difference analysis that contrasts affected and unaffected German districts around the onset of the economic crisis in January, 2009. Here we decompose all contributed content into content contributed in working time (9am - 6pm on worling days) and in leisure time of the week (6pm - 9am and weekends). The columns contain different measures of contribution activity to Wikipedia during working and leisure time: first, in the panel of results for working time, the number of (1) bytes added and (2) deleted by anonymous editors, as well as (3) anonymous edits; in the second panel the results for leisure time include the same set of content contribution measures. The variable of interest, which captures the treatment effect, *Treated districts after T* (reatment), 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. Observations range from 6 Months before to 6 Months after the crisis started. All specifications include time period (month-year) dummies. Standard errors, clustered by districts, in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 19   Robustness: Difference-in-Difference-in-Differences for German Districts with Different Intensity of Manufacturing.

	Anonymous Content			Registered Content		
	(1)	(2)	(3)	(4)	(5)	(6)
	Log (Added (KB))	Log (Deleted (KB))	Log (# Edits)	Log (Added (KB))	Log (Deleted (KB))	Log (# Edits)
Treated Districts After T.	0.099	-0.222*	0.060**	-0.031	-0.070	-0.029
	(0.064)	(0.131)	(0.030)	(0.199)	(0.193)	(0.102)
Treated with Manufacturing after T	0.034	0.249	-0.003	0.064	0.110	-0.040
	(0.080)	(0.176)	(0.050)	(0.356)	(0.276)	(0.164)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3120	3120	3120	3120	3120	3120
Number of Districts	240	240	240	240	240	240
R^2	0.06	0.16	0.44	0.01	0.01	0.00

NOTES: The table shows the results of our main difference-in-difference analysis that contrasts affected and unaffected German districts around the onset of the European Financial Crisis in 2009. The columns contain different measures of contribution activity to Wikipedia: (1) the total length of anonymously added content (in KB), (2) the total length of anonymously deleted content (in KB), (3) the number of edits (revisions) by anonymous editors, (4) the total length of content added by registered users (in KB) (5) the total length of content deleted by registered users (in KB), (6) the number of edits (revisions) by registered editors. The variable of interest, which captures the treatment effect, *Treated districts after T* (reatment), 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. Observations range from 6 Months before to 6 Months after the crisis started. All specifications include time period (month-year) dummies. Standard errors, clustered by districts, are in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 20   Robustness: Difference-in-Difference-in-Differences for German Districts with Different Availability of High-Speed Internet.

	Anonymous Content			Registered Content		
	(1)	(2)	(3)	(4)	(5)	(6)
	Log (Added (KB))	Log (Deleted (KB))	Log (# Edits)	Log (Added (KB))	Log (Deleted (KB))	Log (# Edits)
Treated Districts After T.	0.112*	-0.163	0.048	-0.005	-0.081	-0.053
	(0.058)	(0.118)	(0.029)	(0.206)	(0.194)	(0.107)
Treated with Internet after T	-0.018	0.130	0.083*	-0.045	0.315	0.083
	(0.091)	(0.261)	(0.042)	(0.392)	(0.298)	(0.139)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3120	3120	3120	3120	3120	3120
Number of Districts	240	240	240	240	240	240
R^2	0.06	0.16	0.44	0.00	0.01	0.00

NOTES: The table shows the results of our main difference-in-difference around the *end* of the European Financial Crisis in 2009. The analysis contrasts affected and unaffected German districts, and the columns contain different measures of contribution activity to Wikipedia: (1) the total length of anonymously added content (in KB), (2) the total length of anonymously deleted content (in KB), (3) the number of edits (revisions) by anonymous editors, (4) the total length of content added by registered users (in KB) (5) the total length of content deleted by registered users (in KB), (6) the number of edits (revisions) by registered editors. The two variables of interest, which capture the treatment effect, are *Treated districts after T*(reatment), an interaction term between dummies for the districts that are affected by the crisis with the time dummy indicating the period after the crisis, and *Treated with Internet after T*(reatment), the interaction between *Treated districts after T*(reatment) and the dummy for districts with the share of high-speed Internet access above the median. All specifications include time period (month-year) dummies. Standard errors, clustered by districts, are in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

B.4 Robustness of Unemployment vs. Part-Time Labor (Kurzarbeit)

In this section we exploit an additional source of variation in the number of hours worked. This variation was generated by the "Kurzarbeit"-program, which allowed firms to temporarily reduce the work hours of their employees by half while the government stepped in to compensate workers for a large part of their income loss. It is important to note that this reduction was understood to be temporary and the workers would be guaranteed to return to full time hours after the program's phase-out (which occurred in 2011). We observe the number of firms which used this program and the resulting number of workers in a district who worked fewer hours. Table 21 shows that these two shocks were not perfectly correlated, which means that we have two subgroups of treated districts: those exposed to only an increase in the unemployment rate, and those faced an increase in the number of workers with reduced weekly hours (Kurzarbeit). Moreover we have an additional group of control districts that did experience increased levels of reduced weekly hours.

In Table 22 we use the additional treatment dummy to shed light on the drivers of our main effect, by distinguishing these districts. The table repeats the regression in our main specification, but adds the dummy for an increase in the workers that reduced weekly working hours and a dummy for districts that saw both, increased unemployment rates and increased workers with a temporary reduction in weekly working hours. Moreover we control for GDP/Capita and the share of workers in Manufacturing and Services, in addition to the district fixed effect that we employed. The regressions show that the effect is mainly driven by districts that were subject to both shocks, which reinforces our argument, that the crisis has resulted in a relative increase of contributions by new users.

Table 21 Correlations of Unemployment and Kurzarbeit.

	Treated (UR)	Treated (KA)	Treated (KA+UR) X After
Treated (UR)	1		
Treated (KA)	0.185***	1	
Treated (KA+UR) X After	0.424***	0.439***	1
Observations	3120		

$^{*}\,p<0.05,\,^{**}\,p<0.01,\,^{***}\,p<0.001$

Table 22   Contrasting unemployment with part time labor (Kurzarbeit) for German districts - Robustness Check.

	Anonymous Content			Registered Content		
	(1) Log (Added (KB))	(2) Log (Deleted (KB))	(3) Log(# Edits)	(4) Log (Added (KB))	(5) Log (Deleted (KB))	(6) Log(# Edits)
Treated (KA) X After	-0.039	-0.015	0.003	0.276	-0.092	0.057
	(0.070)	(0.142)	(0.034)	(0.313)	(0.321)	(0.180)
Treated (UR) X After	-0.021	-0.095	0.047	0.157	0.084	0.073
	(0.088)	(0.182)	(0.043)	(0.282)	(0.263)	(0.138)
Treated (KA+UR) X After	0.166**	-0.180	0.073*	0.210	0.023	0.004
	(0.078)	(0.155)	(0.041)	(0.268)	(0.287)	(0.153)
GDP/Capita	0.012	-0.025	0.003	-0.007	-0.099*	-0.040
	(0.015)	(0.033)	(0.008)	(0.060)	(0.059)	(0.033)
Share Empl. Manuf.	0.012	-0.228	-0.016	-0.004	-0.516	-0.178
	(0.195)	(0.464)	(0.111)	(0.993)	(0.655)	(0.391)
Share Empl. Serv.	0.034	-0.290	-0.004	0.039	-0.549	-0.181
	(0.192)	(0.439)	(0.110)	(0.985)	(0.642)	(0.386)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3120	3120	3120	3120	3120	3120
Number of Districts	240	240	240	240	240	240
R^2	0.06	0.16	0.44	0.01	0.01	0.01

NOTES: The table shows the results of our main difference-in-difference around the *end* of the European Financial Crisis in 2009. The analysis contrasts affected and unaffected German districts, and the columns contain different measures of contribution activity to Wikipedia: (1) content added by *'anonymous edits'* (length in KB) (2) content deleted by anoymous users (in KB) (3) number of edits (revisions) by anonymous editors (4) content added by registered editors from the district (length in KB) (5) content deleted by 'registered edits' (in KB) (6) number of edits (revisions) by registered editors. The variable of interest, which captures the treatment effect, *Treated districts after T* (reatment), 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 time period (month-year) dummies. Standard errors, clustered by districts, in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

B.5 Robustness: Local Interest (District-speciftc) Content Analysis

Table 23   Local Knowledge Edits about German Districts before/after Crisis.

	Anonymous Content			Registered Content		
	(1)	(2)	(3)	(4)	(5)	(6)
	Log (Added (KB))	Log (Deleted (KB))	Log (# Edits)	Log (Added (KB))	Log (Deleted (KB))	Log (# Edits)
Treated (KA+UR) after T.	0.101**	0.055	0.063*	0.149**	0.064	0.024
	(0.047)	(0.057)	(0.037)	(0.066)	(0.070)	(0.042)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes	Yes	Yes
Single Criterion Dummies	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3029	3029	3029	3029	3029	3029
Number of Districts	233	233	233	233	233	233
R^2	0.08	0.09	0.32	0.04	0.05	0.11

NOTES: Difference-in-difference analysis that contrasts content generated *about* affected and unaffected German districts around the onset of the European Financial Crisis in 2009 (based on : 6 months before/after the crisis start). Columns 1-3 show the results for three measures of contribution activity by anonymous editors to Wikipedia: (1) added content (in KB), (2) deleted content (in KB), and (3) shows the number of edits by anonymous editors. Columns 4-6 show the results for same measures of contribution activity to Wikipedia but for registered editors: (4) added content (in KB), (5) deleted content (in KB), and (6) the number of edits by 'registered editors.' The variable of interest, which captures the treatment effect, *Treated (KA+UR) after T* (reatment), 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 time period (month-year) dummies. Standard errors, clustered by districts, in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 24   Local Knowledge Edits about German Districts before/after Crisis.

	Anonymous Content			Registered Content		
	(1)	(2)	(3)	(4)	(5)	(6)
	Log (Added (KB))	Log (Deleted (KB))	Log (# Edits)	Log (Added (KB))	Log (Deleted (KB))	Log (# Edits)
Treated (KA) X After	0.046	-0.047	0.024	0.132*	0.069	0.096**
	(0.049)	(0.061)	(0.033)	(0.074)	(0.066)	(0.044)
Treated (UR) X After	0.026	0.147***	0.037	0.122	-0.055	0.037
	(0.052)	(0.051)	(0.045)	(0.076)	(0.067)	(0.046)
Treated (KA+UR) after T.	0.101**	0.055	0.063*	0.149**	0.064	0.024
	(0.047)	(0.057)	(0.037)	(0.066)	(0.070)	(0.042)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3029	3029	3029	3029	3029	3029
Number of Districts	233	233	233	233	233	233
R^2	0.08	0.09	0.32	0.04	0.05	0.11

NOTES: Difference-in-difference analysis that contrasts content generated *about* affected and unaffected German districts around the onset of the European Financial Crisis in 2009 (based on : 6 months before/after the crisis start). Columns 1-3 show the results for three measures of contribution activity by anonymous editors to Wikipedia: (1) added content (in KB), (2) deleted content (in KB), and (3) shows the number of edits by anonymous editors. Columns 4-6 show the results for same measures of contribution activity to Wikipedia but for registered editors: (4) added content (in KB), (5) deleted content (in KB), and (6) the number of edits by 'registered editors.' The variable of interest, which captures the treatment effect, *Treated (KA+UR) after T* (reatment), 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 time period (month-year) dummies. Standard errors, clustered by districts, in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

B.6 Controlling for District-Speciftc Pre-Treatment Trends

Table 25 **Local Knowledge Edits about German Districts before/after Crisis (incl. Pretrends).**

	Anonymous Content			Registered Content		
	(1)	(2)	(3)	(4)	(5)	(6)
	Log (Added (KB))	Log (Deleted (KB))	Log (# Edits)	Log (Added (KB))	Log (Deleted (KB))	Log (# Edits)
Treated (KA+UR) after T.	0.118**	0.069	0.073*	0.117*	0.069	0.038
	(0.049)	(0.058)	(0.038)	(0.068)	(0.072)	(0.043)
District pre-treatment trend	0.104***	0.090***	0.050**	0.108***	0.119***	0.145***
	(0.021)	(0.022)	(0.025)	(0.027)	(0.026)	(0.035)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes	Yes	Yes
Single Criterion Dummies	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3029	3029	3029	3029	3029	3029
Number of Districts	233	233	233	233	233	233
R^2	0.09	0.09	0.32	0.05	0.07	0.12

NOTES: Difference-in-difference analysis that contrasts content generated *about* affected and unaffected German districts around the onset of the European Financial Crisis in 2009 (based on : 6 months before/after the crisis start). Columns 1-3 show the results for three measures of contribution activity by anonymous editors to Wikipedia: (1) added content (in KB), (2) deleted content (in KB), and (3) shows the number of edits by anonymous editors. Columns 4-6 show the results for same measures of contribution activity to Wikipedia but for registered editors: (4) added content (in KB), (5) deleted content (in KB), and (6) the number of edits by 'registered editors.' The variable of interest, which captures the treatment effect, *Treated (KA+UR) after T* (reatment), 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 time period (month-year) dummies. Standard errors, clustered by districts, in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

B.7 Count Data Models for the Number of Edits

Table 26   **Negative Binomial Regressions for the Number of Edits to Wikipedia in German Districts.**

	Anonymous Content	Registered Content
	(1)	(2)
	# Edits	# Edits
Treated Districts After T.	0.041**	-0.212
	(0.020)	(0.133)
District FE	Yes	Yes
Time FE	Yes	Yes
Observations	3120	1989
Number of Districts	240	153

NOTES: The table shows the results of an alternative econometric specification, a count data model, which contrasts affected and unaffected German districts around the onset of the European Financial Crisis in January 2009. The columns contain the measures of contribution activity on Wikipedia that follow negative binomial distribution: (1) the number of edits (revisions) by anonymous editors and (2) the number of edits (revisions) by registered editors. The variable of interest, which captures the treatment effect, *Treated districts after T*(reatment), 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. Observations range from 6 months before to 6 months after the onset of the crisis. All specifications include time period (month-year) dummies. Standard errors (in parentheses) are bootstrapped and indicate: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

C Country-Level Analysis

C.1 Country-Level Descriptive Statistics

Table 27   Wikipedia Key Variables Within the Period of 12 Months Before and 12 Months After the Crisis.

	Language Speakers (mln.)	In main country (%)	Views per Speaker	Wikipedians (%)	Active Editors with 5-100 edits (%)	Active Editors with > 100 edits (%)
Bulgarian	8.16	86.05	2	0.02	10.7	3.5
Catalan	4.08	.	3	0.07	13.0	3.8
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

NOTES: *Source: stats.wikimedia.org*. This table gives an overview over the languages that were analyzed in the Country-level study. Column 1 gives the number of speakers. Column 2 shows the percentage of speakers that live inside the associated country. Columns (3) - (6) are means of values for the interval 12 months before to 12 months after crisis.

Table 28   Crisis Indicators: Unemployment Rates (%).

	Affected by Crisis	Crisis Start	Unemployment Rate (%)	Change in Unemployment (%)
Bulgarian	yes	Oct 2008	6.1	0.6
Catalan	yes	Oct 2008	11.5	7.5
Czech	yes	Oct 2008	5.3	2.7
Danish	no	.	4.4	3.2
Dutch	no	.	3.3	0.6
Finnish	no	.	7.1	2.1
German	no	.	7.7	0.1
Greek	yes	June 2009	9.5	2.7
Hungarian	yes	March 2009	9.3	2.1
Icelandic	yes	Oct 2008	4.6	6.9
Italian	yes	May 2009	7.6	1.2
Japanese	no	.	4.4	1.4
Norwegian	no	.	2.8	0.7
Polish	no	.	7.6	-0.2
Romanian	yes	Oct 2008	6.2	0.8
Russian	yes	Oct 2008	7.0	2.3
Slovakian	yes	Oct 2008	10.5	2.9
Slovene	yes	Oct 2008	5.0	1.1
Swedish	no	.	6.9	1.7
Turkish	yes	Oct 2008	10.8	3.7

NOTES: This table shows how countries' unemployment rates were affected during the crisis. Affected countries were identified using mentions in reports by the EU and OECD (for central and eastern Europe see e.g.: *https://www.ecb.europa.eu/pub/pdf/other/art3_mb201007en_pp85-96en.pdf*). These reports typically define countries as strongly affected by the crisis based on a sharp decrease in GDP and the time after the crisis taken for the recovery. As the table shows, such changes typically (but not always) coincided either with a sharp increase in unemployment or a decrease in the hours worked in the economy.

C.2   Speciftcation Tests and Alternative Speciftcations

Table 29   DID Regression for the period of 12 months before and 12 months after the crisis.

	(1) Ln (Views)	(2) Ln (Active Editors (5-100e))	(3) Ln (Very Act. Editors (≤ 100 e))	(4) Ln(Avg.Edits per Article)	(5) Ln (New Words)	(6) Ln (# Wikilinks)	(7) Ln (# Ext. Links)
After treatment	0.129	0.017	-0.005	-0.005	0.054	-0.076**	-0.008
	(0.095)	(0.034)	(0.042)	(0.017)	(0.051)	(0.035)	(0.059)
Treated countries after T	0.153*	0.097**	0.041	-0.001	0.115*	0.006	0.020
	(0.078)	(0.034)	(0.036)	(0.012)	(0.058)	(0.012)	(0.019)
Pretrend	0.007	0.369**	0.386***	0.718***	0.210**	0.923***	0.755***
	(0.018)	(0.131)	(0.091)	(0.081)	(0.088)	(0.093)	(0.113)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	432	480	480	480	480	480	480
Number of Countries	20	20	20	20	20	20	20
R^2	0.39	0.33	0.45	0.96	0.15	0.99	0.98

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) hyperlinks to external sources. All measures of contributions to Wikipedia are in logs, and the month of the estimated crisis onset was omitted from the regressions. 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 time (month - year) fixed effects. Standard errors, clustered by countries, are in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 30   Country Level: Difference-in-Differences Regressions.

	(1) Ln (Views)	(2) Ln (Active Editors (5-100e))	(3) Ln (Very Act. Editors (≤ 100 e))	(4) Ln(Avg.Edits per Article)	(5) Ln (New Words)	(6) Ln (# Wikilinks)	(7) Ln (# Ext. Links)
Treated 3 months before	0.042	0.078	0.113**	0.032**	0.008	0.019	0.061**
	(0.079)	(0.056)	(0.046)	(0.012)	(0.073)	(0.016)	(0.028)
Treated 2 months before	-0.014	0.044	0.090*	0.034**	0.026	0.023	0.079**
	(0.086)	(0.057)	(0.048)	(0.014)	(0.087)	(0.019)	(0.032)
Treated 1 month before	-0.043	0.070	0.056	0.038**	0.061	0.027	0.091**
	(0.065)	(0.052)	(0.049)	(0.016)	(0.063)	(0.021)	(0.035)
Treated in month 0	-0.003	0.033	0.025	0.044**	-0.120	0.028	0.099**
	(0.079)	(0.061)	(0.053)	(0.018)	(0.090)	(0.024)	(0.039)
Treated in month 1	0.017	0.038	0.053	0.051**	-0.088	0.029	0.107**
	(0.082)	(0.051)	(0.041)	(0.020)	(0.114)	(0.028)	(0.043)
Treated in month 2	0.091	0.167***	0.080	0.055**	0.078	0.033	0.117**
	(0.071)	(0.056)	(0.054)	(0.022)	(0.071)	(0.030)	(0.046)
Treated in month 3	0.086	0.158**	0.133*	0.060**	0.201**	0.041	0.128**
	(0.094)	(0.072)	(0.064)	(0.024)	(0.090)	(0.033)	(0.050)
Treated in month 4	0.120	0.150**	0.143**	0.063**	0.178*	0.049	0.137**
	(0.086)	(0.061)	(0.063)	(0.025)	(0.099)	(0.036)	(0.054)
Treated 5-6 months later	0.186	0.151**	0.146**	0.064**	0.109	0.058	0.147**
	(0.109)	(0.059)	(0.055)	(0.028)	(0.069)	(0.040)	(0.061)
Treated 7-8 months later	0.190**	0.196***	0.170***	0.066**	0.213**	0.069	0.172**
	(0.085)	(0.059)	(0.059)	(0.031)	(0.079)	(0.044)	(0.065)
Treated 9-10 months later	0.164	0.171**	0.188***	0.069*	0.183*	0.081	0.189**
	(0.102)	(0.074)	(0.053)	(0.034)	(0.092)	(0.049)	(0.070)
Treated 11-12 months later	0.170*	0.176**	0.184***	0.075*	0.157	0.090	0.207**
	(0.089)	(0.064)	(0.056)	(0.037)	(0.106)	(0.052)	(0.077)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	452	500	500	500	500	500	500
Number of Countries	20	20	20	20	20	20	20
R^2	0.39	0.22	0.32	0.87	0.14	0.93	0.93

NOTES: The table shows the relationship of the European economic crisis with different measures of contributions to Wikipedia (over time). In each column we show a different measure of activity on Wikipedia: (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 measures of contributions to Wikipedia are in logs. The variable of interest is an interaction terms between an indicator for countries that are affected by the crisis with the time dummy indicating the period after the crisis. Observations ranging from 12 months before to 12 months after the crisis. The excluded periods are the months 4-12 prior to treatment. All specifications include time (month - year) fixed effects. Standard errors, clustered by countries, are in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 31   Country Level: Fixed Effects Regressions for the Period of 12 Months Before the Crisis.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Ln (Views)	Ln (Active Editors (5-100e))	Ln (Very Act. Editors (≤ 100 e))	Ln(Avg.Edits per Article)	Ln (New Words)	Ln (# Wikilinks)	Ln (# Ext. Links)
Unemployment Rate (%)	0.030	-0.005	-0.014	-0.000	0.005	0.009	0.014
	(0.028)	(0.019)	(0.017)	(0.004)	(0.026)	(0.008)	(0.016)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	192	240	240	240	240	240	240
Number of Countries	20	20	20	20	20	20	20
R^2	0.31	0.12	0.24	0.83	0.08	0.94	0.90

NOTES: This table shows the relationship between unemployment and contributions to Wikipedia before the moment the economic crisis hits the country. 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 measures of contributions to Wikipedia are in logs. The independent variable of interest, *Unemployment*, is the monthly unemployment rate. All specifications include time (month - year) fixed effects. Standard errors, clustered by countries are in parentheses: *** $p < $ **0.01** , ** $p < $ **0.05** , * $p < $ **0.1** .

Table 32   Country Level: Content Contributions and Youth Unemployment.

	(1) Ln (Views)	(2) Ln (Active Editors (5-100e))	(3) Ln (Very Act. Editors ($\leq$ 100 e))	(4) Ln(Avg.Edits per Article)	(5) Ln (New Words)	(6) Ln (# Wikilinks)	(7) Ln (# Ext. Links)
Youth Unemployment (%)	0.010**	0.007***	0.010***	0.001	0.022***	0.012***	0.019***
	(0.004)	(0.002)	(0.002)	(0.002)	(0.006)	(0.003)	(0.006)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	411	456	456	456	456	456	456
Number of Countries	19	19	19	19	19	19	19
R^2	0.39	0.15	0.25	0.84	0.16	0.97	0.94

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 measures of contributions to Wikipedia are in logs. The independent variable of interest, *Youth unemployment*, is the monthly unemployment rate. Fixed Effects Regressions for the Period of 12 Months Before and 12 Months After the Crisis. All specifications include time (month - year) fixed effects and exclude the estimated period of the crisis onset. Standard errors, clustered by countries are in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.