Spillovers in Networks of User Generated Content *

- Evidence from Natural Experiments on Wikipedia

Michael E. Kummer†

Preliminary and obviously incomplete - Please don't quote or circulate

June 16, 2013

Abstract

Endogeneity in the formation of networks has been a continued obstacle for scientists, who were trying to understand peer effects, externalities or the role of social networks in generating spillovers. This paper suggests a method that aims at overcoming the usually required, but generally strong, assumptions of exogenous link formation or exogeneity of the observed characteristics. Identification is based on exploiting exogenous but local shocks or randomized treatments on a relatively small number of nodes in the network. The method has the additional advantage of providing identification even in very small (village- or classroom) networks, where the researcher cannot observe open triads. The suggested method is applied to data from the German Wikipedia in order to measure how attention to articles spills across links and how additional attention results in new content generation. This application is of interest in its own right. Knowing whether spillovers in content networks exist and how exactly they function would allow important insights into how humans share content or knowledge and how they allocate effort in peer production settings. Understanding this input to the production process of users who privately provide a public good is also relevant to newly created content platforms that are eager to acquire a minimum necessary level of content that guarantees the platform's survival.

Keywords: Social Media, Information, Knowledge, Large-scale Networks.

JEL Classification Numbers:

---

*The author is grateful the Wikimedia Foundation for granting access to the Wikipedia data, Thorsten Doherr and Manfred Knobloch for support with the data processing, George Giorgidze for help with the programming of DSH and Frédéric Schütz for providing us with the data on page views. This work benefitted from discussions with Irene Bertschek, Luis Cabral, Ulrike Cress, Avi Goldfarb, Shane Greenstein, Francois Laisney, Jose Luis Moraga-Gonzalez, Kathy Nosal, Martin Peitz, Marianne Saam, Philipp Schmidt-Dengler, Olga Slivko, Yossi Spiegel, Michael Ward and Andrea Weber. I thank the participants of the MaCCI-IO day at University Mannheim and the IIOC in Boston for valuable input. Timo Schäffer and Daniel Bergman provided outstanding research assistance. Financial support from the WissenschaftsCampus Tübingen is gratefully acknowledged.

†Centre for European Economic Research (ZEW); L 7, 1; 68181 Mannheim; Germany; Email: Kummer@zew.de.

1 Introduction

The role of spillovers in networks is an issue that has interested economists more and more in recent decades. Particularly knowledge and technology spillovers have far reaching implications for innovation and welfare. A special form, spillovers of attention and content generation through links, might be of central importance for the provision of user generated content and the contribution to public goods. Understanding them, might thus help us to take a first step of understanding a potentially important input to the production function of user generated content and explaining its high quality, which has puzzled researchers from both economics and other fields. Knowing whether spillovers of attention and content generation are mediated through links is also relevant for administrators of newly created wikis, who are eager to acquire a minimum necessary level of content that guarantees the platform's survival.

Yet, important as they may be, measuring spillovers in a network is very difficult. While correlations between agents abound, exogeneous sources of variation that allow to pin down the cause and to distinguish it from the effect are usually hard to observe. Previous research in the context of peer production has attempted to analyze the correlation of a node's position in a network and the outcomes of interest (Fershtman and Gandal (2011), Claussen et al. (2012) or Kummer et al. (2012)). However, the network position of a node is frequently driven by many observed and unobserved factors, and the problems induced by this fact persist even when observing the agents in the network over time. Moreover, often the outcome variable is itself an important determinant of the network position, thus giving rise to the classic endogeneity problem. Consequently, in the absence of an exogenous source of variation in the network position, it is hardly ever possible to come up with a design that cleanly separates causes and effects in such a networked setting.

This paper pursues a different strategy. Looking at a very short time window of only several days, I can utilize a source of exogenous variation to attention and interest in a known set of nodes, to analyze spillovers of attention and content generation that are mediated through the links in the network. In other words, this paper does not look at the variation or changes in the network, but at observable changes of attention and activity that affect only a relatively small set of nodes of the network, but not the rest. Thus changing the approach I analyze spillovers of attention and content generation in one of the world's biggest citation networks, namely the network formed by links between articles on the German Wikipedia. I exploit two sources of large variation: (i) exogenous and unpredictable large scale media events, such as the outbreak of political upheaval, earthquakes or plane crashes and (ii) articles that were chosen to be featured on Wikipedia's startpage and were thus highly visible for 24 hours. Such events typically triggered changes in the attention that certain pages received at a known (ex-post) point in time.

I access the revision data of the German Wikipedia around nine large scale events and I track the pages that are most relevant to each of them. analogously I identify 36 articles that were featured and when this happened. I identify both all the pages that received a direct link from the shocked pages and, for large events also those pages that received an indirect link (a page that was two clicks away), a week before the event occurred. I then measure how the attention and content generation at the pages changed at the time of the event. I compare it to the changes of attention and content that occurred in a comparison group of which none of the pages was related to the event.

I use data of Wikipedia, the world's largest and best known wiki. I have access to a database that was put together in a joint effort of the University of Tübingen, the IWM Tübingen and the ZEW Mannheim. It is based on data from the German project, which (currently) has roughly 1.4M articles and thus provides us with a very large number of articles to observe. At the heart of the database are the publicly available data dumps, provided by the Wikimedia foundation. These were augmented with data on the link structure between articles, data on the download frequency of pages and information on major media events, which occurred during our period of observation. We thus obtained a unique database for the time span from Dec 2007 to Dec 2010. Starting from nine large scale media events, I obtain a data set of more than 35,000 articles with an indirect link from one of the nine start pages, that were affected by the exogenous shock. Additionally I extracted a dataset that is based on 36 pages that were featured on Wikipedia's main page for 24 hours. Moreover I obtain the information on clicks and content provision for the neighboring pages. I analyse these two treatments separately.

For large events, I find substantial spillovers of attention even to pages that are two clicks away from the disaster page, but relatively modest and not necessarily robust effects on content provision. The rise in the number of clicks on the pages that were in the neighborhood of the shocked area (2 clicks away) amounts to roughly 35 more clicks (a 100% increase) on average, and thus indicates that attention indeed gets transmitted via the link network. The evidence whether this increase also translates to increased provision of content may be positive but it is small in absolute terms. Furthermore, the evidence is less robust and requires further confirmation. In short, my first results suggest that links matter for for the attention that a node in a citation network receives, but not necessarily to the content that is generated on such nodes. While these patterns might indicate, that the spillover effect is actually a "lookup"-effect, rather than an effect that translates to new provision of content, such a conclusion needs to be confirmed by further research, before it can actually be drawn.

For featured articles, I do not find any effects on the pages that were two clicks away, but the effects on neighboring pages are substantial, both for attention and content generation. I observe an increase of almost 100 percent, both in viewership and editing activity.

However, given the small baseline activity the activity triggered by having a featured article in the neighborhood is small in absolute terms. Aggregated over all neighbors I find that 1 click on the treated page translates to 1 click on one of the neighbors, but that it takes 1000 clicks before an additional revision occurs. Despite the fact that my design allows to infer results already from the reduced form, it is necessary to point out that the data data on large events, and hence also my results are only preliminary. Ongoing work aims at extracting data on more events and improving the details of the design and, eventually, estimating a structural parameter of attention spillovers.

Understanding the existence and the magnitude of attention spillovers that are mediated through links sheds light onto how humans gather information and what they pay attention to when navigating a citation network of interlinked knowledge. Although not all the insights from the present application will carry over to other citation networks such as scientific literature or patent licensing, the use of attention spills in the aftermath of large scale media events have proven to be a useful focal lens to better understand some of the underlying dynamics that is common to all these networks. While the same applies to my results concerning content generation, they are also relevant to the Wikipedia community and to the administrators of other platforms that hope to harvest user generated content. What is more, understanding whether links matter for channeling the flows of content provision on such platforms, might be an important ingredient to solving the conundrum of the sometimes astonishingly high quality of user generated content. Finally, the results will also be relevant when considering Wikipedia as a large public good and for understanding the behavior of the agents who contribute to it.

1.1 Literature

This paper builds on two important streams of literatures: first the literature on social effects, peer effects or spillovers and second on the literature that uses pseudo-treatments to causally identify economic effects.

Social effects, such as peer effects or spillovers in a network are generally difficult to identify, mostly, because they are frequently confounded with other individual specific characteristics or network dynamics. One quite prominent subbranch of this literature pursues the strategy of investigating the relationship of a node's position in a network and its performance.

A well known and important branch in this literature is dedicated to the identification of peer effects. As has been shown, these are extremely difficult to identify in a setting where both the peers average characteristics and their average performance influence the individual's outcome Manski (1993). One of the most widely known approaches to disentangle these effects is based on the structure of networks, or more precisely on the existence of open triads, where a peer is connected to two other peers, who themselves are not connected to each other. In such a situation the outcome of the peer, who is connected to both nodes is instrumented with the performance of one peer before analyzing it's influence on the other. (De Giorgi et al. (2010) and Bramoullé et al. (2009))

Another series of papers has focused on knowledge spillovers in production through social networks. Fershtman and Gandal (2011) investigate indirect or direct knowledge spillovers in the production of open source software and Claussen et al. (2012) pursue a similar question in the electronic gaming industry. Both studies focus on the relationship between developers' network position and the success of the project they are working on. (Kummer et al. (2012)), borrowing borrow from the approach used by Halatchliyski et al. (2010), who analyze authors' contributions in two related knowledge domains, consider a different network in a similar context, namely the hyperlink network of articles. The strategy in all of these papers is based on exploiting variations in the link network, be it between or within nodes in a network, and relating them to the outcome of interest. If this relationship is found to be positive, this is believed to be evidence for spillover effects of knowledge. However, a common criticism of this strategy is, that the variation in the network position might not be exogenous or that it is at least very difficult to identify sources of exogenous variation to the network. The strategy pursued in this paper approaches the problem from a different vector of attack, because it no longer attempts to measure spillover effects by looking at variation in the link structure, but it looks how shocks are transmitted in a given link structure. The underlying reasoning is now based on the observation that pages who happen to be linked to a shocked page receive a spillover, while similar pages who are not linked in such a way do not. I will argue further below that such an approach can be successful under certain conditions, but that it might fail in situations, where these conditions are not satisfied.

As far as the second stream of literature on treatment effects is concerned, it is well established that social effects play an important, though usually not constructive, role for the causal identification of such effects. More precisely, it can be difficult to identify the causal effect of a treatment in the presence of social effects, not least because such effects might lead to a violation of the Stable Unit Treatment Value Assumption (SUTVA) and hence raise doubts about the validity of the control group. (Ferracci et al. (2012)) . The present study proposes to use the treatment of peers in a network to identify social effects and asks under which circumstances it may be possible to causally identify spillovers or peer effects when treatment of peers can be observed. It is worth emphasizing, that the analysis in this paper might be somewhat unusual for readers, who are well familiar with this literature, because it is not particularly interested in the effect of treatment itself, but, instead, aims at exploiting treatments to identify *the spillover effect*. Hence, it is not only not concerned about a possible violation of crucial assumptions such as SUTVA, but rather aims at *exploiting* a situation where it is not satisfied to identify something else (i.e. the spillover effect). This idea is not new, but has been used more and more often in recent studies. Imberman et al. (2009), for example, exploit variation due to a natural disaster in their analysis of the peer-effects of evacuee inflow on Houston's and Louisiana's incumbent school children in the aftermath of the hurricanes Katrina and Rita. Their identification strategy is based on the large variation in peer groups and random allocation of evacuees after the event. (They find small peer effects on average, but they also show that inflow of high-achieving peers has positive efects of achievement.) In a recent study in the realm of e-commerce Carmi et al. (2009) analyse the effect of book recommendations by Oprah Winfrey (as external shocks) on the product network of books on Amazon. They find, that the recommendation does not only trigger a spike in sales of the recommended book, but also of books that are adjacent to the recommended books in Amazon's recommendation network. They measure demand in terms of the products' sales ranks and use a Differences in Differences strategy, where they obtain the control group by exploiting temporal variations in the recommendation network. They find a significant and positive effect on the recommended books' neighbors and the neighbors of neighbors. The current study proposes to apply very similar ideas to estimate the spillover effects in a content network. Berge (2011) compares peers of treated and non-treated agents in a field experiment to measure information and knowledge spillovers from a business training program in Tanzania. Using in depth interviews he finds that "indirectly-treated" male clients become more "business minded" (i.e. they discuss more business, increase their loans and become more risk averse).

Finally, several papers have been dedicated to natural disasters or other sources of exogenous variations on Wikipedia, or they have exploited treatments, similar to the one I use. The well known paper by Zhang and Zhu (2011) exploits an exogenous shock that occurred, when the Chinese government blocked Wikipedia in mainland China to measure the effect on the incentives to contribute. Keegan et al. (2013) analyze the structure and dynamics of Wikipedia's coverage of breaking news events. They show that the coverage of breaking news events is an increasingly important phenomenon on Wikipedia, which makes up an increasing share of edits and, they hypothesize, might become one of Wikipedia's most important sources of new contributors. They contrast the evolution of author networks breaking news event-articles with the genesis of non breaking news (and also "historical" articles) and they find that breaking news articles emerge into well connected collaborations more rapidly than non-breaking articles.

2 The empirical model

2.1 Reduced Form Analysis

In a first step it is useful to apply reduced form regressions to understand the impact of the local treatment on both the treated pages and their neighbors. These are very similar in spirit to the analysis in (Carmi et al. (2009)). The main idea is to simply compare pages grouped by their distance to the page which experiences treatment to their analogue in the control group and to redefine treatment for each set of pages accordingly. This results in the following three regression equations, each for a different set of pages (and without time-dummies):

L0.) Diff in Diff specification at level L0:

$$y_{it} = \alpha_i + \sum_{s \in S} \alpha_s \lambda_s + \sum_{s \in S} \beta_s (\lambda_s * treat_{L0,i}) + \xi_{it} \tag{1}$$

...$treat_{L0}$: treatment on the very page; $S = \{-14, ..., 14\}$

L1.) At level L1 ($treat_{L1}$ featured (in theory) 1 click away):

$$y_{it} = \alpha_i + \sum_{s \in S} \alpha_s \lambda_s + \sum_{s \in S} \beta_s (\lambda_s * treat_{L1,i}) + \xi_{it} \tag{2}$$

L2.) At level L2 ($treat_{L2}$ featured (in theory) 2 clicks away):

$$y_{it} = \alpha_i + \sum_{s \in S} \alpha_s \lambda_s + \sum_{s \in S} \beta_s (\lambda_s * treat_{L2,i}) + \xi_{it} \tag{3}$$

In words, I run the same diff in diff on three levels (on L0, L1 and L2 (shown only for featured)), but, for the present purpose, change the definition of a treatment: while $treat_{L0,i}$ is an indicator variable for a page that is (going to be) featured on Wikipedia's main page, $treat_{L2,i}$ takes the value of 1 for pages that are two clicks away from pages that are (going to be) affected by such a shock. In the regressions above, the cross terms correspond to this indicator variable multiplied with the time dummies. Thus, a cross term captures whether treatment has occurred at a given point in time or not. For an observation in the control-group this variable will always take the value of 0, for an observation in the treated group, this variable will take the value of 1, if the observation corresponds to the event time the time-dummy aims to capture. Hence, if the treatment is effective, the coefficients of the cross terms are expected to be 0 before treatment occurs and positive for the periods after the treatment.

Note that there is a specificity to the major events condition, which stems from the fact that the page of the event itself (the directly shocked node) cannot exist before the shock. I deal with this fact by specifying the L1 set as the set of pages that (will) have a reciprocal link from the L0-page once it is created. Hence, at the time of the shock, these are the pages that are very close at the epicenter. The L2 group is then defined as before, i.e. as the set of pages that received a link from an L1-page one week before the shock actually ocurred.

Other than the cross terms I also include page fixed effects and another full set of time dummies (event time) to control for general (e.g. weekday-specific) activity patterns in Wikipedia. Note that I run each regression twice to take advantage of my two comparison groups: first I contrast the treated pages against the control group and then I contrast it with the placebo treatment, i.e. with the treated articles themselves, but simulating a (placebo) treatment 42 days (i.e. 7 weeks) before the real shock.

This procedure is crude, because it does not consider several important factors, such as how well neighbors are linked among each other or how large the peak of interest is on the originally shocked page. Yet, it is a useful check, since even the results from such a reduced form analysis will provide guidance, as to whether attention-spills exist at all, how far they carry over, and whether they result in increased production.

2.2 Structural Form Analysis

[This section is under construction, hence the present section and the Appendix it refers to, still suffer from limitations.:]

- I have written down a model, but it is not carved in stone as it is.
- As of now I am augmenting a very important linear peer-effects model, which goes back to Manski (to have a contribution of broader interest?)
- If interested, cf. Appendix B and C for what I have so far, or simply stick to the body of the paper and read only the results based on the reduced form.
- I show an extension of the peer effects model in matrix notation as formulated by Bramoullé et al. (2009). – It is not polished and might still have mistakes. Also I probably fail to make the assumptions in a standard way, but I believe it is in principle right.

In this section, I extend the well known linear peer effects model, as it is formulated in Manski (1993), with exogenous shocks. Departing from the version that was used by Bramoullé et al. (2009) to show identification of peer effects in social networks, I show how the availability of exogenous shocks can be exploited to identify spillovers (or the peer effect) in this model, even if the nodes characteristics are endogenous.

Since the derivations involve quite heavy notation, but are otherwise relatively straight forward, the details and derivations can be found in the appendix. Here I only provide the point of departure and the main results. A well known form of the linear model has been formulated in Manski (1993)

$$y_{it} = \alpha \frac{\sum_{j \in P_{it}} y_{jt}}{N_{P_{it}}} + X_{it}\beta + \gamma \frac{\sum_{j \in P_{it}} X_{jt}}{N_{P_{it}}} + \epsilon_{it}$$

where y_{it} denotes the outcome of interest in period t and X_{it} are i's observed characteristics. P_{it} is the set of i's peers and $N_{P_{it}}$ represents the number of i's peers. α is the coefficient of interest. In the present context it measures how the clicks on page A are influenced the clicks on the adjacent pages.. Bramoullé et al. (2009) suggest a more succinct notation based on vector and matrix notation:

$$\mathbf{y_t} = \alpha\mathbf{G}\mathbf{y_t} + \beta\mathbf{x_t} + \gamma\mathbf{G}\mathbf{x_t} + \epsilon_\mathbf{t}$$
$$E[\epsilon_t|\mathbf{x_t}] = 0$$

I augment this model by including a vector of treatment, which, for simplicity, is assumed to take the value of 1 for only one treated node and the value of 0 otherwise. This captures the notion of a "local treatment condition", under which only one node is exposed to treatment.

$$\mathbf{y_t} = \alpha\mathbf{G}\mathbf{y_t} + \mathbf{X_t}\beta + \gamma\mathbf{G}\mathbf{X_t} + \delta_1\mathbf{D_t} + \epsilon_t \qquad E[\epsilon_\mathbf{t}|\mathbf{D_t}] = \mathbf{0} \tag{4}$$

- $\mathbf{G}$ is NxN
- $G_{ij} = \frac{1}{N_{P_i} - 1}$ if i receives a link from j and $G_{ij} = 0$ otherwise
- treated side: $D_t = e_{l0}$; i.e.: a vector with a 1 in the coordinate that corresponds to the treated node and 0's elsewhere.
- untreated side: $D_t = \mathbf{0}$, a vector of zeros.
- Note that I DO NOT assume $E[\epsilon_t|\mathbf{X_t}] = 0$, but only that treatment be exogenous, since weaker assumptions will suffice.

Define the set of observations in the network if treatment occurs in t by the subscript ℓ, and a comparison group in which no node is treated by subscript c. If these sets of nodes can also be observed one period earlier a Difference in Differences (DiD) can be computed. Let this DiD be denoted as $\mathbf{E}[\mathbf{\Delta}\ell - \mathbf{\Delta c}|\mathbf{D_t}]$.

Result: The DiD contains the following quantity:

$$\begin{aligned} \mathbf{E}[\mathbf{\Delta}\ell - \mathbf{\Delta c}|\mathbf{D_t}] &= \mathbf{I} * \delta_\mathbf{1}\mathbf{D_t} + \alpha\mathbf{G} * \delta_\mathbf{1}\mathbf{D_t} + \alpha^\mathbf{2}\mathbf{G^2} * \delta_\mathbf{1}\mathbf{D_t} + \alpha^\mathbf{3}\mathbf{G^3} * \delta_\mathbf{1}\mathbf{D_t} + ... \\ &= \delta_1\mathbf{D_t}(\mathbf{I} + \alpha\mathbf{G} + \alpha^\mathbf{2}\mathbf{G^2} + \alpha^\mathbf{3}\mathbf{G^3} + ...) \end{aligned} \tag{5}$$

...which is nothing else than the direct effect of the treatment plus how it spills through the network. The proof of this result can be found in Appendix C.

Clearly, the DiD alone will not directly reveal α, the parameter of interest, but merely a quantity that is tightly linked to α and δ_1. Yet, the result also highlights that computing the parameters is not necessarily feasible, e.g. because it involves the knowledge of the complete link structure of the nodes. Luckily, a closer look at the nodes independently reveals that already limited information about the link structure can suffice to acquire additional information about the parameters. In particular, it is possible to separately compare the directly treated nodes and their counterparts in the control group and their neighbors (in treated and comparison group). Under quite rigorous, but not uncommon assumptions it shall thus be able to obtain a lower bound estimate for the coefficient α. (cf. Appendix C)

Specifically, if we ignore higher order spillovers[1], we can obtain such an estimate from applying the the Diff in Diff estimator on the level of directly treated nodes and a suitable comparison group and then move on to estimate α based on combining it with a second Diff in Diff at the neighbor level. Let $\Delta l_a - \Delta c_a$ denote such a Diff in Diff ($a \in \{0, 1\}$ denotes whether the nodes are in the center of the network (0) or the neighbors of the start nodes (1):

$$\hat{\delta}_1 = \hat{\Delta l}0 - \hat{\Delta c}0 \tag{6}$$

- $\hat{\Delta l}0 := \frac{1}{NP_l0} * \sum_i (y_{i,l0,t=0} - y_{i,l0,t=1})$
- $\hat{\Delta c}0 := \frac{1}{NP_c0} * \sum_i (y_{i,c0,t=0} - y_{i,c0,t=1})$

$$\hat{\alpha} = \frac{\hat{\Delta l}1 - \hat{\Delta c}1}{\hat{\Delta l}0 - \hat{\Delta c}0} * NP_{l1} \tag{7}$$

with the definition of $\hat{\Delta l}1$ and $\hat{\Delta c}1$ paralleling the definition of $\hat{\Delta l}0$ and $\hat{\Delta c}0$. This estimator has the advantage of being readily available with well known properties. However, as it is, it would only be suitable under the (potentially quite strong) assumption that higher order spillovers are negligible. Whether this is true or not will depend on the size of the spillover effect, but to a very large extent also on the network structure and the number of nodes. In future research I shall proceed to illustrate how second order spillovers affect the estimator in equation 7 and I will show how to derive an upper bound to the size of the problem.

[1] Or maintain the assumption that we can observe the nodes' performance before any higher order spillovers arrive at the treated node

3 Data

3.1 Preparation of the data and definition of the treated and control group

The dataset is based on a full-text dump of the German Wikipedia from the Wikimedia toolserver. To construct the history of the articles' hyperlink network for the entire encyclopedia, it was necessary to parse the data and identify the links. From the resulting tables, we constructed the time-varying graph of the article network, which provides the foundation for how I sample articles in our analysis. Furthermore information about the articles, such as the number of authors who contributed up to a particular point in time or the existence of a section with literature references was added. Hence, the data we use in our analysis, are based on 153 weeks of the the entire German Wikipedia's revision history between December 2007 and December 2010. Since the data are in the order of magnitude of terabytes, it would be not possible to conduct the data analysis using only in-memory processing. We therefore stored the data in a relational database (disk-based) and queried the data using Database Supported Haskell (DSH) (Giorgidze et al. (2010)). This is a novel high-level language which allows to write and efficiently execute queries on nested and ordered collections of data.

To identify major events, we consulted the corresponding page on Wikipedia and selected the 25 events that had the largest impact and for which a site was created after it occurred. For each of these events we identified the page that corresponds to the event, and which are considered to be in the set "L0" (sometimes also called "start pages"). Note that this page is created after the event occured [2]. We then exploited the data on the link structure to identify the set of pages that shared a reciprocal link with the start page and were hence very closely related to the event. After the disaster page existed they were only 1 click away from the event-page (set "L1"). Next, we identified those pages that received a link from an L1 page (unidirectional) (2 clicks away set "L2").[3] Having fixed the set of pages to observe, I extracted daily information about the current state of the articles (page visits, number of revisions, number of distinct authors that contributed, page length, number of external links etc.) 14 days before the event occurred (on a neighboring page) and 14 days after the shock (giving a total of 29 observations per page).

The "featured articles" were found by consulting the German Wikipedia's archive of pages that were selected to be "Seite des Tages" between December 2007 and December

[2] Usually it takes up to two days until the event receives its own page

[3] Note, that I evaluate the L2 pages in the network a week before the shock actually occurred, thus ensuring that we only include pages that had a link before it was known that the start page will be hit. I furthermore exclude pages that receive their indirect (L2) link via a page that has more than 100 links, since such pages are very likely either pure "link pages" very general pages (such as pages about a year), that bare only a very weak relationship to the shocked site.

2010. To reduce the computational burden and to avoid the risk of temporal overlaps of different treatments, we focus on pages that were selected on the 10th of a month. Similar to the procedure for disasters we identified all the pages that received a direct link (L1) and an indirect link (L2) from such a featured article a week before the page was featured.

I am most interested in attention-spills and content, which are not directly related to the event but rather a consequence of the peak in interest and resulting improvements also on the linked pages. Hence, I will not focus on the treated pages directly, but on the set L1, the pages that are "one click away" from a treated page, in my analysis of the "featured articles"[4] However, for disasters it is natural to focus on the pages that received a link from the neighborhood of the event (the indirectly linked set of pages (L2)) in the analysis below. First, because the shock is very large and second, because the event page actually does not exist at the time of the shock so that I cannot be sure, that the L1 pages are not actually treated themselves[5].

Since the approach I take in this paper hinges on the availability of a valid group for comparison, I also need to identify a set of observations, against which I can contrast treated pages and their neighbors. To obtain such observations I pursue two distinct strategies. First, I identified pages, which are similar but unlikely to be affected by the treatment. For a first comparison I focus on the network around older catastrophes, that occurred at a different point in time and were not from exactly the same domain (to avoid overlaps in the link network). [6] Given such a similar page, I, again, identified the set of pages which are one click away and which are two clicks away when the event occurs on the treated page. This gives me a set $L2_{control}$ which is both similar in size and also in the characteristics of the sampled pages (before the shock). Yet, since the choice of the start-pages in the comparison group is somewhat arbitrary. To approach this issue I also sampled the treated pages, but now only 42 days before the disaster or event occurred. For this group I simulate a treatment, by setting their $t = 0$ when no actual treatment occurred (I will refer to this group as the "placebo" group ($L2_{placebo}$) and

[4]Effects on the pages that are 2 clicks away were to small too be measured.

[5]Some of the consequences of major events, such as earthquakes, might change the state of the world and thus trigger a change in content, which is due to the event (e.g. destruction of an important monument) but not merely a consequence of the peak in interest and resulting improvements. Therefore I do not emphasize the change in activity on the pages that are only one click away for disasters. Moreover, to be certain I do not mix up directly and indirectly linked pages, I exclude any pages that were at any later point directly linked to the event page.

[6]This approach is not satisfactory in many ways. In ongoing work I control by using similar events that occur at a different point in time and to reduce the possible overlap. I also plan to select the control groups based on matching procedures. Note however, that my approach is generally quite robust independently of how I specify the control group. Alternatively I tried the following control group: for a region which was affected by an earthquake is compared to a region of similar size and relative importance in a similar, but remote, geographic space or the page of an airline, which lost one of its planes in an air crash is compared to an airline of similar importance but in a different region of the world. Such a change in the specification of the control group does not affect my results. (available upon request).

to their treatment as "placebo-treatment"). The obvious advantage of this comparison group over the control group described above stems from the fact, that it consists of the treated articles and their neighbors themselves. This comes at the cost of observing the pages at a different point in time. A third control group of "unrelated" observations results from the combination applying a placebo to the control group. Although this set of observations actually emerged as an artefact from the data extraction it provides yet another group that can be compared to the treated group.

An example of a natural disaster in the dataset is the "Sichuan Earthquake", which took place on May, 12th 2008 in the Province of Sichuan, PRC. The main consequence of this event were more than 60,000 dead and the region also suffered substantial economic loss. Suitable control pages could be pages about similar regions in far away places, or pages about other regions or countries, which were hit by large natural catastrophes, but at a different point in time. The placebo-control would be the same set of pages (on Sichuan and surrounding pages), but evaluated 7 weeks before the event. Table 1 shows which events were included in the data. These include both Natural Disasters as well as technical or economic catastrophes. Since the main focus of this study lies on the pages that are two clicks away, the table also shows the number of observations that are associated with each event.[7]

A representative Şfeatured articleŤ (ŞSeite des TagesŤ) might be the page about *Banjo-Kazooie*, which was featured on June 10th, 2010. It describes *Banjo-Kazooie*, which appears to be a highly commendable Nintendo-64, *JumpŠnŠrun* video game, which I am admittedly not familiar with. Table 4 shows which featured articles that were chosen by my procedure and were included in the data. In general, the variety of topics that are covered by the articles is much wider than in the other sample. They cover topics as varied as innovations (e.g. the CCD-sensor), places (Helgoland), soccer clubs (Werder Bremen) and art historical topics (Karolingische Buchmalerei - book-illustrations in the carolingian period). For the featured articles treatment the focus of interest lies on the pages that are one click away. As before, the table also shows the number of observations that received a link from the article before it was featured. For one featured article the number of associated observations ranges from a 1,334 to 37,642. Control observations were articles that were featured either later or earlier in time (and neighbors) or, as before, the same set of pages, seven weeks before (after) ŞfeaturingŤ.

3.2 A closer look at the dataset

Summary statistics for the data on large events are shown in Table 2. The data contains 498,916 observations from 17,179 pages on the main variables. From the table it

[7]Note, that each page shows up 29 times in the raw data and was sampled twice (placebo and real treatment), so that the number of corresponding pages (treatment or control) can be inferred by dividing the number of observations by 58.

can be seen, that the average page contains 5512 bytes of content and has undergone 80 revisions. However, the median is substantially lower (3778 bytes and only 37 revisions). Also, the summary statistics of the first differences (variables starting with "del_" reveal, that, on a typical day, nothing happens on a given page on Wikipedia. This highlights the necessity to use major events as a focal lense for analyzing activity on Wikipedia.[8] This is confirmed by the visual inspection of the direct and indirect effect of treatments.

In Figure 1 I plot the average clicks (left column) and the average number of added revisions (right columns) for the three groups of pages (zero clicks away (upper row), one click away (middle row) and two clicks away (lower row)). The two lower rows in this figure contains four lines. The first represents the treated group (or it's neighbors) when they were actually treated (hence flag_treated = 1 and placebo_state = 0). The second line represents the same group but during the placebo treatment at an earlier point in time. The third line (flag_treated = 0 and placebo_state = 0) shows the control group at the time when the real shock occurred and the fourth line represents the "unrelated" observations, which are never treated and taken in the placebo period. [9] The upper row contains only two lines, three lines, showing the control group and the directly treated nodes, which are created only after the onset of the event and have no placebo condition available. It shows, that the directly affected pages experience a very large spike of 15,000 clicks per day on average. Also the number of additional revisions peaks on the first days of treatment, when the pages are created: an average of almost 60 revisions are added to a page on the first day. Also on the pages that are to share a reciprocal a link from the treated page the effect is quite pronounced: Yet, while the clicks on the average L1 page increase by 2,500, the absolute value of the average increase in revision activity is already no more than 5. When I look at pages that are two click away, the effects are much smaller (especially for revisions) but quite pronounced. The clicks on the average adjacent page go up by 35 and the absolute value of the average increase in revision activity is already no more than 0.04.

The data from "featured articles" are shown in Table 5. The data contains 355,917 observations from 12,273 pages on the main variables. Note, that this corresponds to a much smaller number of pages per treatment, which is due to the fact that I focus on the directly linked pages in this condition. The table shows, that the median page contains 4703 bytes of content and has undergone 49 revisions. Also in this sample, the mean is substantially higher (6644 bytes and only 95 revisions). As before, the summary statistics of the first differences show clearly how little activity occurs on a normal day

[8]Further descriptive analyses that compare treated and control groups before and during treatment show that the groups are very similar in their activity levels before the shocks occurred and that the control group did not change it's behavior during treatment. These tables and their description were omitted for reasons of brevity. They are available from the author upon request.

[9]For greater ease of representation I included a graphical representation of only two variables. The summary statistics for these groups before and after treatment are also available as tables upon request.

on any given page on Wikipedia.

Figure 2 plots the aggregate dynamics around the event and corresponds to 1 for the large event condition. I plot the average clicks (left column) and the average number of added revisions (right columns), but now only for the treated pages and the direct neighbors. As before, each of the four figures contains four lines, one for each conditions that can be obtained from combining treatment (yes/no) and placebo (yes/no). The major difference to the large events condition is the brevity of the treatment. Attention rises from the typical levels (below 50 views) to more than 4200 (on average) views but it immediately returns to the old levels the day after treatment was administered. A very similar pattern can be observed for the neighbors where attention is almost twice as high as on a usual day and then falls back to the old levels. A similar pattern can be observed for the number of revisions, but, other than for large events, it can be observed that activity rises already before t=0. Nevertheless, on the day of treatment the spike of activity is pronounced also for the neighbors.

4 Estimation results

In what follows I present estimation results for these equations for both groups and discuss their interpretation. Before I proceed with presenting the details of my estimations, it is worth recalling a few important facts. First of all, recall, that the main focus in this paper lies on the estimation of the equation corresponding to equation 3 for large events and 2 for featured articles in section 2.1. This is due to two reasons: first, the two conditions differ in how local the treatment is, that I exploit for estimation and second, only the “featured articles” exist at treatment, while the page at the center of the treatment does not yet exist for the large events and will only be created during the days to follow.

Moreover, recall that I deal with potentially endogenous link formation that might arise as a result of the treatment by considering only links that had been in place *a week before the treatment*. Moreover, when a page was sampled to lie in both treatment and control group it was excluded from estimation, whenever identifiable. Yet, note that including such pages will bias the estimated coefficients towards zero. Also extremely broad pages with a very large number of links (e.g. pages that correspond to years) were excluded from estimation to avoid biases from oversampling them. Finally, I use the 7 observations from two weeks before treatment (days -14 through to -8) as the reference group in the estimations and I use only flow variables (clicks, new revisions, new authors etc.) to guarantee that my results are not driven by any anticipation effects[10]

[10] Anticipation effects are impossible for disasters but cannot be entirely ruled out in the featured articles condition, where sophisticated users, who can obtain the information about pages that are going to be presented soon. In fact the editors of the daily featured article, have to edit the article in the week

4.1 Large Events

For this group the estimation concerns the set of L2 pages, the pages that are two clicks away from the epicenter (the future page about the disaster). This is not because closer pages are uninteresting, but because the shock of the analyzed events is very big and very likely directly affected a page that shall be directly and bidirectionally linked. If, for example, a city in the province under consideration was hit by the earthquake, the added content on that page might simply cover this very fact. In such a case, this is not an improvement that arose from the increased attention that results from the adjacent event, but a change that is directly caused by the treatment. As was already explained above, this is not the effect I am primarily interested in. Consequently I focused on pages that receive an indirect link, because these are no longer likely to be directly affected by the treatment on the page two clicks away.[11] Moreover, to make sure that also my L2 pages are not directly related to the event, I checked, whether a page that was in L2 when I evaluated the network (a week before the shock) was going to be linked to the page of the disaster at any later point in time. If this was the case, I concluded the page might have been affected by the shock, despite having been in L2 before the shock and eliminated it from the sample. Thus I can ensure that only pages that were indirectly linked at the time of the shock and that also never got directly linked enter the sample.

The results for the estimation of the model for L2 nodes are shown in Table 3.[12] The table shows the results for clicks in the first three columns and the results for the number of added revisions in columns 4,5 and 6. All the specifications are OLS panel regressions, which include a fixed effect for the page and standard errors are clustered on the event level (20 clusters). For ease of representation the table only shows the coefficients for the cross terms from 2 periods before the shock until 4 periods after the shock. As was explained before, until the onset of the event (periods -2 to 0), we would effect insignificant effects for the cross terms and after the event has occurred a positive effect would imply that some form of spillover can be measured. Very much in line with the visual evidence, the average increase in click, relative to the control group (column 1), amounts to up to 40 more clicks on average. For the placebo treatment (column 2) this effect is almost equal. This is somewhat different for the number of revisions (as the graphical analysis had already suggested), since the effects are much smaller. A small effect is revealed from the first day after the treatment. This effect is small in absolute terms, since roughly one in 30 to 40 pages gets an additional revision. Yet, given the low

before it goes online, to make sure it fits on Wikipedia's start page, which invariably results in increased activity during the week before treatment.

[11]The results for the L1 group are included in the appendix. The effects are very large and statistically significant. The estimated coefficients for the L0 group (not reported) are in the 10,000s for clicks and between 60-35 for revisions. However, due to the lack of sufficient observations, even these very large coefficient estimates are not statistically different from zero.

[12]Non-parametric comparisons of the coefficients of each group taken separately confirm the results from the panel regressions and are thus not reported. They are available upon request.

levels in average activity on a given page on a given day, this is still a noteworthy effect.

4.2 Neighbors of Featured Articles

Table 6 shows the results for the featured articles. For this reduced form estimation I consider the model for L1 nodes (equation 2). This is the relevant group here, because the treatment takes place entirely inside Wikipedia (usually no media coverage or anything of the like) and it is "completely local", in the sense that no two articles can be featured at the same time. Hence, the different nature of the treatment guarantees that only the treated page is directly affected and any variation in the neighbors is almost certainly a result of the processes that take place inside Wikipedia.

The first three columns of the table show the results with clicks as the dependent variable. The estimation is the same as in Table 7 and also the clustering is implemented on the level of events (like before). The main insight of this table is that it confirms the statistical significant of the effect and provides a quantification of its size. The size of the effect is estimated to be 37 to 38 additional clicks on the average neighbor page on the day of treatment. Also on the revisions (columns 4-6) I observe an important effect of about 0.035 additional revisions one day after the treatment of the neighbor page. Note two things here: Firstly, the effect is very small in absolute terms and corresponds to one additional edit per 30 pages. Secondly however, this is an increase in contribution activity by 80 to 100 per cent.

I tested the robustness of my results by excluding the first third of the "featured articles".[13] Table 8 shows the result of the check and adds a new dependent variable, the change in the number of editors (in columns 5 to 6). In general, results reveal the same pattern as table 7, but the significance levels might be lower. The number of authors moves largely in parallel with the number of revisions, indicating that twice as many new authors as usual edit the article due to the treatment of their neighbor. Yet, while this is a large effect in relative terms it means that only one in 70 articles is edited by a new author.

Another way of understanding the meaning of these point estimates, consists in aggregating the changes in clicks and revisions over all neighboring articles and then averaging over the 36 different featured articles. This is done in figures 3 and 4 in order to summarize and illustrate the insights from the "featured articles" condition. I find that, on average, there are 4000 clicks on all neighbors taken together (Figure 3). Given that also the average treated articles received an additional 4000 clicks this corresponds to a one for one conversion of clicks on the treated page to clicks on one of the neighbors. In other words, the average visitor clicks on exactly one of the links. The total number of

[13] This is clearly not final, but splitting the sample is a common and useful first check to test whether the results are robust.

revisions on the neighboring pages (Figure 4) increases from approx. 4.5 to roughly 8.5. This is an additional four changes, which means that the 4000 initial additional clicks are converted in 4000 additional clicks and four new revisions or a ratio of 1000:1000:1.

Finally I report results of an extended analysis, which were omitted here, for reasons of space.[14] I included the number of clicks on the treated page in the regression and, as expected, the number of links on the neighboring pages is positively related to that value. Moreover I split the sample in well connected articles (many links) and poorly connected ones, but I do not find a significant relationship between this variable and the number of visits. The same is true for a variable that captures whether a page is very long or not. I get a positive but insignificant point estimate for page views. However, when I consider only "stubs", i.e. pages that do not exceed a length of only 1500 bytes, I find a much stronger effect in the number of edits. This indicates that the new content that is provided after all, is provided on pages, where the existing content is little.

5 Concluding remarks and further research

In this paper I analyze whether the link network between articles on the German Wikipedia influences how much content is provided to it by users. I use observable exogeneous shocks, such as large scale media events or natural disasters as focal lense to analyze the spillovers through networks of user generated content that are mediated through links. I deploy this strategy to see whether some of the additional attention on the shocked pages is channeled across the links in the link network.

In the analysis I rely on self-generated linkage data from the German Wikipedia that is matched with data on the articles and with information on Wikipedia page views. From the resulting database of more than 35,000 articles that were indirectly linked to shocked pages. Though my results are preliminary, I find substantial spillovers of attention even to pages that are two clicks away, but relatively modest and not necessarily robust effects on content provision.

This relationship suggests, that in network that consist of interlinked nodes of content to be created by users, the links between nodes seem to be an important medium for attention spillovers. How much attention a node will receive can be influenced by the links it receives. Yet, my results indicate also that the spillovers in attention may be mostly for the purposes of looking up information. It remains to be determined, whether links also affect how much content will be created.

The answer to these questions sheds light on whether the links forming the content network can be used to manage contribution flows on platforms that rely on user generated content. Hence, it is important for administrators of both new wikis and of

[14]They are available from the author upon request.

burgeoning platforms for knowledge documentation, who are worried about channeling the flows of content contribution. Note that the link network between articles is a citation network. Thus, our findings allow for a more abstract reading when interpreting Wikipedia as a peer produced tool for the documentation of human knowledge, i.e. a setting of peer production, similar to the production of open source software or scientific research. Viewed under this light, our results suggest that the attention to a certain field or project will be more likely, if it receives links from other areas. However, the analysis of the "featured analysis" suggests that the average visitor clicks on exactly one of the links. The total number of revisions on the neighboring pages (Figure 4) increases from approx. 4.5 to roughly 8.5. This is an additional four changes, which means that the 4000 initial additional clicks are converted in 4000 additional clicks and four new revisions or a ratio of 1000:1000:1.

This project is ongoing work and several improvements are conceived or even under way, but could not be included in the current version of this paper. In **Further research** I hope to exploit the heterogeneity in the direct treatment effects more thoroughly. In particular, I hope to understand, whether the attention (that is currently measured as average effect) is evenly distributed across nodes, or whether the users actually herd in only a few of the directly linked pages.

Moreover, even though my design allows a causal interpretation of the reduced form estimates, I plan to add a structural model of the underlying dynamic with which the clicks on neighboring pages are transmitted to each other. This entails the provision of both an identification strategy and estimates of the parameter of interest, which will then quantify how attention spills between nodes and how it translates to content generation. Another promising area for further analysis would aim at investigating whether new authors are attracted by the events or whether contributions are made only by authors that contributed to the subject before.

References

Berge, Lars Ivar Oppedal, "Measuring Spillover Effects from Business Training," 2011.

Bramoullé, Yann, Habiba Djebbari, and Bernard Fortin, "Identification of peer effects through social networks," *Journal of econometrics*, 2009, *150* (1), 41–55.

Carmi, Eyal, Gal Oestreicher-Singer, and Arun Sundararajan, "Spreading the oprah effect: The diffusion of demand shocks in a recommendation network," 2009.

Claussen, J., O. Falck, and T. Grohsjean, "The strength of direct ties: Evidence from the electronic game industry," *International Journal of Industrial Organization*, 2012, *30* (2), 223–230.

Ferracci, Marc, Grégory Jolivet, and Gerard J van den Berg, "Evidence of Treatment Spillovers Within Markets," 2012.

Fershtman, C. and N. Gandal, "Direct and Indirect Knowledge Spillovers: The 'Social Network' of Open Source Software," *RAND Journal of Economics*, 2011, *42* (1).

Giorgi, Giacomo De, Michele Pellizzari, and Silvia Redaelli, "Identification of social interactions through partially overlapping peer groups," *American Economic Journal: Applied Economics*, 2010, *2* (2), 241–275.

Giorgidze, G., T. Grust, T. Schreiber, and J. Weijers, "Haskell Boards the Ferry: Database-Supported Program Execution for Haskell," in "Revised selected papers of the 22nd international symposium on Implementation and Application of Functional Languages, Alphen aan den Rijn, Netherlands," Vol. 6647 of *Lecture Notes in Computer Science* Springer 2010. Peter Landin Prize for the best paper at IFL 2010.

Halatchliyski, I., J. Moskaliuk, J. Kimmerle, and U. Cress, "Who integrates the networks of knowledge in Wikipedia?," in "Proceedings of the 6th International Symposium on Wikis and Open Collaboration" ACM 2010, p. 1.

Imberman, Scott, Adriana D Kugler, and Bruce Sacerdote, "Katrina's children: Evidence on the structure of peer effects from hurricane evacuees," Technical Report, National Bureau of Economic Research 2009.

Keegan, Brian, Darren Gergle, and Noshir Contractor, "Hot Off the Wiki: Structures and Dynamics of WikipediaŠs Coverage of Breaking News Events," *American Behavioral Scientist*, 2013.

Kummer, Michael, Marianne Saam, Iassen Halatchliyski, and George Giorgidze, "Centrality and Content Creation in Networks–The Case of German

Wikipedia," *ZEW-Centre for European Economic Research Discussion Paper*, 2012, (12-053).

Manski, Charles F, "Identification of endogenous social effects: The reflection problem," *The review of economic studies*, 1993, *60* (3), 531–542.

Zhang, X. and F. Zhu, "Group Size and Incentives to Contribute: A Natural Experiment at Chinese Wikipedia," *The American Economic Review*, 2011, *101*, 1601–1615.

6 Tables and Figures

6.1 Disasters

6.1.1 Descriptive Analysis for Disasters

Table 1: Included disasters and the pages that are associated with them (2 clicks away).

name of event	No.
Air-France-Flug_447	5,887.0
Air-India-Express-Flug_812	21,373.0
Amoklauf_von_Winnenden	3,973.0
Anschläge_am_26._November_2008_in_Mumbai	290.0
Ausbruch_des_Eyjafjallajökull_2010	14,848.0
Buschfeuer_in_Victoria_2009	1,827.0
Erdbeben_in_Haiti_2010	21,866.0
Erdbeben_in_Sichuan_2008	13,253.0
Erdbeben_von_L'Aquila_2009	7,424.0
Flugzeugabsturz_bei_Smolensk	19,343.0
Grubenunglück_von_San_José	6,235.0
Josef_Fritzl	7,308.0
Kaukasuskrieg_2008	19,807.0
Kolontár-Dammbruch	5,800.0
Luftangriff_bei_Kunduz	101,674.0
Northwest-Airlines-Flug_253	65,424.0
Pandemie_H1N1_2009/10	8,903.0
US-Airways-Flug_1549	13,688.0
Unglück_bei_der_Loveparade_2010	28,855.0
Versuchter_Anschlag_am_Times_Square	8,990.0
Wald-_und_Torfbrände_in_Russland_2010	15,863.0
Wohnhausbrand_in_Ludwigshafen_am_Rhein	11,977.0
Zyklon_Nargis	21,982.0
Überschwemmungskatastrophe_in_Pakistan_2010	72,326.0
Total	498,916.0

NOTES: The table shows summary statistics for all pages that are two clicks away from a start page (be it treated, placebo or control). Pages included 17,179

Table 2: Summary statistics: FirstSumstatsPanelL2dis of main variables

	count	mean	sd	min	p10	p50	p90	max
Length of page (in bytes)	498916	5512	6164	16	34	3778	12818	76176
Number of authors	498916	28	33	1	1	17	68	435
Clicks	498916	30	164	0	0	0	62	29865
Number of Revisions	498916	80	127	1	2	37	201	2083
SNwik_degree	498916	117	430	0	5	30	264	27537
Dummy: literature section	498916	.19	.39	0	0	0	1	1
Number of images	498916	1.3	2.7	0	0	0	3	57
Number language links	498916	13	18	0	0	7	37	179
References (footnotes)	498916	1.2	4.1	0	0	0	3	150
Links to further info	498916	2.6	5.2	0	0	1	6	218
time_var	498916	0	8.4	-14	-12	0	12	14
del_num_revisions	481712	.034	.35	0	0	0	0	44
del_page_length	481712	1.5	97	-22416	0	0	0	22416
del_num_authors	481712	.012	.12	0	0	0	0	11
del_SNwik_degree	481712	.046	2.2	-1148	0	0	0	216
del_num_imagelinks	481712	.00042	.077	-27	0	0	0	20
del_num_refs	481712	.0011	.11	-32	0	0	0	26
del_num_extlinks	481712	.0014	.33	-16	0	0	0	214

Figure 1: Contrasting means of clicks vs. number of added revisions over time: looking at all 4 groups in one plot.

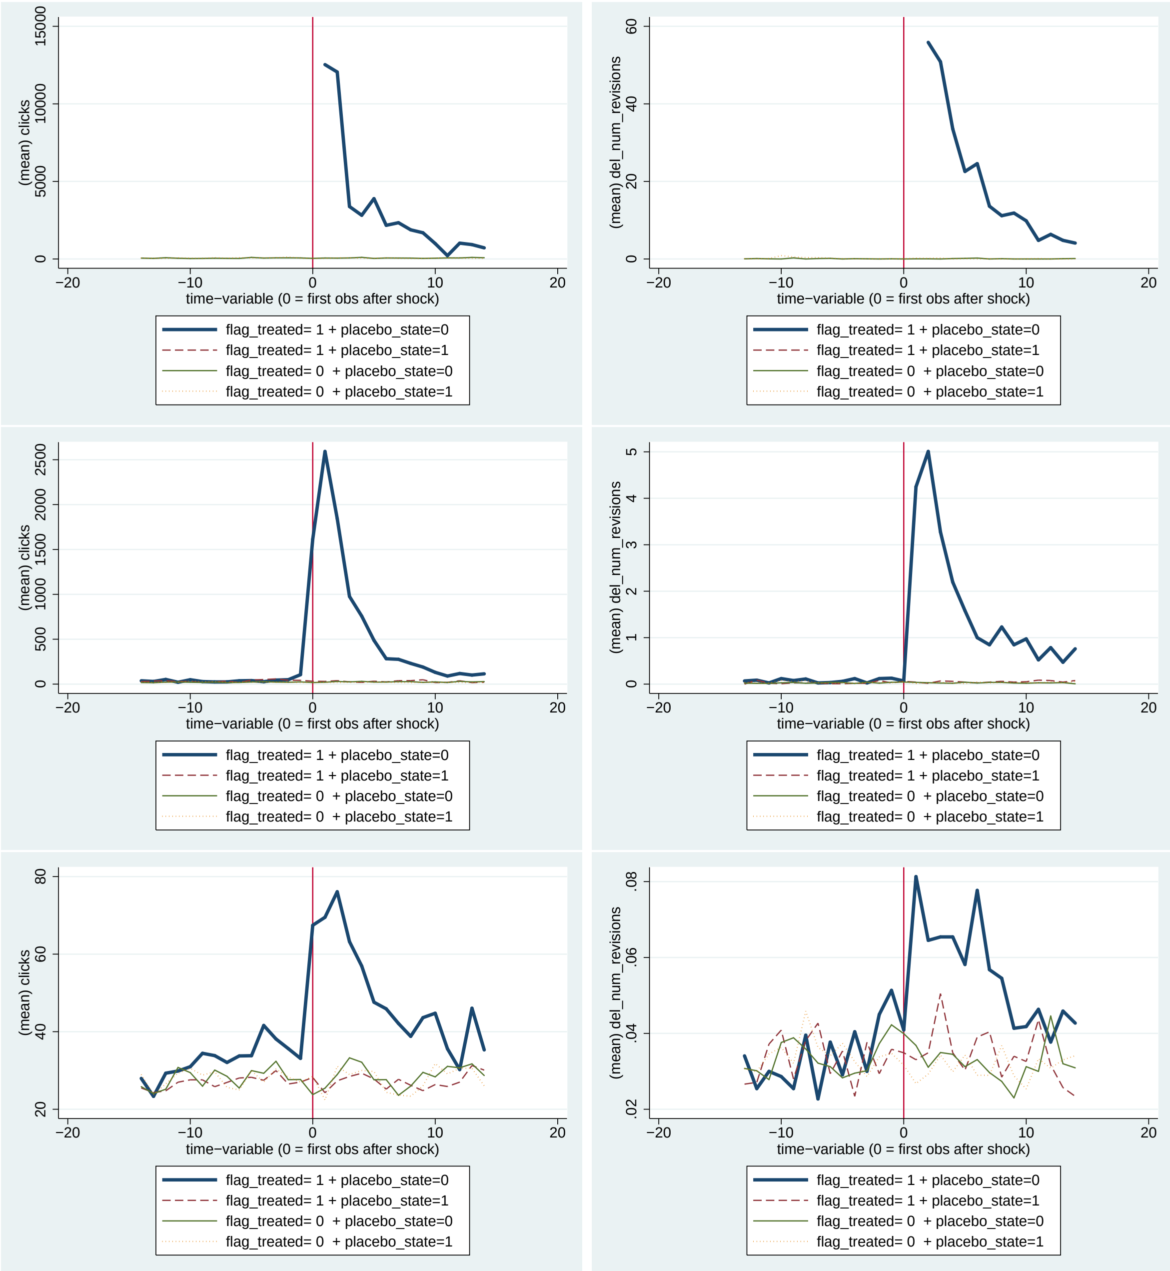

Notes: The upper row shows the average effect on the event pages, the middle row the directly treated pages (L1, with reciprocal link), and the lower row for the pages that are one click away from L1. Directly hit pages received up to 15,000 additional clicks and up to 60 new revisions on average. Pages that will have a reciprocal link received up to approx. 2,500 clicks and up to 5 additional revisions. However, not only the treated pages, but also their neighbors received 40 additional clicks and up to 0.04 additional revisions on average.

Table 3: Relationship of clicks/added revisions and time dummies.

	clicks			del revisions		
	(1) compare control	(2) compare placebo	(3) compare all	(4) compare control	(5) compare placebo	(6) compare all
t = -2	4.637	4.911	4.462	0.008	0.016*	0.012
	(4.734)	(5.138)	(4.568)	(0.009)	(0.009)	(0.008)
t = -1	2.167	1.894	2.460	0.013	0.016	0.018**
	(3.956)	(3.750)	(3.506)	(0.009)	(0.010)	(0.008)
t = 0	24.488***	20.455***	21.991***	0.003	0.009	0.009
	(5.715)	(6.164)	(5.456)	(0.009)	(0.010)	(0.008)
t = 1	40.036***	39.578***	41.395***	0.037*	0.042*	0.043**
	(10.978)	(11.011)	(10.858)	(0.018)	(0.024)	(0.018)
t = 2	39.466***	39.822***	38.764***	0.026**	0.025*	0.028**
	(14.381)	(14.453)	(14.236)	(0.012)	(0.015)	(0.011)
t = 3	24.386***	28.787***	26.941***	0.034***	0.018	0.032***
	(8.682)	(8.408)	(8.249)	(0.011)	(0.013)	(0.010)
t = 4	19.878***	21.984***	21.246***	0.029**	0.030**	0.032***
	(6.730)	(6.872)	(6.625)	(0.012)	(0.013)	(0.012)
Constant	27.945***	27.631***	27.585***	0.032***	0.032***	0.034***
	(0.911)	(1.326)	(0.589)	(0.001)	(0.003)	(0.001)
All cross	Yes	Yes	Yes	Yes	Yes	Yes
Time Dummies	Yes	Yes	Yes	Yes	Yes	Yes
Observations	189728	98582	377938	181104	94101	360759
Number of Pages	8624	4481	17179	8624	4481	17179
Adj. R^2	0.003	0.004	0.002	0.000	0.001	0.000

Standard errors in parentheses

Fixed Effects Panel-Regressions with heteroscedasticity robust standard errors.

Only crossterms closer to treatment are shown, but all were included. Reference group t-14 to t-5

* p<0.10, ** p<0.05, *** p<0.01

6.2 Page of Day

6.2.1 Descriptive Analysis for Page of Day

Table 4: Included disasters and the pages that are associated with them (1 clicks away).

name of event	No.
Afrikaans	8,352.0
Alte_Synagoge_(Heilbronn)	3,741.0
Banjo-Kazooie	7,801.0
Benno_Elkan	9,135.0
Bombardier_Canadair_Regional_Jet	5,568.0
CCD-Sensor	37,642.0
Charles_Sanders_Peirce	18,241.0
Das_Kloster_der_Minne	3,538.0
Deutsche_Bank	21,257.0
Eishockey	10,672.0
Ekel	17,806.0
Fahrbahnmarkierung	2,726.0
Geschichte_Ostfrieslands	14,036.0
Geschichte_der_deutschen_Sozialdemokratie	20,184.0
Ghetto_(Venedig)	1,334.0
Glanzstoff_Austria	16,646.0
Glorious_Revolution	10,904.0
Granitschale_im_Lustgarten	5,713.0
Gustav_Hirschfeld	8,700.0
Hallenhaus	5,452.0
Helgoland	14,384.0
Jaroslawl	21,460.0
Jupiter_und_Antiope_(Watteau)	2,378.0
Karolingische_Buchmalerei	10,266.0
Katholische_Liga_(1538)	3,074.0
Martha_Goldberg	3,364.0
Naturstoffe	18,676.0
Paul_Moder	4,466.0
St._Martin_(Memmingen)	4,060.0
Stabkirche_Borgund	2,668.0
Taiwan	10,498.0
USS_Thresher_(SSN-593)	5,945.0
Visum	3,944.0
Wenegnebti	3,306.0
Werder_Bremen	17,980.0
Total	355,917.0

NOTES: The table shows summary statistics for all pages that are two clicks away from a start page (be it treated, placebo or control).

Table 5: Summary statistics: FirstSumstatsPanelL1PoD of main variables

	count	mean	sd	min	p10	p50	p90	max
Length of page (in bytes)	355917	6644	6722	17	43	4703	15072	81585
Number of authors	355917	33	35	1	2	22	77	324
Clicks	355917	35	142	0	0	0	84	20384
Number of Revisions	355917	95	130	1	3	49	236	1382
SNwik_degree	355917	133	491	0	7	41	305	27687
Dummy: literature section	355917	.3	.46	0	0	0	1	1
Number of images	355917	2.2	7.7	0	0	1	5	319
Number language links	355917	13	18	0	0	6	38	180
References (footnotes)	355917	1.3	4.3	0	0	0	4	182
Links to further info	355917	2.3	4.2	0	0	1	6	155
time_var	355917	0	8.4	-14	-12	0	12	14
del_num_revisions	343644	.042	.38	0	0	0	0	42
del_page_length	343644	2.1	156	-31473	0	0	0	31462
del_num_authors	343644	.014	.13	0	0	0	0	9
del_SNwik_degree	343644	.058	1.1	-90	0	0	0	438
del_num_imagelinks	343644	.001	.26	-50	0	0	0	132
del_num_refs	343644	.0014	.091	-7	0	0	0	18
del_num_extlinks	343644	.0007	.097	-19	0	0	0	16

Figure 2: Page of Day experiment: Contrasting means of clicks vs. number of added revisions over time: looking at all 4 groups in one plot.

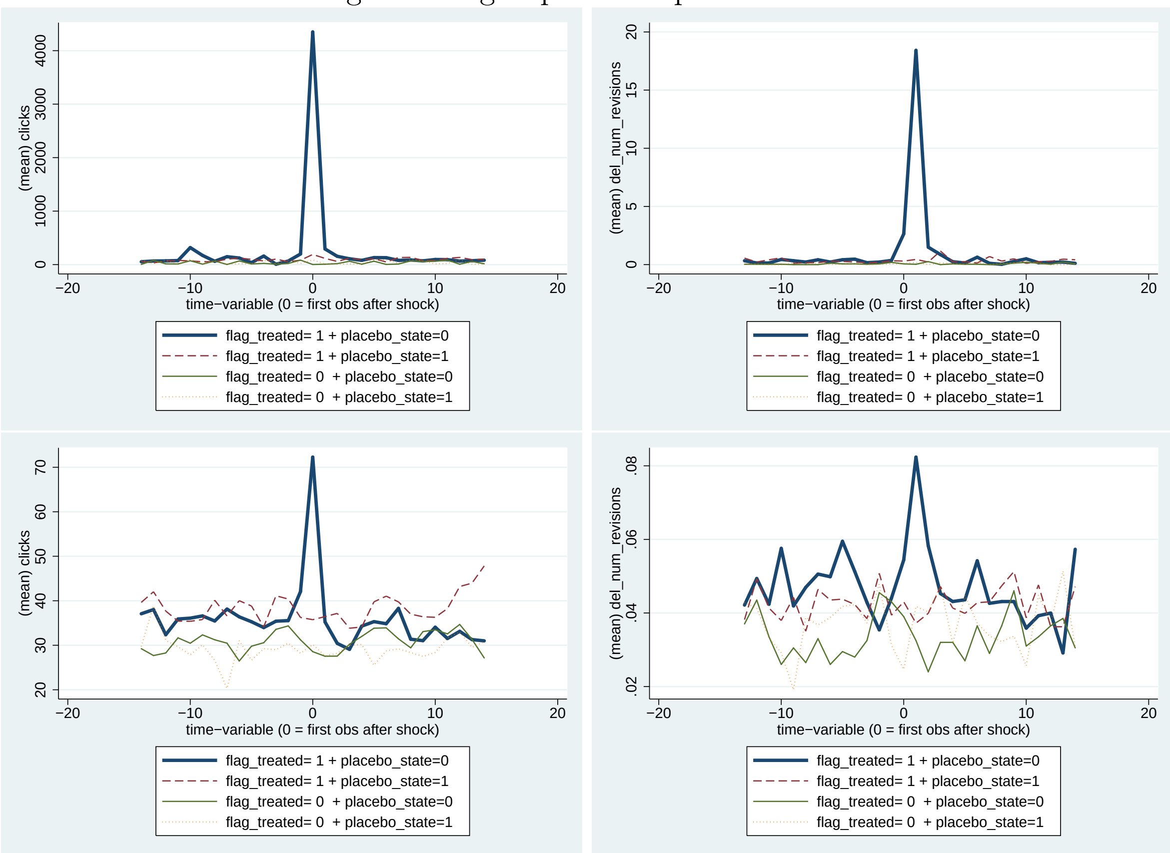

NOTES: The upper row shows the average effect on the directly treated pages, the lower row for the pages one click away.

Figure 3: Figure contrasting the mean of clicks on featured articles, with the aggregated clicks on all neighboring pages.

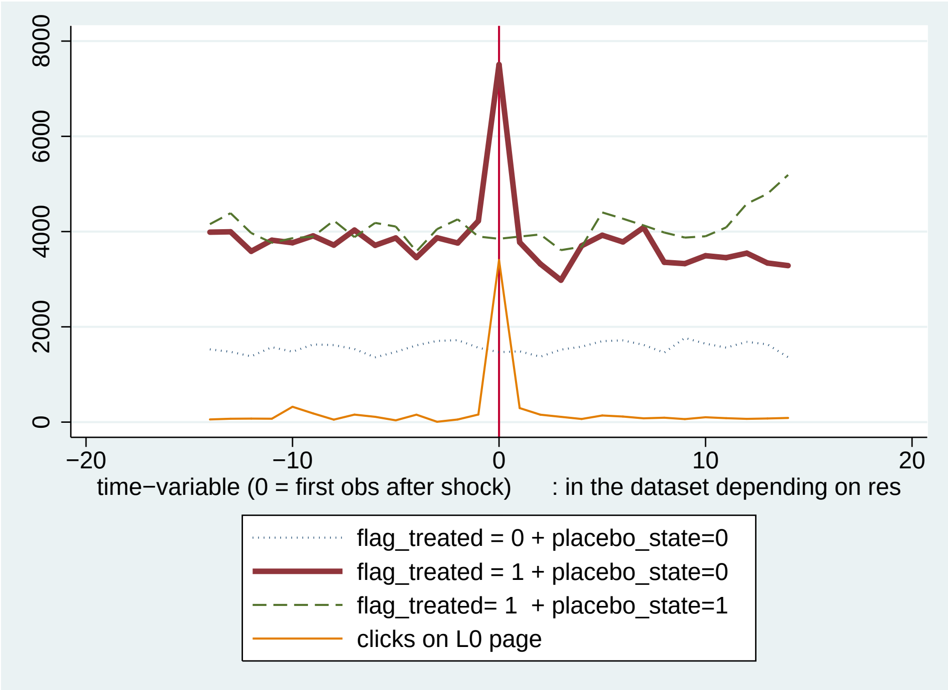

Notes: The figure shows the aggregated effect on the pages that are one click away. The average treated page received up to 4000 additional clicks, all neighbors together received approx. the same number of additional clicks

Figure 4: Figure showing the aggregated new revisions on all neighboring pages.

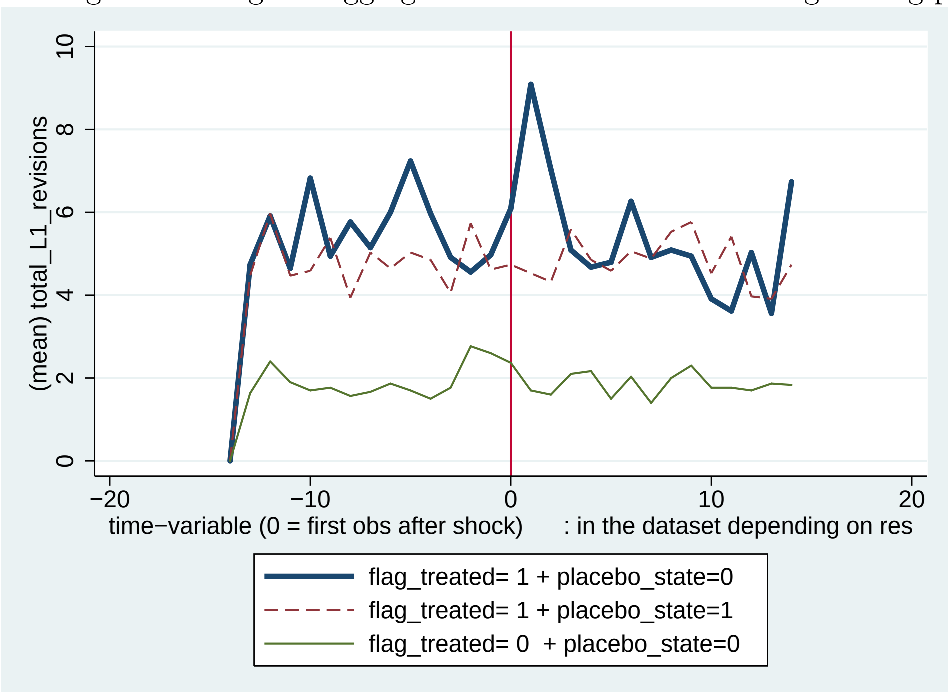

Notes: The figure shows the aggregated effect on the pages that are one click away. All neighbors of treated articles together received approx. four additional revisions.

6.2.2 Regression Tables for Page of the Day

Table 6: Featured Articles: Relationship of clicks/added revisions and time dummies at the neighbors.

	clicks			del revisions		
	(1) compare control	(2) compare placebo	(3) compare all	(4) compare control	(5) compare placebo	(6) compare all
t = -2	-5.187	-2.836	-3.176	-0.027**	-0.022**	-0.024***
	(3.752)	(3.509)	(3.015)	(0.010)	(0.010)	(0.008)
t = -1	4.478	7.836	6.784	-0.016	-0.002	-0.005
	(5.374)	(5.887)	(5.087)	(0.010)	(0.008)	(0.007)
t = 0	37.353***	38.560***	37.438***	-0.002	0.005	0.006
	(10.743)	(10.995)	(10.691)	(0.009)	(0.008)	(0.007)
t = 1	1.349	0.835	0.926	0.033***	0.038***	0.034***
	(2.735)	(3.309)	(2.416)	(0.012)	(0.013)	(0.012)
t = 2	-3.514	-4.724	-4.423	0.017**	0.012	0.011
	(3.371)	(3.603)	(3.126)	(0.007)	(0.010)	(0.008)
t = 3	-7.524**	-2.729	-5.164	-0.004	-0.009	-0.009
	(3.614)	(5.804)	(4.049)	(0.009)	(0.012)	(0.010)
t = 4	-4.056	2.207	-0.566	-0.006	-0.005	-0.004
	(3.280)	(6.080)	(3.875)	(0.011)	(0.009)	(0.008)
Constant	34.246***	37.018***	34.652***	0.043***	0.045***	0.041***
	(0.776)	(0.797)	(0.582)	(0.002)	(0.002)	(0.001)
All cross	Yes	Yes	Yes	Yes	Yes	Yes
Time Dummies	Yes	Yes	Yes	Yes	Yes	Yes
Observations	135322	182336	270006	129171	174048	257733
Number of Pages	6151	8288	12273	6151	8288	12273
Adj. R^2	0.004	0.003	0.002	0.000	0.000	0.000

Standard errors in parentheses
Fixed Effects Panel-Regressions with heteroscedasticity robust standard errors.
Only crossterms closer to treatment are shown, but all were included. Reference group t-14 to t-5
* p<0.10, ** p<0.05, *** p<0.01

A Table-Appendix

Table 7: Relationship of clicks/added revisions and time dummies.

	clicks		del revisions	
	(1) compare control	(2) compare placebo	(3) compare control	(4) compare placebo
t = -2	-5.329**	-3.627	0.033	0.044*
	(2.526)	(5.724)	(0.023)	(0.025)
t = -1	6.802**	-2.482	0.016	0.037
	(2.722)	(5.710)	(0.025)	(0.026)
t = 0	18.739	15.547	0.000	-0.020
	(12.240)	(14.856)	(0.017)	(0.042)
t = 1	113.499**	115.780**	0.062	0.081*
	(44.580)	(44.941)	(0.040)	(0.042)
t = 2	127.670***	137.607***	0.225***	0.240***
	(47.978)	(48.167)	(0.084)	(0.085)
t = 3	69.339***	73.648***	0.146***	0.175***
	(21.774)	(22.271)	(0.046)	(0.046)
t = 4	46.821***	57.460***	0.107***	0.131***
	(16.231)	(16.757)	(0.036)	(0.037)
Constant	19.890***	32.316***	0.032***	0.049***
	(4.649)	(5.972)	(0.009)	(0.013)
All cross	Yes	Yes	Yes	Yes
Time Dummies	Yes	Yes	Yes	Yes
Observations	22440	18240	22440	18240
Number of Pages	1496	1216	1496	1216
Adj. R^2	0.010	0.009	0.006	0.006

Standard errors in parentheses

Fixed Effects Panel-Regressions with heteroscedasticity robust standard errors.

Only crossterms closer to treatment are shown, but all were included. Reference group t-7 to t-5

* p<0.10, ** p<0.05, *** p<0.01

Table 8: Robustness Check: Relationship of clicks/added revisions and time dummies for only a reduced number of events.

	clicks		del revisions		del authors	
	(1) compare control	(2) compare placebo	(3) compare control	(4) compare placebo	(5) compare control	(6) compare placebo
realtreat_x_period_13	-3.107	3.357	-0.026**	-0.019	-0.014***	-0.006
	(4.242)	(3.994)	(0.011)	(0.014)	(0.005)	(0.005)
realtreat_x_period_14	8.172	17.857**	-0.010	-0.003	-0.002	-0.004
	(7.726)	(8.676)	(0.011)	(0.011)	(0.004)	(0.004)
realtreat_x_period_15	40.980**	47.763***	-0.009	0.010	-0.007	-0.006
	(16.056)	(16.440)	(0.011)	(0.011)	(0.004)	(0.005)
realtreat_x_period_16	-0.763	2.874	0.038**	0.041**	0.014**	0.014*
	(2.753)	(4.633)	(0.017)	(0.019)	(0.006)	(0.007)
realtreat_x_period_17	-4.398	-0.573	0.012	0.005	0.003	-0.004
	(4.792)	(4.476)	(0.008)	(0.013)	(0.004)	(0.005)
realtreat_x_period_18	-7.070	-0.809	-0.006	-0.026	-0.005	-0.007
	(4.774)	(4.811)	(0.011)	(0.016)	(0.005)	(0.005)
realtreat_x_period_19	-1.833	6.907	-0.015	-0.008	-0.007*	-0.001
	(4.043)	(4.254)	(0.013)	(0.012)	(0.004)	(0.004)
_cons	33.467***	36.855***	0.045***	0.047***	0.016***	0.015***
	(1.180)	(1.098)	(0.003)	(0.003)	(0.001)	(0.001)
All cross	Yes	Yes	Yes	Yes	Yes	Yes
Time Dummies	Yes	Yes	Yes	Yes	Yes	Yes
Observations	84194	111298	80367	106239	80367	106239
Number of Pages	3827	5059	3827	5059	3827	5059
Adj. R^2	0.003	0.003	0.000	0.000	0.001	0.001

Standard errors in parentheses
Fixed Effects Panel-Regressions with heteroscedasticity robust standard errors.
Only crossterms closer to treatment are shown, but all were included. Reference group t-14 to t-5
* p<0.10, ** p<0.05, *** p<0.01

B The empirical model and structural identification of the parameter of interest.

B.1 Introductory remarks

[incomplete, needs to be formulated as text]

This section presents the structural model and discusses the coefficients we are interested, the usual problems in identifying them and possible avenues that have been suggested by the previous literature.

The underlying relationship we are interested in:

- What's the role of links in content generation?
- Is an article more likely to be improved because of spillovers through links.

$$y_{it} = \alpha \frac{\sum_{j \in P_{it}} y_{jt}}{N_{P_{it}}} + X_{it}\beta + \gamma \frac{\sum_{j \in P_{it}} X_{jt}}{N_{P_{it}}} + \epsilon_{it} \tag{8}$$

- α is the coefficient of interest. In context it measures for example how the clicks on page A are influenced the clicks on the adjacent pages.
- Yet, it is - generally - very hard to identify, even in stable networks.
- In Wikipedia there might be variation in P_{it} and y_{it}
- THIS RELATIONSHIP IS OF GENERAL INTEREST TO A VERY LARGE LITERATURE ON PEER EFFECTS!
- Bramoullé et al. (2009) suggest a more succinct notation:.

$$\mathbf{y_t} = \alpha \mathbf{G} \mathbf{y_t} + \beta \mathbf{x_t} + \gamma \mathbf{G} \mathbf{x_t} + \epsilon_\mathbf{t} \tag{9}$$

$$E[\epsilon_t | \mathbf{x_t}] = 0 \tag{10}$$

B.2 Setup and Basic Idea

I augment the model in equation 8 by observable treatments (shocks) that are locally applied.

$$y_{it} = \alpha \frac{\sum_{j \in P_{it}} y_{jt}}{N_{P_{it}}} + X_{it}\beta + \gamma \frac{\sum_{j \in P_{it}} X_{jt}}{N_{P_{it}}} + \delta_1 D_{it} + \epsilon_{it} \tag{11}$$

Interpretation of the 2 new coefficients:

- δ_1 ... measures the direct treatment effect if a node(page) is, treated.
- $X_{it}\beta$ may contain an individual fixed effect and an additively separable age-dependent part: $X_{it}\beta = \beta_i + \widetilde{X_{it}}\beta_1 + \beta_2 f(age)$.

To see how local treatments can be used as a source of identification, consider two pairs of nodes.

B.2.1 Local application of treatment

First, consider 2 connected nodes, where one is treated ($l0$) in period t and the neighbors are not treated ($l1 \in L1$). Assume for simplicity that l0 is the only treated node in l1's neighborhood.

$$\ell 0 :: y_{l0t} = \alpha \frac{\sum_{j \in P_{l0t}} y_{jt}}{N_{P_{l0t}}} + X_{l0t}\beta + \gamma \frac{\sum_{j \in P_{l0t}} X_{jt}}{N_{P_{l0t}}} + \delta_1 1 + \epsilon_{l0t} \tag{12}$$

$$\ell 1 \in L1 :: y_{l1t} = \alpha \frac{y_{l0t} + \sum_{j \in P_{l1t}/l0} y_{jt}}{N_{P_{l1t}}} + X_{l1t}\beta + \gamma \frac{\sum_{j \in P_{l1t}} X_{jt}}{N_{P_{l1t}}} + \delta_1 0 + \epsilon_{l1t} \tag{13}$$

B.2.2 Controls in remote part of the network around $c0$

Second, take two remote nodes c0 and $c1 \in C1$, where nothing happens (nobody gets treated).

$$c0 :: y_{c0t} = \alpha \frac{\sum_{j \in P_{c0t}} y_{jt}}{N_{P_{c0t}}} + X_{c0t}\beta + \gamma \frac{\sum_{j \in P_{c0t}} X_{jt}}{N_{P_{c0t}}} + \delta_1 0 + \epsilon_{c0t} \tag{14}$$

$$c1 \in C1 :: y_{c1t} = \alpha \frac{\sum_{j \in P_{c1t}} y_{jt}}{N_{P_{c1t}}} + X_{c1t}\beta + \gamma \frac{\sum_{j \in P_{c1t}} X_{jt}}{N_{P_{c1t}}} + \delta_1 0 + \epsilon_{c1t} \tag{15}$$

From this equation it can easily be seen, how the local treatment will allow to measure the spillover or peer effect. This will be possible despite the richness in other sources of variation, provided (i) the shocks are large enough and (ii) the "control network" allows to credibly infer the dynamics in the "treated network", had no treatment taken place. To formalize this more concretely, I will take a small detour and rewrite the model in the more succinct notation, that was already mentioned above.

B.2.3 Condensed Notation

Following Bramoullé et al. (2009), this can be written in Matrix notation and X might include a time-dependent component (e.g. a linear function of age) as well:

$$\mathbf{y_t} = \alpha \mathbf{G} \mathbf{y_t} + \mathbf{X_t}\beta + \gamma \mathbf{G} \mathbf{X_t} + \delta_1 \mathbf{D_t} + \epsilon_t \qquad E[\epsilon_t | D_t] = 0 \tag{16}$$

- $\mathbf{G}$ is NxN
- $G_{ij} = \frac{1}{N_{P_i}-1}$ if i receives a link from j and $G_{ij} = 0$ otherwise
- treated side: $D_t = e_{l0}$; i.e.: a vector with a 1 in the coordinate that corresponds to the treated node and 0's elsewhere.
- untreated side: $D_t = \mathbf{0}$, a vector of zeros.
- I DO NOT assume $E[\epsilon_t | \mathbf{X_t}] = 0$, since weaker assumptions will suffice.
- I DO NOT require that the structure of the network is exogenous.

The reduced form is given by:

$$\mathbf{y_t} = (\mathbf{I} - \alpha \mathbf{G})^{-1} [\mathbf{X_t}\beta + \gamma \mathbf{G} \mathbf{X_t} + \delta_1 \mathbf{D_t} + \epsilon_t] \tag{17}$$

Clearly, taking the first difference, we obtain a term that depends on the time-dependent component and the effect of any changes in the independent variables.[15]

$$\begin{aligned} \boldsymbol{\Delta} \mathbf{y_t} &= \mathbf{y_t} - \mathbf{y_{t-1}} = \\ &= (\mathbf{I} - \alpha \mathbf{G})^{-1} [\Delta \mathbf{X_t}\beta + \gamma \mathbf{G} \boldsymbol{\Delta} \mathbf{X_t} + \delta_1 \boldsymbol{\Delta} \mathbf{D_t} + \epsilon_t - \epsilon_{t-1}] \end{aligned} \tag{18}$$

I use two control groups in this paper: For one, I use a set of nodes that are remote to the treated node and second, I use the same nodes but only several weeks before the shock. Let's start by looking at the control group formed by the same network, but S periods earlier, i.e. in period $t - S$ then we have.

$$\mathbf{y_{t-S}} = \alpha \mathbf{G} \mathbf{y_{t-S}} + \mathbf{X_{t-S}}\beta + \gamma \mathbf{G} \mathbf{X_{t-S}} + \delta_1 \mathbf{D_{t-S}} + \epsilon_{t-S}$$

[15] If βX_{it} is modelled to contain an additively separable age-dependent part as in our example above, $\Delta X_{it-S}\beta$ would contain $\frac{df(age)}{dt}$. [which is to be eliminated by taking DiffDiff]

Analogously, the first difference of the reduced form will contain a time-dependent component and the effect of any other changes in the independent variables.[16]

$$\begin{aligned}\mathbf{\Delta y_{t-S}} &= \mathbf{y_{t-S}} - \mathbf{y_{t-S-1}} = \\ &= (\mathbf{I} - \alpha\mathbf{G})^{-1}[\Delta\mathbf{X_{t-S}}\beta + \gamma\mathbf{G}\mathbf{\Delta X_{t-S}} + \delta_1\mathbf{\Delta D_{t-S}} + \epsilon_{t-S} - \epsilon_{t-S-1}]\end{aligned}$$

Denoting $\epsilon_t - \epsilon_{t-1}$ as $\Delta\epsilon_t$ proceed to take the Difference in Differences, we obtain:

$$\begin{aligned}&\mathbf{\Delta y_t} - \mathbf{\Delta y_{t-S}} = \\ &= (\mathbf{I} - \alpha\mathbf{G})^{-1} * \quad \{[\Delta\mathbf{X_{t-S}}\beta + \gamma\mathbf{G}\mathbf{\Delta X_t} + \delta_1\mathbf{\Delta D_t} + \Delta\epsilon_t] - \\ &\qquad\qquad -[\Delta\mathbf{X_{t-S}}\beta + \gamma\mathbf{G}\mathbf{\Delta X_{t-S}} + \delta_1\mathbf{\Delta D_{t-S}} + \Delta\epsilon_{t-S}]\}\end{aligned} \tag{19}$$

Rearranging and denoting ν=$\epsilon_t - \epsilon_{t-1} - \epsilon_{t-S} + \epsilon_{t-S-1}$ gives:

$$\begin{aligned}\mathbf{\Delta y_t} - \mathbf{\Delta y_{t-S}} &= (\mathbf{I} - \alpha\mathbf{G})^{-1} * \quad \{(\beta + \gamma\mathbf{G}) \ \ [\Delta\mathbf{X_t} - \Delta\mathbf{X_{t-S}}] + \\ &\qquad +\delta_1[\mathbf{\Delta D_t} - \mathbf{\Delta D_{t-S}}] + \nu\} = \\ &= (\mathbf{I} - \alpha\mathbf{G})^{-1} * \quad \{(\beta + \gamma\mathbf{G}) \ \ [\Delta\mathbf{X_t} - \Delta\mathbf{X_{t-S}}] + \\ &\qquad +\delta_1\mathbf{D_t} + \nu\}\end{aligned}$$

where the second equation holds by construction, because treatment occurs only in period t and hence $\mathbf{D_{t-1}} = \mathbf{D_{t-S}} = \mathbf{D_{t-S-1}} = \mathbf{0}$. We can thus simplify $\mathbf{\Delta D_t} - \mathbf{\Delta D_{t-S}} = \mathbf{D_t}$.
If we now take Conditional Expectations w.r.t. $\mathbf{D_t}$ we have:

$$\begin{aligned}\mathbf{E}[\mathbf{\Delta y_t} - \mathbf{\Delta y_{t-S}}|\mathbf{D_t}] = (\mathbf{I} - \alpha\mathbf{G})^{-1} * \quad &\{(\beta + \gamma\mathbf{G}) \ E[\Delta\mathbf{X_t} - \Delta\mathbf{X_{t-S}}|D_t] + \\ &+\delta_1\mathbf{D_t} + E[\nu|D_t]\}\end{aligned} \tag{20}$$

which reduces to:

$$\mathbf{E}[\mathbf{\Delta y_t} - \mathbf{\Delta y_{t-S}}|\mathbf{D_t}] = (\mathbf{I} - \alpha\mathbf{G})^{-1}\{\delta_1\mathbf{D_t}\} \tag{21}$$

if the following relatively weak identifying assumptions are satisfied:

[16] Note that also here $\Delta X_{it-S}\beta$ would contain $\frac{df(age_{t-S})}{dt}$, if βX_{it} is modelled to contain an additively separable age-dependent term.

- $E[\nu|D_t] = E[\Delta\epsilon_t - \Delta\epsilon_{t-S}|D_t] = 0$.[17]
- $E[\Delta\mathbf{X_t} - \Delta\mathbf{X_{t-S}}|D_t] = 0$, which means that the expected changes of the pages are the same between t-1 and t and between t-S-1 and t-S. This is satisfied if $\Delta X_t|D_t$ is stationary of order one.

Provided $(\mathbf{I} - \alpha\mathbf{G})^{-1}$ is well defined[18], and using the property that $(\mathbf{I} - \alpha\mathbf{G})^{-1} = \sum_{s=0}^{\infty} \alpha^s G^s$, the general impact of a local treatment will be:

$$\begin{aligned} \mathbf{E}[\boldsymbol{\Delta}\mathbf{y_t} - \boldsymbol{\Delta}\mathbf{y_{t-S}}|\mathbf{D_t}] &= \mathbf{I} * \delta_1\mathbf{D_t} + \alpha\mathbf{G} * \delta_{\mathbf{1}}\mathbf{D_t} + \alpha^{\mathbf{2}}\mathbf{G^2} * \delta_{\mathbf{1}}\mathbf{D_t} + \alpha^{\mathbf{3}}\mathbf{G^3} * \delta_{\mathbf{1}}\mathbf{D_t} + ... \\ &= \delta_1\mathbf{D_t}(\mathbf{I} + \alpha\mathbf{G} + \alpha^{\mathbf{2}}\mathbf{G^2} + \alpha^{\mathbf{3}}\mathbf{G^3} + ...) \end{aligned} \tag{22}$$

where $\mathbf{D_t}$ is a vector which is 1 at the treated nodes (if they are *currently* treated) and 0 otherwise. The proof for the control group consisting of remote nodes is analogous

[but needs to be written down. cf. Cameron Trivedi]

[enough of a proof ?]

B.2.4 General Pattern of first and higher order spillovers

[Notation here partly inconsistent with notation from above. Proof also not yet done for both groups]

Above we have shown for both control groups, what is measured by the Difference in Differences. From now on I shall refer to a node in the control condition by c and to a node in the treated condition by ℓ. Hence let us recollect that if $\mathbf{D_t}$ denotes the vector of treatments which is 1 at the treated nodes and 0 otherwise, estimation of the difference in differences returns

$$\begin{aligned} \text{(23)}\ \mathbf{E}[\boldsymbol{\Delta}\ell - \boldsymbol{\Delta}\mathbf{c}|\mathbf{D_t}] &= \mathbf{I} * \delta_{\mathbf{1}}\mathbf{D_t} + \alpha\mathbf{G} * \delta_{\mathbf{1}}\mathbf{D_t} + \alpha^{\mathbf{2}}\mathbf{G^2} * \delta_{\mathbf{1}}\mathbf{D_t} + \alpha^{\mathbf{3}}\mathbf{G^3} * \delta_{\mathbf{1}}\mathbf{D_t} + ... \\ &= \delta_1\mathbf{D_t}(\mathbf{I} + \alpha\mathbf{G} + \alpha^{\mathbf{2}}\mathbf{G^2} + \alpha^{\mathbf{3}}\mathbf{G^3} + ...) \end{aligned}$$

Clearly, the DID alone will not directly reveal α, the parameter of interest, but merely a quantity that is tightly linked to α and δ_1. Yet, the result also highlights that computing the parameters is not necessarily feasible, e.g. because it involves the knowledge of the complete link structure of the nodes. Luckily, a closer look at the nodes independently

[17] Particularly, any time trends or other dynamics, is to be eliminated by the Differences in Differences, if $\frac{df(age)}{dt}$ is the same evaluated at t-S and at t.

[18] This is the case if $\alpha < 1$

reveals that already limited information about the link structure can suffice to acquire additional information about the parameters.

B.2.5 Analysis on the Node Level

- Recall the effect of treatment on the treated network:

$$E[\Delta l - \Delta c | D_t] = \delta_1 D_t (I + \alpha G + \alpha^2 G^2 + \alpha^3 G^3 + ...)$$

- What matters for each focal node j is its own row in this set of equations.
- Note further, that under the local treatment assumption $D = e_i$
- Hence, for each node we need to evaluate its corresponding ji element in the matrix G and it's higher orders.

The higher orders of the adjacency matrix G will contain the same knowledge that comes from the sampling strategy and the knowledge about local treatment. Some nodes (L0) are known to be directly treated, and some (L1) have a direct link so that the entry in G that links them to the treated node is positive. However, for those, who only have an indirect link, the corresponding entry in G takes the value 0 and only the relevant element of G^2 will be greater than 0.

- For a shocked node $l0 \in L0$, a neighbor $l1 \in L1$ and the indirect neighbors (2 clicks away, 3 clicks away etc.) we have:

$$\begin{aligned} l0: \; E[\Delta y_{i,t} - \Delta y_{c,t} | ..] &= \delta_1 (1 + 0 + \alpha^2 G^2_{ii} + \alpha^3 G^3_{ii} + ...) \\ l1: \; E[\Delta y_{j,t} - \Delta y_{c,t} | ..] &= \delta_1 (0 + \alpha G_{ij} + \alpha^2 G^2_{ij} + \alpha^3 G^3_{ij} + ...) \\ l2: \; E[\Delta y_{k,t} - \Delta y_{c,t} | ..] &= \delta_1 (0 + 0 + \alpha^2 G^2_{ik} + \alpha^3 G^3_{ik} + ...) \\ & \qquad \qquad \qquad \qquad \qquad \qquad etc. \end{aligned}$$

- The basic idea of this paper is to back out the point estimates for α and δ_1 from the sequence of reduced form Diff in Diff estimates for increasingly large link-distances.
- Two parameters and as many equations as can reasonably be traced
- The precise estimates are based only on the higher orders of G...
- ... which is the (unsolved) computational challenge...

[maybe here only summarize the results of the next two subsubsections and shift them out of sight (e.g. to a separate Appendix)]

B.3 A first estimator that ignores higher order spillovers

Note, that if we now were to neglect all spillovers of order 2 and higher [19] we would attribute all changes in the l0 node directly to the shock and all changes in the L1-set, directly to the spillovers. Then we could simply write the difference in the shocked node as δ_1 and set any differences in the nodes that are two or more clicks away to 0, which is equivalent to assuming:

$$\mathbf{E}[\mathbf{\Delta l} - \mathbf{\Delta c}] =^{\mathbf{b.A.}} \mathbf{I} * [\mathbf{D}] + \alpha * \mathbf{G} * [\mathbf{D}] + \mathbf{0} + \mathbf{0} + ... \tag{24}$$

which is equivalent to having[20]:

$$\begin{aligned} E[y_{igt}|D_{it} = 0] - E[y_{igt}|\mathbf{D_{it} = 1}] &= \delta_1 && for\ treated\ L0 - nodes \\ E[y_{it}|D_{it} = 0, \mathbf{D_{l0} = 0}] - E[y_{it}|D_{it} = 0, \mathbf{D_{l0} = 1}] &= 0 && for\ L2\ and\ further \end{aligned}$$

...but also implies, that there are no “multiplication-effects” or “feedback-loops” between the nodes. In the light of the formalization presented here, this is obviously a heroic assumption, yet note that in the impact evaluation literature with fixed and stable classroom sizes or villages, this assumption is implicitely, but very commonly taken. (cf. DE GIORGI, DUFLO, etc. etc.). The Diff in Diff for the neighbors of the treated nodes[21] would simply reduce to:

$$\Delta l1 - \Delta c1 \quad = \quad \frac{\alpha}{NP_{l1}}\delta_1 \tag{25}$$

Given the necessary assumptions, it is obvious that a consistent estimator of δ_1 and the observed difference in difference will be enough to estimate α. Specifically, if we (for now) maintain the assumption that we can observe the nodes’ performance before any higher order spillovers arrive at the treated node, we can obtain such an estimate from applying the the Diff in Diff estimator on the level of directly treated nodes and a suitable comparison group and then move on to estimate α:

[19] Neglecting higher-order spillovers is like implicitely introducing a temporal structure where a spillover takes time to occur and taking a snapshot after the first order effect had just enough time to spill onto it’s neighbors, but not yet enough time for any second and higher order spillovers. This is possible if, for example, spillovers are slow and the temporal structure of the available data is fine grained enough. Formalizing these higher order spillovers is quite involved and depends on the specific structure of the network and the nature of the links. Hence, treating them explicitely is tackled in the next section.

[20] D_{l0} denotes the value of D at the central node, that is related to the focal node.

[21] Which corresponds to an Indirect Treatment Effect or an “Externality”

$$\hat{\delta}_1 = \Delta\hat{l}0 - \Delta\hat{c}0 \tag{26}$$

- $\Delta\hat{l}0 := \frac{1}{NP_l0} * \sum_i (y_{i,l0,t=0} - y_{i,l0,t=1})$
- $\Delta\hat{c}0 := \frac{1}{NP_c0} * \sum_i (y_{i,c0,t=0} - y_{i,c0,t=1})$

$$\hat{\alpha} = \frac{\Delta\hat{l}1 - \Delta\hat{c}1}{\Delta\hat{l}0 - \Delta\hat{c}0} * NP_{l1} \tag{27}$$

with the definition of $\Delta\hat{l}1$ and $\Delta\hat{c}1$ paralleling the definition of $\Delta\hat{l}0$ and $\Delta\hat{c}0$. This estimator has the advantage of being readily available with well known properties. However, as it is, it would only be suitable under the (potentially quite strong) assumption that higher order spillovers are negligible. Whether this is true or not will depend on the size of the spillover effect, but to a very large extent also on the network structure and the number of nodes. In what follows I shall proceed to illustrate how second order spillovers affect the estimator in equation 27 and I will show how to derive an upper bound to the size of the problem.

B.4 Considering higher order spillovers

As was just pointed out, it will often be the case, that nodes will have a feedback effect on each other, so that the neighbors change in performance (due to the original impulse) will affect the neighbors' neighbours, but also feed back on the treated neighbor. The differences between period 0 and 1 will then also include the second order spillovers. Obviously, the diff in diff estimators will then also observe the changes in outcome at the end of this process, when all higher order spills have taken place. The real structural relationship without explicit characterization of the higher order spills can be thought of as follows:

$$\Delta l0 - \Delta c0 = \delta_1 + HO_{l0} \tag{28}$$

$$\Delta l1 - \Delta c1 = \frac{\alpha}{NP_{l1}}\delta_1 + HO_{l1} \tag{29}$$

where HO_{l0} and HO_{l1} typically depend on the underlying network of peers and need to be characterized from scratch, taking into account the network structure.

B.5 Benchmarks without using the information on the network structure.

Luckily, even if the information on G is not available, it is possibly to derive benchmarks (upper and lower bound estimate) which can provide useful information on their own.

[to be completed]

B.6 Precise Estimator that exploits the information on the network structure.

[to be completed]

C idea in a nutshell

[such a section shall be useful, to get a clear overview...]

D Aside: Reaction to treatment of the neighbor

Everything that was derived above was derived under the assumption that the nodes do not observe or at least do not react to the local treatment of their neighbors. In general however, the subjects of treatment and their neighbors might observe each other and react to these observations.

D.1 Setup with “observing neighbors”

An Example of such a setting could be children in a class at school, who get annoyed or jealous when they observe that their peer was treated in a nice way and they were not. Also economic agents in a village, who observe that their neighbor was refused a social service for failure to comply with the requirement of sending their kids to school, might adapt their behavior in reaction to this observation. Another such situation could be commuters in a city, who observe when their friends got caught after the local transport authority increases the frequency of controls and the punishment for failure to present a valid ticket. In such situations the students/villagers might react to *merely observing* the treatment of their neighbors and they might select a different value for the outcome variable.

To model such a situation we need to further augment the model in equation 8 by both the observable treatments (shocks) that are locally applied, and a term that captures the possible reaction to the treatment of the neighbor.

$$y_{it} = \alpha \frac{\sum_{j \in P_{it}} y_{jt}}{N_{P_{it}}} + X_{it}\beta + \gamma \frac{\sum_{j \in P_{it}} X_{jt}}{N_{P_{it}}} + \delta_1 D_{it} + \delta_2 \frac{\sum_{j \in D_{jt}}}{N_{P_{it}}} + \epsilon_{it} \tag{30}$$

Interpretation of the 2 new coefficients:

- δ_1 ... measures the direct treatment effect if a node(page) is, ITSELF, treated.
- δ_2 ... in general: measures reactions of the node, when it "observes" treatment of one (or several) of its peers.

D.1.1 Local application of treatment

consider 2 connected nodes, where one is treated ($l0$) in period t and the neighbors are not treated ($l1 \in L1$). Assume for simplicity that l0 is the only treated node in l1's neighborhood.

$$l0 :: y_{l0t} = \alpha \frac{\sum_{j \in P_{l0t}} y_{jt}}{N_{P_{l0t}}} + X_{l0t}\beta + \gamma \frac{\sum_{j \in P_{l0t}} X_{jt}}{N_{P_{l0t}}} + \delta_1 1 + \delta_2 \frac{\sum_{j \in P_{l0t}} 0}{N_{P_{l0t}}} + \epsilon_{l0t} \tag{31}$$

(32)

$$l1 \in L1 :: y_{l1t} = \alpha \frac{y_{l0t} + \sum_{j \in P_{l1t}/l0} y_{jt}}{N_{P_{l1t}}} + X_{l1t}\beta + \gamma \frac{\sum_{j \in P_{l1t}} X_{jt}}{N_{P_{l1t}}} + \delta_1 0 + \delta_2 \frac{1 + \sum_{j \in P_{l1t}/l0} D_{jt}}{N_{P_{l1t}}} + \epsilon_{l1t}$$

D.1.2 Assumptions for identification (conjecture)

NOTE: Assuming we have such a local treatment available, we get two types of spillover effects:

- δ_2 ... "behavior change" of the node, when it "observes" treatment of its peer.
- the "pure" spillover α, that we observe, because treatment will affect the outcome of $l0$
- $\rightarrow \alpha$ is only identified if δ_2 is believed to be 0
- $\rightarrow$ otherwise only the total "treatment-of-peer"-effect can be measured. (but that can also be interesting)

D.1.3 Controls in remote part of the network around $c0$

take two remote nodes c0 and $c1 \in C1$, where nothing happens (nobody gets treated).

$$c0 :: y_{c0t} = \alpha \frac{\sum_{j \in P_{c0t}} y_{jt}}{N_{P_{c0t}}} + X_{c0t}\beta + \gamma \frac{\sum_{j \in P_{c0t}} X_{jt}}{N_{P_{c0t}}} + \delta_1 0 + \delta_2 \frac{\sum_{j \in P_{c0t}} 0}{N_{P_{c0t}}} + \epsilon_{c0t} \tag{33}$$

$$c1 \in C1 :: y_{c1t} = \alpha \frac{\sum_{j \in P_{c1t}} y_{jt}}{N_{P_{c1t}}} + X_{c1t}\beta + \gamma \frac{\sum_{j \in P_{c1t}} X_{jt}}{N_{P_{c1t}}} + \delta_1 0 + \delta_2 \frac{\sum_{j \in P_{c1t}} 0}{N_{P_{c1t}}} + \epsilon_{c1t} \tag{34}$$