Discussion Paper No. 13-098

Spillovers in Networks of User Generated Content

– Evidence from 23 Natural Experiments on Wikipedia

Michael E. Kummer

ZEW

Zentrum für Europäische
Wirtschaftsforschung GmbH

Centre for European
Economic Research

Discussion Paper No. 13-098

Spillovers in Networks of User Generated Content

– Evidence from 23 Natural Experiments on Wikipedia

Michael E. Kummer

Spillovers in Networks of User Generated Content *

- Evidence from 23 Natural Experiments on Wikipedia

Michael E. Kummer†

December 1, 2013

Abstract

Endogeneity in network formation hinders the identification of the role that social networks play in generating spillovers, peer effects and other externalities. This paper tackles this problem and investigates how the link network between articles on the German Wikipedia influences the attention and content generation individual articles receive. Identification exploits local exogenous shocks on a small number of nodes in the network. It can thus avoid the usually required, but strong, assumptions of exogenous observed characteristics and link structure in networks.

Exogenous variation is generated by natural and technical disasters or by articles being featured on the German Wikipedia's start page. The effects on neighboring pages are substantial; I observe an increase of almost 100 percent in terms of both views and content generation. The aggregate effect over all neighbors is also large: I find that a view on a treated article converts one for one into a view on a neighboring article. However, the resulting content generation is small in absolute terms.

My approach also applies if, due to a lack of network data, identification through partial overlaps in the network structure fails (e.g. in classrooms). It helps bridge the gap between the experimental and social network literatures on peer effects.

Keywords: Social Media, Information, Knowledge, Spillovers, Networks, Natural Experiment

JEL Classification Numbers: L17, D62, D85, D29

*The author is grateful to the Wikimedia Foundation and Frédéric Schütz for access to the Wikipedia data. Thanks to Thorsten Doherr and Manfred Knobloch for support with data processing and to George Giorgidze for help with DSH. I benefitted from discussions with Irene Bertschek, Christoph Breunig, Luis Cabral, Ulrike Cress, Neil Gandal, Avi Goldfarb, Shane Greenstein, Thorsten Grosjean, Maximilian Kasy, Francois Laisney, Jose Luis Moraga-Gonzalez, Kathy Nosal, Gal Oestreicher-Singer, Martin Peitz, Marianne Saam, Philipp Schmidt-Dengler, Olga Slivko, Yossi Spiegel, Joel Waldfogel, Michael Ward, Andrea Weber and the participants in several seminars at ZEW and U Mannheim, U Linz, U Kaiserslautern and WU Wien and of the 11^{th} ICT-Conference (ZEW, Mannheim), the EARIE 2013 (Évora), the 11^{th} conference of Media Economics (Tel Aviv) and the IIIrd ICTCM (Munich) for valuable input. James Binfield, Timo Schäffer and Daniel Bergman provided outstanding research assistance. Financial support from the WissenschaftsCampus Tübingen is gratefully acknowledged.

†Centre for European Economic Research (ZEW);L7 1;68181 Mannheim, Germany; Kummer@zew.de.

1 Introduction

This paper measures the spillovers of attention and contribution effort transmitted through links in German Wikipedia. Understanding how links channel human attention is important in a wide range of settings, such as peer production, advertisement or public decision making. To give an example, such knowledge can be useful in peer production to more effectively discover unnoticed mistakes, bugs or biases. Analogously citations and their effect on attention play an important role in scientific production and innovation. Yet, endogeneity in the formation of networks has been a constant obstacle for those who try to measure peer effects or the role of social networks in generating spillovers or other externalities. While correlations between nodes' network positions and their outcomes abound, exogenous sources of variation that allow us to pin down causes and distinguish them from effects are usually hard to observe. I address this issue by exploiting local exogenous shocks on a small number of nodes in Wikipedia's article network.

This paper contributes to the literature in two ways. Firstly, I measure spillovers of attention and how they convert to content production in the German Wikipedia by exploiting the fact that sudden spikes in attention affect not only the shocked nodes in a network, but are also transmitted to their neighbors (Carmi et al. (2012)). I exploit large-scale events like natural disasters and accidents or the advertisement of featured articles on Wikipedia's start page as exogeneous sources of sudden spikes in the attention an article receives. Previous research in the field of peer production has analyzed the correlation between 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 outcome variable might itself drive network position, thus giving rise to the classic endogeneity problem. Moreover, nodes are likely to be peers (Bramoullé et al. (2009)) and researchers interested in measuring interactions between them face the reflection problem laid out by Manski (1993). This paper overcomes these two problems by exploiting local exogenous treatments of single nodes in Wikipedia's article network.

Secondly, I modify the framework of Bramoullé et al. (2009) to allow the incorporation of local and randomized treatments and their spillover effects after shocks in a network setting. I show how it is possible to use exogenous treatments of individual nodes in networks, which are the focus of both Carmi et al. (2012) and the present paper, as a new and complementary source of identification beside partial overlaps in the network structure (Bramoullé et al. (2009), De Giorgi et al. (2010)). The suggested formalization is quite general and nests not only exogenous treatments of single nodes in networks, but also partial population treatments (Moffitt (2001), Dahl et al. (2012)). The identification strategy uses a combination of exogenous shocks and an estimator based on comparing differences. After obtaining reduced form estimates based on minimal assumptions, I exploit knowledge of the network structure to back out the structural parameters describing the spillover effects. Finally, some of my insights carry over to impact evaluation studies based on a two-stage randomization over sub-populations (villages) and then individuals inside sub-populations (Angelucci and De Giorgi (2009), Kuhn et al. (2011), Crépon et al. (2013)). It becomes clear why the reduced form coefficients are boundary estimates if detailed link information is not available. I show under which conditions it is possible to derive both the upper and the lower bound estimates of the parameter of interest. This second (typically lower) bound is a second contribution to the literature. My contribution to the literature is discussed in Section 2. Detailed information about the formal framework is presented in Section 3.

In the application, I exploit two sources of variation that trigger substantial changes in the attention that certain pages receive at a known (ex-post) point in time: (i) exogenous and unpredictable large-scale events such as earthquakes or plane crashes and (ii) articles that are chosen to be advertised on Wikipedia's homepage and are thus highly visible for 24 hours. To obtain my dataset I augment the publicly available data dumps provided by the Wikimedia Foundation[1] 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. I use 23 large-scale events, 34 articles that were featured on Wikipedia's main page, and all their respective network neighbors. The resulting dataset contains information on views and content generation of almost 13,000 articles, 14 days before and after the events (more than 750,000 observations). Details about the data are provided in Section 4.

I document substantial effects for neighboring pages of featured articles for both attention (= views) and content generation (= editing activity). Articles in the neighborhood of the shocked area were viewed 35 more times on average. For large events, I even observe spillovers of attention and content generation to pages that are two clicks away from the (originally non-existent) disaster page. The effect corresponds to an increase of almost 100 percent in terms of views and also editing activity also almost doubled. On the aggregate, the effects are very large. Over all neighbors I find that each view on the treated page translates to a view on one of the neighbors. However, given the small baseline editing activity, the content generation triggered by having a featured article (or even a natural disaster) in the neighborhood is small in absolute terms. It takes one thousand views before an additional revision occurs. In short, links matter for the attention that a node in a citation network receives, but much less for the content that is generated on such nodes. This may be justified given the maturity that the German Wikipedia had reached by 2007. More results can be found in Section 5.

My extension of the framework of Bramoullé et al. (2009) is the first formalization

[1] 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.

that allows exogenous treatment of single nodes in networks that I am aware of. The model helps us to understand how frequently measured reduced form estimates of an ITE can be related to the "structural" social or spillover effects. It also highlights how to back out the parameter of interest if detailed link information is available or estimate the upper and the lower bounds of the parameter of interest, if it is not. I apply these techniques to my data and obtain an interval estimator for the spillover effect of interest, finding that an average increase of ten views on the neighboring pages results in an increase of 2.22 to 2.92 views on the page in the center. These results suggest that placing links has an effect, but that it is small. For more details please refer to the results in Section 5.

The hyperlink network between articles can be interpreted as a citation network. Thus, my findings allow for a more abstract reading when interpreting Wikipedia as a peer production tool for the documentation of human knowledge. Consequently their relevance extends to other settings of peer production including open source software or scientific research. While it is true that my strategy requires a lot from the data[2], recent advances in data handling techniques and the increasing availability of data on social interactions will provide further applications where this strategy can be used.

There are two sets of results in this paper. Firstly, Section 3 discusses the structural peer effects model and identification through local treatments in networks. Detailed derivations of the estimator and the bounds are in Appendix B. Secondly, empirical results and how to relate my reduced form estimates to the structural model are described in Section 5. The relevant literature and this paper's contributions are discussed in the next section (2). Section 4 discusses the data collection and the relevant variables. Concluding remarks, limitations and avenues for further research are offered in Section 6. The Appendix contains summary statistics, robustness checks, additional figures and a discussion of why network neighbors should not react to their neighbor's treatment.

2 Literature

This paper builds on two important strands of the literature: firstly that on social effects, such as peer effects or spillovers (cf. Manski (1993)) in networks and secondly that which uses pseudo-treatments to causally identify economic effects (typically not in a network context).

Social effects[3] or spillovers in a network are generally difficult to identify because they

[2]Exogenous treatments of individuals in in networks (or groups) could rarely be observed in previous studies. Researchers often have the network structure and no exogenous source of identification, or exogenous variation yet no information on the network structure. However, such data are increasingly available from field experiments or online sources.

[3]In what follows I use the term "social effects" in the broad sense that Manski (1993) uses it in order to subsume "social norms, peer influences, neighborhood effects, conformity, imitation, contagion, epidemics, bandwagons, herd behavior, social interaction or independent preferences." (cf. ibid.)

are frequently confounded with other individual-specific characteristics or network dynamics. One prominent subbranch of this literature investigates the relationship between a node's network position and its performance.

A series of papers is dedicated to the identification of peer effects. These are difficult to identify when both the peers' average characteristics and their average performance influence individual outcomes (Manski 1993). One of the most well-known approaches to disentangling these effects is to exogenously vary the composition of peer groups (Sacerdote (2001), Imberman et al. (2009)). Imberman et al. (2009), for example, exploit variation in the peer groups of Houston and Louisiana's incumbent school children due to evacuee inflow in the aftermath of the hurricanes Katrina and Rita. Their identification strategy is based on the large variation in peer groups and the random allocation of evacuees after the event.

Other approaches use network structure or, more precisely, the existence of partial overlaps (open triads) wherein a peer is connected to two other peers, who themselves are not connected to each other (De Giorgi et al. (2010) and Bramoullé et al. (2009)). 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 its influence on the other. I extend that formal framework to allow the inclusion of exogenous treatments and combine it with a simple difference-in-differences estimator (rather than using an IV approach in a cross-section). I thus use a completely different source of identification to measure peer effects. Using exogenous sources of variation for identification has the advantage of not requiring the network structure to be independent of observables. Furthermore the strategy also works in the absence of open triads in the network structure,[4] which renders the notation compatible with a two-layered setting with randomization across subpopulations (villages, classrooms, etc.) and subsequent treatment of randomly selected individuals within subpopulations. Ballester et al. (2006), finally, propose a strategy for identifying the "key player" of a group or a network, which is a very different purpose from the one in the present study. Nevertheless, since my analysis of the network shares features of the one in their paper, I was able to draw upon some of their insights and their formalization of the network structure.

Another relevant branch of this literature has focused on the effect of knowledge spillovers on production in social networks. Fershtman and Gandal (2011) investigate indirect and direct knowledge spillovers in the production of open source software. Claussen et al. (2012) using panel data to control for unobserved heterogeneity, consider the electronic gaming industry. Oestreicher-Singer and Sundararajan (2012) contribute by exploiting Amazon's link network and showing that items connected through a visible link influence each other three times more than when the link is invisible. Kummer

[4]Formally open triads are equivalent with linear independence of the adjacency matrix of the graph that represents the network (typically denoted by $\mathbf{G}$) and its own square.

et al. (2012), following Halatchliyski et al. (2010) who analyze authors' contributions in two related knowledge domains, considers the network of linked articles on the German Wikipedia. The strategy of the above mentioned papers exploits variations in the link network, between or within nodes, and relates them to the outcome of interest. If a positive relationship is found, it is taken as prima facie evidence for the existence of spillover effects of knowledge or economic success. However, this strategy could be criticized if the network position is not exogenous; it is often difficult to identify exogenous sources of variation in a network. The strategy pursued in this paper, as in Carmi et al. (2012), approaches the problem from another angle. I do not attempt to measure spillover effects by looking at variation in the link structure. Instead I consider how shocks are transmitted within a given link structure. The estimates are based on the observation that articles that are linked to a shocked article receive a spillover, while similar articles that are not linked to a shocked article do not. I will show below the conditions under which such an approach can be successful (Section 3). However, this approach will fail if the conditions are not satisfied (cf. Appendix C).

As for the second stream of literature, on treatment effects, it is well established that social effects or spillovers play an important role in the causal identification of treatment effects. More precisely, it can be difficult to identify the causal effect of a treatment in the presence of social effects or spillovers, because they might lead to a violation of the Stable Unit Treatment Value Assumption (SUTVA) and hence compromise the validity of the control group (Ferracci et al. (2012)). Since such externalities threaten the identification of treatment effects, researchers have come to understand the importance of adding a second layer of randomization at the level of classrooms, villages, districts etc. (Miguel and Kremer (2003), Angelucci and De Giorgi (2009), Kuhn et al. (2011) and many more). Since such randomization immediately lends itself to computing indirect treatment effects, many of the aforementioned papers put special emphasis on them. Crépon et al. (2013) are concerned about the possibility that labor market programs have a negative impact on the non-eligible. They test their hypothesis by both randomizing over treated populations and varying the treatment intensity. There is a close relationship between the idea of randomizing across subunits and the Partial Population Experiment that Moffitt (2001) suggested as a solution to the Reflection Problem Manski (1993). Dahl et al. (2012) provide an example of such an experiment, exploiting the introduction of a new policy that made it easier for some fathers to leave their jobs and spend time at home with their babies and measuring how the increased take-up of treated fathers impacted the probability that their (old-regime) peers' also decided to stay at home.

This study uses the treatment of *peers in a network* to identify spillovers and asks under which circumstances it is, in general, possible to causally identify spillovers or peer effects when treatment of peers can be observed. This paper might be somewhat unusual for readers who are familiar with this literature, because it is not focused on the effect of treatment itself but instead uses treatments to identify *the spillover effect*. Hence, it exploits the violation of the SUTVA to measure the indirect treatment effects and identify the spillover effect. Moreover, in order to identify spillovers I use the fact that exogenous treatment sometimes affects only a single node and use the local network formed by nearby neighbors as an analogue to "villages" over which nature has randomized. Such local treatments are analogous to the Partial Population Treatment coined by Moffitt (2001) and might be dubbed "Mini Population Treatments". This idea is not new; it has been used increasingly often in recent studies. In a widely-quoted paper Aral and Walker (2011) develop an approach based on hazard modeling and use randomized treatments of individuals in networks to measure contagion. Most closely related to my approach is a study in the realm of e-commerce. Using the same method as this paper, Carmi et al. (2012) analyze the effect of the external shocks of recommendations by Oprah Winfrey on the product network of books on Amazon. They find that a recommendation not only triggers a spike in sales of the recommended book but also of books adjacent in Amazon's recommendation network. They measure demand in terms of the products' sales ranks and, like this paper, use a difference-in-differences strategy. They obtain a control group by exploiting the fact that Amazon's algorithm chooses different books to highlight in the recommendation network at different points in time. Their data structure allows analysis of which characteristics of the linked items can predict a higher spillover, which is beyond the scope of this paper. Although applying a similar estimation method, my contribution adds in two ways. Firstly, I provide a formalization of why this source of exogenous variation ensures identification of the social/spillover parameter and of how to relate the reduced form estimators to the structural parameter of interest. Secondly, the German Wikipedia is a quite different type of network, as it is a citation network of articles that are created in a peer production process. This network is formed by a very large number of links (edges) that are not set to an a priori fixed number of items. Moreover, the links are placed by humans rather than an algorithm. Another difference is that my data structure allows the direct observation of visits (attention) to the linked sites. Rather than sales I analyze the decision to contribute information to the article. The insights I obtain are hence complementary and my results suggest that more studies of the many different types of networks are needed before all phenomena can be understood.

In a somewhat more qualitative field experiment "offline", Berge (2011) compares peers of individually treated and non-treated agents to measure knowledge spillovers from a business training program in Tanzania. Using in-depth interviews he finds that "indirectly-treated" male clients become more "business minded", discussing business more, increasing their loans and becoming more risk averse. Once again we see a close relationship between their work and randomized treatment of villages or small groups. Hence it is not surprising that a paper by Banerjee et al. (2012) exploits very detailed information on village networks to analyze the spread of information about microfinance through the villages.

This paper contributes to the literature by providing a simple formalization that includes exogenous local treatments into an existing framework to analyze peer effects in networks (Bramoullé et al. (2009)). It provides a formalization of why this approach guarantees identification and shows how to relate the reduced form estimators to the structural parameter of interest. As in the approach of Bramoullé et al. (2009), local treatments as source of exogenous variation ensure identification of the social parameter even in the famously underidentified model by Manski (1993). Yet the extension in this paper guarantees identification based on a complementary source of exogeneity. An additional advantage is that the formalization applies to the randomized treatment of subpopulations (Partial Population Treatment) in general. Moreover, all of my insights can be applied directly if only one member of the subpopulation (one pupil in a class, one villager in a community) is treated (Mini Population Treatment). I also derive an upper and a lower bound coefficient of interest that can be obtained even without any information on the network structure itself.

On the one hand, this paper shows how treatments diffuse across networks if the agents are linked and how average and indirect treatment effects can be related to a structural parameter that quantifies spillovers. On the other hand it formalizes a relatively new way to identify peer effects and spillovers in networks, based on truly exogenous variation. It thus helps bridge the gap between the experimental and social network literatures.

Finally, a source of exogenous variation in this paper is natural disasters, accidents and large scale events. Several papers other than Imberman et al. (2009) have been dedicated to natural disasters or other sources of exogenous variations on Wikipedia. Ashenfelter and Greenstone (2004) exploit the effect of a change in mandated speed limits on the number of fatal car accidents to estimate the value of a statistical life. A well-known paper by Zhang and Zhu (2011) uses the blocking of Wikipedia by the government 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 contrast the evolution of articles on breaking news events with the genesis of non-breaking news (and "historical" articles) and find that breaking news articles emerge into well-connected collaborations more rapidly than non-breaking news articles. For this reason they hypothesize that breaking news articles may become an important source of new content contributors.

This paper's application illustrates a new way of looking at content networks such as the one formed by Wikipedia articles. It provides new insights into the dynamics of user activity in the world's largest knowledge repository, measures how users allocate their attention and shows how attention is converted into contributions.

3 The Empirical Model

In what follows I will discuss the empirical model. I first give a basic and informal intuition of my estimation approach (subsection 3.1). Next I discuss the assumptions made to identify the effect of the exogenous treatments I use (subsection 3.2) and the reduced form estimation used in the regressions (subsection 3.3). The last and most extensive subsection (3.4) describes the structural model. There I discuss how and under which assumptions the researcher can identify the parameter that measures spillovers from the reduced form estimates if she observes the network information. In the same section I also show how to compute an upper and a lower bound for the coefficient when the network information is not available. An important case where my arguments do not apply are situations where the neighbors of the treated nodes/individuals observe the treatment and adjust their outcome as a reaction.[5] Appendix C shows, how the structural model would have to be extended to include such a possibility and which cahllenges to identification of the spillover parameter would emerge as a result.

3.1 Basic Intuition - Throwing Stones into a Pond

This subsection provides an intuitive explanation of the data structure and the estimation approach. The basic idea of the research approach can be imagined as "throwing stones into a pond and tracing out the ripples". The design of this paper uses the fact that certain nodes were affected by a large increase of attention, that this was exogeneous, and that ex-post it is known to the researcher when exactly the pseudo-experiment occurred. Moreover, since the link structure is also known, it is possible to observe what happens to the directly or indirectly neighboring nodes. As in a pond, we would expect the largest effect on the directly hit node and a decreasing amount of attention the further away an article is from the epicenter.

The schematic representation in Figure 1 shows how the data is structured. Wikipedia articles are the nodes of the network. They are represented by a circle with a letter inside. Each circle represents a different article in the German Wikipedia. Articles are connected to each other via links, which are visible on Wikipedia as highlighted blue text. Clicking on such text forwards the reader to the next article and these links form the edges of my directed network. In Figure 1 and 2, they are represented by a line between two nodes. Moreover, an important aspect of my identification strategy requires the observation of two disconnected subnetworks at the same time. This is represented by showing them as network L and network C facing each other in both figures. I will maintain this notation also in all derivations that follow. I focus on subnetworks around a start node. These start nodes are denoted by subscript 0. Hence, the start node of the two networks are

[5] E.g. classmates, that react with protest to an unfair punishment of their peer.

Figure 1: Schematic representation of a start node and its direct and indirect neighbors in two subsections of the network.

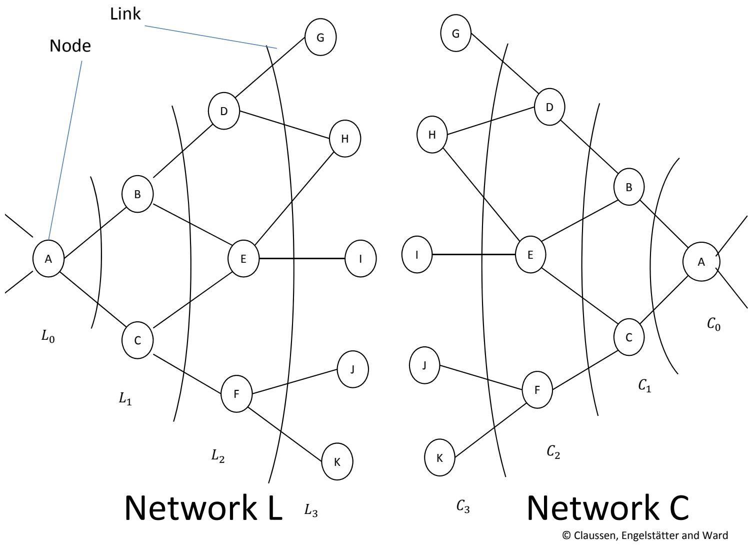

NOTES: The Figure illustrates the structure of the data. Wikipedia articles are the nodes of the network. Each circle with a letter inside represents a different article in the German Wikipedia. Articles are connected to each other via links, which are represented as lines. In general the network may be directed or undirected (Wikipedia articles are directed). The left side of the figure draws on a representation in a working paper on network formation by Claussen, Engelstaetter and Ward.

denoted by $\ell 0$ and $c0$. The nodes that receive a direct link from a start node (direct neighbors) in network L form the set of direct neighbors L_1 and a focal node from that set is sometimes denoted as $\ell 1$[6]. The set of indirect neighbors[7] in the network L forms L_2 and so on. Analogously the set C_1 is made up of direct neighbors of the start node in network C and C_2 are the indirect neighbors of node $c0$

In a typical network in which the outcome of the individual nodes depends on the outcome of their neighbors we would observe many correlations and cross influences, but it would be difficult to discern where they originate from or whether they are due to underlying and unobserved background factors which merely affect the nodes in similar ways. The schematic representation of Figure 2 illustrates the mechanism of local exogenous shocks ("the stone in a pond"). The shocked node is colored in dark blue, the direct neighbors are colored in light blue and so on. As I will show formally in the next sections, identification of the spillover hinges on the ability to observe a valid second network from which it is possible to infer what the outcomes would have been if no treatment had taken place. If this is possible, we can use these outcomes for comparing the size of the outcomes layer by layer. More information about how the layers are identified and

[6]While the set L_0 consists only of one node ($L_0 = \{\ell 0\}$), set L_1 consists of multiple nodes.

[7]Indirect neighbors are defined as receiving at least one link from a node in set L_1 without themselves being in L_1. Hence the shortest path from the start node to an indirect neighbor is via two clicks.

Figure 2: Schematic representation of a local treatment, which affects only one of the two subnetworks and there only a single node directly.

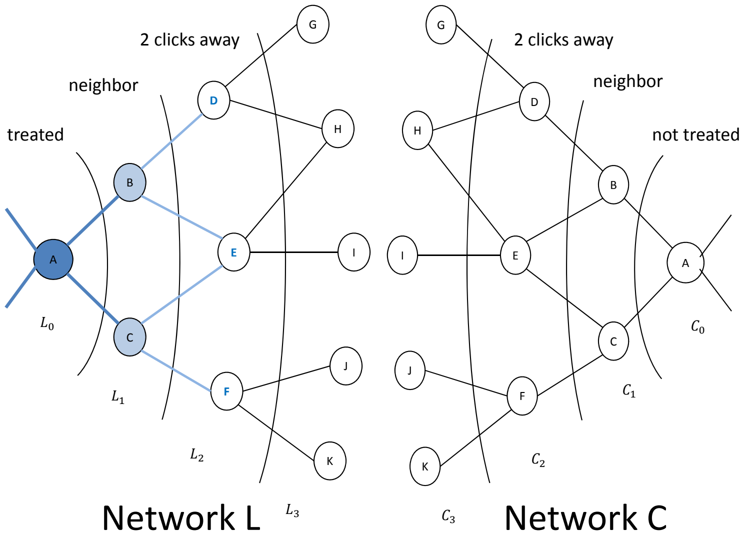

NOTES: The design of this paper uses the fact that certain nodes were affected by a large and exogenous increase of attention, and that it is known to the researcher when the pseudo-experiment occurred. The Figure illustrates the effect of a large local shock on Wikipedia, which affects only subnetwork L. The shocked node is colored in dark blue, the direct neighbors are colored in light blue and so on. If we observe a valid second network from which it is possible to infer what the outcomes would have been if no treatment had taken place, we can use these outcomes for comparing the size of the outcomes layer by layer.

obtained is provided in Section 4.

3.2 Identifying Assumptions for the Treatment Effects

Uncovering the spillover parameter, will be based on estimating Difference in Differences for each layer seperately. To clarify the assumptions in the reduced form estimation by layer, I use the control-treatment notation from impact evaluations (cf. Angrist and Pischke (2008)). Doing so serves two purposes: first, it highlights the parallels between the methodology based on ideas similar to the Partial Population Treatment (cf. Moffitt (2001)). Second, it provides a better intuition of the crucial assumptions made. Terminology and notation are inspired by Kuhn et al. (2011). Readers who are familiar with Average and Indirect Treatment Effects and the assumptions of a Difference-in-Differences strategy might prefer to merely browse the formulas or skip this section.

Even though it is a core concept in impact evaluation I briefly discuss the Average Treatment Effect (cf. Angrist and Pischke (2008) it in the context of my application (first part of this subsection). I revisit the concept of Indirect Treatment Effects as analyzed in Kuhn et al. (2011) in light of my content network with equidistant layers of articles around a shock in the second part of this subsection. The identifying assumption of the reduced form analysis will be as follws: Absent treatment, the control observations have a similar rate of change across time as the treated subnetworks, i.e. they grow at similar rates and are affected similarly by any dynamics that affect the entire Wikipedia (weekday dynamics etc.).

3.2.1 Average Effect of Direct Treatment (ATE)

Consider node in network $i \in \{\ell, c\}$ in period t. The ATE measures the effect of treatment on the treated. We would like to compare the observed outcome after treatment to the (unobservable) outcome of the same individual if we had not treated them.

$$\mathbf{E}[y^1_{\ell 0,t}|d_{\ell 0,t} = 1] - \mathbf{E}[y^0_{\ell 0,t}|d_{\ell 0,t} = 1] \tag{1}$$

ℓ denotes the subnetwork which is treated in period t and the subscript c the subnetwork that is not. $d_{i,t}$ indicates if node i itself was directly treated or not. Superscript 1 denotes the outcome of a treated observation and superscript 0 the outcome of the untreated counterpart. One of these cannot be observed and hence is counterfactual. $\mathbf{E}[y^0_{\ell 0,t}|d_{\ell 0,t} = 1]$ denotes the (counterfactual) outcome that we would observe for $\ell 0$ in period t, had it not been treated.

The challenge lies in the fact that the second term in this difference cannot be observed. We estimate the counterfactual observation of the treated using a comparable node/individual[8] in a period where it is not treated. I take two approaches to obtain such an observation: (i) a simple approach compares the observation "before and after" the treatment, and attributes all observed changes in outcomes to the treatment. This is equivalent to making the assumption that, absent treatment, the node/individual would have the same outcome as in the previous period.[9]

Assumption ATE-before-after:

$$\mathbf{E}[y^0_{\ell 0,t}|d_{\ell 0,t} = 1] = \mathbf{E}[y^0_{\ell 0,t-1}|d_{\ell 0,t-1} = 0] \tag{2}$$

Using this as counterfactual observation has the advantage of being as close to the treated observation as possible. However, it will fail to capture any period-specific effects that would have affected all nodes even without any treatment. Any such effects (weekday fluctuations, shocks etc.) will simply be attributed to the treatment. (ii) Alternatively, "difference-in-differences" uses a distinct comparison group. This could be individuals in the same populations, which were not eligible for treatment. The (unobservable) counterfactual outcomes of the treated nodes are assumed to be the treated nodes'/individuals' pre-treatment outcome *plus* the *change* of the non-treated control observation.

[8] A node, which is believed to be affected by treatment in similar ways.

[9] If the object/individual was observed more than once before treatment it might be possible to further improve this approach by accounting for trends in the outcomes etc.

Assumption ATE-DiD:

$$(3) \qquad \mathbf{E}[y^0_{\ell 0,t}|d_{\ell 0,t}=1] = \mathbf{E}[y^0_{\ell 0,t-1}|d_{\ell 0,t-1}=0] + + \{\mathbf{E}[y^0_{c0,t}|d_{c0,t}=0] - \mathbf{E}[y^0_{c0,t-1}|d_{c0,t-1}=0]\}$$

It is worth stressing this fact in the context of Wikipedia articles. The crucial assumption is not that the articles *are* very similar, but merely that they *evolve* in a similar way, i.e. on average they have similar growth in readership and edits and they are subject to similar intertemporal fluctuations.

3.2.2 Indirect Treatment Effect (ITE)

While the ATE can be defined on a single group of eligible nodes, the "indirect treatment effect" or ITE requires the introduction of at least one additional group of nodes/individuals that are not eligible for treatment.[10] The ITE measures the (spillover or externality) effect of treatment of eligible objects/individuals on the outcomes of non-eligibles. As for the ATE, we cannot observe the outcome of the non-eligibles had the eligibles not been treated. Well known papers that estimate ITEs are Angelucci and De Giorgi (2009), Kuhn et al. (2011) or Crépon et al. (2013), to name a few.

Since the distance to the epicenter of the treatment is known in my application, we can measure several ITEs and compare the nodes of the subnetworks by layer. ITE_1 refers to direct neighbors of the eligible nodes in a treated subnetwork. Analogously, ITE_2 refers to nodes that are two steps away, ITE_3 to three steps away and so on. Miguel and Kremer (2003) is a well-known example where distance layers were included in the estimation to incorporate a similar notion of distance to treatment in a real world setup.

The richness of my dataset will require an even more involved notation, since I have to differentiate along four dimensions (treatment, time, distance and subnetwork). To capture the notion of layers in the estimation, I use $D_{xr,t}$ as shorthand that takes the value 1 if both of the two following conditions are simultaneously satisfied: (i) the subnetwork x was treated and (ii) there exists a treated node with a shortest distance of r steps.[11] For direct neighbors we have:

$$(4) \qquad ITE_1 = \mathbf{E}[y^{\mathbf{1}}_{\ell 1,t}|D^{\mathbf{1}}_{\ell 1,t}, d^{\mathbf{0}}_{i,t}] - \mathbf{E}[y^{\mathbf{0}}_{\ell 1,t}|D^{\mathbf{1}}_{\ell 1,t}, d^{\mathbf{0}}_{i,t}]$$

As before, $d^mathbf1_{i,t}$ indicates if node i was directly treated or not in period t.[12]

[10] A good example would be an intervention to foster the reintegration into the jobmarket after paternity/maternity leave, for which people without children would not be eligible.

[11] In an estimation by layers the minimum distance to the treated node is exactly r steps.

[12] To save space treatment status is indicated by superscripts, $d^mathbf0_{i,t}$ otherwise. Notation has to be more involved here, because it is no longer possible to talk of a single node, as the treated nodes can have many different neighbors.

$y^{\mathbf{1}}_{\ell r,t}$ is now the outcome if *some neighbor* in D_{xr} was treated in t, and $y^{\mathbf{0}}_{\ell r,t}$ denotes the outcome if nobody in that set was treated.

The object of interest is the ITE_r, Generally, for any range r:

$$ITE_r = \mathbf{E}[y^{\mathbf{1}}_{\ell r,t}|D^{\mathbf{1}}_{\ell r,t}, d^{\mathbf{0}}_{i,t}] - \mathbf{E}[y^{\mathbf{0}}_{\ell r,t}|D^{\mathbf{1}}_{\ell r,t}, d^{\mathbf{0}}_{i,t}] \tag{5}$$

As for the ATE, the ITE has to be estimated, since the counterfactual outcome in the absence of treatment cannot be observed. We again estimate the counterfactual outcome by using two methods: (i) a comparison of the same individual before and after treatment and (ii) a Difference in Differences between neighbors in the compared subnetworks.

Assumption ITE_r-before-after:

$$\mathbf{E}[y^{\mathbf{0}}_{\ell r,t}|D^{\mathbf{1}}_{\ell r,t}, d^{\mathbf{0}}_{i,t}] = \mathbf{E}[y^{\mathbf{0}}_{\ell r,t-1}|D^{\mathbf{0}}_{\ell r,t-1}, d^{\mathbf{0}}_{i,t-1}] \tag{6}$$

Estimating an ITE_r from a before-after estimation has the same advantages and drawbacks as the ATE. Analogously, the drawbacks can be accounted for by computing a difference-in-differences estimator. In the context of an ITE, we need to observe comparable, but untreated, subpopulations (e.g. villages, classrooms, or here, subnetworks) in which we have information about which individuals/nodes are eligible for treatment and which are non-eligible. Ideally we would like to observe a random selection of the subpopulations in which any treatments are to be administered, and in the second step we administer treatment to the eligible nodes. Moreover, we observe both subpopulations before the treatment of one takes place.

Assumption ITE_r-DiD:

$$\begin{aligned}\mathbf{E}[y^{\mathbf{0}}_{\ell r,t}|D^{\mathbf{1}}_{\ell r,t}, d^{\mathbf{0}}_{i,t}] &= \mathbf{E}[y^{\mathbf{0}}_{\ell r,t-1}|D^{\mathbf{0}}_{\ell r,t-1}, d^{\mathbf{0}}_{i,t-1}] + \\ &+ \{\mathbf{E}[y^{\mathbf{0}}_{cr,t}|D^{\mathbf{0}}_{cr,t}, d^{\mathbf{0}}_{i,t}] - \mathbf{E}[y^{\mathbf{0}}_{cr,t-1}|D^{\mathbf{0}}_{cr,t-1}, d^{\mathbf{0}}_{i,t-1}]\}\end{aligned} \tag{7}$$

In words, this means that the counterfactual outcome of the neighbors of the treated can be obtained by computing their pre-treatment outcome *plus* the average *change* of the neighbors of the eligible node in a non-treated control group. As before, the crucial assumption is not that the nodes in the control group *are* very similar to those of the treated subnetwork, but rather that they *grow* similarly and that the way they are affected by Wikipedia-wide fluctuations is the same as long as no treatment occurs. Before moving on to the econometric specification, I conclude this section by summarizing the identification result in terms of the Difference-in-Differences estimator:

Conclusion ITE_r DiD: If Assumption ITE_r-DiD holds, the difference below identifies the ITE_r.

$$
\begin{aligned}
(8) \qquad ITE_r &= \mathbf{E}[y^{\mathbf{1}}_{\ell r,t}|D^{\mathbf{1}}_{\ell r,t}, d^{\mathbf{0}}_{i,t}] - \{\mathbf{E}[y^{\mathbf{0}}_{\ell r,t-1}|D^{\mathbf{0}}_{\ell r,t-1}, d^{\mathbf{0}}_{i,t-1}] + \\
&+ (\mathbf{E}[y^{\mathbf{0}}_{cr,t}|D^{\mathbf{0}}_{cr,t}, d^{\mathbf{0}}_{i,t}] - \mathbf{E}[y^{\mathbf{0}}_{cr,t-1}|D^{\mathbf{0}}_{cr,t-1}, d^{\mathbf{0}}_{i,t-1}])\}
\end{aligned}
$$

Hence, our estimator of the ITE_1 is based on the pre-treatment outcomes and comparing the *change* in the outcomes of direct neighbors of the eligible nodes in a *treated* subnetwork to the direct neighbors of the eligible nodes in the *non-treated* subnetwork. Note that this conclusion also applies to the ATE, when setting r to 0.

3.3 Reduced Form Analysis

To obtain the ITEs for each layer, I apply reduced form regressions which allow the understanding of 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. (2012). The idea is to compare pages grouped by their distance to the page which experiences treatment to their analogue in the control group (L_0 to C_0, L_1 to C_1,...). I denote all reduced form coefficients by ϕ. Furthermore, I define "treatment" for each set of pages along the lines of the indirect treatment effects (ITE_r) in the previous section.[13] I let s indicate the day relative to day 0, the day when the treatment is administered. Hence s runs from -14 to 14. λ_s is an indicator, which takes the value 1 if $t = s$ and 0 otherwise. Each set of pages that corresponds to one layer in the network is regressed seperately. So if I focus on the treated nodes, the neighbors and the indirect neighbors, it results in the following system of fixed effect regression equations, which all are based only on dummy variables:

L_0.) Diff-in-Diffs specification at level L_0:

$$
(9) \qquad y_{it} = \phi_i^{L_0} + \sum_{s \in S} \phi_{1,s}^{L_0} \lambda_s + \sum_{s \in S} \phi_{2,s}^{L_0} (\lambda_s * treat_{L_0,i}) + \xi_{it}
$$

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

L_1.) At level L_1 ($treat_{L1}$ featured (in theory) 1 click away):

$$
(10) \qquad y_{it} = \phi_i^{L_1} + \sum_{s \in S} \phi_{1,s}^{L_1} \lambda_s + \sum_{s \in S} \phi_{2,s}^{L_1} (\lambda_s * treat_{L_1,i}) + \xi_{it}
$$

[13] The dummy in the regression for the neighbors (sets L_1 and C_1) takes the value 1, not if the node was itself treated, but if the corresponding start node ($\ell 0$) was treated in t (and 0 otherwise).

L_2.) At level L_2 ($treat_{L2}$ featured (in theory) 2 clicks away):

$$y_{it} = \phi_i^{L_2} + \sum_{s \in S} \phi_{1,s}^{L_2} \lambda_s + \sum_{s \in S} \phi_{2,s}^{L_2} (\lambda_s * treat_{L_2,i}) + \xi_{it} \tag{11}$$

In words, I run the same difference-in-differences on three levels (on L_0, L_1 and L_2 (shown only for large events)). $treat_{L0,i}$ 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. 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 a value of 0, while for an observation in the treated group this variable will take a value of 1 if it 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. The $ITEs$ from the previous subsection are captured by the ϕ_2 coefficient that corresponds to day 0 in the regressions above. I look at $\phi_{2,0}^{L_1}$ for the ITE_1, which corresponds to L_1 and analogously at $\phi_{2,0}^{L_0}$ for L_0 and $\phi_{2,0}^{L_2}$ for L_2.

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. 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 useful, since the results from the reduced form analysis are based on minimal assumptions and provide guidance as to whether attention spillovers exist at all. They also allow us to see, how far they carry over, and whether they result in increased production. Finally, they allow me to provide a lower bound and an upper bound estimate of the aggregate spillover effects to be expected.

3.4 Structural Form Analysis and Bounds

Beyond measuring the size of the ITEs, I am interested in quantifying the size of the spillovers of attention that exist between Wikipedia articles on normal days. In this section, I augment the well known linear-in-means model for peer effects, as formulated in Manski (1993), with exogenous shocks. Departing from the version that was used by Bramoullé et al. (2009)[14], I show how exogenous shocks can be exploited to identify spillovers (or peer effects). This is possible in my modification of the model, even if the nodes characteristics or the network structure are endogenous. In other words, exogenous shocks are used as a focal lens to identify the spillovers, which is usually very challenging.

[14] They show how identification of peer effects can be achieved in social networks, using an IV-strategy.

Recall that the underlying relationship of interest is the role of links in content generation and whether an article is more likely to be improved because of spillovers through links. The mechanism we have in mind, is that attention from article A can be diverted to article B if a link exists and that some of the users who get to see B start to edit it. Thus, the first thing to show is, how much attention spills via links, which can be modeled using the well known linear-in-means model of the type discussed in Manski (1993), who shows that the coefficient of interest is generally very hard to identify.

I start from the same form of the model.[15] In this section I provide only the point of departure and the main results, the details and derivations can be found in the appendix.[16]

$$y_{it} = \alpha \frac{\sum_{j \in P_{it}} y_{jt}}{N_{P_{it}}} + X_{it-1}\beta + \gamma \frac{\sum_{j \in P_{it}} X_{jt-1}}{N_{P_{it}}} + \epsilon_{it}$$

y_{it} denotes the outcome of interest in period t and X_{it-1} are i's observed characteristics at the end of period $t-1$.[17] P_{it} is the set of i's peers and $N_{P_{it}}$ the number of i's peers. α is the coefficient of interest: It captures the effect of the performance of i's peers and in the present context it measures how the views of an article are influenced by the views of the adjacent articles. The coefficient vector β accounts for the impact of i's own characteristics and γ measures the effect of the peers' average characteristics on i's performance. In the setting of this paper β accounts for how the page's own length or quality might affect how often it is viewed and γ captures how length and quality of neighboring pages affect views of page i. Bramoullé et al. (2009) suggest a more succinct representation based on vector and matrix notation:

$$\mathbf{y_t} = \alpha \mathbf{G} \mathbf{y_t} + \beta \mathbf{X_{t-1}} + \gamma \mathbf{G} \mathbf{X_{t-1}} + \epsilon_\mathbf{t} \quad \mathbf{E}[\epsilon_\mathbf{t} | \mathbf{X_{t-1}}] = 0$$

Clearly this model and, specifically, measuring the social parameter α is of interest to a very large literature. To incorporate exogeneous variation, I augment this model by including a vector of treatments, 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 (a "Mini Population Treatment").

$$\mathbf{y_t} = \alpha \mathbf{G} \mathbf{y_t} + \mathbf{X_{t-1}}\beta + \gamma \mathbf{G} \mathbf{X_{t-1}} + \delta_1 \mathbf{D_t} + \epsilon_\mathbf{t} \qquad \mathbf{E}[\epsilon_\mathbf{t} | \mathbf{D_t}] = 0 \tag{12}$$

A few remarks: $\mathbf{G}$ is a NxN matrix. $G_{ij} = \frac{1}{N_{P_i}-1}$ if i receives a link from j and $G_{ij} = 0$ otherwise. For the treated side $\mathbf{D_t}$ is a vector consisting of zeros and ones that indicates

[15] Note that it is easy to add a fixed effect to the model, but that it will be eliminated when taking differences. Consequently, I omit it for ease of notation.

[16] The derivations involve quite heavy notation, but are otherwise relatively straightforward.

[17] Note, that I can observe the current state of a Wikipedia article once a day at a fixed time.

which nodes are treated. In some of the proofs and in my application I will assume a local treatment that affects only a single node. Formally this is written as $\mathbf{D_t} = \mathbf{e_{\ell 0}}$; that is, a vector of zeros and a unique one in the coordinate that corresponds to the treated node. On the untreated subnetwork we have $\mathbf{D_t} = \mathbf{0}$, a vector of zeros.

Note that I do not require that the structure of the network ($\mathbf{G}$) is exogenous, but rather which node gets treated. It is worth stressing that my setup is fundamentally different from Bramoullé et al. (2009) because it will use an entirely different source of identification. Moreover, there will be no requirements needed concerning the linear independence of $\mathbf{G}$ and $\mathbf{G^2}$.

In this model, the reduced form expectation conditional on "treatment" is given by:

$$\mathbf{E}[\mathbf{y_t}|\mathbf{D_t}] \quad = \quad (\mathbf{I} - \alpha\mathbf{G})^{-1}[(\beta + \gamma G)\mathbf{E}[\mathbf{X_{t-1}}|\mathbf{D_t}] + \delta_1\mathbf{D_t}] \tag{13}$$

Define the set of observations in the subnetwork where treatment occurs in t by the subscript ℓ and a comparison group in which no node is treated by the subscript c. If these sets of nodes can also be observed one period earlier, a difference-in-differences estimator can be computed.

Result 1: Denote the difference in difference estimator as

$$\mathbf{DiD} := [\mathbf{E}[\mathbf{y_{\ell,t}}|\mathbf{D_{\ell,t}}] - \mathbf{E}[\mathbf{y_{\ell,t-1}}|\mathbf{D_{\ell,t-1}}]] - [\mathbf{E}[\mathbf{y_{c,t}}|\mathbf{D_{c,t}}] - \mathbf{E}[\mathbf{y_{c,t-1}}|\mathbf{D_{c,t-1}}]]$$

and assume that the treatment affects only the contemporary outcome of the treated node and not its exogenous characteristics.[18] Then the DiD contains the following quantity:

$$\mathbf{DiD} \quad = \quad \delta_1\mathbf{D_t}(\mathbf{I} + \alpha\mathbf{G} + \alpha^{\mathbf{2}}\mathbf{G^2} + \alpha^{\mathbf{3}}\mathbf{G^3} + ...)$$

In words, this result means that the node is affected by both treatment and second and higher order spillovers, the positive feedback loop that ensues as the neighbors increase their performance in sync with their peers. Instances of a higher order effects[19] are $\alpha^2\delta_1$ in the second round or $\alpha^3\delta_1$ in the third round and so on. The other important factor is whether and how often spillovers of a given order q arrive. This depends on the number of indirect paths of length q that go from the shocked node $\ell 0$ to any focal node j.

Proof: For a proof please refer to Appendix B.3.

My result shows that the difference-in-differences approach alone will not directly reveal α, the social parameter of interest, because nodes might have a feedback effect on each other. The neighbor's change in performance (due to the original impulse) will

[18]The independent characteristics X should not be immediately affected by treatment because this would threaten the identification of the spillover. However, they may adjust over time. As long as we can observe one period where only the outcome is affected, but not the characteristics, the result holds.

[19]Note that I am considering the homogeneous network, so all spillovers have the same magnitude.

affect the neighbors' neighbors, but also feed back on the treated neighbor. The estimator will also observe all the changes in outcome at the end of this process, when all higher order spills have taken place. In some applications this will be the object of interest to the researcher, however in the present context, the research is motivated by the desire to know the effect of the link structure and not of the treatment per se. Consequently it is warranted to dig deeper in order to understand the structural parameters.

Computing the parameters is not necessarily feasible, because it requires the knowledge of the complete link structure. However, a closer look at the nodes independently reveals that limited information about the link structure suffices to acquire additional information about the parameters. In the following two subsections I show how to get the point estimate for the peer effects coefficient if the network is known and I show how to derive upper and lower bound estimates for the parameter if no information about the network is available.

3.4.1 Estimator of the Peer Effects Parameter if the Network Structure can be Observed

If the network structure can be observed, the peer effect parameter α can be backed out by computing the higher orders of the network graph (G-matrix). To know how many spillovers arrive in each round, it suffices to focus on the entries $G_{ij}, G_{ij}^2, G_{ij}^3, ...$ that document the number of paths via 1, 2, 3,... links from the treated node to the neighboring node in question. With this information it is straightforward to compute by how much the observed effect at the node in question has to be discounted and to use this information to compute the true average effect.

3.4.2 Upper and Lower Bound Estimates of the Peer Effects Parameter if the Network Structure is Unobserved

If the network structure cannot be observed, it is possible to obtain boundary estimates for the peer effects based merely on the separate comparison of the directly treated nodes and their counterparts (L_0 vs. C_0) in the control group and their neighbors (L_1 vs. C_1). While randomization and information on the network together are rarely available, a separate comparison of eligible and non-eligible nodes in randomly treated communities or networks is frequently available in empirical settings. Yet, even in such situations it shall be possible to obtain a lower bound estimate for the coefficient α, if the researcher is willing to make some rigorous, but not uncommon, assumptions. In what follows I briefly show how to obtain the bounds. The idea behind this derivation is to select two specific "extreme" types of networks which either minimize or maximize second and higher order spillovers. These benchmark networks are schematically represented in Figure 3. I use a directed network with only "outward bound" links emanating from $\ell 0$ to $\ell 1 \in L1$

Figure 3: Schematic representation of the two extreme networks, used to compute the upper and lower bound estimates of the parameters of interest.

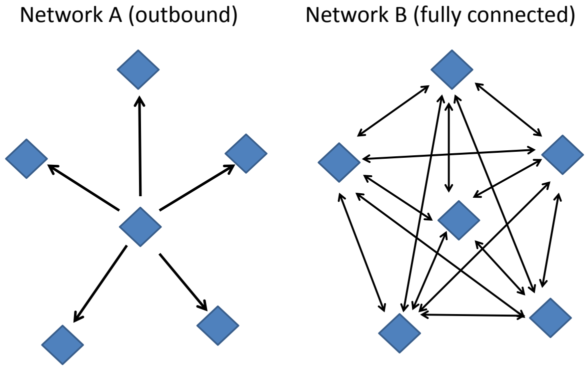

NOTES: The " outbound network" (left) is used to obtain the upper bound estimate. It is a directed network with only "outward bound" links. Holding the number of nodes and the observed ITEs fixed, the social parameter will be estimated to be largest in this type of network. The fully connected network (right), is the benchmark case from which the lower bound of the social parameter can be estimated.

to obtain the upper bound estimate of the social/spillover parameter α. The opposed benchmark is a fully connected network, where every node is the direct neighbor of every one of its peers. From there I obtain the lower bound estimate of the social parameter. A more detailed account is provided in Appendix B.4.

If we ignore higher order spillovers,[20] we can obtain an upper bound estimate for the ATE ($\overline{\delta_1}$) by applying the difference-in-differences estimator on the level of directly treated nodes (L_0) and a suitable comparison group (C_0) . After that I can move on to estimate the upper bound for the parameters for spillovers ($\overline{\alpha}$) by combining it with a second difference-in-differences estimator at the neighbor level. Let $DiD_{(\ell a - ca)}$ denote such a difference-in-differences ($a \in \{0, 1\}$) whether the nodes are in the center of the network (L_0 or C_0) or are the neighbors of the start nodes (L_1 vs. C_1):

$$\begin{aligned} \hat{\overline{\delta_1}} &= \widehat{DiD_0} = \Delta\hat{\ell}0 - \Delta\hat{c}0 \\ \hat{\overline{\alpha}} &= \frac{\widehat{DiD_1}}{\widehat{DiD_0}} NP_{\ell 1} \end{aligned} \tag{14}$$

- $\Delta\hat{\ell}0 := \frac{1}{NP_{\ell 0}} * \sum_i (y_{i,\ell 0,t=1} - y_{i,\ell 0,t=0})$
- $\Delta\hat{c}0 := \frac{1}{NP_{c0}} * \sum_i (y_{i,c0,t=1} - y_{i,c0,t=0})$

with $\widehat{DiD_1} = \Delta\hat{\ell}1 - \Delta\hat{c}1$ and the definitions of $\Delta\hat{\ell}1$ and $\Delta\hat{c}1$ paralleling those of $\Delta\hat{\ell}0$ and $\Delta\hat{c}0$. In my application's reduced-form estimations of the previous section DiD_1

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

corresponds to $\phi_{2,0}^{L_1}$ and DiD_0 is estimated by $\phi_{2,0}^{L_0}$. This upper bound estimator would be suitable under the potentially quite strong assumption that higher order spillovers are negligible. I proceed to show how to compute the lower bound estimates under the assumption of *maximal* second order spillovers. The lower bound gives an idea of the maximal size of the problem that might result from trusting the easily computed upper bound estimates.

It is also possible to compute a lower bound estimate for α and δ_1. This bound can be obtained by imagining that the network is fully connected, i.e. every node links to every other node, assuming that all effects are of the same sign, strictly ordered and (w.l.o.g) positive.[21] Further computations in Appendix B show that in a network with N nodes, the lower bound of the estimator for α is characterized by the solution to the following quadratic equation:

$$\underline{\alpha}^2 - [\frac{DiD_0}{DiD_1} + (N-1)]\underline{\alpha} + (N-1) = 0 \tag{15}$$

This equation has two solutions, one of which lies between 0 and 1. The closed form solution for $\underline{\alpha}$ is hence given by:

$$\underline{\alpha} = \frac{1}{2}[\frac{DiD_0}{DiD_1} + (N-1)] - \sqrt{\frac{1}{4}[\frac{DiD_0}{DiD_1} + (N-1)]^2 - (N-1)} \tag{16}$$

Recall that all the quantities required are readily available from the reduced form estimations. DiD_1 corresponds to $\phi_{2,0}^{L_1}$ and DiD_0 is estimated by $\phi_{2,0}^{L_0}$. In Appendix B.4 I provide a proof for my claims and explain how this bound is derived. Which of the estimates is more accurate will depend on the size of the spillover effect, but to a very large extent also on the real network structure and the number of nodes. The upper bound estimator would be quite suitable if the researcher assumes (potentially quite strong) that higher order spillovers are negligible. It would also be appropriate in networks with very sparse connections among its members. The lower bound estimator might be more suitable if the researcher believes the network to be highly connected and expects the spillover coefficient to be relatively large. The bounds have several limitations (cf. Appendix B.4) and for some applications the bounds might turn out to be to wide to be actually informative. Still, taken together, the bounds can provide a useful first characterization of the spillover parameters in question.

4 Data

This section gives detailed information about the dataset. Subsection 4.1 explains how the database was put together and the procedure I used to extract the dataset that

[21] The precise assumption is $DiD_0 > DiD_1 > HO^B > 0$, as stated and explained in Lemma 1

I use. Subsection 4.2 describes the dataset itself.

4.1 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, I constructed a time-varying graph of the article network, which provided the foundation for how I sample articles in my 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 I use 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 was not be 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)).[22]

"Featured articles" were found by consulting the German Wikipedia's archive of pages that were selected to be advertised on Wikipedia's main page ("Seite des Tages") between December 2007 and December 2010. To reduce the computational burden and to avoid the risk of temporal overlaps of different treatments, I focus on pages that were selected on the 10^{th} of a month. I identified all the pages that received a direct link (L_1) or an indirect link (L_2) from such a featured article more than a week before treatment. I evalutated links with this time gap before the shock actually occurred to make sure that the results are not driven by endogeneous link formation.[23] Having fixed the set of pages to observe, I extracted daily information on the contemporary state of the articles (page visits, number of revisions, number of distinct authors that contributed, page length, number of external links etc.). I determine these variables on a daily basis, 14 days before the event occurred (on a neighboring page) and 14 days after the shock (giving a total of 29 observations per page).

To identify major events, I consulted the corresponding page on Wikipedia and selected the 26 largest events with spontaneous onset. For each of these events we identified the page that corresponds to the event, which are considered to be in the set "L_0" (sometimes also called "start pages"). Note that this page is typically created *after the event*

[22]This is a novel high-level language allowing the writing and efficient execution of queries on nested and ordered collections of data.

[23]I thus 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 (L_2) 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.

occurred[24], which obliges me to identify the pages, that user will most likely turn to until the disaster's page is in place. To achieve this, I used the link data to identify the set of pages that later shared a reciprocal link with the start page. Such a reciprocal link indicates that they were closely related to the event. After the event page came in to existence they were only one click away (set "L_1"). Next, we identified those pages that received a link from an L_1 page (unidirectional) (2 clicks away set "L_2")

I am most interested in attention spillovers and content provision, which are not directly related to the events but rather a consequence of the spike in interest and the resulting improvements to the linked pages. Hence, I will not focus on the treated pages directly, but on the set L_1 that are "one click away", in my analysis of the "featured articles".[25] For disasters the shock is very large and the event page usually does not exist at the time of the shock, so the L_1 pages might have been treated themselves.[26] Hence, I focus on the indirectly linked set of pages (L_2) in the analysis below.

The approach I take in this paper hinges on the availability of a valid control group. To obtain such observations I pursue two distinct strategies. The first approach uses pages which are similar but unlikely to be affected by the treatment. For a first comparison I selected articles and neighbors thereof that were featured either later or earlier in time. Given such a similar page, I identified their direct and indirect neighbors when the event occurred on the treated page. This gives me a set $L2_{control}$ which is similar in both size and characteristics to the sampled pages (before the shock). Yet, the choice of the start pages in the comparison group is somewhat arbitrary.[27] I address this issue by simulating a treatment on the treated pages 42 days before the disaster or event occurred. I refer to the articles in this "placebo-treatment" as $L2_{placebo}$, because for them $t = 0$ when no actual treatment occurred. By design, this comparison group consists of the same set of articles (treated and their neighbors). This comes at the cost of observing the articles at a different point in time. A third control group of "unrelated" observations results from applying a placebo to the control group.[28]

For disasters I proceeded along similar lines. I focused on the network around older catastrophes that occurred at a different point in time and were not from exactly the

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

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

[26]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 merely *due to the event* (e.g. destruction of an important monument). Consequently, I do not emphasize the change in activity on the pages that are only one click away for disasters. I also exclude pages if they were later directly linked to the event page.

[27]Ideally the selection of comparison pages should be based on matching procedures, which is unfortunately not possible without computing the characteristics of all the 1,000,000 nodes. My approach is however quite robust independently of how I specify the control group. I also compared to the neighbors around articles of similar size and relative importance, about similar topics, but in a remote geographic space or technical domain. Such a change in the specification of the control group does not affect my results. (available upon request).

[28]This set of observations actually emerged as an artifact from the data extraction. Nevertheless it provides yet another group that can be compared to the treated group.

same domain, to avoid overlaps in the link network ($L_{control}$). Alternatively, I observe the same set of pages seven weeks before the disaster ($L_{placebo}$).

In Table 4 shows which featured articles were chosen by my procedure and included in the data. In general, they cover various topics such as innovations (e.g. the CCD-sensor), places (Helgoland), soccer clubs (Werder Bremen) or art history topics (Carolingian book-illustrations). I show the number of observations that received a link from an article before it was featured, separated by whether or not they belong to a time-series with actually treated observations.[29] It ranges from 2,088 to 33,872. Table 5 shows the number of articles that belong to each featured article. Table 7 shows which events were included in the data and shows the associated number of observations for each of them. The dataset includes both natural disasters as well as technical or economic catastrophes.

4.2 A Closer Look at the Dataset

Summary statistics for the data on large events are shown in Table 6. The data contains 425,981 observations from 7,379. From the table it can be seen that the average page contains 5658 bytes of content and has undergone 84 revisions. However, the median is substantially lower at 3885 bytes and only 40 revisions. Also, the summary statistics of the first differences (variables starting with "Delta:" 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[30], which is confirmed by the visual inspection of the direct and indirect effect of treatments.

In Figure 4 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 its 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.[31] The upper row contains four lines showing the control group and the directly treated nodes, which are created only after the onset of the event. Most of these 23 pages did not exist at all before

[29]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.

[30]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 its behavior during treatment. These tables and their description were omitted for reasons of brevity. They are available from the author upon request.

[31]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.

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

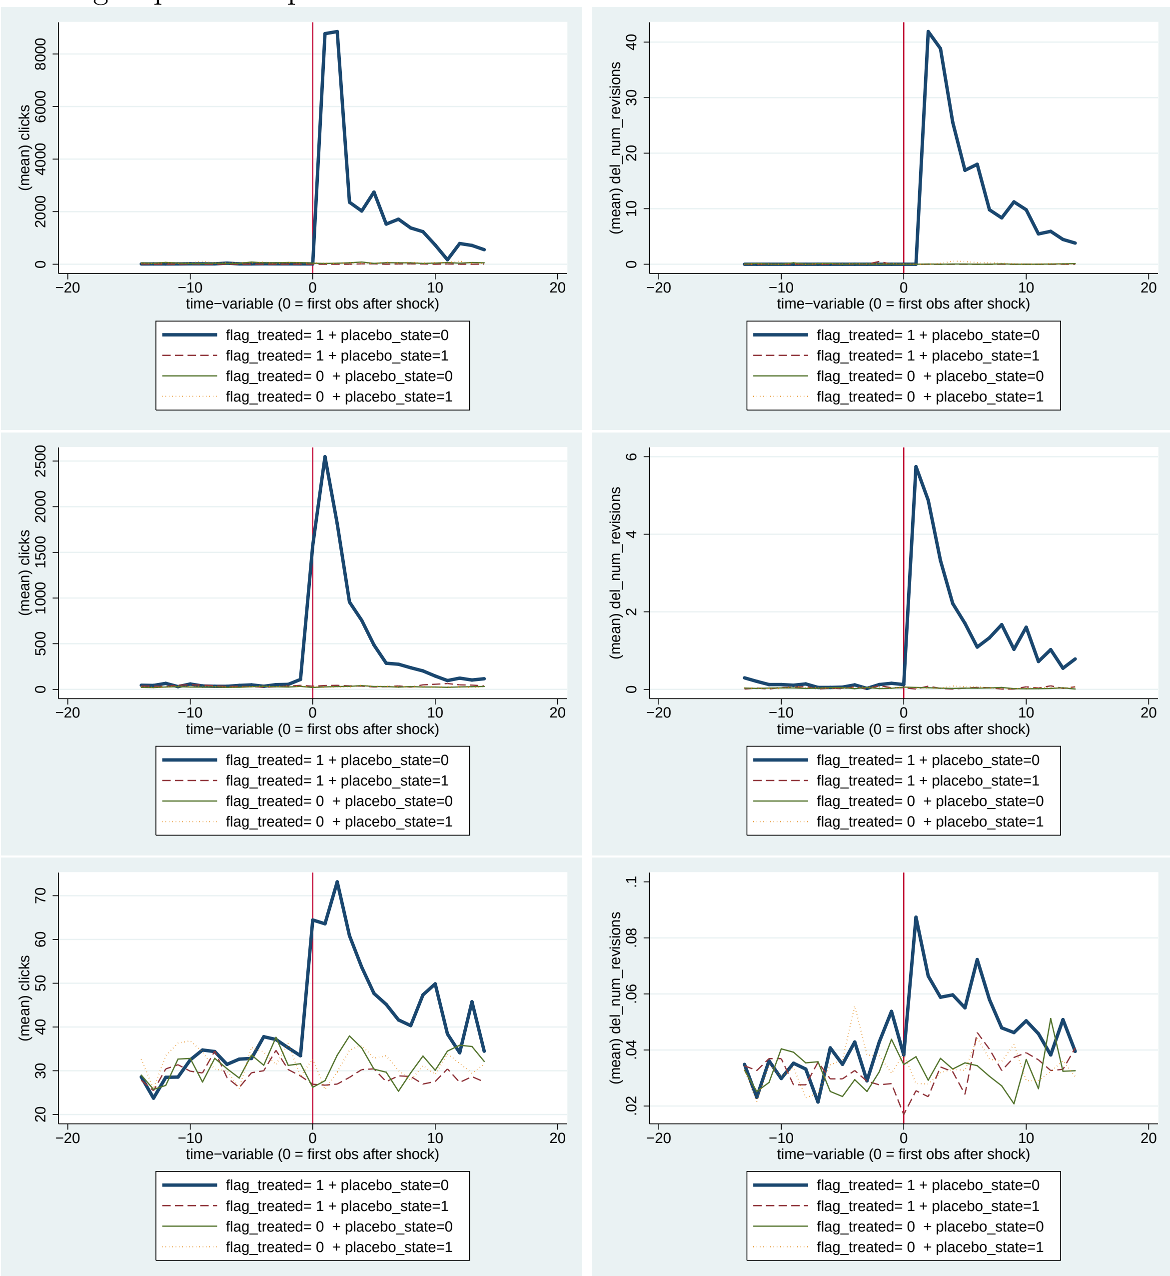

Notes: The upper row shows the average effect on the event pages (which by definition were created after the event), the middle row the directly treated pages (L1, with reciprocal link), and the lower row for the pages that are one click away from L_1. The left column shows the average number of clicks the right column shows the average number of edits. The outcomes are shown for the treated articles and the control groups separately. Directly hit pages received up to 8,500 additional clicks and up to 40 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 35 additional clicks and up to 0.04 additional revisions on average.

the onset of the event and therefore only a few have a placebo condition available. The row shows that the directly affected pages experience a very large spike of 8,500 clicks per day on average. The number of additional revisions peaks on the first days of treatment, when the pages are created: an average of almost forty revisions are added to a page on the first day. On the pages that are to share a reciprocal a link from the treated page the effect is quite pronounced: while the clicks on the average L_1 page increase by 2,500, the absolute value of the average increase in revision activity is only five. When I look at pages that are two clicks 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.

A summary of the data from "featured articles" are shown in Table 3. The data contains 317,550 observations from 5,489 pages[32] 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 4833 bytes of content and has undergone 48 revisions. In this sample, the mean is substantially higher at 6794 bytes and 95 revisions. As before, the summary statistics of the first differences show clearly how little activity occurs on a normal day on any given page on Wikipedia.

Figure 5 plots the aggregate dynamics around the day when the start node was shown on Wikipedia's main page and corresponds to Figure 4 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 direct neighbors. As before, each of the four figures contains four lines, one for each condition that can be obtained by 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 typical levels, below 50 views, to more than 4200 views on average, but immediately returns to the old levels the day after treatment is 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. Excepting large events, activity rises already before $t = 0$. Nevertheless, on the day of treatment the spike of activity is also pronounced for the neighbors.

[32]Since pages were observed also in the placebo condition, each page is sampled twice, and hence I observe 10,950 distinct time series.

Figure 5: "Today's featured articles": Contrasting means of clicks vs. number of added revisions over time: looking at all 4 groups in one plot.

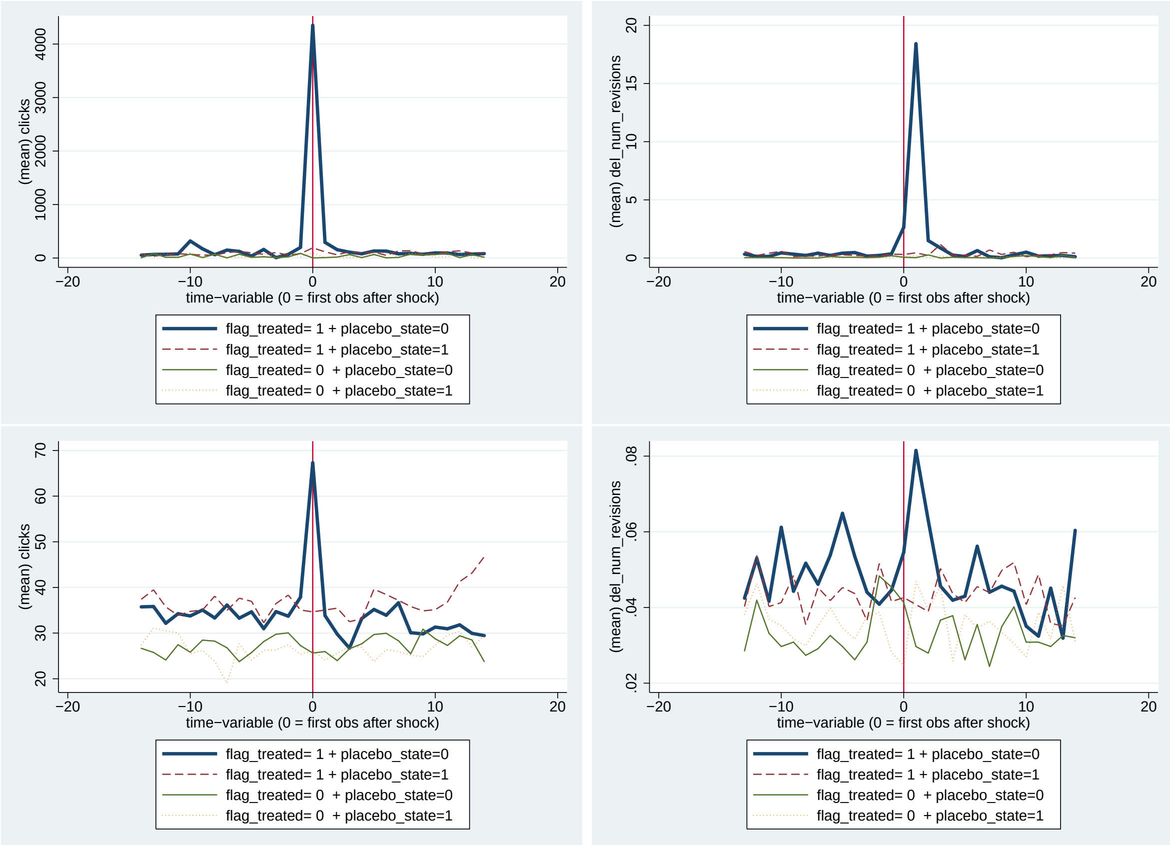

NOTES: The upper row shows the average outcome on the directly treated pages (set L_0 containing 63 pages total), the lower row for the pages one click away (set L_1, which contains 5,489 pages). The left column shows the average number of clicks the right column shows the average number of edits. The outcomes are shown for the treated articles and the control groups separately.

5 Estimation Results

In what follows I present my estimation results and discuss their interpretation. Before I proceed with the details of my estimations, it is worth recalling a few important facts. The point of departure of the estimations in this paper is estimating Equation 11 (Section 3.3) for large events and Equation 10 for "featured articles" . This is due to two reasons: first, the two conditions differ in how local the treatment I exploit for estimation is. Second, only the "featured articles" exist at the time of treatment, while the page at the center of a large event treatment typically does not exist and will instead be created in the following days.

Moreover, I avoid potentially endogenous link formation during treatment by considering only links that had been in place *a week before the treatment*. When a page is found to lie in both the treatment and control groups it is excluded from the estimation, because including such pages will bias the estimated coefficients towards zero. 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. Finally, I use the seven observations from two weeks before treatment (days -14 through to -8) as the reference group in the estimations and I include only flow variables such as views, new revisions, new authors etc. to guarantee that my results are not driven by any anticipation effects.[33] The following two subsections report the results for both conditions.

5.1 Large Events

For this group the estimation concerns the set of L_2 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 likely to directly affect pages that will eventually 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 explained above, this is not the effect I am primarily interested in. Consequently I focus on pages that were indirectly linked at the time of the shock and that never became directly linked enter the sample. These articles are no longer likely to be directly affected by the treatment on the page two clicks away.[34] Moreover, to make sure that my L_2 pages are not directly related to the event I checked whether a page that was in L_2 when I evaluated the network a week before the shock was going to be linked to the page of the disaster later. Since this indicates a potential direct relationship, I eliminated such pages from the sample.

The results of the estimation of the model for L_2 nodes are shown in Table 1. 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 (23 clusters). 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

[33]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 before it is advertised, to make sure it fits into the corresponding box on Wikipedia's main page. This alone results in increased activity during the week before treatment. After carefully studying this process, I am not very concerned about this feature of the data, because the magnitude of the day-0 effect suggests that the vast majority of attention influx is due to readers who do not anticipate which page is to be advertised.

[34]The results for the L_1 group are included in the appendix. The effects are very large and statistically significant. The estimated coefficients for the L_0 group (not reported) are close to 4,500 for clicks and between 20 and 25 for revisions. However, due to the lack of sufficient observations, even these very large coefficient estimates are not statistically different from zero.

Table 1: Relationship of clicks/added revisions and time dummies for indirect neighbors of shocked articles (2 clicks away from epicenter) in the large events condition.

	clicks			del revisions		
	(1) compare control	(2) compare placebo	(3) compare all	(4) compare control	(5) compare placebo	(6) compare all
t = -2	3.172	3.487	2.065	0.010	0.015	0.008
	(4.709)	(4.545)	(4.431)	(0.009)	(0.010)	(0.008)
t = -1	0.978	3.144	3.450	0.010	0.026***	0.017**
	(3.993)	(3.742)	(3.866)	(0.010)	(0.008)	(0.008)
t = 0	37.391**	36.047**	35.394**	0.003	0.021*	0.004
	(14.421)	(14.386)	(14.287)	(0.011)	(0.011)	(0.010)
t = 1	35.020***	35.397***	36.757***	0.049**	0.062**	0.055**
	(11.098)	(11.113)	(11.082)	(0.024)	(0.023)	(0.023)
t = 2	38.767***	44.730***	42.319***	0.037**	0.043***	0.038***
	(13.650)	(13.589)	(13.523)	(0.014)	(0.012)	(0.012)
t = 3	22.069**	30.895***	25.945***	0.021*	0.025**	0.024**
	(9.168)	(8.730)	(8.448)	(0.011)	(0.011)	(0.010)
t = 4	17.601**	21.918***	19.063***	0.026**	0.028**	0.026**
	(7.065)	(6.917)	(6.838)	(0.012)	(0.013)	(0.012)
Constant	29.900***	29.994***	30.735***	0.033***	0.032***	0.031***
	(1.001)	(1.289)	(0.664)	(0.002)	(0.002)	(0.001)
All cross	Yes	Yes	Yes	Yes	Yes	Yes
Time Dummies	Yes	Yes	Yes	Yes	Yes	Yes
Observations	162338	104214	323158	154959	99477	308469
Number of Pages	7379	4737	14689	7379	4737	14689
Adj. R^2	0.003	0.003	0.002	0.001	0.001	0.000

NOTES: The table shows the results of the reduced form regressions to estimate the ITE in the large events condition. Columns (1)-(3) show the results for clicks and Columns (4-6) for new edits to the articles . Specification (1) and (4) contrast treated and comparison group; (2) and (5) show the comparison of treated articles with themselves but seven weeks earlier (placebo treatment). In Columns (3) and (6) I juxtapose the treated subnetworks with all available comparison groups at the same time. Fixed Effects Panel-Regressions with heteroscedasticity robust standard errors. The unit of observations is the outcome of a page i on day t. The time variable is normalized and runs from -14 to 14.; Only crossterms closer to treatment are shown, but all were included. Reference group t-14 to t-5; standard errors in parentheses: *** p<0.01, ** p<0.05, * p<0.1; no. of obs. = 323158; no. of clusters = 44; no. of articles = 7379.

simulating a (placebo) treatment 42 days (i.e. 7 weeks) before the real shock. The third column compares the shocked pages against all the available comparison groups at once.

For ease of representation the table only shows the coefficients for the cross terms from two periods before the shock until four periods after the shock. As explained above, until the onset of the event (periods -2 to 0), we would expect insignificant effects for the cross terms and after the event has occurred a positive effect would imply that some form of spillover is present. Very much in line with the visual evidence, the average increase in clicks relative to the control group (Column 1) amounts to 35-38.7 more clicks on average. For the placebo treatment (Column 2) this effect is almost equal, but a bit larger from the second day onwards.

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 consistently revealed from the first day after the treatment. This effect is small in absolute terms, since roughly one in twenty to thirty pages gets an additional revision. Yet, given the low levels in average activity on a given page on a given day, this is still a noteworthy effect. Moreover, when comparing the pages with the placebo treatment I observe a small increase in editing activity before the onset of the event, which is however not confirmed by the comparison with the control group. The size of the effect still more than doubles after day 1, at which point the comparison with the control group suggests a drastic increase in editing activity.

5.2 Neighbors of Featured Articles

Table 2 shows the results for the "featured articles". For this reduced form estimation I consider the model for L_1 nodes (Equation 10) in Section 3.3. This is the relevant group here because the treatment takes place entirely inside Wikipedia[35] 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 1 and the clustering is implemented on the level of events as 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 33.1 to 34.6 additional clicks on the average neighbor page on the day of treatment. In terms of revisions (Columns 4-6), I observe an important effect of about 0.032 additional revisions one day after the treatment of the neighbor page. Note two things here: first, the effect is very small in absolute terms and corresponds to one additional edit per thirty pages. Second however, this is an increase in contribution activity of eighty to one hundred per cent.

I test the robustness of my results by excluding the first third of the "featured articles".[36] Table 8 shows the result of the test and adds a new dependent variable, the change in the number of editors (in Columns 5 to 6). Results reveal the same patterns as Table 1, but at lower significance levels. 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 seventy articles is edited by a new author.

Another way of understanding the meaning of these point estimates is to aggregate the changes in clicks and revisions over all neighboring articles and then averaging over

[35] Unlike in the disaster case, when an article is advertised on German Wikipedia's start page this is usually not covered by media or anything of the like.

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

Table 2: Relationship of clicks/added revisions and time dummies for direct neighbors of shocked articles in the 'featured articles' condition.

	clicks			del revisions		
	(1) compare control	(2) compare placebo	(3) compare all	(4) compare control	(5) compare placebo	(6) compare all
t = -2	-5.064	-2.629	-3.139	-0.028**	-0.018*	-0.019**
	(4.051)	(3.477)	(3.062)	(0.011)	(0.011)	(0.008)
t = -1	2.149	4.792	3.957	-0.021*	-0.004	-0.007
	(3.082)	(4.187)	(3.242)	(0.012)	(0.008)	(0.007)
t = 0	33.128***	34.638***	34.008***	-0.006	0.004	0.004
	(9.162)	(9.294)	(9.082)	(0.009)	(0.008)	(0.007)
t = 1	-0.158	0.773	0.645	0.032**	0.033**	0.030**
	(2.266)	(3.214)	(2.346)	(0.012)	(0.014)	(0.012)
t = 2	-2.523	-3.700	-3.438	0.015*	0.017	0.014
	(2.965)	(3.144)	(2.758)	(0.008)	(0.011)	(0.009)
t = 3	-8.373**	-3.807	-5.864	-0.011	-0.012	-0.013
	(3.371)	(5.435)	(3.949)	(0.012)	(0.013)	(0.011)
t = 4	-2.557	2.038	0.054	-0.016	-0.009	-0.008
	(2.766)	(5.615)	(3.535)	(0.014)	(0.009)	(0.008)
Constant	31.982***	35.354***	32.534***	0.043***	0.046***	0.042***
	(0.816)	(0.768)	(0.580)	(0.002)	(0.002)	(0.001)
All cross	Yes	Yes	Yes	Yes	Yes	Yes
Time Dummies	Yes	Yes	Yes	Yes	Yes	Yes
Observations	120758	166518	240900	115269	158949	229950
Number of Pages	5489	7569	10950	5489	7569	10950
Adj. R^2	0.004	0.003	0.002	0.000	0.000	0.000

NOTES: The table shows the results of the reduced form regressions to estimate the ITE in the 'featured articles' condition. Columns (1)-(3) show the results for clicks and Columns (4-6) for new edits to the articles. Specification (1) and (4) contrast treated and comparison group; (2) and (5) show the comparison of treated articles with themselves but seven weeks earlier (placebo treatment). In Columns (3) and (6) I juxtapose the treated subnetworks with all available comparison groups at the same time. Fixed Effects Panel-Regressions with heteroscedasticity robust standard errors. The unit of observations is the outcome of a page i on day t. The time variable is normalized and runs from -14 to 14.; Only crossterms closer to treatment are shown, but all were included. Reference group t-14 to t-5; standard errors in parentheses: *** p<0.01, ** p<0.05, * p<0.1; no. of obs. = 240900; no. of clusters = 63; no. of articles = 5489.

the 34 different "featured article" clusters. This is done in Figures 6 and 7 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 6). Given that the average treated articles received an additional 4000 clicks this corresponds to a one to 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 revisions on the neighboring pages (Figure 7) increases approximately from 4.5 to 8.5. This is an additional four changes, which means that the 4000 initial additional clicks are converted in 4000 additional clicks on neighbors and four new revisions or a ratio of 1000:1000:1.

Finally I report results of an extended analysis which is omitted here, for reasons of space.[37] I include 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

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

value. Moreover, I split the sample into well-connected articles with many links and poorly connected ones with few, 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 1500 bytes, I find a much stronger relative effect in the number of edits. This indicates that new content is provided on pages where the existing content is limited.

5.3 Bounds for the Structural Estimator

Unfortunately I cannot compute the precise structural estimator because the full matrix $\mathbf{G}$ formed by the German Wikipedia is too large to be computed in memory. Hence I cannot solve for $\mathbf{G}^2$ and higher orders of the link matrix.[38] However, it is possible to present upper and lower bound estimates of the structural parameters that are discussed in Subsection 3.4 and derived formally in Appendix B.

To compute these values the researcher has to decide where to evaluate the number of peers. I choose to evaluate the coefficients at the median which is 31 for indirect neighbors of disaster pages and 36 for neighbors of "featured articles". This is a crude first evaluation which primarily serves to highlight how easy it is to retrieve the structural parameters once this decision is made. The rest reduces to a back of the envelope calculation.

To compute the upper bound of the social/spillover parameter α and the shock δ_1 I use Equation 40. My preferred estimates are taken from the "featured article" condition. Estimates from the disaster condition are reported in parentheses. $\overline{\delta_1}$ is directly estimated to be 4,200 (2,440 in the disaster condition) The estimate of $\overline{\alpha}$ is 0.292 (34/4,190*N=36) based on "featured articles" (based on disasters: (38/2,440*N=31) = 0.483).

Computing the lower bound estimates is slightly more involved since it involves plugging the estimates and the number of nodes into the closed form solution given in Equation 46. This gives the point estimator for the lower bound of α, which is estimated to be $\underline{\hat{\alpha}} = 0.222$ for "featured articles" and $\underline{\hat{\alpha}} = 0.320$ for disasters.

To conclude this section I attempt to quantify the meaning of these results: literally they mean that if the average clicks on the neighboring pages are increased by ten, this alone would result in an increase of 2.215 to 2.92 clicks on the page, which all come from the neighbors. Even though caution is needed to make the following claim, the results suggest that placing links has an effect, but that it is small. Provided this out of equilibrium thought experiment is warranted: creating additional links from neighbors that increase aggregate viewership of the neighbors by 200 is predicted to result in 1.61 additional views on the target page.[39] While this absolute effect in clicks is very small, the

[38] Ongoing work is attempting to solve this issue. If these efforts are fruitful, the results might be included in a revised version of this paper.

[39] As before I use the median number of neighbors for these thought experiments. Consequently 200

conversion to content is even smaller than that since even huge shocks did not generate many revisions on neighboring articles. This suggests that placing links strategically will only generate large effects, if the pages that link out are very frequented. However, for the normal traffic on a typical Wikipedia page we would expect very small effects.

6 Conclusions, Limitations and Further Research

This paper investigates how the network of links between articles on the German Wikipedia influences the attention and content generation individual articles receive. I use large scale media events and natural disasters as observable exogenous shocks to analyze the spillovers of attention and content generation mediated through links. I find substantial spillovers of attention in terms of both views and editing activity. Articles in the neighborhood of shocked articles received 35 more visits on average - an increase of almost 100 percent. The findings indicate that links that point to an article influence how much attention a node will receive. My structural estimates suggest that an article will receive 30 percent of the number of average views on neighboring articles. Hence, by placing links to oft frequented nodes and thus increasing the average daily views on their neighbors by ten, one could obtain three additional daily visits to an article.

My results also indicate that the spillovers in attention may be driven to a large extent by users who only look up information. The analysis of the "featured articles" suggests that the average visitor clicks on exactly one of the links. Yet, while increased content provision is large in relative terms, it is modest in absolute terms. The total number of daily revisions on the neighboring pages (Figure 7) increases approximately from 4.5 to 8.5, which is a small increase given the size of the treatment. Hence, my estimates suggest that using the link network is probably an expensive and inefficient strategy for channeling contribution flows.

My results may be interesting for wiki administrators charged with channeling flows of content contribution, be it when setting up a firm wiki or when realizing the Wikimedia Foundation's vision of a world in which "every single human being can freely share in the sum of all knowledge" (Wikimedia-Foundation (2013)). A promising area for further analysis would investigate whether new authors are attracted by the events or whether contributions are made only by authors that previously contributed to the subject. Future studies should also investigate if these estimates apply only to mature wikis like the German Wikipedia or whether the small size of the effects will pertain on younger wikis with less content. This question can only be analyzed via studying smaller or younger Wikis, which would be a promising endeavor. However, when interpreting my results

aggregate view correspond to five more views on average. The quantification is based on the upper bound estimates of α in the "featured article" condition (and would be 3.31 for disasters).

on a more abstract level and considering Wikipedia as a setting of peer production or a citation network that documents human knowledge, it is probably advantageous that this paper uses data from a mature wiki. The findings suggest that the attention to a certain field or project will be higher if it receives links from other articles in other areas.

This paper suffers from several limitations. Most importantly, the strategy of exploiting local exogenous treatments will not allow the identification of the social spillover parameter if neighbors of the treated nodes observe the treatment and adjust their outcome as a reaction to the mere fact that their neighbor was treated. An example would be a teacher who selectively punishes or favors a single student: if the other pupils react to the special treatment, by changing their motivation to study for the subject, then their changed performance will reflect the sum of the spillover and their behavioral adjustment. In Appendix C, I outline that case and illustrate formally why the spillover parameter can no longer be identified. Moreover, future research should thoroughly exploit the heterogeneity in intensity of direct treatment effects. In particular, I hope to understand how attention, currently measured as average effect, is distributed across neighbors. Is it evenly distributed or do users herd to only a few of the linked pages? Another promising area would use the methodology based on exogenous local treatments alongside that based on the network structure and the exploitation of open triads (Bramoullé et al. (2009), De Giorgi et al. (2010)). The approaches are complementary; research along these lines will result in valuable insights. Finally, my evaluation of the structural parameters of the underlying dynamic with which attention is transmitted between neighboring pages is based on several assumptions. Most importantly it was not yet possible to surmount the computational hurdle of exploiting the detailed network information when obtaining the structural estimates. Future research should include this information and investigate which population parameter should be optimally included for relating reduced form and structural parameters.

References

Angelucci, Manuela and Giacomo De Giorgi, "Indirect Effects of an Aid program: How do Cash Transfers Affect Ineligibles' Consumption?," *The American Economic Review*, 2009, *99* (1), 486–508.

Angrist, Joshua D and Jörn-Steffen Pischke, *Mostly Harmless Econometrics: An Empiricist's Companion*, Princeton University Press, 2008.

Aral, Sinan and Dylan Walker, "Creating Social Contagion through Viral Product Design: A Randomized Trial of Peer Influence in Networks," *Management Science*, 2011, *57* (9), 1623–39.

Ashenfelter, Orley and Michael Greenstone, "Using Mandated Speed Limits to Measure the Value of a Statistical Life," *Journal of Political Economy*, 2004, *112* (S1), S226–S267.

Ballester, Coralio, Antoni Calvo-Armengol, and Yves Zenou, "Who's Who in Networks. Wanted: The Key Player," *Econometrica*, 09 2006, *74* (5), 1403–17.

Banerjee, Abhijit, Arun G Chandrasekhar, Esther Duflo, and Matthew O Jackson, "The Diffusion of Microfinance," Technical Report, National Bureau of Economic Research 2012.

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, "Is Oprah Contagious? Identifying Demand Spillovers in Online Networks," *NET Institute Working Paper.*, 2012, *No. 10-18.*, Available at SSRN: http://ssrn.com/abstract=1694308 or http://dx.doi.org/10.2139/ssrn.1694308 (August 3, 2012).

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), 223–30.

Comola, Margherita and Silvia Prina, "The Impact of Exogenous Interventions on Networks: A Dynamic Peer Effect Model," 2013.

Crépon, Bruno, Esther Duflo, Marc Gurgand, Roland Rathelot, and Philippe Zamora, "Do Labor Market Policies have Displacement Effects? Evidence from a Clustered Randomized Experiment," *The Quarterly Journal of Economics*, 2013, *128* (2), 531–80.

Dahl, Gordon B, Katrine V Løken, and Magne Mogstad, "Peer Effects in Program Participation," Technical Report, National Bureau of Economic Research 2012.

Ferracci, Marc, Grégory Jolivet, and G. J. van den Berg, "Evidence of Treatment Spillovers within Markets," online 2012. last retrieved on Nov. 11, 2013 at http:www.efm.bris.ac.ukecgrjeval_resubmit.pdf.

Fershtman, C. and N. Gandal, "Direct and Indirect Knowledge Spillovers: The "Social Network" of Open-source Projects," *RAND Journal of Economics*, 2011, *42* (1), 70–91.

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), 241–75.

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, Iassen, Johannes Moskaliuk, Joachim Kimmerle, and Ulrike Cress, "Who Integrates the Networks of Knowledge in Wikipedia?," *Proceedings of the 6th International Symposium on Wikis and Open Collaboration*, 2010, pp. 1–10.

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 Wikipedia's Coverage of Breaking News Events," *American Behavioral Scientist*, 2013, *57* (5), 595–622.

Kuhn, Peter, Peter Kooreman, Adriaan Soetevent, and Arie Kapteyn, "The Effects of Lottery Prizes on Winners and their Neighbors: Evidence from the Dutch Postcode Lottery," *The American Economic Review*, 2011, *101* (5), 2226–47.

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), 531–42.

Miguel, Edward and Michael Kremer, "Worms: Identifying Impacts on Education and Health in the Presence of Treatment Externalities," *Econometrica*, 2003, *72* (1), 159–217.

Moffitt, Robert A, "Policy Interventions, Low-level Equilibria, and Social Interactions," *Social Dynamics*, 2001, pp. 45–82.

Oestreicher-Singer, Gal and Arun Sundararajan, "The Visible Hand? Demand Effects of Recommendation Networks in Electronic Markets," *Management Science*, 2012, *58* (11), 1963–81.

Sacerdote, Bruce, "Peer Effects with Random Assignment: Results for Dartmouth Roommates," *The Quarterly Journal of Economics*, 2001, *116* (2), 681–704.

Wikimedia-Foundation, "Frequently Asked Questions," online 2013. last retrieved on Sept. 9, 2013 at http://wikimediafoundation.org/wiki/FAQ/en.

Zhang, X. and F. Zhu, "Group Size and Incentives to Contribute: A Natural Experiment at Chinese Wikipedia," *The American Economic Review*, 2011, *101*, 1601–15.

A Data-Appendix

A.1 Summary Statistics for "Today's Featured Articles"

Table 3: Summary statistics: direct neighbors of shocked articles in the 'featured articles' condition

	mean	sd	min	p10	p50	p90	max
Length of page (in bytes)	6794	6784	17	51	4833	15262	81585
Number of authors	33	35	1	2	21	77	324
Clicks	33	131	0	0	0	77	20384
Number of Revisions	95	130	1	3	48	237	1382
Links from Wikipedia	118	301	0	6	36	286	9484
Dummy: literature section	.3	.46	0	0	0	1	1
Number of images	2.3	8.1	0	0	1	5	319
Number language links	13	18	0	0	6	37	180
References (footnotes)	1.3	4.5	0	0	0	4	182
Links to further info	2.3	4.2	0	0	1	6	155
time variable (normalized)	0	8.4	-14	-12	0	12	14
Delta: Number of Revisions	.042	.39	0	0	0	0	42
Delta: Length of page	2.1	159	-31473	0	0	0	31462
Delta: Number of authors	.015	.13	0	0	0	0	9
Delta: Links from Wikipedia	.054	1.1	-90	0	0	0	438
Delta: Number of images	.00099	.27	-50	0	0	0	132
Delta: References	.0014	.097	-18	0	0	0	18
Delta: Links further info	.00078	.1	-19	0	0	0	16

NOTES: The table shows the distribution of the main variables. The unit of observations is the outcome of a page i on day t. The time variable is normalized and runs from -14 to 14.; no. of obs. = 317550; no. of start pages = 63; no. of articles = 5489.

Table 4: Included "featured articles" and the number of observations that are associated with them.

name of event	flag_real_treatment		
	0	1	Total
	No.	No.	No.
Afrikaans	5,481.0	1,943.0	7,424.0
Alte_Synagoge_(Heilbronn)	1,885.0	1,131.0	3,016.0
Banjo-Kazooie	5,191.0	2,030.0	7,221.0
Benno_Elkan	5,133.0	2,900.0	8,033.0
Bombardier_Canadair_Regional_Jet	4,205.0	1,073.0	5,278.0
CCD-Sensor	31,001.0	2,871.0	33,872.0
Charles_Sanders_Peirce	11,716.0	3,219.0	14,935.0
Das_Kloster_der_Minne	1,827.0	1,102.0	2,929.0
Deutsche_Bank	10,005.0	9,860.0	19,865.0
Eishockey	4,698.0	4,698.0	9,396.0
Ekel	10,295.0	5,336.0	15,631.0
Fahrbahnmarkierung	1,276.0	1,276.0	2,552.0
Geschichte_Ostfrieslands	7,453.0	6,177.0	13,630.0
Geschichte_der_deutschen_Sozialdemokratie	9,599.0	8,033.0	17,632.0
Glanzstoff_Austria	14,094.0	1,537.0	15,631.0
Glorious_Revolution	6,206.0	2,668.0	8,874.0
Granitschale_im_Lustgarten	3,857.0	928.0	4,785.0
Gustav_Hirschfeld	6,438.0	1,740.0	8,178.0
Hallenhaus	2,117.0	2,001.0	4,118.0
Helgoland	8,120.0	5,104.0	13,224.0
Jaroslawl	12,789.0	5,829.0	18,618.0
Jupiter_und_Antiope_(Watteau)	1,160.0	928.0	2,088.0
Karolingische_Buchmalerei	4,843.0	4,553.0	9,396.0
Katholische_Liga_(1538)	1,682.0	464.0	2,146.0
Martha_Goldberg	1,595.0	1,595.0	3,190.0
Naturstoffe	9,338.0	9,222.0	18,560.0
Paul_Moder	1,798.0	1,682.0	3,480.0
St._Martin_(Memmingen)	1,653.0	1,711.0	3,364.0
Stabkirche_Borgund	1,421.0	899.0	2,320.0
Taiwan	5,017.0	4,669.0	9,686.0
USS_Thresher_(SSN-593)	3,712.0	1,479.0	5,191.0
Visum	1,624.0	1,624.0	3,248.0
Wenegnebti	1,798.0	1,363.0	3,161.0
Werder_Bremen	8,555.0	8,323.0	16,878.0
Total	207,582.0	109,968.0	317,550.0

NOTES: For each "featured article", the table shows the number of observations associated with all articles that are one clicks away from a start page. Observations associated with actually "featured articles" are shown separately from control observations. Pages included 5,489.

Table 5: Included “featured articles” that were advertised on German Wikipedia’s start page and the number of articles that are associated with them (1 clicks away).

name of event	No.
Afrikaans	128.0
Alte_Synagoge_(Heilbronn)	52.0
Banjo-Kazooie	125.0
Benno_Elkan	139.0
Bombardier_Canadair_Regional_Jet	92.0
CCD-Sensor	586.0
Charles_Sanders_Peirce	258.0
Das_Kloster_der_Minne	51.0
Deutsche_Bank	343.0
Eishockey	162.0
Ekel	270.0
Fahrbahnmarkierung	44.0
Geschichte_Ostfrieslands	235.0
Geschichte_der_deutschen_Sozialdemokratie	306.0
Glanzstoff_Austria	270.0
Glorious_Revolution	153.0
Granitschale_im_Lustgarten	83.0
Gustav_Hirschfeld	142.0
Hallenhaus	71.0
Helgoland	228.0
Jaroslawl	321.0
Jupiter_und_Antiope_(Watteau)	36.0
Karolingische_Buchmalerei	162.0
Katholische_Liga_(1538)	37.0
Martha_Goldberg	55.0
Naturstoffe	320.0
Paul_Moder	61.0
St._Martin_(Memmingen)	59.0
Stabkirche_Borgund	40.0
Taiwan	167.0
USS_Thresher_(SSN-593)	90.0
Visum	56.0
Wenegnebti	55.0
Werder_Bremen	292.0
Total	5,489.0

NOTES: For all “featured articles”, the table shows the number of associated articles that are two clicks away from one of the corresponding start pages (be it treated or control).

A.2 Summary Statistics for Disasters

Table 6: Summary statistics: indirect neighbors of shocked articles (2 clicks away from the epicenter) in the large events condition

	mean	sd	min	p10	p50	p90	max
Length of page (in bytes)	5658	6287	16	33	3885	13210	76176
Number of authors	29	34	1	1	18	71	435
Clicks	33	174	0	0	0	70	29865
Number of Revisions	84	133	1	2	40	211	2083
Links from Wikipedia	123	447	0	5	31	269	27611
Dummy: literature section	.2	.4	0	0	0	1	1
Number of images	1.3	2.4	0	0	0	4	57
Number language links	13	18	0	0	7	37	179
References (footnotes)	1.3	4.2	0	0	0	4	150
Links to further info	2.7	5.1	0	0	1	7	130
time variable (normalized)	0	8.4	-14	-12	0	12	14
Delta: Number of Revisions	.035	.35	0	0	0	0	44
Delta: Length of page	1.8	106	-22416	0	0	0	27500
Delta: Number of authors	.013	.12	0	0	0	0	11
Delta: Links from Wikipedia	.049	2.5	-1148	0	0	0	216
Delta: Number of images	.00047	.078	-27	0	0	0	20
Delta: References	.0014	.13	-32	0	0	0	29
Delta: Links further info	.0011	.12	-15	0	0	0	31

NOTES: The table shows the distribution of the main variables. The unit of observations is the outcome of an article i on day t. The time variable is normalized and runs from -14 to 14.; no. of obs. = 425981; no. of start pages = 44; no. of articles = 7379.

Table 7: Included disasters, associated observations and the associated number of pages (2 clicks away).

name of event	flag_real_treatment 0 No.	1 No.	Total No.
Air-France-Flug_447	4,495.0	1,392.0	5,887.0
Air-India-Express-Flug_812	19,662.0	1,711.0	21,373.0
Amoklauf_von_Winnenden	2,088.0	2,146.0	4,234.0
Bahnunfall_von_Halle_(Belgien)	2,436.0	580.0	3,016.0
British-Airways-Flug_38	6,699.0	1,624.0	8,323.0
Buschfeuer_in_Victoria_2009	928.0	957.0	1,885.0
Deepwater_Horizon	8,178.0	3,596.0	11,774.0
Erdbeben_in_Haiti_2010	15,602.0	6,322.0	21,924.0
Erdbeben_in_Sichuan_2008	11,571.0	1,508.0	13,079.0
Erdbeben_von_LâĂŹAquila_2009	3,654.0	1,885.0	5,539.0
Flugzeugabsturz_bei_Smolensk	12,412.0	8,758.0	21,170.0
GrubenunglÃijck_von_San_JosÃľ	8,033.0	551.0	8,584.0
Josef_Fritzl	6,264.0	1,044.0	7,308.0
Kaukasuskrieg_2008	18,705.0	1,276.0	19,981.0
KolontÃąr-Dammbruch	4,669.0	1,073.0	5,742.0
Luftangriff_bei_Kunduz	113,767.0	7,772.0	121,539.0
Northwest-Airlines-Flug_253	65,279.0	1,276.0	66,555.0
Sumatra-Erdbeben_vom_September_2009	4,002.0	2,726.0	6,728.0
US-Airways-Flug_1549	7,888.0	5,220.0	13,108.0
UnglÃijck_bei_der_Loveparade_2010	15,283.0	13,572.0	28,855.0
Versuchter_Anschlag_am_Times_Square	10,353.0	1,334.0	11,687.0
Wald-_und_TorfbrÃďnde_in_Russland_2010	13,485.0	2,204.0	15,689.0
ZugunglÃijck_von_Castelldefels	1,508.0	493.0	2,001.0
Total	356,961.0	69,020.0	425,981.0

NOTES: For each event, the table shows the number of observations associated with all articles that are two clicks away from a start . Observations associated with actually "featured articles" are shown separately from control observations. Pages included 7,379

name of event	No.
Air-France-Flug_447	102.0
Air-India-Express-Flug_812	369.0
Amoklauf_von_Winnenden	74.0
Bahnunfall_von_Halle_(Belgien)	52.0
British-Airways-Flug_38	144.0
Buschfeuer_in_Victoria_2009	33.0
Deepwater_Horizon	203.0
Erdbeben_in_Haiti_2010	379.0
Erdbeben_in_Sichuan_2008	227.0
Erdbeben_von_LâĂŹAquila_2009	96.0
Flugzeugabsturz_bei_Smolensk	368.0
GrubenunglÃijck_von_San_JosÃľ	149.0
Josef_Fritzl	129.0
Kaukasuskrieg_2008	346.0
KolontÃąr-Dammbruch	99.0
Luftangriff_bei_Kunduz	2,107.0
Northwest-Airlines-Flug_253	1,151.0
Sumatra-Erdbeben_vom_September_2009	116.0
US-Airways-Flug_1549	226.0
UnglÃijck_bei_der_Loveparade_2010	499.0
Versuchter_Anschlag_am_Times_Square	202.0
Wald-_und_TorfbrÃďnde_in_Russland_2010	273.0
ZugunglÃijck_von_Castelldefels	35.0
Total	7,379.0

NOTES: For each event in the data, the table shows the number of pages that are two clicks away from one of the two associated start pages (be it treated or control).

A.3 Additional Regression and Figures

Figure 6: Figure contrasting the mean of clicks on featured articles, with the aggregated clicks on all neighboring pages.

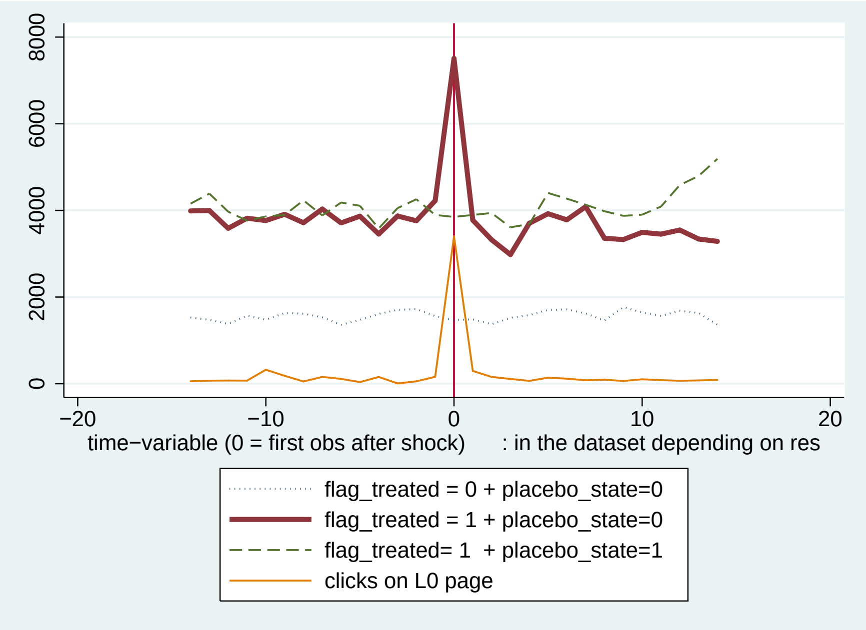

Notes: The figure shows the aggregated effect on the pages that are one click away. The average treated page received up to 4000 additional clicks, all neighbors together received approx. the same number of additional clicks

Figure 7: Figure showing the aggregated new revisions on all neighboring pages.

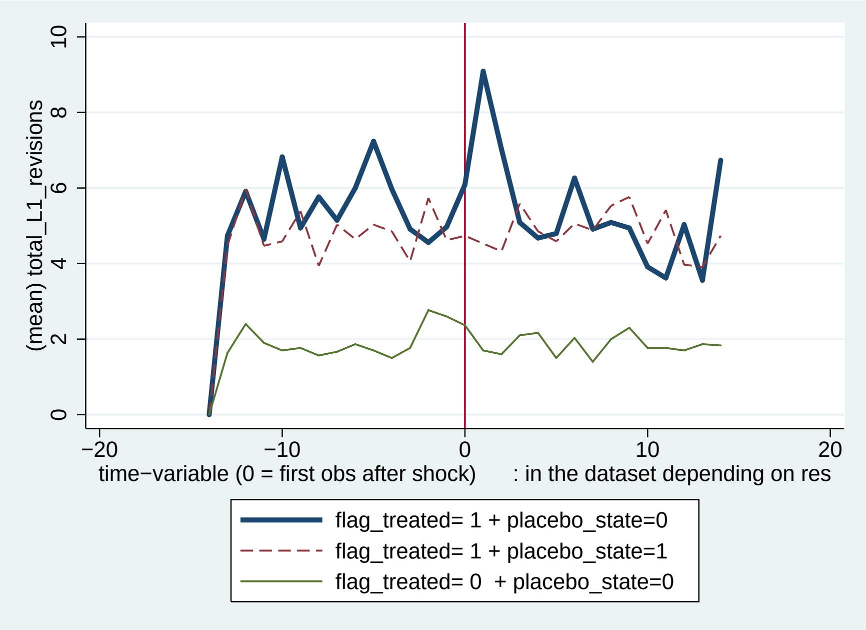

Notes: The figure shows the aggregated effect on the pages that are one click away. All neighbors of treated articles together received approx. four additional revisions.

Table 8: Robustness Check: Relationship of clicks/added revisions and time dummies for direct neighbors of shocked articles in the 'featured articles' condition 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	-2.117	4.304	-0.026**	-0.020	-0.014**	-0.009*
	(5.554)	(4.147)	(0.012)	(0.015)	(0.006)	(0.005)
realtreat_x_period_14	2.953	11.074*	-0.012	-0.004	-0.002	-0.005
	(4.448)	(5.974)	(0.013)	(0.011)	(0.005)	(0.004)
realtreat_x_period_15	34.625**	40.149***	-0.013	0.004	-0.006	-0.008
	(13.296)	(13.572)	(0.011)	(0.011)	(0.005)	(0.005)
realtreat_x_period_16	-1.463	2.145	0.037*	0.033	0.015**	0.012
	(2.685)	(4.649)	(0.018)	(0.021)	(0.007)	(0.008)
realtreat_x_period_17	-3.262	-0.427	0.012	0.005	0.002	-0.006
	(4.308)	(4.076)	(0.010)	(0.015)	(0.004)	(0.006)
realtreat_x_period_18	-10.195**	-3.046	-0.009	-0.033*	-0.005	-0.010*
	(4.874)	(4.977)	(0.012)	(0.019)	(0.006)	(0.005)
realtreat_x_period_19	-3.023	5.034	-0.031*	-0.019	-0.009	-0.003
	(4.074)	(3.968)	(0.016)	(0.012)	(0.005)	(0.005)
_cons	30.930***	34.526***	0.047***	0.050***	0.016***	0.016***
	(1.291)	(1.135)	(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	73084	98252	69762	93786	69762	93786
Number of Pages	3322	4466	3322	4466	3322	4466
Adj. R^2	0.004	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

This section presents the structural model and discusses the parameters of interest, the challenges in identifying them and the approach taken to tackle them.

I depart from the well known linear-in-means model as formulated by Manski (1993)[40] This model is based on the idea, that agent i's performance depends not only on her own characteristics, but also on both the performance and he characteristics of i's peers:

$$y_{it} = \alpha \frac{\sum_{j \in P_{it}} y_{jt}}{N_{P_{it}}} + X_{it-1}\beta + \gamma \frac{\sum_{j \in P_{it}} X_{jt-1}}{N_{P_{it}}} + \epsilon_{it} \tag{17}$$

where y_{it} denotes the outcome of interest in period t and X_{it-1} are i's observed characteristics at the end of period $t-1$ (beginning of period t)[41] P_{it} is the set of i's peers and $N_{P_{it}}$ represents the number of i's peers. β measures the effect of i's own characteristics and γ accounts for how i's performance is affected by the peers' average characteristics. 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:

$$\mathbf{y_t} = \alpha \mathbf{G} \mathbf{y_t} + \beta \mathbf{X_{t-1}} + \gamma \mathbf{G} \mathbf{X_{t-1}} + \epsilon_\mathbf{t} \quad \mathbf{E}[\epsilon_\mathbf{t} | \mathbf{X_{t-1}}] = \mathbf{0}$$

Clearly this model and specifically measuring α is of general interest to a very large literature. Moreover the linear in means model provides the weakest basis for identification. Hence, I conjecture that the insights carry over to less weakly identified models.

B.2 Setup and Basic Intuition

Augment the model (eq. 17) by observable and locally applied treatments (shocks):

$$y_{it} = \alpha \frac{\sum_{j \in P_{it}} y_{jt}}{N_{P_{it}}} + X_{i,t-1}\beta + \gamma \frac{\sum_{j \in P_{it}} X_{j,t-1}}{N_{P_{it}}} + \delta_1 D_{it} + \epsilon_{it} \tag{18}$$

where the new coefficient δ_1 measures the direct effect if a node(page) is treated.

Note that $X_{it-1}\beta$ may contain an individual fixed effect and an additively separable age-dependent part: $X_{it-1}\beta = \beta_i + \widetilde{X_{i,t-1}}\beta_1 + \beta_2 f(age)$. To see how local treatments can

[40] Note that it is easy to add a fixed effect to the model, but that it will be eliminated when taking differences. Consequently, I omit it for ease of notation.

[41] The choice of the temporal structure depends on the application that the researcher has in mind. In the present application many independent variables are stock variables (articles' characteristics such as page length), while the dependent variables are typically flows (clicks or new revisions).

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 ($\ell 0$) in period t and the neighbors are not treated ($\ell 1 \in L1$). Assume for simplicity that $\ell 0$ is the only treated node in $\ell 1$'s neighborhood.

$$\ell 0 :: \boldsymbol{y_{\ell 0t}} = \alpha \frac{\sum_{j \in P_{\ell 0t}} y_{jt}}{N_{P_{\ell 0t}}} + X_{\ell 0t-1}\beta + \gamma \frac{\sum_{j \in P_{\ell 0t}} X_{jt-1}}{N_{P_{\ell 0t}}} + \delta_1 \boldsymbol{1} + \epsilon_{\ell 0t} \tag{19}$$

$$\ell 1 \in L1 :: y_{\ell 1t} = \alpha \frac{\boldsymbol{y_{\ell 0t}} + \sum_{j \in P_{\ell 1t}/\ell 0} y_{jt}}{N_{P_{\ell 1t}}} + X_{\ell 1t-1}\beta + \gamma \frac{\sum_{j \in P_{\ell 1t}} X_{jt-1}}{N_{P_{\ell 1t}}} + \delta_1 0 + \epsilon_{\ell 1t} \tag{20}$$

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

Second, take two remote nodes c0 and $c1 \in C1$, where nobody gets treated:

$$c0 :: y_{c0t} = \alpha \frac{\sum_{j \in P_{c0t}} y_{jt}}{N_{P_{c0t}}} + X_{c0t-1}\beta + \gamma \frac{\sum_{j \in P_{c0t}} X_{jt-1}}{N_{P_{c0t}}} + \delta_1 0 + \epsilon_{c0t} \tag{21}$$

$$c1 \in C1 :: y_{c1t} = \alpha \frac{\sum_{j \in P_{c1t}} y_{jt}}{N_{P_{c1t}}} + X_{c1t-1}\beta + \gamma \frac{\sum_{j \in P_{c1t}} X_{jt-1}}{N_{P_{c1t}}} + \delta_1 0 + \epsilon_{c1t} \tag{22}$$

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

I use the notation suggested by Bramoullé et al. (2009) and incorporate the newly proposed vector of treatments. The equations above can be written in Matrix notation and X might include a time-dependent component (e.g. a linear function of age) as well:

$$\mathbf{y_t} = \alpha \mathbf{G y_t} + \mathbf{X_{t-1}}\beta + \gamma \mathbf{G X_{t-1}} + \delta_1 \mathbf{D_t} + \epsilon_\mathbf{t} \qquad \mathbf{E}[\epsilon_\mathbf{t} | \mathbf{X_{t-1}}] = \mathbf{0} \tag{23}$$

$\mathbf{G}$ is a NxN matrix, which captures the link structure in the network. $G_{ij} = \frac{1}{N_{P_i}-1}$ if i receives a link from j and $G_{ij} = 0$ otherwise. Note that I do not require $\mathbf{G}$ to be exogenously given, but only $\mathbf{D_t}$, a vector which is 1 at the treated nodes (if they are *currently* treated) and 0 otherwise. In some of the proofs and in my application I will assume a local treatment that affects only a single node. Formally this is written as $\mathbf{D_t} = \mathbf{e_{\ell 0}}$; that is, a vector of zeros and a unique one in the coordinate that corresponds to the treated node. On the untreated subnetwork we have $\mathbf{D_t} = \mathbf{0}$, a vector of zeros.

It is worth stressing that my setup is fundamentally different from Bramoullé et al. (2009), because it will not be based on finding an instrument for $\mathbf{Gy}$. Instead, I use an entirely different source of identification based on differences and exogenous shocks that affect only one part of the network. Hence, there will be no requirements on the linear independence of $\mathbf{G}$ and $\mathbf{G^2}$.

B.3 Proof of Result 1

I shall now proceed to provide the formal argument for Result 1. To increase the readability I will make a few assumptions to keep things simple. Most importantly I assume the network $\mathbf{G}$ to be stable over time but I allow $\mathbf{X_t}$ to change dynamically. I set the comparison group (which was indexed by c) to be the group itself S periods earlier, which results in an S-period Difference in Differences. This setting is reasonably close to comparing the evolution of nodes in a very stable network during a post and a pre-treatment stage. Importantly the nodes in the network have to be observed over time and have to evolve in a stable fashion, to ensure that the first differences are the same at t and $t-S$.[42] At the end of the formal derivations I will discuss the consequences of relaxing the requirement of a stable network or the consequences of adding the assumption that $\mathbf{X_t}$ does not change between the periods of observation.

Result 1: A Difference in Differences estimator contains the following quantity:

$$\mathbf{DiD} \quad = \quad \delta_1 \mathbf{D_t}(\mathbf{I} + \alpha\mathbf{G} + \alpha^{\mathbf{2}}\mathbf{G^2} + \alpha^{\mathbf{3}}\mathbf{G^3} + ...)$$

Proof.

The reduced form corresponding to equation 23 is given by:

$$\mathbf{y_t} = (\mathbf{I} - \alpha\mathbf{G})^{-1}[\mathbf{X_{t-1}}\beta + \gamma\mathbf{G}\mathbf{X_{t-1}} + \delta_1\mathbf{D_t} + \epsilon_\mathbf{t}] \tag{24}$$

and the expectation conditional on the "treatment" is:

$$\begin{aligned} \mathbf{E}[\mathbf{y_t}|\mathbf{D_t}] \quad &= \quad (\mathbf{I} - \alpha\mathbf{G})^{-1}[(\beta + \gamma\mathbf{G})\mathbf{E}[\mathbf{X_{t-1}}|\mathbf{D_t}] + \delta_1\mathbf{D_t} + \mathbf{E}[\epsilon_\mathbf{t}|\mathbf{D_t}]] =^{b.A.} \\ &=^{b.A.} \quad (\mathbf{I} - \alpha\mathbf{G})^{-1}[(\beta + \gamma\mathbf{G})\mathbf{E}[\mathbf{X_{t-1}}|\mathbf{D_t}] + \delta_1\mathbf{D_t}] \end{aligned} \tag{25}$$

[42] The setting is reasonably close to the "placebo condition" of my application below.

Taking the first difference, we obtain:

$$\begin{aligned}(26)\quad \mathbf{\Delta_t E[y|D]} = & \quad \mathbf{E[y_t|D_t] - E[y_{t-1}|D_{t-1}]} = \\ = & \quad \mathbf{(I-\alpha G)^{-1}}[(\beta+\gamma\mathbf{G})\{\mathbf{E[X_{t-1}|D_t] - E[X_{t-2}|D_{t-1}]}\} + \delta_1\mathbf{\Delta D_t}] = \\ = & \quad \mathbf{(I-\alpha G)^{-1}}[(\beta+\gamma\mathbf{G})\{\mathbf{E[X_{t-1}|D_t] - E[X_{t-2}|D_{t-1}]}\} + \delta_1\mathbf{D_t}]\end{aligned}$$

...where $\mathbf{\Delta D_t = D_t - D_{t-1}}$ and the second equality holds, because treatments are assumed to start in period t, but not before. That difference contains the time-dependent component and the effect of any changes in the independent variables.[43]

Now consider the control group formed by the same network, but S periods earlier:

$$\mathbf{y_{t-S}} = \alpha\mathbf{G y_{t-S}} + \mathbf{X_{t-S-1}}\beta + \gamma\mathbf{G X_{t-S-1}} + \delta_1\mathbf{D_{t-S}} + \epsilon_{\mathbf{t-S}}$$

The first difference of the reduced form's conditional expectations are:

$$\begin{aligned}\mathbf{\Delta_{t-S} E[y|D]} = & \quad \mathbf{E[y_{t-S}|D_{t-S}] - E[y_{t-S-1}|D_{t-S-1}]} = \\ & = \mathbf{(I-\alpha G)^{-1}}[(\beta+\gamma\mathbf{G})\{\mathbf{E[X_{t-S-1}|D_{t-S}] - E[X_{t-S-2}|D_{t-S-1}]}\} + \delta_1\mathbf{\Delta D_{t-S}}] = \\ & = \mathbf{(I-\alpha G)^{-1}}[(\beta+\gamma\mathbf{G})\{\mathbf{E[X_{t-S-1}|D_{t-S}] - E[X_{t-S-2}|D_{t-S-1}]}\} + 0]\end{aligned}$$

with $\mathbf{\Delta D_{t-S}} = 0$, since treatments are assumed to start in period t, but not earlier. Proceeding to take the Difference in Differences, we obtain:

$$\begin{aligned}\mathbf{DiD} := & \quad \mathbf{\Delta y_t E[y|D]} - \quad \mathbf{\Delta y_{t-S} E[y|D]} = \\ = & \quad \mathbf{(I-\alpha G)^{-1}} \quad [(\beta+\gamma\mathbf{G})\{\mathbf{E[X_{t-1}|D_t] - E[X_{t-2}|D_{t-1}]}\} + \delta_1\mathbf{D_t}] - \\ & \qquad - \quad (\beta+\gamma\mathbf{G})\{\mathbf{E[X_{t-S-1}|D_{t-S}] - E[X_{t-S-2}|D_{t-S-1}]}\}]\end{aligned}$$

Denoting the change in the expectation of $\mathbf{X_{t-1}}$ conditional on $\mathbf{D_t}$ more concisely by $\{\mathbf{E[X_{t-1}|D_t] - E[X_{t-2}|D_{t-1}]}\} = \mathbf{\Delta_t(E[X|D])}$ and rearranging gives:

$$(27) \qquad \mathbf{DiD} = \mathbf{(I-\alpha G)^{-1}} \quad [(\beta+\gamma G)\{\mathbf{\Delta_t(E[X|D]) - \Delta_{t-S}(E[X|D])}\} + \delta_1\mathbf{D_t}]$$

which reduces to:

$$(28) \qquad \mathbf{DiD} = \mathbf{(I-\alpha G)^{-1}}\{\delta_1\mathbf{D_t}\}$$

if $\mathbf{\Delta_t(E[X|D]) = \Delta_{t-S}(E[X|D])}$. Thus, the identifying assumption is that the expected changes of the pages between $t-1$ and t are the same as from $t-S-1$ and $t-S$.

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

This is satisfied if $\Delta X_t | D_t$ is stationary of order one.

Provided $(\mathbf{I} - \alpha\mathbf{G})^{-1}$ is invertible we can use the property that $(\mathbf{I} - \alpha\mathbf{G})^{-1} = \sum_{s=0}^{\infty} \alpha^s G^s$[44], the general impact of a local treatment is:

$$\mathbf{DiD} = \delta_1 \mathbf{D_t}(\mathbf{I} + \alpha\mathbf{G} + \alpha^{\mathbf{2}}\mathbf{G^2} + \alpha^{\mathbf{3}}\mathbf{G^3} + ...) \tag{29}$$

which completes the proof. ■

Discussion of the assumptions used:

1. $\mathbf{E}[\epsilon_{\mathbf{t}} | \mathbf{D_t}] = \mathbf{0}$

2. α is smaller than the norm of the inverse of the largest eigenvalue of $\mathbf{G}$. A regularity condition to ensures that the expression $(\mathbf{I} - \alpha\mathbf{G})^{-1} = \sum_{s=0}^{\infty} \alpha^s G^s$ is well defined.

3. I assumed the network to be stable over time and used it's earlier state as control observation. Formally this is written as $\mathbf{G_{\ell,t}} = \mathbf{G_{\ell,t-1}} = \mathbf{G}$ and $\mathbf{G_{c,t}} = \mathbf{G_{\ell,t-S}} = \mathbf{G}$. This assumption could be relaxed, but only at the expense of strengthening the following assumption.

4. $\mathbf{\Delta_t}(\mathbf{E}[\mathbf{X}|\mathbf{D}]) - \mathbf{\Delta_{t-S}}(\mathbf{E}[\mathbf{X}|\mathbf{D}])$, which means that the expected changes of the pages between $t-1$ and t are the same as from $t-S-1$ and $t-S$[45]. This is the analogue of the well known common trends assumption.

5. SUTVA on the level of subnetworks: the non-treated subnetwork is not affected by treatment of the treated subnetwork. In the present context SUTVA holds for my placebo condition and, given the size of the Wikipedia network, it is also plausibly satisfied for the control group formed by a remote part of the network.

The proof for the control group consisting of remote nodes is analogous. It relaxes the third assumption and requires a more general formulation of the fourth. The qualitative meaning of the generalized assumption will be the same: Absent treatment the treated *network* and the control *network* must "*evolve* in the same way." To be more precise, the link formation and the way in which the characteristics of the nodes change over time have to be the same (common trends) in both networks in order to guarantee that the counterfactual outcome of the treated network can be inferred from its own past and the evolution in the control network.[46] However, I have to maintain the assumption that the network formation *process* is not affected by the treatment. If this is the case, all

[44]$\mathbf{G}$ is invertible if $\alpha < 1$ (Bramoullé et al. (2009)) and the infinite sum is well defined if α is smaller than the norm of the inverse of the largest eigenvalue of $\mathbf{G}$ (Ballester et al. (2006)).

[45]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.

[46]The derivations require a lot of notational overhead and the resulting conditions are quite unwieldy. Assumption 4 would refer not only to $\mathbf{\Delta X}$, but to $\mathbf{\Delta GX}$, in order to allow for relaxing Assumption 3.

estimates of indirect treatment effects, will reflect a sum of the treatment on the existing network and new spillovers due to the changes in the link network (cf. Comola and Prina (2013)), which will lead to upward biases if not accounted for.

B.3.1 Estimating α: Analysis on the Node Level

Above we have shown 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 identifies:

$$\mathbf{DiD} = \delta_1 \mathbf{D_t}(\mathbf{I} + \alpha\mathbf{G} + \alpha^{\mathbf{2}}\mathbf{G^2} + \alpha^{\mathbf{3}}\mathbf{G^3} + ...) \tag{30}$$

When taking the analysis back from the level of treated networks and look at the nodes individually it is worth noting that for each focal node j its own row in this set of equations is all that matters. To simplify this analysis I will now begin to use the local treatment assumption, which exploits the fact that only a single node in the network is treated. This is like a partial population treatment Moffitt (2001) with only one single node (a mini population) being treated.

Local Treatment Assumption: Under the local treatment assumption $\mathbf{D_t} = \mathbf{e_i}$, where $\mathbf{e_i}$ is an elementary vector with node i being the only treated node.

Node that if only one node is treated, the spillover dynamic is greatly simplified. With $D = e_i$ the only factor to be evaluated for each node is its corresponding ji element in the matrix $\mathbf{G}$, $\mathbf{G^2}$ and it's higher orders.

The information that's contained in the higher orders of the adjacency matrix $\mathbf{G}$ will be the same as the information from the sampling strategy in combination with knowing who was affected by the local treatment. Some nodes (L0) are known to be directly treated, and some (L1) have a direct link so that the entry in $\mathbf{G}$ that links them to the treated node is positive. However, for those who only have an indirect link, the corresponding entry in $\mathbf{G}$ takes the value 0 and only the relevant element of $\mathbf{G^2}$ will be greater than 0.

If only one node in the network is treated, we distinguish a shocked node $\ell 0 \in L0$, a neighbor $\ell 1 \in L1$ and the indirect neighbors (2 clicks away, 3 clicks away etc.) as follows:

$$\begin{aligned} \ell 0: \ \mathbf{DiD_0} &= \delta_1(1 + 0 + \alpha^2 G_{ii}^2 + \alpha^3 G_{ii}^3 + ...) \\ \ell 1: \ \mathbf{DiD_1} &= \delta_1(0 + \alpha G_{ij} + \alpha^2 G_{ij}^2 + \alpha^3 G_{ij}^3 + ...) \\ \ell 2: \ \mathbf{DiD_2} &= \delta_1(0 + 0 + \alpha^2 G_{ik}^2 + \alpha^3 G_{ik}^3 + ...) \\ &\qquad\qquad\qquad\qquad\qquad\qquad \textit{etc.} \end{aligned} \tag{31}$$

Differentiating the nodes with respect to their distance from $\ell 0$ and estimating these strata separately results in as many estimation equations as can reasonably be traced and two parameters to be estimated. This fact is the basic idea of this paper, because it enables the researcher to back out the estimates for the structural parameters α and δ_1. All that is needed is a sequence of reduced form Differences in Differences estimates for increasingly large link-distances. If the precise information on $\mathbf{G}$ and its higher orders is available the parameters can be directly estimated.[47] If not, it is possible to compute an upper and a lower bound for the parameters α and δ_1. In the next subsection I proceed to show how the boundary estimates can be computed.

B.4 Estimating Bounds for the Parameters of Interest

If the researcher lacks information on $\mathbf{G}$ it is possible to compute an upper and a lower bound for the parameters α, the parameter that accounts for the social effect or spillover, and δ_1, the treatment effect (net of spillovers). This is useful, since the precise information on $\mathbf{G}$ is often not easy to obtain or computing its higher orders might confront the researcher with substantial computational challenges. In what follows I will show how to obtain these bounds. In Subsection B.4.1, I will give and intuitive account underlying the bounds and in Subsection B.4.2, I will set up the preliminaries, including a Lemma that will be used for obtaining both bounds. Subsection B.4.3 obtains the upper bound and Subsection B.4.4, finally, provides the proof for the lower bound.

B.4.1 Intuition for obtaining Bounds

The goal in this section is to back out a lower and an upper bound estimate for α and δ_1, that is based only on the estimated DiD's and the number of nodes. In my proofs I use the local treatment assumption (only one individual in the network is treated), for both ease of notation and understanding, which applies to "today's featured articles".[48]

To see why we can bound the parameter, even without knowing the details of the network structure, we can select two 'specific 'extreme" types of networks which either minimize or maximize the higher order effects. For greater convenience, I repeat the illustration of such networks in Figure 8. The network that minimizes higher order spillovers is a directed network with only "outward bound" links from $\ell 0$ to $\ell 1 \in L1$ [49]. This implies no links between the nodes in $L1$ and will serve as upper bound.

[47] To do this use all the ij values that correspond to each individual focal node j as weights for α, α^2, α^3, etc. and minimize a quadratic loss function. Unfortunately I cannot show this here, because the full matrix $\mathbf{G}$ formed by the German Wikipedia is too large to be computed in memory.

[48] I conjecture that extending the proof to partial population or randomized treatments will be straight forward. It merely means taking into account that more than one node gets treated and that the effects from the treated can also spill to the other treated, which will render the formulas quite unwieldy.

[49] and possibly further on to $\ell 2 \in L2$, $\ell 3 \in L3$ and so on.

Figure 8: Schematic representation of the two extreme networks, used to compute the upper and lower bound estimates of the parameters of interest.

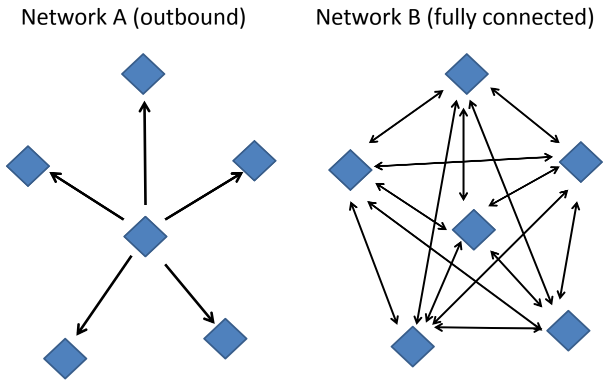

NOTES: The " outbound network" (left) is used to obtain the upper bound estimate. It is a directed network with only "outward bound" links. Holding the number of nodes and the observed ITEs fixed, the social parameter will be estimated to be largest in this type of network. The fully connected network (right), is the benchmark case from which the lower bound of the social parameter can be estimated.

The opposite type of network is a network, where every node is the direct neighbor of every one of its peers[50]. Considering the fully connected network is useful for two reasons: First, the fully linked structure implies that there are only tow types of nodes (treated or not) and that higher order spillovers are the same for every node of the same type. Second, given α and N, the fully connected network has the greatest second and higher order spillovers. Every node affects every other node via a direct link and everybody will get second and higher round spillovers from *every* other node. This allows to derive a closed form solution for the lower bounds of the relevant parameters.

B.4.2 Preliminaries

Before I proceed to characterize the bounds of the coefficient, it is useful to point out a fact that will be important in the argument that follows. First, note that the formulas in equation 31 can be rewritten without explicit characterization of the higher order spills:

$$DiD_0 = \delta_1 + HO_{\ell 0} \tag{32}$$

$$DiD_1 = \frac{\alpha}{NP_{\ell 1}}\delta_1 + HO_{\ell 1} \tag{33}$$

where $HO_{\ell 0} = \delta_1(\alpha^2 G^2_{ii} + \alpha^3 G^3_{ii} + ...)$ and $HO_{\ell 1} = \delta_1(\alpha^2 G^2_{ij} + \alpha^3 G^3_{ij} + ...)$. These effects typically depend on the underlying network of peers and need to take into account the network structure. However, I can use a simple insight concerning the size of the higher

[50] I will sometimes refer to this network as "classroom" network

order effects.

Lemma 1 *Given the total effect, larger higher order effects, imply smaller coefficients, i.e. for* $DiD_0 > DiD_1 > HO^B > HO^A \geq 0$*: for any* $HO^A < HO^B$, $\alpha^A > \alpha^B$ *and* $\delta_1^A > \delta_1^B$.[51]

Proof. We have to make the following two comparisons:

$$\begin{aligned} DiD_0 &= \delta_1^A + HO^A \quad & vs. \quad & DiD_0 = \delta_1^B + HO^B \\ DiD_1 &= \frac{\alpha^A}{NP_{\ell 1}}\delta_1^A + HO^A \quad & vs. \quad & DiD_1 = \frac{\alpha^B}{NP_{\ell 1}}\delta_1^B + HO^B \end{aligned}$$

This can be transformed as follows:

$$\delta_1^A = DiD_0 - HO^A \quad vs. \quad \delta_1^B = DiD_0 - HO^B \tag{34}$$

$$\alpha^A = \frac{(DiD_1 - HO^A)}{\delta_1^A}NP_{\ell 1} \quad vs. \quad \alpha^B = \frac{(DiD_1 - HO^B)}{\delta_1^B}NP_{\ell 1} \tag{35}$$

From equation 34 it is immediately obvious that $HO^A < HO^B$ implies $\delta_1^A > \delta_1^B$. For comparing α substitute the corresponding δ_1 from 34 into 35, define $HO^A := HO^B - \varepsilon$ (for $\varepsilon > 0$) and rewrite equation 35 as

$$\alpha^A = \frac{a}{b}NP_{\ell 1} \quad vs. \quad \alpha^B = \frac{a-\varepsilon}{b-\varepsilon}NP_{\ell 1} \tag{36}$$

defining $a = (DiD_1 - HO^A)$ and $b = DiD_0 - HO^A$. Comparing α^A $vs.$ α^B is equivalent to comparing $\frac{a}{b}$ $vs.$ $\frac{a-\varepsilon}{b-\varepsilon}$. Since we have $a, b, \varepsilon > 0$, $\varepsilon < b$ and $\varepsilon < a$:

$$\begin{aligned} \frac{a}{b} - \frac{a-\varepsilon}{b-\varepsilon} > 0 \quad &\Leftrightarrow \quad a(b-\varepsilon) - b(a-\varepsilon) > 0 \\ &\Leftrightarrow \quad a\varepsilon < b\varepsilon \\ &\Leftrightarrow^{b.A.} \quad a < b \end{aligned}$$

The last inequality holds by the initial assumptions, which completes the proof. ■

With this lemma in hand we can now proceed to derive benchmarks (upper and lower bound estimates) for the parameters of interest.

[51] Note that the requirement $DiD_1 > HO^B$ has bite, since it implies $\alpha < 0.5$. This assumption need not be satisfied in all applications, but it applies well to settings where the spills dissipate quickly and to settings where the direct effect on the treated is much larger than on the neighbors ($DiD_0 >> DiD_1$). This is the case in most applications and certainly so in the present one.

B.4.3 Upper Bound: Network without higher order spillovers.

In the "outbound" network higher order spills back to the originating nodes do not exist[52]: $HO_{\ell 0}$ and $HO_{\ell 1}$ would be 0. This is equivalent to assuming:

$$\mathbf{DiD} =^{\mathbf{b.A.}} \overline{\delta_{\mathbf{1}}}\mathbf{D_t}(\mathbf{I} + \overline{\alpha}\mathbf{G} + \mathbf{0} + \mathbf{0} + ...) \tag{37}$$

which is equivalent to having[53]:

$$\begin{aligned} DiD_0 &= \overline{\delta_1} && for\ treated\ L0 - nodes \\ DiD_2 &= 0 && for\ L2 \\ & && ...analogously\ for\ L3\ and\ higher \end{aligned} \tag{38}$$

By Lemma 1 this assumption leads to an upper bound of both the coefficients If all effects are of the same sign and $DiD_0 > DiD_1 > HO > 0$.[54] The difference in differences for a node $\ell 1 \in L1$[55] would simply reduce to:

$$DID_1 = \frac{\overline{\alpha}}{NP_{\ell 1}}\overline{\delta_1} \tag{39}$$

A consistent estimator of $\overline{\delta_1}$ and the observed difference in difference will be enough to estimate $\overline{\alpha}$. In the "outbound network", we can obtain such an estimate from applying the the Difference in Differences estimator on the level of directly treated nodes and then move on to estimate $\overline{\alpha}$:

$$\begin{aligned} \hat{\overline{\delta_1}} &= \widehat{DiD_0} = \Delta\hat{\ell}0 - \Delta\hat{c}0 \\ \hat{\overline{\alpha}} &= \frac{\widehat{DiD_1}}{\widehat{DiD_0}} NP_{\ell 1} \end{aligned} \tag{40}$$

- $\Delta\hat{\ell}0 := \frac{1}{NP_{\ell 0}} * \sum_i (y_{i,\ell 0,t=0} - y_{i,\ell 0,t=1})$
- $\Delta\hat{c}0 := \frac{1}{NP_{c0}} * \sum_i (y_{i,c0,t=0} - y_{i,c0,t=1})$

with the definition of $\widehat{DiD_1}$ and the underlying $\Delta\hat{\ell}1$ and $\Delta\hat{c}1$ paralleling the definition of $\Delta\hat{\ell}0$ and $\Delta\hat{c}0$.

Discussion: The assumption in equation 37 implies no "multiplication-effects" or

[52] Admittedly, in such a network, endogeneity would not be a problem in the first place.

[53] $D_{\ell 0}$ denotes the value of D at the central node, that is related to the focal node.

[54] DiD_0 (DiD_1) denotes the Difference in Differences for treated nodes (neighbors). For the reverse relationships ($DiD_0 < DiD_1 < HO < 0$) the estimate based on assuming an "outward bound" network gives a lower bound, if the effects go in opposite directions, my claims do not necessarily hold and will have to verified by the researcher. Slightly more involved assumptions will be needed.

[55] Which corresponds to an Indirect Treatment Effect or an "Externality"

"feedback-loops" between the nodes.[56] In the light of the formalization presented here, this is a strong assumption. However, in the impact evaluation literature with fixed and stable classroom sizes or villages, this assumption is almost taken implicitely, whenever the researchers report merely the ATE and ITEs. (cf. Angelucci and De Giorgi (2009), Carmi et al. (2012), Dahl et al. (2012), etc. etc.).

Having said that, the upper bound estimator is quite suitable if higher order spillovers are negligible. In what follows I compute the lower bound estimates under the assumption of maximal higher order spillovers. This will give a sense of the maximal size of the bias that might result from assuming away the higher order complexities of a network.

B.4.4 Lower Bound: Network with maximum higher order spillovers.

In this subsection I derive the lower bound estimates under the assumption of a fully connected network. Formally, consider the matrix $\underline{\mathbf{G}}$, that corresponds to a fully connected network:

$$\underline{\mathbf{G}} = \begin{pmatrix} 0 & \frac{1}{N-1} & \frac{1}{N-1} & \cdots & \frac{1}{N-1} \\ \frac{1}{N-1} & 0 & \frac{1}{N-1} & \cdots & \frac{1}{N-1} \\ \frac{1}{N-1} & \frac{1}{N-1} & 0 & \cdots & \frac{1}{N-1} \\ . & . & . & & . \\ . & . & . & & . \\ \frac{1}{N-1} & \frac{1}{N-1} & \frac{1}{N-1} & \cdots & 0 \end{pmatrix}$$

First, observe that all nodes are direct neighbors, i.e. $NP_{\ell 0} = NP_{\ell 1} = NP_{\ell} = N-1$. Next, note that there are only two types of nodes: Directly treated nodes and neighbors. Let us now characterize the higher order spillovers that arrive at the treated node. From equation 31 we know that the spillovers that arrive at a node in L0 are given by:

$$\ell 0: \; DiD_0 = \delta_1 (1 + 0 + \alpha^2 G_{ii}^2 + \alpha^3 G_{ii}^3 + ...)$$

The formula above points out that no spillovers of order 1 arrive at the treated node, since i does not link on to himself.[57] But in a network characterized by $\underline{\mathbf{G}}$, (and maintaining local treatment) the second order spillovers arrive from *every* neighbor, i.e. NP_{ℓ} times, third order spillovers arrive $(N-1)^2 - (N-1)$ times etc.[58] The number of channels for

[56] 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. This is possible if, for example, spillovers are slow or if the temporal structure of the available data is fine grained enough.

[57] Note that this is precisely the point where the local treatment assumption is most useful, because had we treated T > 1 nodes, then we would have to count T-1 direct spillovers that arrive at i, which obviously would render the following considerations less tractable.

[58] Counting the number of channels for third and higher order spillovers is a matter of combinatorics: The number of channels for higher order increases at an almost exponential rate, leading to potentially very large effects, that are moderated only by the decrease of the primary effects during transmission.

spillovers of order S is given by:

$$\begin{aligned}\#channels_{ii,S} &= (N-1)^{S-1} - (N-1)^{S-2} + (N-1)^{S-3} + ... \\ &= \sum_{s=1}^{S-1} (N-1)^s (-1)^{(S-1)-s} \qquad S \geq 2\end{aligned}$$

The sum of second and higher order spillovers arriving at the treated node is:

$$\begin{aligned}HO_{ii} &= \sum_{S=2}^{\inf} \delta_1 \frac{\alpha^S}{(N-1)^S} \#channels_{ii,S} \\ &= \sum_{S=2}^{\inf} \delta_1 \frac{\alpha^S}{(N-1)^S} \sum_{s=1}^{S-1} (N-1)^s (-1)^{(S-1)-s}\end{aligned}$$

All non-treated neighbors are the same and the number of channels for spillovers of order S from node i to node j is computed almost[59] in the same way:

$$\begin{aligned}\#channels_{ij,S} &= (N-1)^{S-1} - (N-1)^{S-2} + (N-1)^{S-3} + ... \\ &= \sum_{s=0}^{S-1} (N-1)^s (-1)^{(S-1)-s} \qquad S \geq 2\end{aligned}$$

Again the sum of second and higher order spillovers at the neighboring nodes is:

$$\begin{aligned}HO_{ij} &= \sum_{S=2}^{\inf} \delta_1 \frac{\alpha^S}{(N-1)^S} \#channels_{ij,S} \\ &= \sum_{S=2}^{\inf} \delta_1 \frac{\alpha^S}{(N-1)^S} \sum_{s=0}^{S-1} (N-1)^s (-1)^{(S-1)-s}\end{aligned} \tag{41}$$

Before we can move on to derive the lower bound estimates, note that we have $\sum_{s=1}^{S-1}(N-1)^s(-1)^{(S-1)-s} < (N-1)^{S-1}$ which will be a convenient fact for simplifying the estimation of the lower bound.

$$\begin{aligned}HO_{ii} &= \sum_{S=2}^{\inf} \delta_1 \frac{\alpha^S}{(N-1)^S} \sum_{s=1}^{S-1} (N-1)^s (-1)^{(S-1)-s} < \\ &< \sum_{S=2}^{\inf} \frac{\alpha^S}{(N-1)^S} (N-1)^{S-1} = \\ &= \frac{1}{(N-1)} \sum_{S=2}^{\inf} \alpha^S = \frac{\alpha^2}{(N-1)} \frac{1}{1-\alpha}\end{aligned} \tag{42}$$

Let us call this expression $\overline{HO_{ii}}$. Analogously we obtain $\overline{HO_{ij}} = \frac{\alpha^2}{(N-1)} \frac{1}{1-\alpha}$. These values can now be used in the equations 32 and 33 from above (and rewritten here for

[59] s now starts at 0.

convenience):

$$
\begin{aligned}
DiD_0 &= \delta_1 + HO_{\ell 0} \\
DiD_1 &= \frac{\alpha}{NP_{\ell 1}}\delta_1 + HO_{\ell 1}
\end{aligned}
$$

With lemma 1 in hand we can plug in the upper bounds that we derived for higher order effects in a fully connected network to back out the lower bounds of the coefficients α and δ_1, which are characterized by the following two equations.

$$
DiD_0 = \underline{\widehat{\delta_1}} + \overline{HO_{\ell 0}} \tag{43}
$$

$$
DiD_1 = \frac{\underline{\widehat{\alpha}}}{NP_{\ell 1}}\underline{\widehat{\delta_1}} + \overline{HO_{\ell 1}} \tag{44}
$$

It is somewhat tedious, but straight forward to show, that solving this system of equations results in a quadratic equation for $\underline{\widehat{\alpha}}$:

$$
\underline{\widehat{\alpha}}^2 - [\frac{DiD_0}{DiD_1} + (N-1)]\underline{\widehat{\alpha}} + (N-1) = 0 \tag{45}
$$

The closed form solution for $\underline{\widehat{\alpha}}$ is hence given by:

$$
\underline{\widehat{\alpha_{1/2}}} = \frac{1}{2}[\frac{DiD_0}{DiD_1} + (N-1)] + / - \sqrt{\frac{1}{4}[\frac{DiD_0}{DiD_1} + (N-1)]^2 - (N-1)} \tag{46}
$$

It is easy to see that under weak regularity conditions[60] one solution is above 1 and another one between 0 and 1. The latter one is the solution for $\underline{\widehat{\alpha}}$ and it can easily be used to retrieve $\underline{\widehat{\delta_1}}$ from equation 32

Discussion: Note that this closed form solution requires only the number of nodes, and the two estimates from the Difference in Differences (for treated nodes and neighbors). It can be computed when nothing is known about the network, except how many agents and who was treated. It is thus as readily available as the upper bound estimators.

Clearly, one would immediately wish for more[61]. Also, having more information about the network structure or even the link strength between nodes is certainly desirable and, generally, will allow for more interesting additional results. Finally, while the proof here advantageously uses the local treatment assumption, I conjecture, that it is straightforward to extend it to treatments of more than one node.

[60] $DiD_0 > DiD_1$, which is to be expected for most treatments and follows from $\alpha < 0.5$ and $N > 1$

[61] Note that if there is reason to believe that α is greater than 0.5 an analogue of Lemma 1 that relaxes my assumption of $\alpha < 0.5$ is required.

C Aside: Reaction to treatment of the neighbor

Everything above was derived under the assumption that nodes do not observe or at least do not react to the local treatment of their neighbors. This is appropriate for neighbors of Wikipedia articles that get advertised on the start page.[62] In general however, subjects might observe treatment of their neighbors and react to the fact.

An example are children at school, who get annoyed or jealous when their peer was treated in a nice way and they were not.[63] In such situations the students/villagers might react to *merely observing* the treatment of their neighbors by selecting a different value for the outcome variable. To model such a situation we need to further augment the model in equation 18 by both the observable treatments (shocks) that are locally applied, and a term that captures the possible reaction to the treatment of the neighbor.

$$y_{it} = \alpha \frac{\sum_{j \in P_{it}} y_{jt}}{N_{P_{it}}} + X_{it}\beta + \gamma \frac{\sum_{j \in P_{it}} X_{jt}}{N_{P_{it}}} + \delta_1 D_{it} + \delta_2 \frac{\sum_{j \in P_{it}} D_{jt}}{N_{P_{it}}} + \epsilon_{it} \tag{47}$$

Where δ_1 measures the direct treatment effect and the new coefficient δ_2 measures reactions of the node, when it "observes" treatment of one (or several) of its peers. Consider again two connected nodes, where one is treated ($\ell 0$) in period t and the neighbors are not treated ($\ell 1 \in L1$). Assume for simplicity that $\ell 0$ is the only treated node in $\ell 1$'s neighborhood. Similarly, but different, we have:

$$\ell 0 :: y_{\ell 0t} = \alpha \frac{\sum_{j \in P_{\ell 0t}} y_{jt}}{N_{P_{\ell 0t}}} + X_{\ell 0t}\beta + \gamma \frac{\sum_{j \in P_{\ell 0t}} X_{jt}}{N_{P_{\ell 0t}}} + \delta_1 1 + \delta_2 \frac{\sum_{j \in P_{\ell 0t}} 0}{N_{P_{\ell 0t}}} + \epsilon_{\ell 0t} \tag{48}$$

$$\ell 1 \in L1 :: y_{\ell 1t} = \alpha \frac{y_{\ell 0t} + \sum_{j \in P_{\ell 1t}/\ell 0} y_{jt}}{N_{P_{\ell 1t}}} + X_{\ell 1t}\beta + \gamma \frac{\sum_{j \in P_{\ell 1t}} X_{jt}}{N_{P_{\ell 1t}}} + \delta_1 0 + \delta_2 \frac{1 + \sum_{j \in P_{\ell 1t}/\ell 0} D_{jt}}{N_{P_{\ell 1t}}} + \epsilon_{\ell 1t} \tag{49}$$

Now we get two types of spillover effects in this model: First the "pure spillover" α, due to the effect of treatment on the outcome of $\ell 0$. But second, also the "behavior change" of the node, δ_2, when it "observes" treatment of its peer kicks in.

Applying a Difference in Differences strategy alone will measure the joint effect of these two "spillovers". It will not identify α seperately, unless δ_2 is believed to be 0. If this assumption is not warranted only the total "treatment-of-peer"-effect can be measured. Depending on the application we might care about the effect of treatments, in which case this aggregate effect will be interesting. It is simply important to be aware that it is not possible to identify the pure spillover effect in such a setting.

[62] For two reasons: (i) Wikipedia articles cannot react and (ii) the advertisement is not associated with any changes in the real world, so there is no reason for any updates.

[63] Other examples entail economic agents in a village, who observe that their neighbor was refused a social service for failure to comply with a requirement (e.g. sending their kids to school) or commuters in a city, who observe when their friends got caught (after the local transport authority increased the frequency of controls and the punishment for failure to present a valid ticket).