Tilburg University

Will they stay or will they go? How network properties of WebICs predict dropout rates of valuable Wikipedians

Lerner, J.; Kenis, P.N.; van Raaij, D.P.A.M.; Brandes, U.

Published in:
European Management Journal

Publication date:
2011

Document Version
Publisher's PDF, also known as Version of record

[Link to publication in Tilburg University Research Portal](#)

Citation for published version (APA):
Lerner, J., Kenis, P. N., van Raaij, D. P. A. M., & Brandes, U. (2011). Will they stay or will they go? How network properties of WebICs predict dropout rates of valuable Wikipedians. *European Management Journal*, 29(5), 404-413.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

journal homepage: www.elsevier.com/locate/emj

Will they stay or will they go? How network properties of WebICs predict dropout rates of valuable Wikipedians

Jürgen Lerner ^c, Patrick Kenis ^{a,*}, Denise van Raaij ^b, Ulrik Brandes ^c

^a Antwerp Management School, Sint-Jacobsmarkt 9-13, 2000 Antwerpen, Belgium

^b Department of Organization Studies, Tilburg University, P.O. Box 90153, NL-5000 LE Tilburg, The Netherlands

^c Department of Computer and Information Science, University of Konstanz, P.O. Box 67, 78457 Konstanz, Germany

KEYWORDS

Online communities;
Governance;
Dropouts;
Network-analysis;
Reputation;
Controversies

Summary This paper contributes to our understanding of an increasingly prevalent work system, web-based internet communities (WebICs). We are particularly interested in how WebICs are governed given the fact how different they are compared to more classical forms of organization. We study the governance of a WebIC by studying the structure and dynamics of their edit network. Given the fact that the edit network is a relational structure, social network analysis is key to understanding these work systems. We demonstrate that characteristics of the edit network contribute to predicting the dropout hazard of valuable WebIC members. Since WebICs exist only thanks to the activity of their contributors, predicting drop-outs becomes crucial. The results show that reputation and controversy have different effects for different types of Wikipedians; i.e., an actor's reputation decreases the dropout hazard of *active* Wikipedians, while participation on controversial pages decreases the dropout hazard of *highly active* Wikipedians.

© 2011 Elsevier Ltd. All rights reserved.

Introduction

Although work systems that are not based on bureaucratic authority have been around a long time (Coleman, 1982), recently several scholars have started to study alternative types of work systems (Marsden, 2005; Sinha & Van de Ven, 2005). One such form that receives increased attention by organization scholars is what we call web-based information communities (WebICs). WebICs comprises work systems

facilitated by the Internet infrastructure and composed of actors voluntarily attempting to create a product or service that is distributed by some type of open source license (Van Raaij, Brandes, Kenis, & Lerner, 2008). Although organization scholars have so far been particularly interested in WebICs that produce software, for instance the Linux kernel (e.g., Kuk, 2006; Lee & Cole, 2003; O'Mahony & Ferraro, 2007), WebICs often produce a wide array of other types of products, such as information services (e.g., Psychnet; Wikipedia), portals & management systems (e.g., Sahana)

* Corresponding author. Tel.: +32 3 2654734.

E-mail address: Patrick.Kenis@ams.ac.be (P. Kenis).

and tools used in plant genetic engineering that can contribute to for instance drugs development (e.g., Cambia). Consequently, WebICs can potentially create societal value (Benkler, 2003; Watson et al., 2005).

Given the fact that such work systems are more and more prevalent and given the fact that they function different from classical forms of organizations an important question becomes how they are actually governed. They are neither governed by hierarchical mechanisms nor by market mechanism (Kenis & Provan 2006, p. 233). They are also not governed by scientific experts (e.g., Giles, 2005), or by the members of the WebICs themselves (e.g., Brandes, Kenis, Lerner, & van Raaij, 2009a; Stein & Hess, 2007). It appears that WebICs are, in the first place, governed by those who decide to contribute and by their acts of contribution. WebICs live, and are thus governed, by the activity of those who contribute to the work system. In order to better understand the governance of WebICs, the question then becomes how to operationalize this type of governance and how to study whether the governance structure and dynamics relates to a vital characteristic of a WebIC. That is what we propose to contribute in this paper. First, we will argue that network analysis is a promising way to describe the governance structure and dynamics in a WebIC and secondly, we will demonstrate that these governance characteristics are capable of explaining a vital phenomenon in WebICs, i.e. the dropout hazard of valuable WebIC members. Being able to explain dropout hazard of WebICs members is crucial given the fact that WebICs exist only thanks to the activity of their contributors.

The research field we have chosen for identifying and explaining the quality of WebICs is the English version of Wikipedia (http://en.wikipedia.org/wiki/Main_Page). In particular, we study the relational structure of the *edit-networks* associated to *individual* Wikipedia entries (i.e., a single page referring to a specific encyclopedic topic, e.g., *gun politics* [http://en.wikipedia.org/wiki/Gun_politics]). The edit-network associated with a Wikipedia page has as actors the authors of the page and encodes how actors edit the page and how they respond to the edits of others. Among other things, this information can be used to determine the role of authors (e.g., to discriminate between those who provide content and those who delete content), to determine authors that erase or defend each other's content most actively, or to determine whether the community decomposes in different groups of opinion. Consequently, we explicitly conceptualize the work system of a WebIC, or in this particular case, a singular Wikipedia entry as a relational production system and thus we can describe and analyze it with the help of social network analysis (for the importance of studying the relational structure see Ganley & Lampe, 2009; Monge & Contractor, 2003). We will not just describe the structural properties of WebICs but we will also spell out its governance function by demonstrating that the structure helps to explain dropout rates. To date, organizational scholars have hardly studied the organizing processes that take place at page level. In contrast, more formal governance that is executed by *Wikipedia administrators* (e.g., formal rules, punishments and the development of a two-tier governance structure distinguishing 'stable' and 'live' versions of entries) have been studied in order to better understand how WebICs differ from other

types of governance structures (Garud, Jain, & Tuertscher, 2008; Giles, 2005) such as inter organizational networks or hierarchies (Demil & LeCocq, 2006). In contrast we will study the relational and network properties of WebICs in order to better understand their functioning.

Theoretical considerations

The value of studying drop out from a relational perspective has already been demonstrated by others. For example, Krackhardt and Porter have emphasized the impact of relational factors (e.g., colleagues leaving) on turnover (e.g., Krackhardt & Porter, 1985, 1986), and currently, the effect of networks of people leaving one company for another, is taken into account in popular and scientific literature (e.g., Dess & Shaw, 2001). Looking at it from the other side, i.e., why employees' stay, Mitchell and colleagues introduced the construct job embeddedness (2001). Results suggested that non-attitudinal factors such as how well one fits and is linked with ones' job and the external community, and which sacrifices are perceived when leaving a job, determine staying or leaving. Moreover, it shows that on-the-job external activities (e.g., relationships with colleagues) and off-the-job external activities (e.g., commitment to community, work-life balance) influence people to stay.

Also computer scientists report on the relationship between relational features of the work system and continued activity. Lento, Welser, Gu, and Smith (2006) showed that interacting with active actors is a strong and significant predictor of continued user activity in Weblogs. Also earlier work on the causes of dropout from Wikipedia suggests the relevance of relational explanations; i.e., the process of getting feedback on user talk pages¹ and replying on feedback as a Wikipedian predicts dropout hazards. Getting feedback increased the hazard to dropout, while replying to feedback attenuated this effect (Brandes, Kenis, Lerner, & van Raaij, 2009b).

These relational predictors of turnover are closely related to the more informal, and self-regulating nature of WebIC governance, which is characterized by discretionary feedback, peer review and community norms (Moon & Sproull, 2000), and substitutes economic means to influence employees levels of commitment or satisfaction that are appropriate in hierarchical organizations (Ganley & Lampe, 2009; Moon & Sproull, 2008). Based on these results from the employee turnover literature and empirical studies in computer science, we will explore whether dropout is empirically related to the following three relational properties: *actor reputation*, *actor reputation increase* and *participation in controversial pages*. Before introducing the rational for these relationships we will introduce the definition of the different concepts used in this study.

Dropout Hazard of valuable Wikipedians: In WebICs many 'eyeballs' are needed for contributing and improving content (e.g., Garud et al., 2008; Raymond, 1999). This is similar to the observation in organizations in which employee turnover is considered an issue affecting every type of organization (Beadles, Lowery, Petty, & Ezel, 2000), having an (negative) impact on organizational performance (Dess & Shaw, 2001;

¹ User talk pages are discussion pages related to an editor (see http://en.wikipedia.org/wiki/Talk_pages).

Mitchell, Holtom, Lee, & Erez, 2001; Morrell, Loan-Clarke, & Wilkinson, 2001), particularly if key people with scarce skills and critical knowledge dropout (Dess & Shaw, 2001; Mitchell & Lee, 2001). Yet, whereas the original organizational turnover research started from the assumption that turnover is negative (Beadles et al., 2000, p. 331), others have questioned this (e.g., Abelson & Baysinger, 1984; Beadles et al., 2000; Siebert & Zubanov, 2009), and suggested to distinguish between functional and dysfunctional turnover (Dalton, Todor, & Krackhardt, 1982). This idea was tested by several others, who conclude that: "the basic premise of turnover functionality is the explicit recognition that different employees are of different value to the organization" (Beadles et al., 2000, p. 332). It seems appropriate also in the case of WebICs to distinguish between functional and dysfunctional turnover (Morrell et al., 2001). Consequently, we assume that when highly valued Wikipedians quit an entry (i.e., a single Wikipedia page) the quality of this entry will decrease, while in case hardly valued Wikipedians leave the quality will increase.

The challenge is thus to distinguish valuable from low valuable actors, which by definition is a subjective judgment (Dalton et al., 1982). For hierarchical organizations, so-called 'hard' measures (e.g., productivity, sales) and 'soft' measures (e.g., supervisory appraisals, self-perceptions) have been applied. However, these type of data are unavailable for WebICs and are, moreover, unlikely to be appropriate since WebICs differ from hierarchies (Ganley & Lampe, 2009; O'Mahony & Bechky, 2008)². Also consistent with our two starting points (i.e. using data readily available in Wikipedia and relational data) we distinguish valuable from less valuable WebIC members by looking at their activity levels; i.e., (highly) active Wikipedians are considered as more valuable compared to less active Wikipedians, because they are more involved in providing and assessing encyclopedic content. Consequently, we only study (highly) active Wikipedians since we assume that only their dropout negatively affects the quality of a Wikipedia entry.

Network properties of a WebICs as predictors of dropout of valuable Wikipedians: Markus defines WebIC governance as follows: "the means of achieving the direction, control, and coordination of wholly or partially autonomous individuals and organizations on behalf of an OSS[Open Source Software; added by authors] project to which they jointly contribute" (Markus, 2007, p. 152). Her summary of appropriate governance mechanisms illustrates the importance of relational or network properties, e.g., peer reviewing as a monitoring mechanism and reputation signaling practices rather than economic incentives. It is from such a perspective that we analyze whether the factors actor reputation, actor reputation increase and participation in controversial pages are related to drop-out. We certainly do not claim that these are the only three factors which are related to the quality of WebICs but consider them for the moment being helpful in assessing whether dropout hazard

of Wikipedians is a good indicator of the quality for WebICs. In any case, a clear advantage is that they all can be calculated from information which is readably available online.

Actor reputation and actor reputation growth

According to Campbell (1960 in: Lee & Cole, 2003) reputation mechanisms contribute to preserving and reproducing selected variations by keeping actors motivated to continue to contribute. This is in line with former research on participation to WebICs. Also in this body of research reputation is referred to as a motivating factor (Kuk, 2006, see also Ganley & Lampe, 2009; Hertel, Niedner, & Herrmann, 2003; Lakhani & von Hippel, 2003; Shah, 2006; von Hippel, 2005) and contributes to the management of intellectual capital. Lee and Cole argue that an actor's reputation in a network is based on (1) the appropriateness of the solutions added by an actor and (2) the quality of critical evaluations provided by the actor. What is emphasized here is the quality of contributions. Reputation is clearly a relational construct because reputation and reputation growth depend on the behavior and attitudes of others (e.g., Ganley & Lampe, 2009). Based on earlier results that demonstrate that WebIC members start to participate in order to increase their reputation, we propose that as a node's reputation on the one hand, and the growth of a node's reputation (i.e., developing in a positive direction from low to higher or vice versa in a negative direction) on the other hand decreases the chance the node drops out. Thus:

H1. The higher the reputation of a WebIC member, the lower the chance the WebIC member drops out.

H2. The higher the reputation growth of a WebIC member, the lower the chance the WebIC member drops out.

Controversies

Controversies within (online) work groups can have a positive or a negative effect on their quality. For one thing, literature on the development of knowledge suggests that creative solutions are built from the recombination of existing ideas (Hargadon & Bechky, 2006), and that variation is a necessary condition for the quality of knowledge creation (Campbell in: Lee & Cole, 2003). This also implies that, at least to some degree, controversies are required. Moreover, challenging the ideas of others, also referred to as problem solving dialogues between diverse subgroups or experts who bring in exclusive knowledge (Alavi & Leidner, 2001) is suggested to be an important feature of WebICs that are able to produce high quality outcomes. Yet, the question remains whether controversy motivates or demotivates voluntary WebIC participants? Research on why people are motivated to contribute for free demonstrates that users participate because they enjoy working together (Lerner & Tirole, 2005; Roberts, Hann, & Slaughter, 2006; Shah, 2005) i.e., the process of working together (Füller, Bartl, Ernst, & Mühlbacher, 2006a), the fun experienced when working together (Franke & Shah, 2003; Füller, Jawecki, & Mühlbacher, 2006b; Hienerth, 2006; Shah, 2006; von Krogh & Speath, 2007), and the enjoyment of the work itself (Lakhani & von Hippel, 2003) stimulates people to volunteer. Additionally,

² For instance, the fact that WebIC participants are volunteers without formal labor contracts and that WebICs are open systems in which open source licenses are used (e.g., GNU General Public License) (Lerner & Tirole, 2001; Schweick, Evans, & Grove, 2005; von Krogh & von Hippel, 2003).

Lettl, Herstatt, and Gemuenden (2006) showed that embeddedness in a supportive environment has a positive effect on user participation. An effect of low levels of controversies on user participation can also be deduced from results showing that rivalry between users demotivates WebIC members (Franke & Shah, 2003; Hienerth, 2006; von Hippel, 2002). Therefore, we propose that if others disapprove a member's contributions, less community support is experienced and the joy and fun of the working process itself decreases. This is particularly the case for WebICs in which a lot of conflict is going on. Therefore, the following hypothesis is proposed:

H3. The more often a WebIC member participates in WebICs characterized by high levels of controversies, the higher the chance the WebIC member drops out.

Methods of analysis

In order to better understand whether a WebIC's network characteristics predicts Wikipedians to drop out, a network-analytic and longitudinal perspective will be applied. We will study how the edit-network evolves over time in terms of actor reputation, actor reputation growth, and participation in controversial pages. To explore whether these patterns of edit-network evolution estimate the chance to dropout, lifetime-analysis is applied. Earlier (Brandes et al., 2009b) a statistical model for the dropout hazard (i.e., the conditional probability of users dropping out at time t , under the precondition that they survived up to t) was developed. The dropout rate was modeled in the following functional form:

$$h_u(t) = h_u(W_t; \theta) = \exp\left(\sum_{i=1}^k \theta_i \cdot s_i(u; W_t)\right)$$

Here the dropout hazard of user u at time t is assumed to be explained by statistics $s(u, W)$ that describe specific aspects of the Wikipedia history W at time t from the point of view of user u . (As an example, a specific statistic $s(u, W)$ might encode the reputation of u at time t .) The estimated parameters θ give information about whether a hypothetical cause of dropout (e.g., reputation) shows the predicted effect. If the parameter associated with the reputation statistic for example, is significantly negative, then users with higher reputation would have lower chance to drop out – pointing to the fact that users might be motivated by reputation. A significantly positive parameter would indicate that the associated statistic encodes a cause of increased dropout. Details of how the parameters can be estimated from an observed set of dropouts and survivors are given in Brandes et al. (2009b).

Sample strategy valuable and very valuable users

Determining the “lifetimes” of Wikipedians is not straightforward – it is even not obvious how to decide whether a certain Wikipedian “died” or whether she or he will resume editing at a later time. We restricted the analysis to active users, i.e., users with a specified minimum number of edits, in order to concentrate on Wikipedia's most valuable members, and since it is unclear whether a user with only, say,

two edits is a dropout or just rarely involved. To determine active users and divide them into dropouts and survivors we used the revision history of the English Wikipedia dating October 10, 2008, see <http://download.wikimedia.org/>. A user was called active if he did at least 1000 edits before July 1, 2008, and a user was called highly active if he did at least 10,000 edits before July 1, 2008. Out of several million registered users (excluding Bots; i.e., software programs that perform routine tasks), 17,714 qualified as active and 2463 as highly active. The time period from July 1, 2008 until October 8, 2008 was used to split (highly) active users into dropouts and survivors. A user was called a survivor if he did at least one edit after July 1, otherwise he was classified as dropout. This criterion resulted in 13,126 active survivors, 4588 active user dropouts, 2130 highly active survivors, and 333 highly active users who dropped out. The definition of active users (i.e. a requirement of 1000 edits) was taken from the standard set by Wikipedia (see the Missing Wikipedian page: http://en.wikipedia.org/wiki/Wikipedia:Missing_Wikipedians).

Edit-network structure

Indicators for an actor's reputation, an actor's reputation increase and participation in controversial pages can be derived from the edit-history of a Wikipedia page or what we call the edit-network (see Figure 1). In Brandes et al. (2009a) it is demonstrated that the edit-network associated with a Wikipedia page has as nodes (what we called actors before) the authors of the page and encodes how authors contributed to the page and how authors interacted with each other while editing the Wikipedia page. This information is computed from the complete history of the specific Wikipedia page; i.e., from the sequence of its revisions, by determining which part of the text has been added, has been deleted, or remained unchanged when going from one version of the page to the next.

Actor reputation

A Wikipedian's reputation refers here to the appropriateness of her or his solution added perceived by a page's co-authors. This idea is somewhat similar with the work of Adler and Alfaro (2007), who for developing a reputation system for Wikipedia, considered whether text edited by a user persists on the page or gets deleted afterwards. We will measure the reputation of user u based on the ratio of persistent text over all text added (in network analysis this can

Figure 1 Illustrative edit-network resulting from actor A adding new text to a page which is later deleted by B and subsequently restored by C. Note: The deletion of A's text by B, as well as the undoing of B's edit by C, are interpreted as disagreements (red ties); the restoring of A's text by C is interpreted as agreement (blue tie). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Figure 2 Aggregated network structure of selected pages. The left example showing a low level of bipolarity, while the right example shows a high level of bipolarity (the ellipses denoting actors and the ties showing disagreement relationships).

be considered a centrality measure). Since the number of words that are added by u , as well as the number of persistent words, can change from one edit to the next we get a time-varying measure of reputation for each time point t .

$$reputation(u, t) = [a(u, t) - d(u, t) + r(u, t)] / a(u, t)$$

$a(u, t)$ being the number of words added by user u before or at t ; $d(u, t)$ being the number of words added by u and deleted afterwards (by any user) before or at t ; $r(u, t)$ being the number of words added by u , deleted afterwards and restored later (by any user) before or at t .

Consequently $a(u, t) - d(u, t) + r(u, t)$ is the number of words added that are either not deleted or, if deleted, then restored and $reputation(u, t) = [a(u, t) - d(u, t) + r(u, t)] / a(u, t)$ is the ratio at time t of those persistent words over the number of all added words being a centrality measure with a possible score between 0 and 1.

Actor reputation increase

It might be the case that not only the absolute reputation is a predictor for dropout, but also the increase/decrease of reputation. For instance, if two users u and v both have on a particular day³ a reputation score of (say) 0.5 but u 's reputation increased recently from 0.3 to 0.5 whereas v 's reputation decreased from 0.7 to 0.5, then u might perceive his reputation score as satisfying/motivating while v considers the same value as a setback. To test whether such effects can be observed we define the reputation-increase of user u at time t with respect to a given time lag dt to be the difference.

$$reputation\ increase(u, t) = reputation(u, t) - reputation(u, t - dt)$$

Reputation increase is thus normalized to the interval $[-1.0, 1.0]$; it is positive if reputation increased and negative if reputation decreased.

³ It would not make a substantial difference for this particular analysis whether we operationalize reputation increase based on the increase during one day, a week (or a month) due to technical reasons: the reputation increase within one week or month is the sum of the daily increase within this week. So the weekly increase is high if and only if it is high for (some of) the days of that week. From another point of view, the weekly increase is the average of the daily increases – up to a constant factor that changes nothing.

Participation in controversial pages

First, we operationalized controversial pages in terms of the edit-networks levels of bipolarity. Bipolarity characterizes the global collaboration structure and estimates whether the author community of an individual Wikipedia entry decomposes into two groups of opinion that mutually undo the edits of each other (see Figure 2 for two examples). We define bipolarity as follows (see Brandes et al., 2009a, also for visualization on high and low bipolar edit-networks):

$$Bipolarity = \frac{w - c}{w + c}$$

The author community was partitioned into two groups, A and B, such that the aggregated weight of negative edges between the two groups is maximized. Denote this weight with w and let c denote the aggregated weight of negative relationships connecting authors within A or within B. The bipolarity score lies between -1 and $+1$. It is equal to $+1$ if the graph (e.g., edit-network) is bipartite, i.e., edges connect only members from different groups and, therefore, the division into opposing groups is perfect. The bipolarity score equals 0 if all pairs of actors are connected, and all edges, including loops have the same weight-indicating no opposition at all. Negative weights of the bipolarity are theoretically possible but unlikely and have never been observed in our tests. In fact, it would indicate that all authors delete mostly their own edits but rarely the edits of any other author. Since the bipolarity (controversy) level of a page may change over time, we computed the bipolarity of a given page in separate time intervals of half a year. In the following, let $bipolarity(p, t)$ denote the bipolarity of page p in the half-year interval containing time point t . Starting from bipolarity as an operationalization of controversy, the statistic $sum_controversy(u, t)$ is defined to be the sum over the bipolarity of pages edited by user u before or at time t (where the bipolarity of a page p is counted as often as u has edited p).

Number of edits

Since $sum_bipolarity(u, t)$ is, by design, monotonically increasing with the number of edits done by u , a positive (respectively negative) parameter associated with

Table 1 Estimated reputation parameter, standard errors and *t*-ratios for dropout chance of (highly) active users.^a

Statistic	Active users (at least 1000 edits)		Highly active users (at least 10,000 edits)	
	Parameter (s.e.)	<i>t</i> -Ratio	Parameter (s.e.)	<i>t</i> -Ratio
Reputation	−0.19300****(0.015)	12.87	−0.140 (0.147)	0.95

**** = parameter significantly different from zero at 0.1%.

^a The constant just normalizes the model to the empirical time scale in which one unit corresponds to the expected time-to-dropout of a (hypothetical) user for which the effects of all other statistics add up to zero. Since the value of the constant does not provide much information, we do not report on its values.

Table 2 Estimated reputation parameter, reputation increase parameter, standard errors and *t*-ratios for dropout chance of (highly) active users.

Statistic	Active users (at least 1000 edits)		Highly active users (at least 10,000 edits)	
	Parameter (s.e.)	<i>t</i> -Ratio	Parameter (s.e.)	<i>t</i> -Ratio
Reputation	−0.195 (0.015)****	13	−0.154 (0.148)	1.04
Reputation increase	0.093 (0.166)	0.56	1.620 (0.976)	1.65

**** = parameter significantly different from zero at 0.1%.

participation to controversial pages could just be due to the possibility that users wear out over time (respectively get more robust against dropout over time). Thus, we included the number of edits as a control variable when analyzing the effect of participation to controversial pages. To define the statistics encoding how much a particular user u contributed to the edit-network up to a time point t , let $E_{u,t}$ denote the set of revisions that u performed on pages of the main namespace on or before time t . The respected statistic is defined by:

$$\text{Edit}(u; W, t) = \text{Log}(1 + |E_{u,t}|)$$

The logarithmic scaling of the number of revisions has been chosen due to the extremely skewed distribution (there are users who perform more than 100.000 revisions, while most of the selected users have a count of only slightly more than 1000). A significant positive (negative) parameter associated with edits implies that users with higher numbers of revisions to the main namespace have a higher (or lower) hazard to drop out.

Since it is computationally infeasible to compute the edit-networks of all (more than two million) Wikipedia pages, we approximate the reputation scores in the following way. First we randomly select 200 pages in an importance-driven sampling procedure⁴: a page is selected with probability proportional to the number of its edits that are done by active users. Therefore, pages on which active users are highly involved have a higher probability to be selected. The reputation scores that are computed on the selected pages serve as an approximation of the true

reputation scores. Due to the large N (more than 17,000 users) we strongly believe that this approximation does not severely change our parameter estimates – although the approximate reputation of some individual users might be far of the true value.

Results

In the result section we report on three models that explored whether the network properties studied predict the dropout hazard of (highly) active Wikipedians. The first hypothesis was partly confirmed, the second hypothesis was rejected, and for the third hypothesis we found results contrary to what we had expected.

Reputation

The parameter associated with reputation is significantly negative for active users, and is non-significant for highly active users (see Table 1).

This indicates that the dropout hazard of active users decreases with an actor's higher reputation scores, but that reputation is not a good predictor for the hazard of highly active users to dropout. Apparently, for the highly active users, reputation in terms of the persistence of words is not that important. This could be explained by the different roles that are fulfilled by active versus highly active Wikipedians. If a user does 10,000 edits, and sometimes up to more than 100,000 edits, than he cannot add much text as 10,000 paragraphs or even sentences is an incredible amount of work. Hence, *highly active* Wikipedians mostly do revision work: spell-checking, reverting vandalism, or checking formal requirements on text written by others, while in general *active* Wikipedians are content providers. For the latter, reputation in terms of persisted text is more motivating.

⁴ Contrary to an uniformly random sample which would give the same probability to everyone of 2 million pages. Because most pages are very small, containing very little revisions, an uniformly random sample would decrease our chance to learn something about, for instance, the reputation of (*highly*) active users.

Table 3 Estimated participation on controversial pages parameter, edit parameter, standard errors and *t*-ratios for dropout chance of (highly) active.

Statistic	Active users (at least 1000 edits)		Highly active users (at least 10,000 edits)	
	Parameter (s.e.)	<i>t</i> -Ratio	Parameter (s.e.)	<i>t</i> -Ratio
Participation on controversial page	−0.001 (0.001)	1.00	−0.0058****(0.0015)	3.87
Number of edits	0.00048 (0.001)	0.478	0.0048**** (0.0012)	4.00

**** = parameter significantly different from zero at 0.1%.

Reputation increase

If we add the parameter *reputation increase* to the model, the hazard of (highly) active users to drop out is not predicted (see Table 2).

Hence, reputation increase does not predict either the dropout hazard for active or for highly active Wikipedians. Moreover, if we add reputation increase to the reputation model, the predictive value of reputation only slightly improves for active users. Hence, our second hypothesis is rejected. It could be that for active Wikipedians a certain level of reputation is a sufficient motivator; i.e., that increasing of already high reputation or vice versa does not motivate or demotivate. It could also be the case the reputation increase remains unnoticed.

Participation in controversial pages

The data demonstrate that if highly active users participate in controversial pages, the dropout hazard of highly active users decreases. Hence, participation in controversial pages seems to motivate highly active Wikipedians, but by contrast, has no effect on the hazard that active users will leave (see Table 3).

Hence, hypothesis 3 is rejected. Instead, participation in controversial pages has a positive effect. We checked whether our results could be influenced by our operationalization, by developing three alternative statistics; (1) counting number of edits to pages with bipolarity bigger than 0.7; (2) taking the average bipolarity of pages edited by the user, and; (3) taking the average bipolarity of pages edited today by the user. None of these led to significant results.

Moreover, we included the number of edits in the model to check the robustness of the effect of participation on controversial pages on the dropout chance of (highly) active users. Earlier (Brandes et al., 2009b), it was reported that the number of edits has a negative significant effect on the dropout hazard of active Wikipedians. Without taking the number of edits into account, participation in controversial pages had a significantly negative effect (i.e., more bipolarity – less dropout). The *sum_controversy* statistic (i.e., the statistic referring to the sum over the bipolarity of pages edited by user *u*) has the property that it is monotonically increasing with the number of edits (more edits, higher value of this statistic). So it could be the case that users become just less likely to drop out when they accumulate edits. Indeed, in a model with just the number of edits

and a constant, the number of edits is significantly negative as well. The results presented here demonstrate that the effect of participation in controversial pages is robust for highly active users when both statistics are included; but not for active users. This could be explained by the nature of the product produced; encyclopedic information. Our hypotheses were based on results about WebICs producing software and sport equipments. It could be that these developers are indeed demotivated by rejections of their contributions, while highly active Wikipedians reverting what they see as ‘vandalism’ are motivated by rejections; that is to say, are motivated to reject the contribution of the other again.

Discussion

The results suggest that different network features of WebIC have varying impact on different types of Wikipedians, which potentially differ in their value for Wikipedia entries. This suggests that these online data sources can be helpful for explaining and identifying WebIC quality, although better operationalizations are needed for WebIC quality as will be elucidated in the next section. Yet, the preliminary results suggest that different governance mechanisms can be effective under various conditions. While currently organization scholars have mostly described general features of WebICs (Demil & Lecocq, 2006; O'Mahony & Bechky, 2008), our results suggest that not each feature of a WebIC has the same effect and that this effect depends on the type of WebIC. While studies on software WebICs report on the importance of the the absence of controversies for sustaining volunteers, this does not seem to work in encyclopedic WebICs. By contrast, controversy motivates highly active Wikipedians. Reputation is also generally understood as an effective feature of WebIC governance by motivating users to participate (e.g., Hars & Ou, 2002; Hertel, Niedner, & Herrmann, 2003; Lakhani & von Hippel, 2003; Markus, 2007; Shah, 2000, 2006; von Hippel, 2002). Only a couple of studies argue that reputation is not a significant predictor of participation (Jeppesen & Frederiksen, 2006; Franke & Shah, 2003). Also our research provides nuances to these conclusions by suggesting that reputation is not always effective as a mean of preventing dropout.

These early results could also be relevant for studies on hierarchical organizations and employment turnover. A tendency is noted that hierarchies move towards less formal, hierarchical governance structures (e.g., Sinha & Van de Ven, 2005), for instance when they open up R&D

departments, engage in open innovation processes and crowd surfing (Chesbrough, 2006), and start co-producing with WebICs. Under these circumstances, in which contracts and economic incentives are less appropriate (O'Mahony & Bechky, 2008), organizations need alternatives to governance mechanisms appropriate for hierarchies. Better understanding the effects of WebIC governance, among others, the role of issues such as reputation and controversies, could be helpful. Although it is difficult to generalize on the basis of the preliminary and limited study, it is interesting to learn that phenomena which are generally highly valued in more formal organizational contexts (such as reputation) or generally not valued in these contexts (such as controversies) seem to play a different role in the contexts studied here. In any case, it turns out that the network perspective is a very useful perspective in bridging common and traditional theoretical questions about organizing (e.g. who is in charge and who is motivated to contribute) with new practices. This certainly calls for further research.

Generally speaking, the different effects Wikipedia's network structure has on different types of members and the large amount of fine-grained data available on the Internet, suggests it would be worthwhile to further the effect of other characteristics of the edit network and also to distinguish between functional and dysfunctional dropouts. However, in order to do this, we acknowledge that much more work is needed. For instance, in order to come to conclusions about whether the drop out of a member is dysfunctional, more information about the properties of a dropout are needed; e.g., the role the member fulfilled or the reputation the member has. It would be fruitful to compare active and highly active members that fulfill similar roles. Likewise, information on the reputation of a dropout would make it more likely to draw conclusions about the harmful effect of dropout for Wikipedia.

Another pitfall is the use of the number of edits for defining a Wikipedian's activity level. A Wikipedian who has contributed a very limited number of edits, but whose edits are very lengthy is qualified as less valuable than a Wikipedian doing a lot of small spell-checking corrections. In future research, we will deal with this issue by also considering Wikipedians who did less numbers of edits, but at the same time controlling for the role they fulfill.

What we also demonstrated is that available and embedded online data are useful for explaining dropouts. The results suggest that this kind of data can be used to study the governance of WebICs, although also here a number of improvements are possible. Based on the current results, it can be argued that variations in a WebICs relational structure influence WebICs outcomes in terms of sustaining volunteers. In future research we will broaden the range of WebIC relational features by considering the effect of, for instance, peer review or the presence of generalized exchange. Moreover, edit-networks are just one of the many network structures that can be computed from Internet data. In future research we will explore whether networks with different types of transactional content (e.g., the discussion networks related to Wikipedia's discussion pages) are appropriate for measuring features of WebIC governance and are helpful for predicting WebIC quality.

References

- Abelson, M., & Baysinger, B. (1984). Optimal and dysfunctional turnover: toward an organizational level model. *Academy of Management Review*, 9, 331–342.
- Adler, B.T. & de Alfaro, L. (2007). A content-driven reputation system for the Wikipedia. In *Proceedings 16th Intl. Conf. WWW2007*: 261–270.
- Alavi, M., & Leidner, D. E. (2001). Review: Knowledge management and knowledge management systems: Conceptual foundations and research issues. *MIS Quarterly*, 25, 107–136.
- Beadles, C. M., II, Lowery, N. A., Petty, M. M., & Ezel, H. (2000). An examination of the relationship between turnover, functionality, turnover frequency and organizational performance. *Journal of Business and Psychology*, 15, 331–337.
- Benkler, Y. (2003). Coase's Penguin, or, Linux and the nature of the firm. *The Yale Law Journal*, 112, 369–446.
- Brandes, U., Kenis, P., Lerner, J. & van Raaij, D. (2009a). Network analysis of collaboration structure in Wikipedia. In *Proceedings 18th Intl. World Wide Web Conference (WWW2009)*.
- Brandes, U., Kenis, P., Lerner, J. and van Raaij, D. (2009b). Is editing more rewarding than discussion? A statistical framework to estimate the causes of dropout from Wikipedia. In *Proceedings 1st Intl. Workshop Motivation and Incentives on the web (Webcentives '09) at the 18th Intl. World Wide Web Conference (WWW2009)*.
- Chesbrough, H. (2006). *Open Innovation: The new perspective for creating and profiting from technology*. Cambridge: Harvard Business School Press.
- Coleman, J. S. (1982). *The asymmetric society*. Syracuse: Syracuse University Press.
- Dalton, D. R., Todor, W. D., & Krackhardt, D. M. (1982). Turnover overstated: The functional taxonomy. *Academy of Management Review*, 7, 117–123.
- Demil, B., & Lecocq, X. (2006). Neither market nor hierarchy nor network: The emergence of bazaar governance. *Organization Studies*, 27, 1447–1466.
- Dess, G. G., & Shaw, J. D. (2001). Voluntary turnover, social capital, and organizational performance. *Academy of Management Review*, 26, 496–524.
- Franke, N., & Shah, S. (2003). How communities support innovative activities: An exploration of assistance and sharing among end-users. *Research Policy*, 32, 157–178.
- Füller, J., Bartl, M., Ernst, H., & Mühlbacher, H. (2006a). Community based innovation: how to integrate members of virtual communities into new product development? *Electronic Commerce Research*, 6, 57–73.
- Füller, J., Jaweck, G., & Mühlbacher, H. (2006b). Innovation creation by online basketball communities. *Journal of Business Research*, 60, 60–71.
- Ganley, D., & Lampe, C. (2009). The ties that bind: social network principles in online communities. *Decision Support Systems*, 47, 266–274.
- Garud, R., Jain, S., & Tuertscher, P. (2008). Incomplete by design and designing for incompleteness. *Organization Studies*, 29, 351–371.
- Giles, J. (2005). Internet encyclopedias go head to head. *Nature*, 438, 900–901.
- Hars, A., & Ou, S. (2002). Working for free. Motivations for participating in OSS projects. *International Journal of electronic Commerce*, 6, 25–39.
- Hargadon, A. B., & Bechky, B. A. (2006). When collections of creatives become creative collectives. A field study of problem solving at work. *Organization Science*, 17, 484–500.
- Hertel, G., Niedner, S., & Herrmann, S. (2003). Motivation of software developers in open source projects: An internet-based survey of contributors to the Linux kernel. *Research Policy*, 32, 1159–1177.

- Hienert, C. (2006). The commercialization of user innovations: The development of the Rodeo kayak industry. *R&D Management*, 36, 273–293.
- Jeppesen, L. B., & Frederiksen, L. (2006). Why do users contribute to firm-hosted user communities? The case of computer-controlled music instruments. *Organization Science*, 17, 45–63.
- Kenis, P., & Provan, K. G. (2006). The control of public networks. *International Public Management Journal*, 9, 227–247.
- Krackhardt, D., & Porter, L. W. (1985). When friends leave: A structural analysis of the relationship between turnover and stayer's attitudes. *Administrative Science Quarterly*, 30, 242–261.
- Krackhardt, D., & Porter, L. W. (1986). The snowball-effect: Turnover embedded in communication networks. *Journal of Applied Psychology*, 71, 50–55.
- Kuk, G. (2006). Strategic interaction and knowledge sharing in the KDE developer mailing list. *Management Science*, 52, 1031–1042.
- Lakhani, K. R., & von Hippel, E. (2003). How OSS works: "free" user-to-user assistance. *Research Policy*, 32, 923–943.
- Lee, G. K., & Cole, R. (2003). From a firm-based to a community-based model of knowledge creation: The case of the Linux kernel development. *Organization Science*, 14, 633–649.
- Lento, T., Welser, H. T., Gu, L., & Smith, M. (2006). The ties that blog: examining the relationship between social ties and continued participation in the Wallop weblogging system. In *Proc. 3rd Annual Workshop on the Weblogging Ecosystem*.
- Lerner, J., & Tirole, J. (2001). The open source movement: Key research questions. *European Economic Review*, 45, 819–826.
- Lerner, J., & Tirole, J. (2005). The economics of technology sharing: Open source and beyond. *The Journal of Economic Perspectives*, 19, 99–120.
- Lettl, C., Herstatt, C., & Gemuenden, H. G. (2006). User's contributions to radical innovation: Evidence from four cases in the field of medical equipment technology. *R&D Management*, 36, 251–272.
- Markus, M. L. (2007). The Governance of free/open source software projects: Monolithic, multidimensional, or configurational? *Journal of Management and Governance*, 11, 151–163.
- Marsden, P. V. (2005). The sociology of James S. Coleman. *Annual Review of Sociology*, 31, 1–24.
- Mitchell, T. R., Holtom, B. C., Lee, T. W., & Erez, M. (2001). Why people stay: Using job embeddedness to predict voluntary turnover. *Academy of Management Journal*, 44, 1102–1121.
- Mitchell, T. R., & Lee, T. W. (2001). The unfolding model of voluntary turnover and job embeddedness: Foundations for a comprehensive theory of attachment. *Research in organizational behavior*, 23, 189–246.
- Moon, J. Y., & Sproull, L. S. (2000). The essence of distributed work: The case of the Linux kernel. *First Monday*, 5, http://ftp.firstmonday.dk/www/issues/issue5_11/moon/.
- Moon, J. Y., & Sproull, L. S. (2008). The role of feedback in managing the internet-based volunteer work force. *Information Systems Research*, 19, 494–515.
- Monge, P. R., & Contractor, N. S. (2003). *Theories of communication networks*. Oxford: Oxford University Press.
- Morrell, K., Loan-Clarke, J., & Wilkinson, A. (2001). Unweaving leaving: The use of models in the management of employee turnover. *International Journal of Management Review*, 3, 219–244.
- O'Mahony, S., & Bechky, B. (2008). Boundary organizations: Enabling collaboration among unexpected allies. *Administrative Science Quarterly*, 53, 422–459.
- O'Mahony, S., & Ferraro, F. (2007). The emergence of governance in an open source community. *Academy of Management Journal*, 50, 1079–1106.
- Raymond, E. S. (1999). *The Cathedral and the bazaar: Musings on Linux and open source by accidental revolutionary revised edition*. Cambridge: O'Reilly.
- Roberts, J. A., Hann, I. H., & Slaughter, S. A. (2006). Understanding the motivations, participation, and performance of OSS developers: A longitudinal study of the Apache projects. *Management Science*, 52, 984–999.
- Schweick, C., Evans, T., & Grove, J. M. (2005). Open source and open content: A framework for global collaboration in social-ecological research. *Ecology and Society*, 10, 33 (<http://www.ecologyandsociety.org/vol10/iss1/art33/>).
- Siebert, W. S., & Zubanov, N. (2009). Searching for the optimal level of employee turnover: A study of a large U.K. retail organization. *Academy of Management Journal*, 52, 294–313.
- Sinha, K. K., & Van de Ven, A. H. (2005). Designing work within and between organizations. *Organization Science*, 16, 389–408.
- Shah, S. (2005). Open beyond software. In C. Dibona, D. Cooper, & M. Stone (Eds.), *Open sources 2* (pp. 339–360). Sebastopol: O'Reilly.
- Shah, S. (2006). Motivation, governance, and the viability of hybrid forms in open source software development. *Management Science*, 52, 1000–1014.
- Stein, K., & Hess, C. (2007). Does it matter who contributes? A study on featured articles in the German Wikipedia. In *Proc. 18th ACM Conf. Hypertext and Hypermedia (Hypertext 2007)*, 171–174.
- Van Raaij, D., Brandes, U., Kenis, P., & Lerner, J. (2008). *Determining the quality of WebICs by structural parameters*. Amsterdam: Paper prepared for EGOS conference.
- von Hippel, E. (2002). Open source software projects as user innovation networks. *Open source software: Economics, law and policy*, Toulouse, 20–21.
- von Hippel, E. (2005). *Democratizing innovation*. MIT Press.
- von Krogh, G., & Speath, S. (2007). The OSS phenomenon: Characteristics that promote research. *Journal of strategic information systems*, 16, 236–253.
- von Krogh, G., & von Hippel, E. (2003). Special issue on open source software development. *Research Policy*, 32, 1149–1157.
- Watson, R. T., Boudreau, M.-C., Greiner, M., Wynn, D., York, P., & Gus, R. (2005). Governance and global communities. *International Journal of Management*, 11, 125–142.

Jürgen Lerner has a MSc in Mathematics and a PhD in Computer Science from the University of Konstanz where he currently works on a postdoc position. His research area is the development of methods for social network analysis and visualization and his work is characterized by tight collaboration with social scientists. Application areas include political event networks, networks with negative ties, personal networks of immigrants, and social interaction in open collaboration systems, such as Wikipedia.

Patrick Kenis is Academic Dean of Antwerp Management School since 2010 and Professor Tilburg University, the Netherlands. He received his PhD in Social and Political Sciences from the European University Institute in Florence, Italy. His research area is in organisation theory and the design and functioning of multi-organization networks.

Denise van Raaij obtained a Master degree in Policy and Organization Studies and holds an Ph.D in Organization Studies. From 2007 until 2010 she worked as a post doc researcher at the department of Organization Studies at Tilburg University, where she focused on the logic of production and performance of open innovation networks. Currently, she works as a lecturer at the same department.

Ulrik Brandes is a Professor of Computer Science at the University of Konstanz. With a background in algorithmics, his main interests are in network analysis and visualization, with application to social networks in particular. He is a member of the Graph Drawing Steering Committee, the board of directors of the International Network for Social Network Analysis (INSNA), and the editorial board of the Journal of Graph Algorithms and Applications, as well as an associate editor of Social Networks. He was recently awarded a prestigious Reinhart Koselleck-Project by Deutsche Forschungsgemeinschaft (DFG), and co-directs a transfer center for network research methods.