Dominance, Deference, and Hierarchy Formation in Wikipedia Edit-Networks

Jürgen Lerner and Alessandro Lomi

Abstract Does co-editing of Wikipedia articles reveal users dominating others? Do these dyadic dominance orderings (if any) lead to a global linear hierarchy among contributing users? In this article we claim that dominance (respectively deference) is revealed by users undoing (respectively redoing) edits of others. We propose methods to turn the history of Wikipedia pages into a dynamic multiplex network resulting from three types of interaction events: dyadic dominance, dyadic deference, and third-party assigned dominance ties. We analyze various local temporal patterns for the different types of ties on a sample of page histories comprising 12,719 revisions by 7,657 unique users. On the dyad level we analyze whether two users tend to agree on a dominance order among them or whether dominated users tend to fight back. On the neighborhood level we analyze various degree effects including whether dominant users tend to dominate in the future and whether subordinate users tend to get dominated. On the triad level we analyze whether users have a preference for transitive closure over cyclic closure of dominance ties. These dynamic patterns shed light on the micro processes that can foster or impede the emergence of a global linear hierarchy.

1 Introduction

The formation of dominance hierarchies is a universal pattern in many human and non-human societies. For instance, experiments with domestic chicken [15, 5] revealed that interaction among two individuals results with overwhelming probability in a clearly dominant and a clearly subordinate one, that dominant (respectively subordinate) individuals tend to dominate (respectively get dominated by) others, and

Jürgen Lerner
University of Konstanz, Germany, e-mail: juergen.lerner@gmail.com

Alessandro Lomi
Università della Svizzera italiana, Lugano, Switzerland e-mail: alessandro.lomi@usi.ch

that dominance networks of several individuals tend to be transitive and cycle-free. Experiments on dominance among humans (often denoted by terms like *status*, *reputation*, *prestige*, or *power*) have been performed with small groups (compare [6] and references therein) but empirical studies on hierarchy formation in larger and non-artificial human groups are rare. Collaboration in Wikipedia provides an opportunity to study large-scale, longitudinal, and completely observed data on hierarchy formation in task-oriented human groups. Analyzing hierarchy formation is relevant for understanding Wikipedia since, as any production community, it has to solve the problems of coordination and control. Moreover, acquired high or low status might be a primary source of motivation or frustration of users [18]. However, this paper does not attempt to determine the consequences of successful or failed hierarchy formation but rather analyzes the micro-processes that foster or impede the formation of a global linear hierarchy.

Contributions. In this paper we propose methods to turn the histories of Wikipedia pages into sequences of three types of timestamped and weighted interaction events: dyadic dominance, dyadic deference, and third-party assigned dominance ties. Dyadic dominance ties result from undoing edits and are tentatively interpreted as user A claiming: "I (A) dominate you (B)." Dyadic deference ties result from redoing edits and are tentatively interpreted as A claiming: "You (B) have high status." Finally, third-party assigned dominance ties result from user C favoring A's edits over B's edits and are tentatively interpreted as C claiming: "A dominates B." Thus, the difference between dyadic dominance and third-party dominance is whether the dominance from A to B is claimed by A or by a third actor C.

We turn these events into a dynamic multiplex network, encoding past interaction among users. Crucially, we aggregate not only the type and weight of events that are actually observed but normalize by the *potential* for such events. We analyze how a tie's embedding in the network of past events influences the probability of future typed events on it (see Figure 1 for details). This analysis tests the validity of the tentative interpretation of events and reveals which of these types are appropriate or inappropriate for uncovering dominance among users.

2 Background and related work on hierarchy formation and Wikipedia research

Linearity of hierarchies. Dominance hierarchies are universal in groups of many non-human and human species, e. g., [15, 5, 2]. This tendency to form linear hierarchies has often been attributed to advantages in the group's fitness (cf. [2, 17]); an interesting perspective for our topic: can the success or failure of task-oriented online communities be explained by the (in-)ability to form a hierarchy? Whatever the hypothetical causes or consequences of hierarchy formation, empirical tests of these need ways to assess the degree of linearity in the hierarchical structure of a group. Indices for linearity that have been defined for *tournament graphs* (i. e., graphs in which every undirected dyad $\{A,B\}$ has a dominant and a subordinate node), such as Landau's h or Kendall's K, have been shown to be inappropriate for sparse networks [16]. Global hierarchy indicees for sparse graphs exist (e. g., [14]); alternatively, it has been proposed to measure the linearity of sparse dominance graphs via the relative frequencies of small subgraphs, most notably transitive triads (pointing to linearity) and cyclic triads (pointing to non-linearity) [16, 17].

In this paper we will also consider local configurations but we stress two differences to the two last-mentioned papers. First, we are not analyzing networks of stable dominance ties but dynamic networks of relational events. Thus, instead of counting configurations, we model the probability of current events on a dyad (A,B) as a function of how (A,B) is locally embedded into the network of past events. Second, in networks resulting from the co-editing among Wikipedians there is no reason to assume *a priori* that reciprocated dominance ties are rare. This marks a considerable difference to, say, pecking-networks among chicken where dominance ties are rarely reverted [5]. In Wikipedia, anecdotal evidence, such as the term "edit war" or the "three-revert rule[1]", suggests that at least some users do not accept it when their edits are undone but have a tendency to fight back. Therefore we must start our analysis not with analyzing types of triangles or stars but on the lower dyadic level. Figure 1 illustrates the different network effects considered in this paper.

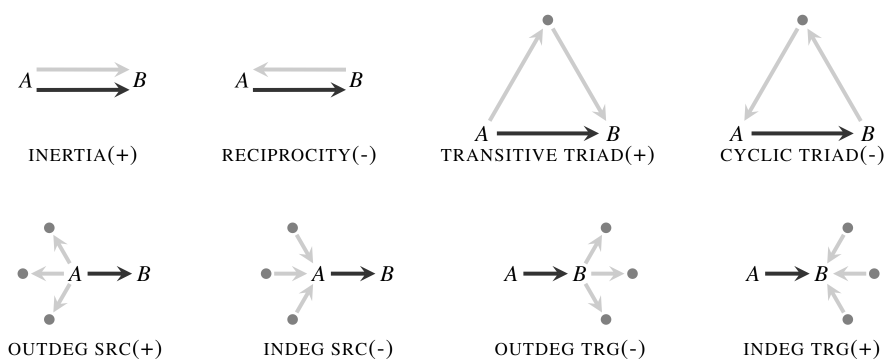

Fig. 1 Local configurations of past dominance events (light gray) explaining future dominance on the tie from A to B (dark gray). A plus sign (+) indicates a hypothetical increase in the probability; a minus sign (-) indicates a hypothetical tendency for decreased dominance probability on (A,B). All of these hypotheses are derived from the assumption that dominance ties point from higher to lower in the hierarchy. Note that the ties are not binary but have weights between zero and one, as explained in Sect. 3.

Wikipedia research. Wikipedia[2] is an open, Web-based project to create a user-generated encyclopedia using wiki software [12]. Launched in 2001, Wikipedia is

[1] https://en.wikipedia.org/wiki/Wikipedia:Edit_warring

[2] www.wikipedia.org

one of the Top-10 most visited websites worldwide[3] and is the largest and most popular general reference work on the internet. Its societal relevance, together with the free availability of its complete database, made Wikipedia also a popular case for empirical research and here we can only discuss some of the most closely related previous work. *Reputation* systems for Wikipedians have been proposed, e. g., in [1, 8]. It has been shown, among others, that contributions of users with low reputation are more likely to be undone in the future; this finding corresponds to the hypothesized effect of INDEGREE TARGET in the notation from Figure 1. Other possible patterns in the Wikipedia edit networks are, however, not tested in these two papers, but the largest difference is that we do not seek to define a global reputation index for users but systematically evaluate dynamic local patterns that can foster or hinder the emergence of a linear dominance hierarchy. *Event sequences* (compare [4]) resulting from co-editing Wikipedia articles are analyzed in [7, 9] but none of these papers is specifically about dominance among users (nor about status or reputation of users). *Signed networks* (that is, networks with positive and negative ties) have been defined resulting from co-editing articles ([3]), from votes for or against requests for adminship ([11]), or from both ([13]). Subsequently, these three papers analyze triadic or global patterns confirming or contradicting balance theory and/or status theory in these signed networks. Adding to these previous papers, we evaluate more systematically the consistency of local dynamic patterns with linear hierarchy formation on the dyad level, the neighborhood level (degree effects), and the triadic level. As it has been argued above and will be empirically shown below, the formation of linear hierarchies can be challenged not only with triads but already at a lower level. Last but not least, to the best of our knowledge our paper is the first that also considers third-party assigned dominance ties in which a user C states a dominance order between two different users A and B. The distinction between dyadic dominance and third-party dominance is highly important, since—as we will show in this paper—the latter type of dominance ties is more consistent with linear hierarchy formation.

3 Dominance, deference, and third-party dominance

Edit events. We propose to compute relational events expressing dominance, deference, and third-party dominance by successively comparing the text of subsequent revisions of the same Wikipedia article in a similar way as in previous work, e. g., [1, 3, 8, 13]. As in these papers, we determine for each revision which part of the text is newly added, which is deleted, and which previously deleted text is restored by reverting a deletion. As in previous work, we do not treat it as a text modification if large parts of the text (complete sentences in our case) are just moved or duplicated. As it is usual, we consider a sequence of consecutive revisions by the same user as one revision whose text is that of the last one in the sequence. Authorship of text is

[3] http://www.alexa.com/topsites

maintained at the word level. Note that the same word can appear in different places in the text and these different instances can have different authors. Augmenting the computation of edit events proposed in [3], we encode the user interaction resulting from it in a more complete way, as explained in the following.

For each word w in the text of each revision we maintain pointers to three potentially different users playing different roles with respect to w:

$$[\texttt{author}(w), \texttt{deleter}(w), \texttt{restorer}(w)] \ .$$

Here $\texttt{author}(w)$ is the author who originally added the word w. This pointer is set at the revision when the word is added and is never changed afterward. The pointer $\texttt{deleter}(w)$ gives the last user who deleted the word. It points to `nil` when the word is originally added (indicating that no one deleted it so far) and is updated whenever the word is deleted. The pointer $\texttt{restorer}(w)$ gives the last user who added or restored the word. It is set to the author when the word is originally added but, in contrast to $\texttt{author}(w)$, the last restorer of a word can change over time when a word is restored after being deleted.

Adding a word, thus, assigns the author of it but creates no interaction events. Interaction events arise when a word is deleted or restored as defined in Figure. 2. Note that we generate a dyadic event only if `active` (i. e., the user who performs the revision) is different from the target of the event and we generate a third-party dominance event only if the active user, the source, and the target are three pairwise different users.

Fig. 2 Edit events resulting from the deletion of a word (*left*) and a word being restored (*right*). Solid lines encode *dyadic deference* events by which the active user re-does the target user's edit. Dashed lines encode *dyadic dominance* events by which the active user makes the target user's edit undone. Dotted lines encode *third-party dominance* assignments by which the active user re-does the source user's edit that has been made undone by the target user. After deleting a word w the user `active` becomes $\texttt{deleter}(w)$ and after restoring w user `active` becomes $\texttt{restorer}(w)$. Note that $\texttt{author}(w)$ is only set when w is originally added and does never change again.

The event potential. While iterating over the revisions of a page we do not only consider events that happen but also the *potential* for such events. More precisely, we keep track for each user B and for each of the dyadic event types x (that is, dyadic dominance and dyadic deference) how many events of type x can have target B. Likewise, for each ordered pair of different users (A, B) we keep track of the potential for third-party dominance events which a user C (different from A and from B) can assign to the dyad (A,B).

The network of past events. While iterating over the sequence of revisions of a page, we successively update six functions (called *dyad-level attributes*) defined on ordered pairs (A,B) of different users. Three of these attributes count events of the three types that actually happened on (A,B) and three of them (the *cumulative potentials*) add up the number of events (of the three types) that could have happened on (A,B) at the edit times.

Finally, to describe the past interaction on dyads (A,B) we consider, separately for the three event types, the ratio of actually observed events divided by the cummulative potential for such events.[4] These ratios are between zero and one (including these borders) and can be interpreted as probabilities: the *past dyadic dominance ratio* on (A,B) is the probability that a randomly chosen word of B that could have been made undone by A during the history of the page is actually undone by A. Similar interpretations apply to *past dyadic deference ratio* and *past third-party dominance ratio*. Henceforward, when we speak of past dominance, deference or third-party dominance, we refer to these ratios.

4 Statistical model

Outcome variables. Whenever a revision r is performed by a user A, then A has a certain potential to initiate events of the two dyadic types to various target users B and A has a certain potential to initiate third-party dominance events on various dyads (B,C).[5] The three outcome variables that we consider are the ratios of the number of events actually performed in r divided by the respective potential for such events. Thus, for each event type we use a binomial model where instances are words that can potentially be changed, a "success" instance is such a word that is actually changed in the revision, and a "failure" instance is such a word that is left unchanged. The probability that a potential change occurs is specified in logistic regression models with explanatory variables introduced below.

Explanatory variables. When modeling the probability of change events that could happen in revision r, we use only information about past interaction resulting from revisions that happened strictly before r. These explanatory variables are defined by combinations of three dyadic attributes (past dyadic dominance ratio, past dyadic deference ratio, and past third-party dominance ratio) on the configurations shown in Figure 1. Specifically, for the degree variables we add up the attribute values of all in-coming respectively out-going dyads incident to source respectively target. For the transitive triad variables we sum over all users C (different from A

[4] Here we resolve $0/0$ to be equal to 0, since no event of that type could have happened so far on such a dyad.

[5] Here we speak of the potential for events in revision r. Note the difference to the cumulative potential used for defining tie-weights in the network of past events.

and B) the product of the attribute value on (A,C) with the value on (C,B) and take the square-root of this sum. For the cyclic triads we consider the dyads (C,A) and (B,C) accordingly.

To obtain better interpretable explanatory variables we divide them by their standard deviation. With this normalization it is easier to compare the effect sizes of the various variables. Since average probabilities are very close to zero (cf. Table1), we can interpret the estimated parameters in the following intuitive (not formally correct) way: if we estimated a parameter θ for the variable x when modeling the dyadic dominance probability p, then (hypothetically) increasing x by one standard deviation (that is by 1) multiplies the probability p by $\exp(\theta)$.

Empirical data. We analyzed the histories of a sample of ten articles from the English-language Wikipedia, randomly chosen from the set of articles that have at least 1000 revisions.[6] In March 2016 there are 56,042 articles (pages in the main namespace that are not redirects) that have at least a thousand revisions. (Altogether there are about 5 million articles; the mean number of revisions per article is just 86.) The ten sampled articles have together 12,719 revisions (disregarding successive revisions by the same user) performed by 7,657 different users. We note that our

Table 1 Number of instances and non-null instances in the analyzed data.

	dyadic dominance	dyadic deference	third-party dominance
no. potential dyads	3,126,047	1,753,160	4,852,052
no. non-null dyads	37,823	21,411	21,335
dyad-density	1.21%	1.22%	0.44%
no. potential words	361,673,769	359,365,077	348,420,292
no. changed words	1,738,728	785,233	783,190
word-change density	0.48%	0.22%	0.23%

number of observations is not just ten since the unit of analysis is not the page but the dyadic event. Table 1 gives the number of dyad-timepoints on which there could have happened an event of the various types, the number of actual dyadic events, the number of words that could have been modified, and the number of actual word modifications. The approach to analyze 10 random pages (rather than just one) has been chosen since it reduces the likelihood of accidentally analyzing a page with an exceptional structure. The restriction to pages with at least a thousand revisions is motivated by the consideration that hierarchy formation takes some time and also a number of users that is not too small. What blows up the runtime of our analysis is that we consider not only the actually occurring events but also those that could have happened. However, we strongly believe that this is necessary since an observation such as “user A deleted 10 of user B’s words” is meaningless if we disregard how many of B’s words user A did not touch and/or if we disregard all the other users

[6] These turned out to be the pages: Balika Vadhu; Ganymede (moon); Greed; Jay Park; List of Hollyoaks locations; Mothra; Pea; Shiv Sena; Swimsuit; and The Third Man.

with which A potentially could have interacted but did not. The results reported in the next section have been estimated to maximize the joint likelihood of all events from all sampled pages.

5 Results and discussion

Dyad-level effects. Table 2 reports logistic regression parameters explaining the probability of dyadic dominance by past interaction on the same and the reverse dyad. In the first model, we observe that past dyadic dominance on (A,B) increases the probability of future dyadic dominance on (A,B); thus, actors continue to dominate their subordinates. However, we see that past dyadic dominance on the reverse dyad (B,A) also increases the probability of dyadic dominance on (A,B); thus, subordinate actors have a tendency to fight back which is a hindrance to hierarchy formation. Likewise, we see that past deference on (A,B) reduces the probability of dyadic dominance on (A,B) (as expected). However, past deference on (B,A) also reduces the probability of dyadic dominance on (A,B); this makes the interpretation that deference goes from lower to higher in the hierarchy questionable.

Table 2 Explaining dyadic dominance by past dyadic dominance and dyadic deference on the same dyad.

	dyad model	dyadic inertia	dyadic reciprocity
(Intercept)	-5.427 (0.001)***	-5.427 (0.001)***	-5.427 (0.001)***
dyadic dominance inertia	0.222 (0.000)***	0.115 (0.000)***	
dyadic deference inertia	-0.288 (0.002)***	-0.071 (0.003)***	
dyadic dominance reciprocity	0.060 (0.000)***		-0.064 (0.000)***
dyadic deference reciprocity	-0.093 (0.000)***		0.030 (0.001)***
undirected dyadic dominance		0.127 (0.000)***	0.262 (0.000)***
undirected dyadic deference		-0.243 (0.001)***	-0.321 (0.003)***
AIC	17,531,308.231	17,531,308.231	17,531,308.231
Num. obs.	3,126,047	3,126,047	3,126,047

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

Looking more closely at the parameter sizes, we see that a dyadic dominance event on (A,B) has two effects. First it increases the future hostility (likelihood of dominance events) on (A,B) and on (B,A). This is consistent with a structural balance interpretation of negative ties (and inconsistent with a status interpretation) and has also been found by Leskovec et al. [11] who analyzed voting behavior of Wikipedians. A second effect, however, is that a dyadic dominance event on (A,B) increases the future dominance on (A,B) more than on (B,A), thereby increasing the relative dominance of (A,B) over (B,A). This second effect becomes more transparent if we control for the increase in dominance activity on both dyads (A,B) and (B,A) by defining a variable *undirected dyadic dominance* which is the sum of dyadic dominance inertia and dyadic dominance reciprocity (normalized to standard deviation one). In the second and third model in Table 2 we see that, controlling for the undirected increase in dominance activity, a dominance event on (A,B) increases the future dominance probability on (A,B) more than expected and that it increases the future dominance probability on (B,A) less than expected. A similar result is obtained for dyadic deference, where a deference event on (A,B) *decreases* the future dominance probability on (A,B) more than expected and that on (B,A) less than expected. We note that the three models in Table 2 are equivalent since their variables are linear transformations of each other.

Summarizing this, a dyadic dominance event on (A,B) has two effects: a structural balance effect increasing the hostility level on the undirected dyad $\{A,B\}$ and a hierarchical effect that shifts the relative dominance towards the direction (A,B). It is likely that the experimentally found anti-reciprocity of dominance events among chicken (e. g., [15, 5]) is due to the small network size. In larger and therefore sparser networks it is likely that reciprocation of acts of dominance, albeit rare, would occur with a higher probability than the low baseline probability of interacting at all.

We make similar findings when estimating the probability of dyadic deference by dyadic effects (with the understanding that deference hypothetically points from lower to higher). These results are not reported in this paper.

Table 3 Explaining third-party dominance by past third-party dominance on the same dyad.

	dyad model	dyadic inertia	dyadic reciprocity
(Intercept)	-6.196 (0.001)***	-6.196 (0.001)***	-6.196 (0.001)***
tp dominance inertia	0.379 (0.000)***	0.391 (0.001)***	
tp dominance reciprocity	-0.022 (0.001)***		-0.740 (0.001)***
undirected tp dominance		-0.025 (0.001)***	0.814 (0.001)***
AIC	9,721,545.325	9,721,545.325	9,721,545.325
Num. obs.	4,852,052	4,852,052	4,852,052

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

Table 3 reports logistic regression parameters explaining the probability of third-party dominance by past interaction on the same and the reverse dyad. In contrast to dyadic dominance, we see that third-party dominance is clearly anti-reciprocal: controlling for the undirected increase in the event probability is here not necessary although it strengthens the anti-reciprocity. This means that if a different user C states that A dominates B, then the probability that C (or any other user different from A and B) later reverses this order decreases. Thus, third-party assigned dominance is more consistent with the hierarchical interpretation than dyadic dominance. Apparently bystanders can judge the dominance order among A and B more reliably than A or B themselves.

Neighborhood-level effects (degree effects). We estimated models explaining the probability of dyadic dominance on (A,B) by past interaction on edges incident to A (source) and B (target). For space limitations, the estimated parameters are not reported in this paper and we will only summarize the main findings. We find some effects consistent with the hierarchical interpretation, such as a positive effect of *dominance outdegree source* and *dominance indegree target*. However, we can also find effects inconsistent with this interpretation, such as a positive effect of *dominance outdegree target* (which implies that dominant users tend to get dominated). As in the case of dyad effects, the effects of the degree variables (for dominance and for deference) become consistent with the hierarchy-interpretation once we control for the *undirected* degrees. We also controlled for the dyadic effects from Table 2 in the degree model which did not change the findings qualitatively. We further estimated the probability of dyadic deference events by degree effects. These findings differ qualitatively from those obtained for the dominance probability (whether or not we control for the undirected degrees) and make the interpretation of dyadic deference pointing from subordinate to dominant more questionable.

We also estimated degree-models for third-party dominance (not reported in this paper). Most effects in this model are consistent with the hierarchical interpretation of third-party dominance ties. The exception is a positive effect of *indegree source* which suggests that subordinates are more likely to dominate in the future. As for dyadic dominance, controlling for the undirected degrees brings all effects in accordance with the hierarchical interpretation. Controlling for the dyadic effects from Table 3 in the degree model yields qualitatively the same findings.

Triad-level effects. Table 4 reports logistic regression parameters explaining the probability of dyadic dominance on (A,B) by past interaction on two-paths of the form $(A,C),(C,B)$, forming a transitive triad, and on two-paths of the form $(B,C),(C,A)$, forming a cyclic triad. The first model reveals that the embedding of (A,B) in a dominance two-path increases the probability of a dominance event on (A,B), irrespective of whether the resulting triad is transitive or cyclic. Controlling for the increase in dominance activity caused by a dominance two-path in any direction (*dominance triplet*) reveals a preference for transitive over cyclic closure of dominance ties—consistent with the formation of a linear hierarchy. Similar effects result from two-paths of deference ties. Controlling for dyad effects and degree effects (not reported in this paper), however, does *not* keep these effects stable.

Table 5 reports logistic regression parameters explaining the probability of third-party dominance on (A,B) by past interaction on two-paths of the form $(A,C),(C,B)$, forming a transitive triad, and on two-paths of the form $(B,C),(C,A)$, forming a cylic triad. The first model reveals that indirect third-party dominance ties decrease the probability of third-party dominance on the dyad (A,B) irrespective of the direction of these two-paths. For transitive triplets, this contradicts the hierarchical interpretation but is consistent with a structural balance interpretation of dominance ties (an enemy of an enemy is not an enemy). When we control for the dominance-reducing effect of undirected two-paths, we find a preference for transitive over cyclic closure (consistent with the hierarchical interpretation). As for dyadic dominance, controlling for dyad and degree effects (not reported in this paper) does *not* keep these triadic effects stable.

Table 4 Explaining dyadic dominance by past dyadic dominance and dyadic deference on transitive and cyclic two-paths.

	triad model	transitive triad	cyclic triad
(Intercept)	-5.264 (0.001)***	-5.264 (0.001)***	-5.264 (0.001)***
transitive dominance triplet	0.149 (0.001)***	0.049 (0.001)***	
transitive deference triplet	-0.362 (0.002)***	-0.045 (0.003)***	
cyclic dominance triplet	0.027 (0.000)***		-0.013 (0.000)***
cyclic deference triplet	-0.135 (0.001)***		0.019 (0.001)***
dominance triplet		0.105 (0.001)***	0.156 (0.001)***
deference triplet		-0.355 (0.002)***	-0.405 (0.003)***
AIC	18,958,234.850	18,958,234.850	18,958,234.850
Num. obs.	3,126,047	3,126,047	3,126,047

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

Table 5 Explaining third-party dominance by past third-party dominance on transitive and cyclic two-paths.

	triad model	transitive triad	cyclic triad
(Intercept)	-5.403 (0.001)***	-5.403 (0.001)***	-5.403 (0.001)***
transitive tp dominance triplet	-0.024 (0.001)***	1.330 (0.002)***	
cyclic tp dominance triplet	-1.792 (0.003)***		-1.760 (0.003)***
tp dominance triplet		-2.259 (0.004)***	-0.040 (0.001)***
AIC	10075300.988	10,075,300.988	10,075,300.988
Num. obs.	4,852,052	4,852,052	4,852,052

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

6 Conclusion

In this paper, we proposed methods to derive three types of interaction events from co-editing Wikipedia articles. We analyzed whether local dynamic patterns for these events are consistent with a linear dominance hierarchy among the users. The analysis in this paper revealed that past events can have two distinct effects on future interaction: on one hand on the frequency of events on the undirected dyad $\{A,B\}$ and on the other hand on the relative dominance of (A,B) over (B,A). The effects on the undirected dyads are often more consistent with a structural balance interpretation of dominance events as revealing negative ties. The effects on the directed dyads are often more consistent with a hierarchical interpretation of dominance events. This finding is similar to one made in [11] where voting behavior among Wikipedians was analyzed. We also showed that the effect on the event frequency can obfuscate effects on the hierarchical ordering. This finding is similar to one made in [10] where effects on the interaction frequency were separated from effects influencing the sign of ties. The analysis in our paper also revealed that the three different types of events show different levels of consistency with linear dominance hierarchies. Most notably, third-party assigned dominance was the only event type that is anti-reciprocal, irrespective of whether we control for a change in the interaction frequency or not. On the other hand, dyadic deference was the most inconsistent with the hierarchical interpretation. A promising approach for future research is to link patterns of (failed or successful) hierarchy formation with properties of the page, such as article quality. This would need a larger sample of separately analyzed pages that show variation in their hierarchical structure and in quality.

Acknowledgements This work has been supported by Swiss National Science Foundation (FNS Project Nr. 100018_150126) and Deutsche Forschungsgemeinschaft (DFG Grant Nr. LE 2237/2-1).

References

1. Adler, B.T., de Alfaro, L.: A content-driven reputation system for the Wikipedia. In: Proc. 16th Intl. Conf. WWW, pp. 261–270 (2007)
2. Boyce, W.T., Obradović, J., Bush, N.R., Stamperdahl, J., Kim, Y.S., Adler, N.: Social stratification, classroom climate, and the behavioral adaptation of kindergarten children. Proceedings of the National Academy of Sciences **109**(Supplement 2), 17168–17173 (2012)
3. Brandes, U., Kenis, P., Lerner, J., van Raaij, D.: Network analysis of collaboration structure in Wikipedia. In: Proc. 18th Intl. Conf. WWW (2009)
4. Butts, C.T.: A relational event framework for social action. Sociological Methodology **38**(1), 155–200 (2008)
5. Chase, I.D.: Dynamics of hierarchy formation: the sequential development of dominance relationships. Behaviour **80**(3), 218–239 (1982)
6. Fişek, M.H., Berger, J., Norman, R.Z.: Participation in heterogeneous and homogeneous groups: A theoretical integration. American Journal of Sociology pp. 114–142 (1991)
7. Iba, T., Nemoto, K., Peters, B., Gloor, P.A.: Analyzing the creative editing behavior of Wikipedia editors: Through dynamic social network analysis. Procedia-Social and Behavioral Sciences **2**(4), 6441–6456 (2010)
8. Javanmardi, S., Lopes, C., Baldi, P.: Modeling user reputation in wikis. Statistical Analysis and Data Mining **3**(2), 126–139 (2010)
9. Keegan, B.C., Lev, S., Arazy, O.: Analyzing organizational routines in online knowledge collaborations: A case for sequence analysis in CSCW. In: Proc. 19th ACM Conf. Computer-Supported Cooperative Work & Social Computing, pp. 1065–1079. ACM (2016)
10. Lerner, J.: Structural balance in signed networks: Separating the probability to interact from the tendency to fight. Social Networks **45**, 66–77 (2016)
11. Leskovec, J., Huttenlocher, D., Kleinberg, J.: Signed networks in social media. In: Proc. SIGCHI Conf. Human Factors in Computing Systems, pp. 1361–1370. ACM (2010)
12. Leuf, B., Cunningham, W.: The Wiki Way. Addison-Wesley (2001)
13. Maniu, S., Cautis, B., Abdessalem, T.: Building a signed network from interactions in Wikipedia. In: Proc. Databases and Social Networks, pp. 19–24. ACM (2011)
14. Mones, E., Vicsek, L., Vicsek, T.: Hierarchy measure for complex networks. PloS one **7**(3), e33799 (2012)
15. Schjelderup-Ebbe, T.: Beiträge zur Sozialpsychologie des Haushuhns. Zeitschrift für Psychologie **88**, 225–252 (1922)
16. Shizuka, D., McDonald, D.B.: A social network perspective on measurements of dominance hierarchies. Animal Behaviour **83**(4), 925–934 (2012)
17. Shizuka, D., McDonald, D.B.: The network motif architecture of dominance hierarchies. Journal of The Royal Society Interface **12**, 20150080 (2015)
18. Willer, R.: Groups reward individual sacrifice: The status solution to the collective action problem. American Sociological Review **74**(1), 23–43 (2009)