2011

Motivation through visibility in open contribution systems

James Howison
Carnegie Mellon University

Jamie F. Olson
Carnegie Mellon University

Aniket Kittur
Carnegie Mellon University

Kathleen M. Carley
Carnegie Mellon University

Motivation through visibility in open contribution systems

James Howison
Carnegie Mellon University
Pittsburgh, PA 15214
jhowison@cs.cmu.edu

Jamie F. Olson
Carnegie Mellon University
Pittsburgh, PA 15214
jfolson@cs.cmu.edu

Aniket Kittur
Carnegie Mellon University
Pittsburgh, PA 15214
nkittur@cs.cmu.edu

Kathleen M. Carley
Carnegie Mellon University
Pittsburgh, PA 15214
kathleen.carley@cs.cmu.edu

ABSTRACT
Open contribution systems such as Wikipedia and Linux have been extraordinarily successful at eliciting contributions from many volunteers, but other projects struggle. While research has examined general motivations for contribution, we know little about what triggers contribution at a specific time. In this paper we mine the history of Wikipedia to understand the contribution process in the context of the many competing demands on users' attention at their computers. In particular, we examine the influence of the visible work of others on the timing and amount of contributions to Wikipedia. Using two different statistical models we show converging evidence of a substantial influence of others' visible work on triggering contribution.

Author Keywords
motivation, attention, contribution, online communities, hawkes process, time

ACM Classification Keywords
H.5 Group and Organization Interfaces: Collaborative computing

INTRODUCTION
Open contribution systems have been surprisingly successful. Wikipedia and Linux are canonical examples, but success has been found in other domains, such as building scientifically useful datasets from the records of hobby birdwatchers (the eBirds project) [27]. The systems that support them have emerged as a major new category of software.

The ability of these projects to draw together contributions from so many people, often volunteers, and aggregate them into valuable resources is a question of great interest, both to academics and to practitioners establishing new projects seeking to replicate their success. Research has focused on understanding the motivations of participants (e.g., [25, 17, 14]) and, more recently, examined how individual contributions are assembled into functional wholes (e.g., [13, 6]). While well-known projects have been successful there remain considerable opportunities for increasing contribution; other projects face a challenge of encouraging any contribution at all [17].

Successful projects represent an aggregation of significant amounts of participants' time. For example, estimates put the amount of work put into a Linux distributions at up to 14,005 person-years [1]. This work was not, of course, undertaken all at once; rather it occurred in many separate slices of participants' time. It is, in fact, one of the shared characteristics of these projects that participants vary rarely work "full-time" on the project; rather work is intertwined with other activities, even within a single workday. Regardless of their overall motivation, to actually make a contribution a participant has to turn from their other activities, pay attention to the project and engage in project work. By this process the general motivations of participants are turned into actual contributions. Yet we know little about what prompts, or triggers, sometime participants to turn to a project and actually contribute.

This paper makes the argument that the visibility of the work of others is central to this triggering process. The information systems that support these projects bring the activity of others into the local context of a potential contributor's networked computer, drawing their attention and creating the opportunity for contribution. In a reinforcing cycle, the system then rapidly conveys that contribution out to other participants, generating yet more opportunities for attention and contribution.

In this paper we examine the influence of the visible work of others on the timing and amount of participation in Wikipedia. First we develop our theory, elaborating a set of mechanisms we believe to operate within this cycle, concluding with a simple hypothesis. We then report on two studies which explore this hypothesis with two different

Working paper submitted for review to CHI 2011.

methods. Finally we conclude with a combined discussion and present theoretical and practical implications.

THEORY

Understanding human behavior, especially as it affects the utility and performance of systems, is at the core of the field of human-computer interaction. Such understandings not only help to explain the emergence and success of new socio-technical phenomena but also help in improving designs for new systems.

Motivation in online contribution projects has been extensively studied especially in the contexts of open source software and online communities [15, 25, 14, 7]. In the HCI field existing work has identified the importance of addressing the problem of contribution by ensuring that system design accommodates and activates potential participants' motivations [17, 14, 23]. In particular authors have pointed to making requests for work (especially from high status individuals), emphasizing the benefits of contribution, providing "social proof" that others have contributed [26] and setting specific or high challenge goals [2] as ways to improve contribution. Systems can be designed to emphasize intrinsic motivations, including enjoyment of other's company [4] or their social presence [16, 19]. Systems can be designed to provide services that enhance activity people already undertake, accomplishing aggregation almost as a side-effect as with birding records on eBird [27].

In general this work has concentrated on general motivations for participation and has not addressed the specific question of how these motivations become activated and result in work at a specific time. Considering this question turns attention to the material environment in which participation takes place. Early research pointed to the possibility of computers "mesmerizing" their users into concentrated "flow" [30], but more recent research, after the advent of ubiquitous networking, has emphasized the many competing demands on a computer user's attention [22, 12], often resulting in interruptions [8]. Online communities compete in this crowded environment, bringing the remote activity of others into a computer user's local environment. Below we theorize about how these visible actions might trigger contribution.

We identify three broad mechanisms through which we think visible actions by others contribute to triggering contribution: 1) attention effects, 2) audience effects and 3) semantic effects.

Attention Effects A visible piece of work by another can, under some circumstances, reach out and draw the attention of a sometime participant to the project. This is particularly true when work in the system generates messages that are pushed out to participants, such as through mailing lists or RSS feeds. These type of communications might be encountered in the participant's normal flow of daily work, functioning as intermittent reminders that the project exists. This may be sufficient to trigger work the participant had in mind. This effect seems likely to be particularly powerful

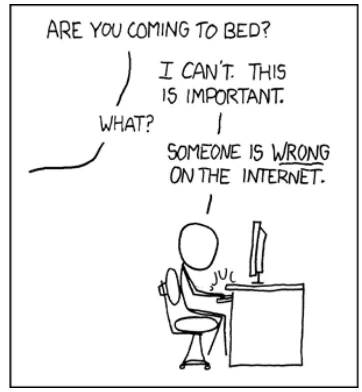

Figure 1. XKCD provides a humorous look at the visible work of others triggering participation, even in the face of other time demands. Reprinted with permission, see http://xkcd.com/about/

if a participant was experiencing frustration in their regular work and was therefore amenable to distraction.

Attention effects are likely to be less powerful when systems do not have push channels, since a participant must themselves reach out to the project to seek updates on activity by others. Nonetheless some systems have hybrids, in which the activity of others is not directly pushed into a participant's local space but is displayed in locations that they are likely to see in the normal course of using that system. For example a logged-in user might see a counter on each page, showing the number of new notifications of work by others. Such a reminder, while it is not reminding them of the existence of the project *per se*, might be enough to shift the users perception of the site as from a passive information source to a venue in which they can contribute.

Audience Effects A visible piece of work by another also functions as a signal that there is a potential audience for contribution. A reminder of the existence of an audience in this sense may be sufficient to trigger contribution on its own. Here the debate between "social loafing" and "social facilitation" is relevant; if participants see that work is being done by others they may not see a need to contribute, but if they seek the approval of others they may be encouraged to participate by the availability of a evaluative audience [17, 31].

A message received in near real-time, however, signals more than the general existence of an audience, it signals that a specific other is paying attention to the project right now; that they are awake, online and engaged. This creates an opportunity for seeking a reaction from that person, but it is a time-limited one, at least if a quick response is needed. Thus the mere appearance of visible work seems likely to trigger rapid response, especially if pushed out quickly enough that contributors can expect a further response from the original actor.

Semantic Effects Beyond the mere notification of others' work the semantics of the content of that work seem likely to play a role in triggering contribution. Many systems, such as collaborative source code management systems like CVS, can be configured to push out a description of that work and even to push out the work itself, showing exactly what changes were made to the shared state of the project. We see three ways that the content of a message can trigger response.

The first is through a negative reaction to the now changed state of the project. Participants are aware that others are seeing these changes; they are the current public face of the project. If they are uncomfortable with that situation then they may well be moved to act rapidly. Vandalism in Wikipedia is a clear example of this but it can also be seen in open source software projects, where a change may be perceived as introducing a memory leak or security hole, or in advice communities where participants can be motivated to respond because they feel the advice of another is wrong or dangerous. Indeed this process has even figured in internet humor, as shown in Figure 1.

Another example of this is common on Wikipedia where controversial topics lead to "edit wars" in which editors with opposing view points fight for control of the public face of a page. These may not always result in simple reversions, however, since participants may feel a need to positively contribute in order to disagree. Indeed the Apache Foundation codifies this by their rule that a negative vote against someone else's desired changes must be accompanied by a full technical explanation and, ideally, a functioning alternative[1].

Secondly, visible work might alter the system in ways that suggest, or even make possible, new work. This effect was observed in open source software development, where the addition of new code made much easier a task that others had deferred because it had been too complex [11]. In Wikipedia the expansion of a "stub" page might makes it obvious where a contributor could add a detail that would not have had a place in the article's previous state. This mechanism is particularly clear in question-and-answer communities where the existence of a question clearly creates the opportunity for an answer (and for additional refinements or extensions of that answer). In some circumstances visible work by others might unblock the local process of a contributor, such as when an answer to a question or the uploading of a missing file makes it possible for the contributor to continue work or to upload work they've done offline.

Reactions to the semantics of other's visible work can operate to trigger contribution even when the visible work is not pushed or conveyed in near real time. Participants returning to a shared space after some time catch up on the work undertaken by others in their absence. Even if the participant had been intending to just do a quick sweep they may be prompted to action by the semantics of what they observe. In this way the patterns of visibility and reaction may be uneven throughout the community; there is no reason to expect that any reaction is a reaction to the exactly previous event.

[1] `http://www.apache.org/foundation/voting.html`

While these mechanisms are distinct, it seems likely that they would work together in a reinforcing manner. Contribution triggered by the semantics of a message might take on more urgency or a different character if the contributor perceives the original author to be still available. A notification drawing the attention of a participant to the project might trigger a visit to the project's shared space where they might come across different content in which the semantics prompt continued participation. A participant periodically reviewing changes might notice a very recent change by a particular person for whom they are reminded they have an immediate question.

If these mechanisms are indeed operating in the way that we theorize then their net effect should be to both synchronize and increase contribution. Since these effects are strongly entangled, for now we focus on the question of synchronization and hypothesize that

> Participants in online work communities are more likely to work soon after others have worked.

EMPIRICAL ENVIRONMENT: WIKIPEDIA

We explore this hypothesis in Wikipedia. Wikipedia is an open, volunteer driven project and thus faces the challenge of attracting and retaining participants and their time. The MediaWiki system on which it runs makes edits by others visible through a particular feature called watchlists. Watchlists allow participants to add pages they are interested in to a monitored list. When others make changes those changes are visible in two ways: through logging in and checking one's watchlist page or by push through RSS feeds. Editors make changes to both the content of Article pages and edit associated Talk pages, which provide a venue for meta-discussion of the Article and its editing. Watchlists show changes to both Talk and Article pages; edits to either could have the effects we discuss above.

In order to bound our inquiry we chose to examine two parts of Wikipedia: pages associated with WikiProject Oregon and WikiProject Math. These are quite distinct areas of wikipedia, one geographically focused and one focused on an abstract area of knowledge. Study 1 examines only WikiProject Oregon, while Study 2 examines both.

Data

We began with the dump of Wikipedia dated March 2008. We then identified pages associated with a WikiProject by finding the template most used by the WikiProject and identifying all pages that carried that template on the date of the dump. For each set of pages we expanded our selection so that it included both Talk and Article pages, even if the relevant template was only on a single page type.

We employed three filters to clean up our dataset. The first was to remove edits made by bots, automated agents undertaking actions that show up as edits. We did this for two reasons: some bot edits are triggered automatically and therefore could overstate the reactivity of edits. Secondly, many of our theorized mechanisms above rely on reactions to other

	Pages	Edits	Editors
Oregon-All	11,242	252,832	25,779
Oregon-Selected	11,066	102,270	216
Math-All	7,174	498,148	36,814
Math-Selected	7,056	277,473	584

Table 1. Summary of datasets before and after applying our filters

humans. We had no way of knowing whether human edits did react to bot edits so removing bot edits is a conservative decision likely to reduce the appearance of reactivity in our dataset. We identified edits by bots if the username was in the bot user group in the Wikipedia database.

The second was to remove all anonymous edits (IP edits), because one cannot be sure that edits represent the same person and because IP edits are responsible for the vast majority of vandalism (97%, according to the WikiProject_Vandalism_studies/Study1 page). While vandalism and revisions are clear examples of visible edits motivating others the net result of a reversion is no change. We did not attempt to filter out the edits made by those doing the reversion (which are known to occur relatively quickly [3]), so our dataset likely includes some portion of these reversion edits, but without their corresponding vandalism edit. This should have the effect of reducing the appearance of responsiveness in the dataset.

Third, the mechanisms we theorize rely on edits being visible and the primary mechanisms for that in Wikipedia is watchlists. The Wikipedia foundation does not provide access to data regarding which editors add which pages to their watchlists, due to privacy concerns. We reasoned that users who are merely experimenting with Wikipedia are less likely to use watchlists than longer term participants. We therefore removed editors with few overall edits by including only those meeting a minimum threshold by picking corners in the distribution of edits counts.

Our final datasets, therefore, included all edits by logged-in editors who had made at least 100 edits to Article pages or at least 50 edits to Talk pages across any of the page pairs within each WikiProject. Table 1 shows the effect of our data cleaning. The thresholds reduced the number of editors examined substantially (while retaining large numbers of edits); we are therefore testing only for responsiveness between edits by relatively experienced editors who are likely to be using watchlists and are thus appropriate for testing our hypotheses about visibility. Figure 2 illustrates the data for one page and four editors; in reality the data is much sparser than this illustration (i.e., much longer periods without any edits).

STUDY 1: LAGGED REGRESSION

Our first approach to testing this hypothesis is to employ a lagged regression. Our main objective is to quantify the effect size of seeing other editors editing in the previous period. To this end we organized the data as shown in Figure 3, imposing a time-window on the data constructed by segmenting continuous time into consecutive periods of equal duration. The length of the appropriate time window should,

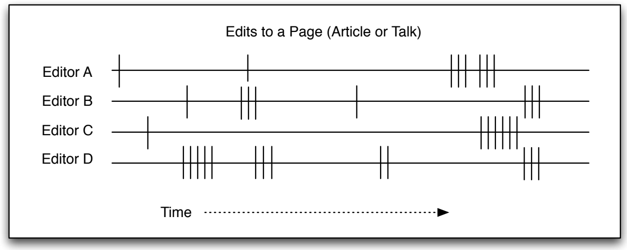

Figure 2. An illustration of our dataset

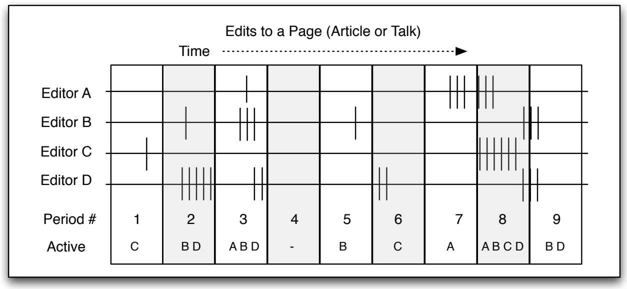

Figure 3. An illustration of applying time windows

in theory, match the expected length of time before a reactive edit becomes visible. We did not have a clear a priori reason to prefer one length over another, so we conducted a sensitivity analysis using different length windows. We discuss the limitations of this approach below.

We are primarily interested in whether people respond to each other, and did not want our results to be skewed by a small number of people undertaking a large number of edits (since contribution follows a power law) so we dichotomized participation for each editor (*self_edit*). The visibility of any other editor is sufficient to activate the mechanisms we theorized above so we created a dichotomized variable indicating whether any editor other than the focal editor was active in the previous period (*lag_other_edit*). Finally we reasoned that editors could be independently active from the previous period and so created a dichotomized variable indicating whether the focal editor had edited in the previous period (*lag_self_edit*). This variable also controls for unobserved factors relevant to self-motivation such as whether the editor was in a good mood or whether they had extra spare time on this day.

We are testing whether visible work makes those who observe it more likely to edit. If a participant didn't see work, even though it happened, then they are not eligible for this effect. As discussed above, in Wikipedia the primary visibility mechanisms is watchlists. Since we do not have access to data as to who has which pages on their watchlists we had to create a proxy. We assumed that the active editors remaining in our dataset keeps a page on their watchlist for 6 months after each edit; if there is a gap of 6 months without an edit that editor is not considered to have the page on the watchlist until they again edit that page. This can be seen in Table 2, derived from Figure 2 in that there is no row in period 2 for editor A, who is assumed not to have the page on their watchlist until after their first edit in period 3. We built this data table for each page pair separately; thus there is a row for each editor on the watchlist of that page pair in that period. Thus we will only find responsiveness if both the triggering event and the triggered event are on the same page. This reduces the possibility that our method would find responsiveness simply due to independent, but temporally overlapping, editing across the many pages in our dataset.

Results

We analyzed the data with a logit regression using the *biglm* package for R statistical environment [24]. Unsurprisingly, given, the extremely large data sets we are working with all results were statistically significant. We measure effect sizes through odds ratios by taking the exponents of the regression coefficients. The odds ratio is the odds that an edit occurs over the odds that an edit does not occur, given the presence of any edits by other editors in the previous period. When the odds ratio is greater than one the probability of an edit is higher than the probability of there being no edit.

Table 3 shows the odds ratios for the predictor variables, given different sized time-windows. As you would expect if an editor was active in the previous period they are far more likely to be active in the next period. The odds ratio results show that, accounting for whether the editor was active in the previous period, having any other editor actor in the previous period makes it more likely that the editor will edit. The effect clearly depends on the size of the window, ranging from 12 times more likely with a window of one hour to just under 3 times more likely that an editor will edit on the day following activity by any other editor.

Interim Discussion

We interpret this analysis as showing clear support for our hypothesis: editors do appear to react to the visible work of other editors on Wikipedia pages, at least amongst those who have performed a reasonably large number of edits. In the context of open online contribution projects where community mangers have few tools for increasing contribution an effect ranging from increasing participation by a factor of 2.9 through a factor of 12 is of potentially great interest.

Period	Editor	self_edit	lag other_edit	lag self_edit
2	B	1	1	0
2	C	0	1	1
2	D	1	1	0
3	A	1	1	0
3	B	1	1	1
3	C	0	1	1
3	D	1	1	1
...	...	...	...	...
5	A	0	0	0
5	B	1	0	0
...	...	...	...	...

Table 2. An illustration of our data tables for Regression, derived from illustration in Figure 3

		Odds Ratios	
Window Size	Row Count	lag self active	lag other active
1 hour	143,574,379	135.179	12.170
2 hours	71,621,741	55.519	9.091
6 hours	23,835,947	16.966	5.359
24 hours	5,970,127	7.828	2.909

Table 3. Log odds ratios for Project Oregon at different window sizes

A weakness of this analysis is that it requires the use of time windows. Time windows require two choices: their size and their location. The right-hand side of Figure 3 makes this clear: a window boundary can easily fall in the middle of a set of edits that extends only over a short period of time (say 10 minutes) but our data transformation would then indicate that the editor was paying attention for two adjacent time-windows; this may explain the extremely high effect sizes for *lag self active*. Our sensitivity analysis goes some way to dealing with both of these issues by varying the length of the windows, which also varies the location of the window breaks. Table 3 shows that while the direction of the effect is robust to the size of window chosen, the size of the effect varies considerably (although it is always substantial).

A further weakness of this approach is that we use only aggregate data and do not have a model of how the editing process actually occurs. This is important because our regression does not rule out the possibility that there are unmeasured variables driving the effects; a continuous-time model would allow greater insight into the event-level generative process.

STUDY 2: HAWKES PROCESS MODEL

Our second analytical approach examines our hypothesis by building and fitting a model. We fit the model to both datasets described above.

We need to be able to model reaction over time but, as discussed above, but can't assume that reaction is to the most recent event. Many models typically used, such as hidden Markov models, Markov random fields and Bayesian networks do, however make this assumption.

An alternative to these models is the mutually-exciting Hawkes process, which is a multidimensional extension of the self-exciting Hawkes process [9]. Although these models have mainly been used to examine temporal reactivity in finance and earthquakes recent work has use Hawkes processes to analyze the dynamics of more social phenomena such as internet traffic and YouTube viewing [18, 5].

Intuitively, a Hawkes process can be conceived of as two stochastic mechanisms. First, there is a continuous background process that generates events with some probability. Second, any generated event has the potential to lead directly to some future event, a process called excitation [10]. In the Wikipedia context this models the understanding presented above: editors have some independent propensity to edit, perhaps deriving from their use of Wikipedia or their long-term goals for improving Wikipedia. They also have the potential to react to visible events.

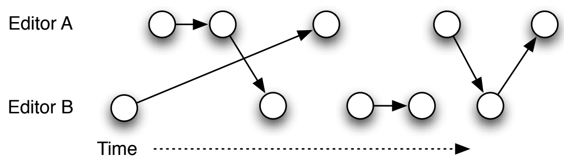

Figure 4. Example of two editors as mutually exciting Hawkes process. Each circle represents an edit event and the arrows indicate direct causal relationships.

Figure 4 illustrates these two mechanisms. Certain edit events occur independently and spontaneously, whereas others have direct causal links to previous edit events in a page's history. For example, a correction or elaboration of another author's works cannot occur unless that author has made their initial edit. In contrast, that author may have simply been reading the article and spontaneously decided to make an addition to it.

Formally, a self-exciting Hawkes process model is part of a class of models including both homogenous and non-homogenous Poisson processes. These type of models are defined in terms of an intensity function, $\Lambda(t|H_t)$, which defines the instantaneously likelihood of an event occurring, shown in general form in equation (1).

$$\lim_{\delta \to 0} \frac{1}{\delta} P(N(t, t+\delta) = 1) = \Lambda(t) \tag{1}$$

For Poisson processes, this intensity is either constant (homogeneous Poisson processes) or a function of external covariate information related to the process (non-homogeneous Poisson processes). Self-exciting Hawkes processes, on the other hand, have an intensity function that depends on the event history of the process itself. This makes the intensity process itself a stochastic process since it depends on the stochastic history of the process. In this way the likelihood of an event can depend directly and specifically on any other previous events in the stream, as depicted in Figure 4.

$$\Lambda(t|H_t) = \mu + \int_{-\infty}^{t} g(t-u)dN(u) \tag{2}$$

The intensity function for the self-exciting Hawkes process is shown in equation (2). The intensity of the background process is represented by v and each editing event contributes some excitation $g(t-u) \geq 0$ where u is time of the exciting event and t is the (potential) time of the excited event.

Since there are multiple editors in the Wikipedia context, each of whom can be excited by edits from any of the others, we need to introduce the idea of separate streams and parameters to model their mutual influence. A stream is a sequence of edits on a page. There are two kinds of sequences which can interact: A sequence of an editor's own edits to a page and a sequence made up of the edits of all others on the page. We allow two types of influence, analogous to the two terms in the regression reported above: the influence of an editor's stream on that editor's own stream and the influence of the stream of edits by all others on that editor's stream. Thus we introduce an additional parameter $\beta_{l \to k}$ to indicate the amount of excitation an event in the stream l contributes to stream k. Below we denote self as 1 and others as 2; the two different types of influence are $\beta_{1 \to 1}$ and $\beta_{2 \to 1}$. By requiring $\int_0^\infty g(v)dv = 1$, these $\beta_{l \to k}$ are equal to the expected number of events in stream k caused by events in stream l. Furthermore, these requirements make $g(v)$ a probability function for $v > 0$, simplifying the interpretation of $g(v)$ as the distribution of inter-arrival times between causally related events.

$$\Lambda_k(t|H_t) = \mu_k + \sum_{l=1}^{K} \int_{-\infty}^{t} \beta_{l \to k} g_k(t-u)dN(u) \tag{3}$$

$$\int_0^\infty g(v)dv = 1 \tag{4}$$

$$g(v) \geq 1, \forall v \geq 0 \tag{5}$$

The Hawkes process requires that we specify parameters for not only the amount of response, but also for the time at which a response occurs. Intuitively, the response time is a mathematical representation of any elements of the editing process which take time. This may be latency of the information technology system, reading of a prior edit or cognitive processing required in constructing a response edit. Modeling this required two decisions: the overall shape of the response curve and whether to allow its parameters to vary for reactions to different characteristics of events.

We reasoned that people do not respond immediately to stimuli, especially given that they may not observe the visible edits for some time. We therefore wanted a distribution that could represent both a skewed distribution with long time-dependencies and model situations where the most likely response time may not be most immediate. The exponential distribution would fail on both these accounts and while the Pareto distribution is positively skewed, it forces the largest probability mass on the most immediate possible response time. We therefore model the time required to respond to an event as a log-normal distribution which can be thought of as the multiplicative effect of many independent positive random variables.

We do not believe that there is sufficient reason to suspect these delays to differ based on the source of the exciting edit event to sacrifice the parsimony of a single distribution of response times. As such, we further assume that each editor has the same distribution of response times without regard to the characteristics of what they are responding to. In particular, whereas we allowed for the possibility that editors might be more likely to follow-up on their own posts than posts of others, we do not allow editors to have a different distribution of response times for their own posts.

Formally, the Hawkes process described above has the following intensity process for an individual editor as in equation (6) where z_i is simply an indicator of whether or not the edit event at time t_i was made by the editor being modeled.

$$\Lambda_1(t|H_t) = \mu_1 + \sum_{i:t_i<t} \left((1 - z_i)\, \beta_{1\to1} + z_i \beta_{2\to1} \right) g_k(t - t_i) \tag{6}$$

$$g(v) = \frac{1}{x\sqrt{2\pi\sigma^2}} \exp^{-\frac{(\ln x - \nu)^2}{2\sigma^2}} \tag{7}$$

Maximum likelihood estimation of Hawkes process models are often computed using numerical optimization [21], but we utilize an expectation-maximization strategy that has been shown to provide more robust estimates in addition to being more computationally efficient [28, 20], implemented in R [24].

Data

The data used for the Hawkes process model were prepared exactly as described for the regression analysis with one modification. The Wikipedia user interface has no means of resolving conflicts created by two simultaneous edits of the same content. This, combined with the potential loss of data due to dropped packets and other limitations of telecommunications technology leads many frequent editors of wikipedia to perform what would be interpreted as a single edit as a sequence of many smaller edits. Whereas in the regression model each edit event merely signals that an editor is active, the Hawkes process model explicitly models the time intervals between edit events. These micro-edits are part of a sequence and so are separated by time intervals much different from those separating distinct edit events. Because a sequence of micro-edits is intuitively a single edit action, we combine any sequence of edits by a single editor to a single page when the edits are separated by less than five minutes, leaving behind a single edit with the date and time of the last micro-edit in the sequence.

Results

The Hawkes process model provides evidence of substantial influence across Wikipedia editors based on the mutually exciting Hawkes process model. Tables 4 summarizes the parameter estimates for the two different communities of Wikipedia editors. We first consider the background activity rates of editors in the two Wikipedia communities. Although there is substantial variability, the overall background activity level is quite low, as is expected across a multi-year time period. Next, Figures 5 shows the distribution of the excitation parameter (β) estimates. Recall from equation (6) that $\beta_{1\to1}$ indicates the amount that an author's edits cause them to produce subsequent edits, while $\beta_{2\to1}$ represents the average amount that an author is caused to edit a page because of each other edit to the page. In both cases these values are the expected number of excited edit events. Although these values appear to be quite small, we will see that they are in fact become rather substantial when compared to the background activity rate.

Variable	1st Quartile	Median	Mean	3rd Quartile
WikiProject Oregon				
$\beta_{1\to1}$	0.177	0.378	0.396	0.617
$\beta_{2\to1}$	1.70e-17	1.44e-2	6.23e-2	6.81e-2
μ	1.67e-7	3.08e-7	7.14e-7	5.97e-7
ν	8.65	11.0	10.3	12.8
σ	0.0851	0.846	1.17	1.66
WikiProject Math				
$\beta_{1\to1}$	0.249	0.467	0.451	0.667
$\beta_{2\to1}$	6.99e-6	3.88e-2	8.45e-2	1.12e-01
μ	1.52e-7	2.68e-7	1.08e-6	6.70e-7
ν	9.65	11.2	11.1	12.7
σ	0.0734	0.741	0.902	1.65

Table 4. Summary of the Hawkes process parameter estimates

We used simulation to construct approximate statistical significance values for these excitation parameters. We simulated 500 sequences of 100 edit events using the Hawkes process specified in equation (6) with no cross-editor excitation ($\beta_{2\to1} = 0$). We then applied the expectation-maximization algorithm to compute a maximum likelihood estimate for the model parameters. The 95th percentile maximum likelihood estimate for the $\beta_{2\to1}$ was 0.042. Any maximum likelihood estimates for $\beta_{2\to1}$ greater than 0.042 are then statistically significant with an approximate p-value of 0.05. Although this ignores the impact of multiple comparisons, we believe this is a useful benchmark and as such have added a vertical line in Figure 5 at 0.042. WikiProjects Oregon and Math had, respectively, 91% and 92% of editors significantly self-exciting, as well as 35% and 48% significantly other-excited, all substantially greater than would be expected if there were no mutual excitement.

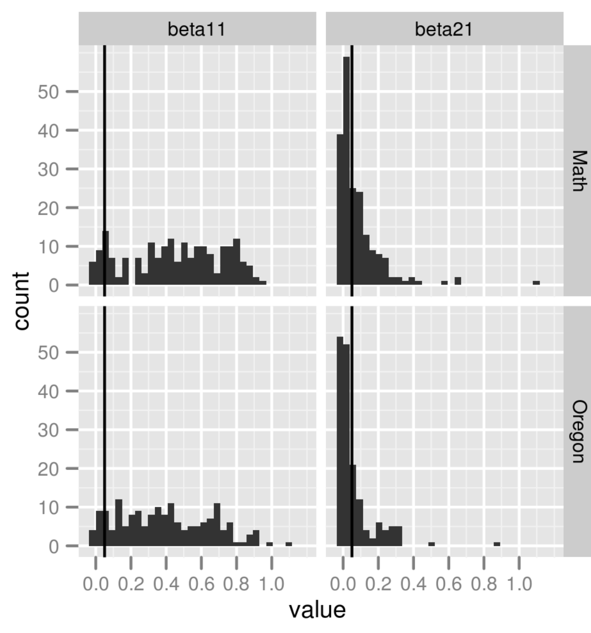

Figure 5. Histogram of the estimated self (beta11, left) and other (beta21, right) excitation quantities for editors in WikiProjects Math (top) and Oregon (bottom)

In addition to computing estimates of the amount of excitement or influence editors had on each other, the Hawkes process model also allows us to model when such caused events occur relative to the causing event. Recall that we model the time delay between excited and exciting events as a log-normal distribution, which means that the log of the time delay follows a normal distribution with mean ν and variance σ^2. This means that the average expected time delay between related edits is $e^{\nu+\sigma^2/2} = 16.8$ hours for WikiProject Oregon and 27.6 hours for WikiProject Math.

Using the average estimated parameter values, Figure 6 shows the impact of several edit events on an editor's intensity process. As you can see, although the absolute impact of each edit event is quite small, the cumulative effect of all the other editing activity on a particular editor's likelihood of producing an edit is quite substantial. For example, at around time 100 two quick edits by others more than double this editor's likelihood of editing. We can also explore the impact of other editors more quantitatively. Again using the average estimated parameter values, we can find that each edit by another editor increases the likelihood of contribution by 54% over the background activity rate.

$$\hat{\beta}_{2\to 1} \int_0^{24\ hours} g_{\nu\sigma}(v)dv = 0.0503 \tag{8}$$

$$\hat{\mu} * 24\ hours = 0.0929 \tag{9}$$

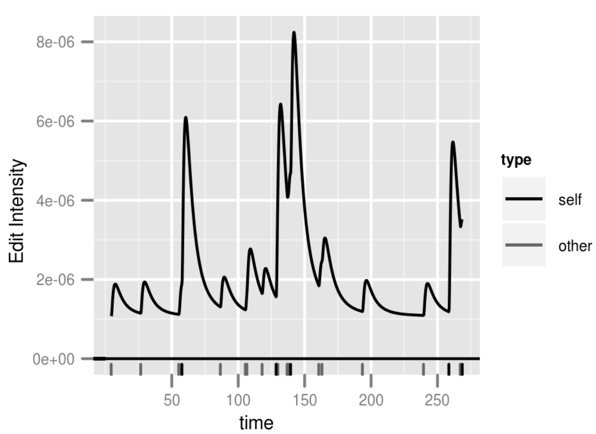

Figure 6. Simulated example of the Hawkes process model using the average estimated parameters. Lighter marks indicate edits by others while darker marks indicate edit by the self author.

DISCUSSION AND LIMITATIONS

Our two studies both found evidence for responsiveness between editors in Wikipedia. The logistic regression results indicate that the likelihood of an editor producing an edit increases by a factor of 2 to 12 when another editor was previously active. The Hawkes process model similarly identified 35-42% of the members of a WikiProject community as significantly influenced to participate in editing based on the activity of other members of the community. Although the observed effect was much smaller than that in the logistic regression, the Hawkes process model indicated an increase of 54% of the baseline daily editing level for each edit made to an article. The cumulative effects of such an increase combined over all edits on all of the WikiProject pages may indeed be comparable to the effect sizes identified in the logistic regression.

The two analyses differ dramatically however, in the temporal quality of response/influence discovered. For the regression model, the largest effect sizes were found when aggregating by 1–2 hour intervals with the effect decreasing substantially for larger time intervals. This would suggest that editors are most strongly influenced by a Wikipedia page's activity levels in the preceding 1–2 hour period. Such an effect would appear to be most consistent with the attentional and audience effects related to our hypothesis.

In contrast, the parameter estimates for the Hawkes process model imply much longer temporal influences across edits. Under the estimated Hawkes process model, a responding edit to an article would be expected to occur 16–27 hours after the original edit. This type of follow-up would seem much more likely to occur from the semantic effects, than either the attentional or audience effects that the logistic regression appears to capture.

There are a number of possible explanations for these divergences. First, the two models were in fact operating on slightly different versions of the edit history. For the Hawkes process model, we deliberately combined what we believe to be "micro-edits" occurring on the same page and by the same editor within 5 minutes in order to limit any possible difference in the responsiveness to one's own edits versus another editor's edits. We did this in order to learn a single model of each editor's response time rather than build additional models depending on who an editor was responding to. This, combined with our decision to permit each editor to have only a single distribution of response-times, may have restricted the Hawkes process model's ability to identify short-term reactions.

Second, limitations of the logistic regression model specified may have limited its ability to model long-term influences. The logistic regression model only models interactions between subsequent time periods, even when such time periods cover a 6 or 24 hour period. This limits its ability to capture longer time-dependencies since as the time interval increases, so to does the volume of edits occurring within any given time period, diluting any influence across time. Recall, too, that fewer than half of the editors in either WikiProject were determined to have been significantly influenced by the edit activity of others. It is possible that heterogeneity among the population of editors also prevented logistic regression from identifying longer time-dependencies.

It is possible that the synchronization we are seeing is resulting not from responses to visible edits but from exogenous factors driving contribution from multiple editors at the same time. For example if a Wikipedia page becomes newsworthy this could alter individual propensities to edit, elevating editors' background processes and producing the appearance of reactivity. Since we do not include data about the relatively popularity of a page at the time of the edits the data in this paper does not rule this out. However this effect is likely to be quite time limited, perhaps only a day or two. The long time frame in our data would seem likely to reduce the impact of such exogenous synchronization. Further, the choice of pages (Oregon and Math) do not appear likely to be relatively high in newsworthiness, certainly not with any regularity.

Despite these concerns, it is also reasonable to conclude that both models did indeed identify real and substantial patterns indicating long and short-term responsiveness between editors, plausibly caused by the attentional, audience and semantic effects we theorize.

Theoretical implications

Studies of motivation in online work have focused at a fairly high level of abstraction, more on generalized membership and post-hoc personal explanations of participation. The results of this paper indicate that processes leading to active contribution, and thus project success, may also be related to the interaction of the local material situation faced by participants with the near real-time state of the project and affordances its systems. In particular our approach suggests that participation in these communities may be less planned and more reactive than current motivation studies assume. This may provide an opportunity to explore the micro-impact of recent interesting findings on contribution, such the finding that volunteer contribution increases when paid participants, who seem likely to be more regular participants, are involved [29].

Such a perspective suggests a focus on how contribution to online work projects intertwines with contributors' daily activities and routines at their networked computers. One could examine affordances of this context that go beyond the visibility we examine in this paper. For example: To what extent does the affordance and experience of context switching play a role in participants' decisions about what to work on at a particular moment? If a participant is frustrated with other temporarily intractable work, does the ability to switch into another, easier, activity lead to an experience of self-efficacy, possibly leading to renewed tenacity on the original frustrating problem? Conversely research might find that context switching leads to low overall productivity due to constant interruption [8].

Practical and Design implications

The results reported in this paper provide strong support for the idea that systems aiming to attract workers and work would be well advised to make the work undertaken by others rapidly visible. Community managers are often looking for ways to increase participation in the community, yet have few tools to do so. This is especially true in volunteer-based communities where external rewards are not available and intrinsic interest is difficult to manipulate. The effect shown in this paper provides a potentially important and relatively easy to implement tool: systems should be designed to reach out to contributors and trigger contribution.

One simple and practical implication of our work is that community managers should consider encouraging any resources they control (even if it is just their own time) to engage in visible work fairly frequently, creating the circumstances for others to be attracted and to attract others in turn. Rapidly visible work may be more effective than spending time on activities like writing a weekly newsletter. Nonetheless there are two issues that suggest caution in approaches attempting to exploit the effect we identify in this paper.

First, attention is a scare resource; extremely noisy and insistent mechanisms are likely to oversaturate a potential contributor's capacity to pay attention; they may simply unsubscribe. Wikipedia watchlists strike this balance somewhat because they are per-page. This is likely to maintain relevance to potential participants, but also keeps the overall flow of notifications in check since activity on individual pages is generally temporally sparse. Action based on our findings should consider the possibility of overloading potential contributors. This suggests an opportunity for designing notification flows such that they consist of the best combination of message type, message content and occur in a temporal pattern most likely to motivate an individual. Experiments could generate variations in a notification flow and dynamically assess their effect on attracting work. In a simple form this might mean throttling notifications, reducing from all edits on a page to just edits near the recipient's historical contributions. In a more complex form that could include filtering out all but the edits most likely to spur added contributions.

The second issue derives from how potential participants perceive efforts to encourage additional work through notification. There is a danger that notifications will be perceived as explicit requests to work, rather than passive notifications or even system side-effects. Participants, especially volunteers, may baulk or even be offended by messages perceived in this way. More subtly, in a version of the Lucas critique, even when notifications are unchanged, common knowledge or even promotion of the effect discussed in this paper may, on its own, reduce its effectiveness, simply because potential participants are actively wondering whether notifications are causing them to turn their attention to the project. On the other hand such concerns may be unwarranted; the effect may be robust to recipient's perceptions or knowledge of it.

CONCLUSION

The similarity of the information systems underlying open contribution projects are not at all coincidental. In a very real sense they bring the project to the participant, generating reinforcing ripples of attention, reaction and contribution. This affordance is central to these projects' ability to compete with other demands on potential participants' time and attention. Online contribution systems offer the benefits of working independently without giving up key benefits of social context, including its ability to motivate and to trigger contribution of others.

REFERENCES

1. J. J. Amor, J. M. Gonzalez-Barahona, G. Robles, and I. Herraiz. Measuring libre software using debian 3.1 (sarge) as a case study: preliminary results. *Upgrade Magazine*, August, 2005.
2. G. Beenen, K. Ling, X. Wang, K. Chang, D. Frankowski, P. Resnick, and R. Kraut. Using social psychology to motivate contributions to online communities. In *Proceedings of the 2004 ACM conference on Computer supported cooperative work*, page 221. ACM, 2004.
3. L. Buriol, C. Castillo, D. Donato, S. Leonardi, and S. Millozzi. Temporal evolution of the wikigraph. In *Proc. Web Intelligence Conference 2006*, pages 45–51. IEEE CS Press, 2006.
4. M. Burke, R. Kraut, and E. Joyce. Membership Claims and Requests: Conversation-Level Newcomer Socialization Strategies in Online Groups. *Small Group Research*, 41(1):4–40, 2010.
5. R. Crane and D. Sornette. Robust dynamic classes revealed by measuring the response function of a social system. *Proceedings of the National Academy of Sciences*, 105(41):15649, 2008.
6. K. Crowston, K. Wei, Q. Li, U. Y. Eseryel, and J. Howison. Coordination of Free/Libre Open source software development. In *ICIS 2005.*, Las Vegas, NV, 2005.
7. S. Faraj and S. L. Johnson. Network exchange patterns in online communities. *Organization Science*, Forthcoming.
8. J. Gluck, A. Bunt, and J. McGrenere. Matching attentional draw with utility in interruption. In *CHI '07*, pages 41–50, 2007.
9. A. G. Hawkes. Spectra of some self-exciting and mutually exciting point processes. *Biometrika*, 58(1):83, 1971.
10. A. G. Hawkes and D. Oakes. A cluster process representation of a self-exciting process. *Journal of Applied Probability*, 11(3):493–503, 1974.
11. J. Howison. *Alone Together*. PhD thesis, Syracuse University, School of Information Studies, 2009. http://james.howison.name/.
12. J. Jin and L. A. Dabbish. Self-interruption on the computer: a typology of discretionary task interleaving. In *CHI '09*, pages 1799–1808, New York, NY, USA, 2009. ACM.
13. A. Kittur, B. Lee, and R. E. Kraut. Coordination in collective intelligence: The role of team structure and task interdependence. In *CHI 2009*, Boston, MA, April 3-9 2009.
14. J. Koh, Y. Kim, B. Butler, and G. Bock. Encouraging participation in virtual communities. *Communications of the ACM*, 50(2):68–73, 2007.
15. K. Lakhani and E. von Hippel. How open source software works: "free" user-to-user assistance. *Research Policy*, 32:923–943, 2003.
16. K. J. Lee. What goes around comes around: an analysis of del.icio.us as social space. In *CSCW '06*, pages 191–194, New York, NY, USA, 2006. ACM.
17. K. Ling, G. Beenen, P. Ludford, X. Wang, K. Chang, X. Li, D. Cosley, D. Frankowski, L. Terveen, A. Rashid, P. Resnick, and R. Kraut. Using social psychology to motivate contributions to online communities. *Journal of Computer-Mediated Communication*, 10(4), 2005.
18. L. Mitchell and M. E. Cates. Hawkes process as a model of social interactions: a view on video dynamics. *Journal of Physics A: Mathematical and Theoretical*, 43(4):045101, 2010.
19. O. Nov and C. Ye. Why do People Tag? Motivations for Collaborative Content Tagging. *Communications of the ACM*, 53(7), 2010.
20. J. F. Olson and K. M. Carley. Exact and Approximate EM Estimation of Mutually Exciting Hawkes Processes. Technical report, Submitted to Statistical Inference for Stochastic Processes, 2011.
21. T. Ozaki. Maximum likelihood estimation of Hawkes' self-exciting point processes. *Annals of the Institute of Statistical Mathematics*, 31(1):145–155, 1979.
22. Pew Foundation. Pew internet and american life project tracking project. Internet, 2010.
23. J. Preece, B. Nonnecke, and D. Andrews. The top five reasons for lurking: improving community experiences for everyone. *Computers in Human Behavior*, 20(2):201–223, 2004.
24. R Development Core Team. R: A language and environment for statistical computing. Technical report, R Foundation for Statistical Computing, Vienna, Austria, 2004.
25. J. A. Roberts, I.-H. Hann, and S. A. Slaughter. Understanding the Motivations, Participation, and Performance of Open Source Software Developers: A Longitudinal Study of the Apache Projects. *Management Science*, 52(7):984–999, 2006.
26. M. Salganik, P. Dodds, and D. Watts. Experimental study of inequality and unpredictability in an artificial cultural market. *Science*, 311(5762):854, 2006.
27. B. L. Sullivan, C. L. Wood, M. J. Iliff, R. E. Bonney, D. Fink, and S. Kelling. ebird: A citizen-based bird observation network in the biological sciences. *Biological Conservation*, 142(10):2282 – 2292, 2009.
28. A. Veen and F. P. Schoenberg. Estimation of Space-Time Branching Process Models in Seismology Using an EM-Type Algorithm. *Journal of the American Statistical Association*, 2006.

29. P. Wagstrom, J. D. Herbsleb, R. E. Kraut, and A. Mockus. The impact of commercial organizations on volunteer participation in an online community. In *Academy of Management Conference*, 2010.

30. J. Webster, L. K. Trevino, and L. Ryan. The dimensionality and correlates of flow in human-computer interactions. *Computers in Human Behavior*, 9(4):411 – 426, 1993.

31. R. Zajonc. Social facilitation. *Science*, 149:269–274, 1965.