

Sustaining Community – Incentive Mechanisms in Online Systems: Final Report of the Group 2005 Workshop

Organizers

Jason Ellis, Christine Halverson, Thomas Erickson
IBM T. J. Watson Research Center
{jasone | krysl | snowfall}@us.ibm.com

Participants

Mark Ackerman, University of Michigan ~ Ana Boa-Ventura, University of Texas ~
Andrea Forte, Georgia Institute of Technology ~ Max Harper, University of Minnesota ~
Jonas Holmström, Swedish School of Economics and Business Administration ~
Sheizaf Rafaeli, Haifa University ~ Leonie Ramondt, Ultralab ~ Bob Sandusky, University of Tennessee
~ Julita Vassileva, University of Saskatchewan ~ Leon Watts, University of Bath

INTRODUCTION

How online communities sustain themselves? Attracting people to an online venue is difficult enough, but once that has been accomplished a number of other questions arise. How do designers ensure that people return? What motivates people to participate? And how does one create an online environment where people not only participate, but participate in such a way that the community, collectively, makes progress towards its goals?

For the purpose of this workshop, we use “community” in the loose sense of a large group of people with generally weak ties. Thus we are concerned with everything from well-established applications such as chat rooms and help desks, to more recent examples of successful online systems like Slashdot, Friendster and Wikipedia.

Open source software provides a good example of what can result from successful answers to these questions. Open source has grown from a tiny experiment to a worldwide phenomenon in a remarkably short time. It is responsible for products – Linux, Apache, Firefox, Thunderbird – that have had broad impact. The sense of community developed in open source has been cited numerous times as a reason for how and why open source succeeds. But what does that mean? One of the assumptions driving this workshop is that at least part of what constitutes this “sense of community” has to do with subtle and entwined incentive systems that motivate and shape the participants’ behaviors.

But what are these incentive mechanisms? Are there different classes or types? What makes incentives work, and how are incentive mechanisms connected to the health of the community overall? The goal of this workshop was to bring together a wide range of participants to explore different community settings, as well as the incentive mechanisms that drive and sustain them.

WORKSHOP CONTENT

In the call for participation, we asked prospective participants to submit position papers that described a particular community that had exhibited sustained activity, and to discuss the ways in which incentives functioned (or failed to function) in supporting sustained activity. In the first section we provide thumbnail sketches of each accepted position paper, along with a link to the full position paper. The subsequent sections attempt to capture some of the issues raised during the workshop, and summarize and synthesize.

The Position Papers: Eight Sustained Communities

Here we provide thumbnail descriptions of the position papers in the order in which they were presented. The thumbnails describe the community and the incentives at play in it (indicated in bold italic typeface).

Comtella, Julita Vassileva

Comtella is a small-scale peer-to-peer online community originally developed for use in the classroom. It is used to share and comment on links to papers and other class resources. Comtella uses an explicit incentive system that rewards users for participation by assigning members a *user status that ranges from “plastic” up to “gold,”* depending on the amount a person has contributed. The amount of participation is also publicly displayed in *star-like visualizations of each participant*. In early versions of the system these incentives produced large numbers of low quality contributions, and so a second incentive mechanism – *c-points* – was introduced to encourage users to rate the papers. *C-points* are awarded for each act of rating papers, and they **can be “invested”** by the users in their own *contributions, to be displayed higher in search list results* (analogous to search engines “sponsored links” mechanism). For more information see “Adaptive Incentive

Mechanism for Sustainable Online Community," by Julita Vassileva.

MovieLens, Max Harper

MovieLens is an online community that uses a recommender system to generate personalized movie recommendations. One incentive for participation is that MovieLens users are told they will begin to receive *personalized recommendations* once they've done 15 ratings; furthermore, they're told that *more ratings will generate better recommendations*. The paper describes a survey of MovieLens users to understand why they rate movies; one of the interesting findings is that (setting aside the most popular reason which of course is to get *movie recommendations*) users say they rate movies because the process of *rating is fun* (48%), and because they like having access to *a list of movies* they've rated (54%). For more information see "User Motivations and Incentive Structures in an Online Recommender System," by Max Harper, Joe Konstan, Xin Li and Yan Chen.

Google Answers, Shiezaf Rafaeli

Google Answers is a fee-based information market where askers pose questions accompanied by a *bid* (from \$2 to \$200 a question), and "researchers" (vetted by Google) can claim a question and have 24 hours to provide an answer. Once the answer is produced, askers may provide an additional payment – a *"tip"* – if they wish, and can give the answer a *rating of up to five stars*; in addition, anyone can see and add *comments on the question/answer*. The paper analyzes behavior in Google Answers, and confirms that monetary incentives had their expected effect: higher priced and better tipped researchers were more likely to participate. However, they also found that if questions/answers that lack comments are excluded from the analysis, the social incentives (ratings and number of comments) were also positively correlated with participation. For more information see "Social and Economic Incentives in Google Answers," by Sheizaf Rafaeli, Daphne Raban and Gilad Ravid.

RePEc, Jonas Holmström

RePEc is a collaborative, all-volunteer digital library that contains meta data about working papers and journal articles in the field of Economics. In addition to supporting searching for and downloading papers (over 300,000 a month), RePEc offers an awareness service called New Economics Papers (NEP) that broadcasts alerts about new additions to RePEc to some 15,000 subscribers; NEP broadcasts are broken down by topic, and NEP editors receive a monthly mailing that details the **number of subscriptions** to each topic list. Another RePEc mechanism allows paper authors to **create their own web pages** that list their papers in RePEc, and provides each author with **download statistics** for their papers. For more information see "Why Does RePEc Persist?" by Jonas Holmström.

Mozilla, Bob Sandusky

Mozilla is one of the largest and most successful open source projects, having produced the Thunderbird email client and Firefox web browser. This paper examines software project management in Mozilla, and views it through the lens of Distributed Collective Information Practices. DCIPs include information compounding (i.e. the creation of bug report networks), negotiation and communal distributed search (e.g. to identify duplicate bug reports). Incentives that motivate Mozilla participants include *enhancing reputation*, *"scratching an itch"* (that is, fixing a problem that bothers the contributor), and working for *pay* (e.g. in projects that have hybrid volunteer/commercial organizational structures). For more information see "Forming and Sustaining Online Communities: The role of Distributed Collective Information Practices," by Bob Sandusky.

Wikipedia, Andrea Forte

Wikipedia is a large open content encyclopedia. In Wikipedia, anyone can begin a new article and anyone can change it; while this might be expected to be a recipe for chaos, in fact it leads to the production of well-formed articles. Drawing on interviews with 22 contributors, the paper explores the question of what motivates them to participate. While Wikipedia does not provide explicit mechanisms for recognition (e.g. there are no bylines on articles), it does provide mechanisms through which participants can and do construct claims for credit. In Wikipedia *registered users can construct their own "user pages,"* and it is here that they can claim credit ('I've started five articles and contributed to 20 others!'). Since Wikipedia includes an edit history for each article, others can verify such claims. Thus, in Wikipedia, the incentives for starting and contributing articles are not built into the system, but rather *the infrastructure supports the assertion and examination of claims*: in short, these incentives are socially mediated rather than being embodied as objects in the system. The paper discusses these and other phenomena, and looks at them in terms of Latour and Woolgar's notion of the "cycle of credit" derived from their studies of scientific communities. For more information see "Why Do People Write for Wikipedia? Incentives to Contribute to Open-Content Publishing," by Andrea Forte and Amy Bruckman.

Archnet, Anne Beamish and Ana Boa-Ventura

ArchNet is a web-based online community made of architects, planners and designers with a focus on the Islamic world. There are over 27,000 members from a remarkable array of countries, and thousands of unique non-member visitors each day. The paper explores the reasons users come to ArchNet – these include *reputation* (it is supported by a well know foundation), *quality of content* (it provides access to a world class collection of images), and *freshness of content* (significant financial investment supports a staff that can enter and update content such as an event calendar and job listings). The

paper also discusses where ArchNet falls short: while its participation in terms of browsing is high (and participants need to explicitly opt in once a year to retain membership), the developers have been disappointed in the level of contribution in terms of content and other forms of participatory use. The paper discusses reasons for the shortfall, and suggests methods for improving participation in ArchNet. For more information see "Building a Culture of Generosity: Activity, Participation and Sustainability in an International Design Community," by Anne Beamish.

Talking Heads, Leonie Ramondt and Leon Watts

Talking Heads is a community for head teachers in the United Kingdom. Its aim is to allow teachers to support one another by sharing knowledge and experience. The paper – which is generally concerned with supporting voluntary communities aimed at supporting social action and transformation – discusses the importance of narrative and storytelling in motivating participation and sustaining group engagement. In the Talking Heads environment, *the ability of participants to identify with one another's problems, and to offer solutions from their own experience* create an intense level of engagement. Other incentives considered in the paper include *the presence and visibility of an audience, the recognition of common experiences and problems, and the possibility effecting change in the outside world.* For more information see "Sustainability through Engagement: Storytelling Strategies as Incentives for Participation," by Leonie Ramondt and Leon Watts.

DISCUSSION

These eight communities, all of which exhibit sustained interaction, provided a wide range of examples of incentive mechanisms that served as a common ground for the workshops discussion. As is typical of workshops, the discussion was wide ranging: we covered a lot of ground, often in a non-linear fashion that is difficult to condense. For the purposes of providing value to others, we offer two sorts of syntheses: a framework for organizing incentive mechanisms, and a list of issues and questions.

Incentive Mechanism Framework

In trying to understand and categorize the incentive mechanisms, we found it helpful to think about incentives using a set of dimensions. We'll illustrate these frameworks using examples mostly drawn from the Comtella, Google Answers, and Wikipedia communities. We choose these three communities as sources of examples because they represent rather different approaches to incentives: Comtella uses deliberately constructed system objects to provide incentives; Google Answers uses monetary incentives; and Wikipedia's incentives are implicit and socially constructed.

Nature of Incentive: Instrumental or Symbolic

Some incentives have value because they allow you to take actions you otherwise could not; other incentives have

value because they are socially appreciated, or, to put it another way, they have symbolic value.

Thus, in Comtella, the c-points a user receives for rating others' contributions are an instrumental incentive because they will cause that user's contributions to appear higher in the lists of search results. In contrast, the Comtella "gold" level – that a user achieves only after making lots of contributions – is a symbolic incentive: while "gold level" users can't do anything they couldn't already do before becoming "gold", they do gain 'bragging rights' amongst their classmates. Note that an incentive may play both roles simultaneously: there is no reason that a higher Comtella user level could not provide the user with additional functionality as well as enhancing their social status.

However, it should be noted that this distinction is not completely clear cut; to see this, let's look more closely at the concept of the "cycle of credit" that was introduced in the Wikipedia position paper. Latour and Woolgar's concept of the "cycle of credit" suggests that, in the scientific community, credit functions in two ways: first, it is given and received in the context of peer-reviewing, publication, etc. as a sort of reward for previous activities; but second, this accrual of credit gives rise to credibility, which is linked with the individual's ability to act in the scientific community, and is manifested through grants, equipment, collaborative opportunities, and so on. To put this in our terms, incentives with symbolic value can, through the mediation of a community, produce instrumental value, resulting in the individual being able to do things that would otherwise be outside of his or her power. This is a nice illustration of why people care about incentives with purely symbolic value – because, in fact, they are more than symbolic. Nevertheless, we think that differentiating between incentives with symbolic value (in that their instrumental impact is mediated by a community), and those with instrumental value (where the incentive is non-socially connected to instrumental power) is a useful distinction.

The Scope of Incentives: Private, Communal, Societal

Another dimension that is useful for thinking about incentives is the scope of incentives. We will distinguish three scopes: private, communal and societal. That is, something may function as an incentive even if only the recipient is aware of it; these seem likely (though not necessarily) to be incentives that have instrumental value. Other things may function as incentives only to the extent that others are aware of them. In this latter case, some incentives may take on value because others in the community are aware of them, whereas others may take on value (or take on additional value) because people beyond the community are aware of them. We will refer to these incentives as having communal and societal scopes, respectively. Let's look at some examples.

In Comtella, users' are awarded status designators – from "plastic" to "gold" – depending on their degree of activity,

and a finer grained picture is signified using graphical star-like shapes that indicate the nature and quantity of their contributions. Thus, a member of a class that is using Comtella may take a private pleasure in having a “gold” user level and being represented by a large, brightly colored star. However, it is likely that the user will take additional pleasure in the fact that his or her high status is also visible to other members of the Comtella community; indeed, users may end up competing with other classmates to have the largest Comtella star (and, in fact, this incentive was sufficiently powerful that it led some participants to game an early version of Comtella). However, for people not in the Comtella community, the user level rating probably carries little weight – it may even seem odd that the user cares about it. Instead, if asked to explain their behavior to outsiders, Comtella users would probably refer to a different sort of incentive: to get a good grade in the class. While in fact a “grade” is at some level analogous to a Comtella “user level” in that it is an item on an arbitrary scale that represents performance, it is nevertheless the case that grades are societally recognized as having value, perhaps because of beliefs about connections with things such as jobs and social standing.

Representation: Designed vs. Constructed vs. Implicit

A third dimension that we can use to classify incentives is how they are represented in the system. Here we distinguish between built-in, constructed and implicit incentive representations.

Built-in incentive representations means that there is some sort of system object – usually associated with a visual indicator – that is created by and under the control of the system. Examples include the users star’s and user levels in Comtella. These appear whether the users are interested or not, and are not intended to be under the direct control of users (though, of course, users can game the system in an attempt to manipulate their incentive representations).

Constructed incentive representations are those that are created by users. Both Wikipedia and RePEc exemplify this approach with their variations on user pages, which are available to registered users of the system. User pages, and similar mechanisms, allow a system’s users to make claims about their contributions; often, the system will also make it easy to provide supporting evidence or for others to view the supporting information (e.g. Wikipedia’s article edit histories allows a Wikipedian’s claims to be checked).

Finally, there may be no incentive representations – that is, the system provides neither a visible representation of how one has contributed, or a means for portraying one’s own contributions. Instead, to the extent people receive kudos for their contributions, it is because other members of the online community observe and track their contributions. Incentives, in situations like these, can either be personal in scope (I know that I am doing something worthwhile), or to the extent my contributions are socially recognized, it is because people are watching and commenting. This type of

incentive may be at play for the Wikipedia contributors who do not construct user pages: those who are registered users may receive the kudos of other Wikipedians who are aware of their activities, while those who contribute anonymously presumably do so out of their own personal motivations.

Summary

Clearly, these dimensions are a rather loose way of classifying incentives. Many incentives do not neatly fall into one category or another – an incentive may be private and communal and societal, for example – nor do they really separate into discrete categories. Nevertheless, they seem like a start at a way of getting a handle on the notion of incentives.

Another issue that arises is what exactly we mean by the term “incentive.” In this workshop report we’ve deliberately used the term quite loosely. Sometimes we’ve referred to internal motivations as incentives (e.g. participants do things to ‘gain the respect of others’ or because ‘it’s fun.’), Sometimes we’ve referred to system objects as incentives: Comtella’s user levels and user star representations. Sometimes we’ve talked about user constructed objects – like Wikipedia or RePEc’s RAS user pages – as incentives. There is no doubt that more conceptual work would be valuable here, perhaps a model of incentives that involves connecting users’ internal motivations to their actions and depictions in the context of a system. However, that takes us beyond the scope of this work.

Issues and Questions

This final section is simply a grab bag of questions and issues that were raised during the workshop. We offer them here as an unordered list, which might provide useful starting points for the work of others.

Cultural and situational issues

Naturally, the role and function of incentive mechanisms vary according to culture and institution. In the case of ArchNet, whose membership is distributed across a diverse array of cultures, discomfort with public participation (as well as discomfort with ‘speaking’ in a non-native language) was hypothesized to be one of the barriers to more active participation. At the other end of the spectrum, a workshop member commented that the emphasis on encouraging participation struck him as odd: ‘in the culture I come from, that’s not a problem – the problem is to get people to shut up!’

Tailoring incentives?

To what extent do incentives need to be tailored for particular people, roles or groups? For example, it may be that newcomers need different incentives from old timers. How might one go about doing this?

The visibility of the incentive structure

To what extent do we want users to understand the overall incentive structure? Do you want them to understand it all? This could support gaming – although it might also discourage gaming because the gaming becomes obvious. Or is it better if users only have a ‘local’ understanding of the incentive structure?

What are the purposes of incentives?

- To encourage people to join the community
- To encourage members to stay in the community
- To encourage participants to be active members
- To encourage members to be active in particular ways
- To address particular problems (fire fighting)
- To encourage participants to act for the “public good.”

What is the role of time in incentive structures?

Should incentives change as the community ages? Should incentives shift according to the phase of activity the community is in? How does community self-regulation play out over time? This seems to be particularly important for designed mechanisms.

How do you control gaming the system?

It seems that mechanisms falling in the designed, communal and instrumental dimensions provide a more likely “target” for gaming.

Implicit vs. explicit incentives

Does making an incentive explicit incent some while disincenting others? How do we make incentives that are implicit repeatable and generalizable?

The dimensions of incentives and their effects

Are some of the dimensions of incentives – nature, scope, representation – predictive of certain effect sizes or shapes? For example, are private incentives necessarily weaker than communal for motivating certain types of contributions?

Incentives and Gaming

How do you control gaming the system? Does the nature of the incentive system play a role in whether users are willing to disrupt it? Are there communities in which gaming the system is more acceptable than others? Instrumental incentive systems in particular seem likely to attract abuses,

simply by virtue of offering community members an immediate reward in the form of new capabilities. On the other hand, symbolic incentive systems may be particularly dangerous when breached, since trust and interpersonal relationships play a role in defining the rewards. The scope of the incentive also has an impact on what repercussions there are if the system is gamed and how critical it is to control it. Private incentives have little impact if gamed; however, societal incentives potentially have a great deal of impact.

How do a community's incentive systems change over time?

Is there an incentive system life cycle? In the workshop, discussions turned to the changing nature of participation in a community over time; both for individual members and in terms of the community character as a whole. This implies that incentives to participate, too, change over time. If we hope to understand the nature of incentive in online communities, we need to be sensitive to the dynamic nature of communities and their norms.

Incentives and Roles

How do you incent people who are playing different roles? How do you incent people to move from one role to another? In a related vein, at any one time, communities sustain diverse memberships wherein one member's needs are met by another's contributions. If incentive systems are to be successful, they must likewise account for complex interrelated forms of participation.

How do participants experience incentive systems?

How do people read incentive systems; how do they individually understand the systems? To what extent are incentive systems understood in a common way among members of a community? To what extent do perceived incentives diverge among members of existing, healthy communities?

CONCLUSION

Online communities, and other networked groups of people, are here to stay. As online communities proliferate, and become a part of ordinary public interaction, a better understanding of incentive mechanisms can inform the design, management and governance of such networked associations of people.