# Authority Models for Collaborative Authoring

**Aaron Krowne and Anil Bazaz**
Department of Computer Science
Virginia Tech
Blacksburg, VA 24061
{*akrowne, abazaz*}*@vt.edu*

***Abstract*— We define an authority model, in the context of CSCW, as the characteristics of the system that determine who has control over objects in the system, and how this control manifests. In this paper we present two natural authority models (free-form and owner-centric) and, employing an empirical study, explore the question "which authority model is best for collaborative, commons-based content authoring?" We discover that the answer varies, depending on whether productivity or user satisfaction is to be the metric of goodness, and likely can not be phrased in absolutes.**

## I. Introduction and Background

During the past two years one of the authors (Krowne) has spear-headed the development of a collaborative online mathematics encyclopedia: PlanetMath [7]. The underlying software of PlanetMath is now being generalized and is called Noosphere [6]. Hence, Noosphere is the software, and PlanetMath is the instance. In this paper we are concerned with Noosphere and systems like it.

### A. The Noosphere System

Noosphere provides a somewhat generic basis for encyclopedic, mathematics-rich content for the commons-based peer production setting [1]. "Commons-based" means that the general expectation is that contributors are doing so by their own volition. This implies that the system lacks any sort of attention to kinds of motivation, such as features for supporting assignments. The second characteristic, which is "peer production," means there is spontaneous revision or review, in addition to just creation (the initiation of some kernel of content). Individuals producing the content are not only concerned with their own corner of the whole product, but are actually involved in improving the general quality by interacting with other authors and objects they have created. As we are interested primarily in this paradigm, we will limit ourselves to it in this paper. However, there are ramifications for CSCW in general, as many of these characteristics are shared in part by various (non-commons-based) CSCW systems.

Noosphere is built to handle the workflow surrounding revisions and corrections. Corrections can be filed by any user to any entry which was not created by them, and there are facilities for accepting or rejecting the corrections. In addition, there is a generic versioning system, with revisions (snapshots) and revision comments. There is a notification and watch system, whereby creators of objects (and other persons associated with them) are notified of messages posted to objects or messages they have created,or other events depending on the object type (for example, correction activity for encyclopedia objects). This system is extremely important for keeping users informed of activity and drawing them back into the site.

This research was partially supported by the Virginia Tech Department of Computer Science, the Systems Research Center, and NSF Grant No. DUE-0121679.

### B. Authority Models in CSCW

We introduce the term *authority model* for CSCW as the characteristics of the CSCW system that determine who has control over objects in the system, and how this control manifests. Previous CSCW work has touched on authority models partially. In [2], "access models" are briefly mentioned. We prefer the term *authority* to *access* because it has a higher-level, inclusive, more social connotation. It conveys the idea of hierarchies, and "who's in charge." Access, on the other hand, conveys little more than a "flat" sense of a Boolean grant or denial. From a technical standpoint, an authority model is always implemented in terms of access rules, but these are simply an encoding of motivating social notions of authority.

A strongly related concept is that of trust. Trust is a notion intimately bound up with authority and access: high access is a consequence of high trust, and the inverse also holds. Complex webs of trust determine access rules and result in the formation of authority structures.

### C. Authority Models in Noosphere

Noosphere has a hybrid authority model: it supports both of the two pure models we describe below, and has features for allowing users to selectively apply the models to all or some of their objects. The fact that Noosphere supports both models was handy for the carrying out of this study.

The first, and default model of authority, we term *owner-centric*. In this paradigm, the creator of an entry (the owner) controls all modifications to the entry, and revisions by others must be done by way of corrections (which are only suggestions as far as the owner is concerned). Owners may chose to globally or selectively expand the notion of ownership using the Access Control List (ACL) system, through which they may create editor groups for their entries and add selected users. Clearly this model resembles the familiar academic environment, where a single author, or a group of authors

(with a first author) creates and revises the document, and is subject to external review from a community of peers. This review takes the form of commentary which is then integrated by the author or authors; the reviewer makes no direct changes. The philosophy, then, is that the content is "hands-off" for everyone but the authors, but at the same time the authors acknowledge that they are not omniscient and would appreciate some input from the community. Some of the explanation for why this "hands-off" paradigm is desired may be the high value placed on uniformity and style of presentation as a distinct, unitary work.

The second model we call *free-form*. This model has gained popularity through the Wiki family of collaborative content management systems. The crux of this paradigm is that anyone may spontaneously make changes to any object, but the creator has the ability to roll back changes to protect against malicious or low-quality edits. This paradigm is supported in Noosphere through the same ACL system which allows granting of edit permission to single users or groups; however some of this complexity is hidden from the user and they have the ability to make the free-form model their global default. The philosophy behind this model is that communal is best: allowing people to instantly make changes to a document they come across will (allegedly) lower the barrier of content creation to its optimal level, raising the net productivity of the content management system. The owner in this paradigm is merely an "administrator," asserting an *a posteriori* sanity check on the content.

The aim of our study was to find out which of the above two models is "best," capturing the notion of best in a variety of ways. There are two key families of metrics to evaluate a productivity system: hard measures of productivity, and the impression the system has left on the user (that is, their satisfaction with it). In this paper we address both, using log analysis results and survey results, respectively.

### D. Related Work

Access control in the CSCW setting has been discussed as early as [4]. Here the facets of operation, role, object-user relationships, and data are discussed. All of these are dimensions that must be considered as attributes of access policies for an authority model. Eillis et al. [2] discuss the need for extending these ideas to real-time groupware.

[3] reviews many different types of trust, showing that trust is not a simple, one-dimensional concept. [5] is particularly relevant as it discusses trust in global virtual teams, which is very reminiscent of the system and scenario of this paper. [9] defines "swift trust," which we think often plays a part in global, common-based collaborative production. We will discuss these ideas more later.

## II. Empirical Study Methods

The general strategy was to test the two supported authority models in Noosphere (owner-centric and free-form) by having people work with them separately, gather some data for each case, and compare the results. To achieve this, modifications to Noosphere were made to allow it to operate in a single-authority model mode rather than the hybrid mode it was designed for, as well as to allow toggling between the two possible modes. This mostly consisted of hiding certain interface elements related to the specific authority model, depending on which mode the system was set to be in, and changing the default ACL parameters. Using this modified system, we were then able to have a set of participants work for a fixed time period in owner-centric mode, and then do the same in free-form mode.

We wanted to be able to draw as "real-world" of a conclusion as possible from this study, which meant that we needed to imitate the commons-based peer authoring setting as closely as possible. However, the study was necessarily artificial in nature, so it was really impossible to fully capture the spontaneous, un-coerced nature of peer production. We had to settle for approximating it. Some users were given writing "assignments," however, these assignments were a free selection of topic, either from a list of available topics or whatever they wanted to write about (depending on how they were recruited). Other users were not given assignments, but were urged to make an attempt at starting an entry. Indeed, many actually did this.

The strongest way we conformed to the real-world situation was the adoption of the remote, asynchronous setting. We never met most participants in person, the participants were anonymous to each other (indeed, some were in another country entirely), everyone worked from their own computer, and the Noosphere system operated asynchronously as usual. A daily reminder e-mail was sent out to each participant for purposes of coordinating the study, to which was also added the usual e-mail notifications from the Noosphere system itself. As these emails were chief elements of CSCW notification, it is clear that the overall interaction style was asynchronous.

The participants of this study were divided into two groups, with one group working first with the free-form model and one working with the owner-centric model. The experiment then progressed to the second part, where the groups worked with the "opposite" authority model. The division into two groups and subdivision into two parts was done to eliminate ordering biases; we wanted to account for initial learning curves with the Noosphere system and/or LaTeX [8], as well as effects coming from fall-offs in motivation.

An important detail is that the participants did not know the names of the authority models they were working with. In this paper, we use the terms "free-form" and "owner-centric," but we felt that using these terms in the experiment could skew subjective user feedback. These terms can be value-laden and may subconsciously affect users, depending on their philosophical leanings. So throughout the course of the study, we used the generics "system 1" and "system 2" to refer to the systems under each "pure" authority model.

Another element of anonymity was participant names. Though many participants didn't know each other, many were recruited from the same sources, and hence would already know a significant fraction of the other users. The

large incidence of this in our experimental population would seriously inhibit the ability to make inferences from the results to the general commons-based setting. We were wary of users partitioning themselves into sub-groups, comfortable only with working with others in these groups, while essentially ignoring the content and activity of the other sub-groups. We sought to eliminate this effect by randomly assigning to participants names of the form "user##" (where "##" represents a number). We think this masked any familiarity biases, and kept people focused on the content and its development.

We ran the study over the course of six days. The first three days made up the first part, and the second three days made up the second part. Before the first day, participants "signed" an online consent form to register for the study and assent to it. Next they were directed to a background survey, which asked questions about the person's experience with mathematics, commons-based peer authoring, and LaTeX. The answers to these questions were used to "evenly" distribute users into the two experimental groups, such that they consisted of about the same numbers of users, numbers of users participating for extra credit, experience in mathematics, knowledge of LaTeX, and self-motivation in mathematics.

The users created entries and/or corrected entries written by other users, as coordinated by the experimental coordination system. Each day they were emailed a reminder which also served as a link to their "assignment" for the day: this assignment was almost always to the effect of "browse the system" and "look for things to revise or correct". The exceptions were at the beginning of the first and second parts of the experiment, when some users were told to create new entries if they could. In addition, on the first day of the study, users were given some background material to read to introduce them to the study, to collaborative authoring and Noosphere, to LaTeX (if they did not already know it), and to the first authority model. At the beginning of the second part, the users were given a primer on the second authority model they would be working under. Care was taken not to refer to any abstract concept of "authority model".

At the end of the study, participants were asked to fill out an exit survey. This survey asked questions about which system the users thought was better for producing content, and how they felt about various aspects of each system. Because answers to the exit survey questions were in the form of "system 1" or "system 2" and these were meaningful only within the context of a particular experimental group, part of our data collection task was to translate these generic identifiers into the more theoretically meaningful handles "free-form" and "owner-centric".

### A. Participants

We recruited 25 participants for this study. Since we needed participants who had a background in mathematics and who either knew LaTeX or would be willing to learn enough of it to get by, we stuck to recruiting from the university setting. The participants were graduate and undergraduate students in the computer science department at Virginia Tech and undergraduates from Dalhousie University taking a linear algebra course. All the participants had good enough mathematical background to make contributions to the content (in principle). They had taken mathematical courses at various levels ranging from second year undergraduate mathematical courses to graduate level mathematical courses, with the number of courses varying from 3 to 15. The participants were experienced in using the Internet, with the extent of this experience ranging from 4 to 10 years. 76% of the participants had used web forums and 64% had looked up mathematics on the Internet. However, 80% of the participants who had never used any collaborative software or used a Wiki-style system. Thus, we had to provide significant guidance throughout the study. In addition, 56% of our participants did not use a computer to write mathematics and 76% had not used LaTeX. Since the participants would have to use LaTeX to write entries and a vast majority of our participants did not know LaTeX, we provided a basic tutorial for them on the markup language.

We must stress that this wide variation in prior experience is actually desirable for out for our study, as it mirrors the real-world situation.

### B. Procedure

We started out by sending a mass-mail (blind-carbon-copied) to all of the users who had expressed interest in participating, telling them how the sign-up procedure would work. When this was ready, we had the coordination system send to each user an email containing a unique hash key, which they'd have to paste into the online consent form to substitute for a signature on paper. The actual text of the consent was customized for the user; users recruited from a classroom who were "in it" for the extra credit were promised this credit for participation. Other participants were promised no compensation. The set of persons who went through this first step became the study participants, and they were subsequently given a link to the background survey for filling out.

Based on the results of the background survey, we divided up the participants into two experimental groups, taking care to evenly distribute experience and motivation in mathematics, those getting extra credit and those who were not, and experience with LaTeX. We also at this point generated random usernames for the participants. All of this was stored in the coordination database.

Soon thereafter we set the coordination system on sending out the daily reminder/assignment messages. The assignment messages for the first day contained links to required and optional reading materials on the study, the Noosphere system, and the first authority model (customized for the user). There was also a link directly into the appropriate Noosphere instance to finalize the creation of a pre-assigned, anonymous account.

The reminder/assignment messages continued daily. The base assignment always instructed the participant to inspect the system, giving special attention to new changes and new activity and checking for completeness, comprehensibility, style, and errors. On the first day of each half of the experiment,

however, the extra-credit participants were given more specific instructions to attempt to write an entry. For the participants recruited from the linear algebra class, we had a special list of topics to select from, which they had covered in their class. The others were able to select as they pleased from the requests list we had loaded into the Noosphere system, which consisted of a variety of topics from all over the mathematical spectrum in terms of subject and difficulty. It was made clear, in addition, that they were not restricted to any of these lists and could just pick something from a text book they had, or anything else they were interested in.

On the last day we sent out another mass mail, telling the participants to wrap up their work, and giving a link to the exit survey.

After we were convinced that all of the responses we were going to get came in, we closed the survey. The results were exported and then loaded into the coordination database. Participation statistics were also generated and loaded into the database (such as score and number of hits to the system). These were used for more in-depth analysis purposes.

## III. Results

In this section we present the distilled, empirical data of our experiment. It is divided into three subsections: participation figures, log results, and survey results.

### A. Participation Figures

In Table I, we present some participation statistics on the two experimental groups, for completeness and perspective. From this table, it can be clearly seen that there was a participation volume disparity between the two groups. As mentioned before, care was taken to form the groups with roughly even distributions of attributes which we thought were likely to influence productivity. However, with such a small sample size, there is a high potential "productivity granularity". In other words, if just one or two persons in one of the groups happens to be more productive, the impact will be large. Later we discuss how this imbalance does not seem to prevent drawing some meaningful initial conclusions from our study.

### B. Survey Results

In Table V we present the overall results of the exit survey. The left column gives a fragment of the actual survey question, which would be of the form "which system do you think had ..." or "which system do you think was ...", etc.. The middle column gives the counts of responses for each possible answer, for all users (the most popular response is in boldface). Note that the questions are divided into category sections for clarity. These sections are "quality", having to do with the quality of the content produced, "economy", having to do with the efficiency of producing the content, "authority", having to do with access and control, and "summary", which are capstone questions meant to capture a general sentiment.

Considering that actual participation varied considerably[1], we thought it prudent to also look at the results considering only the users who accounted for most of the activity of the study. Accordingly, in the third column of Table V we again give the results for the exit survey, but this time having filtered out users who had a Noosphere score of less than 100.

Note that the results for the set of most active participants are generally even more pronounced. Therefore, it is possible that either increased use leads to greater liking of the owner-centric model, or that the most productive tend to like the owner-centric model. We stop short of saying that this is the set of results that matters most, since it is not necessarily the case that the most active users are the only (or most) important users in a collaborative system. That is, the (typically large) class of slightly-active users may collectively have a very high impact in a collaborative production system. The dynamics likely vary from environment to environment.

TABLE I

Results for participation in the two experimental groups.

| | **Group 1** | **Group 2** | **Total** |
|---|---|---|---|
| Total requests[a] | 1402 | 4158 | 5560 |
| Total score[b] | 2139 | 6699 | 8838 |
| Total users | 12 | 13 | 15 |
| Active users[c] | 6 | 6 | 12 |
| Passive users[d] | 2 | 5 | 7 |

[a] A *request* is a web server hit.

[b] *Score* is a Noosphere metric of participation, measured in points. 100 points is given for each entry created, 30 for corrections, 10 for each revision, and so on.

[c] *Active users* are those who had a nonzero score.

[d] *Passive users* are those who had requests but no score; they were "lurkers". We include them because they still were able to provide some valid appraisal of the performance of the authority models, in principle.

### C. Log Analysis Results

In Table II, we present summary statistics of collaborative activity during the experiment, derived from the analysis of logs[2]. Note that numbers for the two experimental groups are summed together when the authority models matched, producing one set of numbers for each authority model.

It is our opinion that the "collaboration" metric here is the most important, but the others are provided for completeness. These figures fairly strongly suggest that the free-form model is the most productive, in quantitative terms.

To further investigate the issue of participation imbalance between groups, we provide Table III and Table IV. These tables break the collaborative activity statistics summarized in Table II down into the individual experiment phases, for each group. The columns in these tables are ordered according to the order the group worked under each authority model.

These latter figures seem to defeat the hypothesis that the participation disparity could have been caused by ordering bias. If ordering bias was present, then Table IV should show a

[1] An exponential participation falloff is observed in live virtual community systems, and is to be expected here

[2] Most of these do not literally come from log files, as there is a significant amount of online participation information within the Noosphere system's database.

decline in participation in the second phase of the experiment, mirroring the decline in Table III. Instead, participation rose, and the *relative* distribution of activity for the two groups matches very closely along authority model lines. Thus we conclude that the participation level disparity is exogenous with respect to authority models and is not in itself of concern in their analysis.

## IV. Discussion

The evidence seems to support the hypothesis that the free-form authority model is the most productive, but that the owner-centric model is preferred by users. However, a significant fraction of users depart from a preference for the owner-centric model (about a third), indicating that there is probably a place for supporting the free-form model as well, even from a pure user-satisfaction standpoint.

Most participants thought that the free-form model was easiest to learn, the least work for authors, the least work for correctors, and resulted in the quickest production, but significantly also thought that the owner-centric model lead to the highest quality of both content and presentation. The overall preference for the owner-centric model seems to indicate that people do not place as high weight on learning curve, turn around time, and revision effort as one might initially postulate.

All of these results could be said to support the claim that the authority model design of the system depends on the application. If the user's feelings are not quite as important as the "bottom line" of production, it may (surprisingly) be best to just support the free-form model. This would seem to justify the growing use of Wiki-flavored systems in corporate and learning environments (though perhaps *not* the academic environment). These environments also care more about the economic realities of production; placing a higher weight on quicker production time and less work for those creating and revising content. These are the *forté* of the free-form model, according to the responses of the participants.

Further, in commons-based settings where users are not *obligated* to contribute, it makes sense to cater more towards their feelings about content control and ownership, providing (and probably defaulting to) the owner-centric model. These settings also tend to stress less the economies of production mentioned above, so we get sort of a "purist" emphasis on correctness and quality that is less efficient to produce.

In general, a preference for greater control over one's own objects is exactly what we would expect to see, considering the issue of trust. We ran the experiment so that users were anonymous to each other, so the resulting setting was equivalent to that of a global virtual team [5]. Hence, participants had no basis upon which to trust each other. This gave us an opportunity to test the "swift trust" theory [9], which states that working teams can quickly build trust based on the quality of each other's work and little more. However, we found no evidence that swift trust was in operation at a higher level of influence than the authority models themselves. If it was, it should have manifested as an ordering bias, with

TABLE II

Collaborative activity analysis ("objective" results).

| | **Free-form** | **Owner-centric** |
|---|---|---|
| Number of revisions | **153** | 114 |
| Number of multiply-edited objects[a] | **15** | 1 |
| Number of objects created | **45** | 28 |
| Co-revisions[b] | **54** | 0 |
| Total corrections[c] | 11 | **28** |
| Total corrected objects | 9 | **16** |
| Collaboration[d] | **65** | 28 |
| Completed collaboration[e] | **62** | 13 |

TABLE III

Collaborative activity analysis (group 1 only).

| | **Free-form** | **Owner-centric** |
|---|---|---|
| Number of revisions | **68** | 11 |
| Number of multiply-edited objects | **5** | 1 |
| Number of objects created | **15** | 2 |
| Co-revisions | **30** | 0 |
| Total corrections | 0 | **4** |
| Total corrected objects | 0 | **2** |
| Collaboration | **30** | 4 |
| Completed collaboration | **30** | 3 |

TABLE IV

Collaborative activity analysis (group 2 only).

| | **Owner-centric** | **Free-form** |
|---|---|---|
| Number of revisions | **103** | 85 |
| Number of multiply-edited objects | 0 | **10** |
| Number of objects created | 26 | **30** |
| Co-revisions | 0 | **24** |
| Total corrections | **24** | 11 |
| Total corrected objects | 14 | **19** |
| Collaboration | 24 | **35** |
| Completed collaboration | 10 | **32** |

[a] A *multiply-edited* object is one that has been edited by more than one person. Note that in both Noosphere authority models this is possible, since in the owner-centric model, the owner can delegate editorship to others.

[b] "Co-revision" means a direct revision of an entry performed by someone other than the creator of that entry.

[c] Note that even in the Noosphere free-form model, corrections (*suggestions* for changes) can be used. In this model, they are still useful as something akin to a "bug tracking" mechanism, which can host discussions relevant to the correction topic.

[d] "Collaboration" in this context means co-revisions + corrections, as the two represent the primary collaborative means of content revision in the two models, and thus the metric facilitates direct comparison.

[e] "Completed collaboration" is the same as collaboration, but only counts closed corrections.

productivity rising in the second half of the experiment, when users had seen some of each other's work and done some work with each other. Yet, we saw that the in first group, which used the owner-centric model second, collaboration dropped precipitously from 30 instances to 4 by the second part of the experiment.

Were swift trust a major factor, it might also have the effect of influencing the survey responses such that users felt more comfortable with the latter authority model, unbeknownst to them, purely because they had more trust for their peers at this time. But there was no such order-sensitive effect in our survey data. These results strongly suggest the particular authority model is the dominating influence on productivity and impression. In other words, there must be certain facts about how people feel about trust in general and how authority models facilitate work that trump situational factors like swift trust.

Admittedly, we were not embarking to prove or disprove swift trust, so we did not optimize our study to isolate this factor. In addition, swift trust was not formulated to deal with a commons-based setting, so perhaps it does not apply to it (though, theoretically speaking, it seems it should). All we can claim here is that it seems to take a back seat to the effects of the choice of authority model.

The above conclusions must be qualified by the limitations of this study. Probably the biggest drawback of this study is the small sample size. The stark difference in productivity between the two groups was less than optimal to work with, though as discussed above, we do not think this effects the general thrust of our conclusions. Also, had we known in advance the rate of actual participation would be so low in each group (6-8 users out of 12-13), we probably would have foregone the division into two groups to ensure that the single group would have a "critical mass" of participation. This may have "activated" the participants in the first group more, though a test for ordering bias would have to come in a later study.

In addition, there is some doubt to be cast on the quality class of exit survey results, as the users had to do a comparison between the two systems from memory based on the impressions they had had over the week of the experiment. It would have been better if the users had been prepped for the exit survey by having them review all of the content that was produced, in a "re-cap" phase. In addition, they could at this point be asked to rate each item for subject-matter and presentation quality. Since these things were not done, it is possible that the difference between participant opinions and our quantitative measurements diverges more than it otherwise would.

This also raises the point that there was no official measure of quality, akin to the productivity counts we determined by log analysis (which measure only quantity). This could be achieved by having experts rate the content in much the same manner as the participants, then using the expert ratings as a yardstick against which to compare participant ratings.

### A. Conclusions and Future Work

In conclusion, we found that users generally *preferred* the owner-centric authority model in a collaborative, commons-based authoring setting. However, we found considerable evidence that the free-form model was actually *the most productive*, quantitatively speaking. Yet, even with the majority preference for the owner-centric model, a significant minority (about one third) preferred the free-form model. It was unknown which model actually yielded the highest quality, but users generally thought that the owner-centric model did.

We think these results are useful in the design of authority models for collaborative commons-based authoring systems, but that they also have some application outside of the commons-based setting to many types of CSCW. For example, we briefly discussed the corporate and academic settings, which would take advantage of different strengths of the two authority models which could effect their selection.

This work is a beginning, and we hope it will spark interest in authority models among the CSCW and related communities. Authority models have the potential to make or break a system, or at least dramatically effect user satisfaction and output productivity. Based on our small investigation and the modest existing literature, we think much work remains to be done, and many new avenues could be explored.

## V. Acknowledgements

We'd like to thank Mary Beth Rosson for her guidance in this study, and Dan Dunlap and Marry Beth Rosson for comments on early versions of this paper. We'd also like to thank Solomon Gifford and Robert Milson for allowing us to recruit participants from their classes and taking the time to coordinate the assignment of extra credit for these students.

## References

[1] Benkler, Y. (2002) Coase's Penguin, or Linux and the Nature of the Firm. *Yale Law Journal*, vol. 112. <http://www.benkler.org/CoasesPenguin.html>.

[2] Eillis, C. A., Gibbs, S. J., Rein, G. L. (1991) Groupware Some issues and experiences. *Communications of the ACM*, 34, 1, 38-58.

[3] Ford, D. P. (2002) Trust and Knowledge Management: The Seeds of Success, in *Handbook on Knowledge Management*, Holsapple, C., Springer-Verlag.

[4] Greif, S., Sarin, S. (1986) Data Sharing in group work. in *Proceedings of the First Conference on Computer supported Cooperative work* (Austin, Tex., Dec. 3-5). ACM, New York, 1986, pp. 175-183.

[5] Javenpaa, S. L., Leidner., D. E., (1998) Communication and Trust in Global Virtual Teams. *Journal of Computer Mediated Communication*, 3, 4.

[6] Krowne, A. P., et al. (2001) Noosphere: a collaborative content creation system for math and science. <http://aux.planetmath.org/noosphere/>.

[7] Krowne, A. P., et al. (2001) PlanetMath: math for the people, by the people. <http://planetmath.org>.

[8] Lamport, L., et al. (1994) LaTeX: A document preparation system. <http://www.latex-project.org/>.

[9] Meyerson, D., Weick, K. E., Kramer, R. M. (1996). Swift trust and temporary groups. in *Trust in organizations: Frontiers of theory and research*, eds. Kramer, R.M., Tyler, T.R., Thousand Oaks, CA: Sage Publications, pp. 166-195.

TABLE V

EXIT SURVEY ("SUBJECTIVE") RESULTS.

| **Question** | **Response Breakdown (all users)** | **Response Breakdown (most active users)** |
|---|---|---|
| | *Quality* | |
| Best content overall? | **Owner-centric - 41% (7)**<br>Free-form - 29% (5)<br>Don't know - 29% (5) | **Owner-centric - 54% (6)**<br>Free-form - 27% (3)<br>Don't know - 18% (2) |
| Best presentation quality? | **Owner-centric - 47% (8)**<br>Free-form - 29% (5)<br>Don't know - 23% (4) | **Owner-centric - 36% (4)**<br>**Free-form - 36% (4)**<br>Don't know - 27% (3) |
| | *Economy* | |
| Fastest corrections? | Owner-centric - 17% (3)<br>**Free-form - 64% (11)**<br>Don't know - 17% (3) | Owner-centric - 18% (2)<br>**Free-form - 72% (8)**<br>Don't know - 9% (1) |
| Least work for authors? | Owner-centric - 23% (4)<br>**Free-form - 41% (7)**<br>Neither - 17% (3)<br>Don't know - 17% (3) | Owner-centric - 9% (1)<br>**Free-form - 54% (6)**<br>Neither - 27% (3)<br>Don't know - 9% (1) |
| Least work for correctors? | Owner-centric - 29% (5)<br>**Free-form - 41% (7)**<br>Neither - 5% (1)<br>Don't know - 23% (4) | **Owner-centric - 45% (5)**<br>Free-form - 27% (3)<br>Don't know - 18% (2)<br>Neither - 9% (1) |
| Most difficult to learn? | Owner-centric - 25% (4)<br>Free-form - 6% (1)<br>**Neither - 56% (9)**<br>Don't know - 12% (2) | Owner-centric - 10% (1)<br>Free-form - 10% (1)<br>**Neither - 70% (7)**<br>Don't know - 10% (1) |
| Most difficult to use? | Owner-centric - 23% (4)<br>Free-form - 11% (2)<br>**Neither - 58% (10)** | Owner-centric - 18% (2)<br>Free-form - 18% (2)<br>**Neither - 63% (7)** |
| | *Authority* | |
| Best control? | **Owner-centric - 64% (11)**<br>Free-form - 29% (5)<br>Don't know - 5% (1) | **Owner-centric - 81% (9)**<br>Free-form - 18% (2) |
| Most comfortable with others' level of access to your entries? | **Owner-centric - 58% (10)**<br>Free-form - 29% (5)<br>Don't know - 11% (2) | **Owner-centric - 81% (9)**<br>Free-form - 18% (2) |
| | *Summary* | |
| Overall best system? | **Owner-centric - 52% (9)**<br>Free-form - 35% (6)<br>Neither- 5% (1)<br>Don't know - 5% (1) | **Owner-centric - 72% (8)**<br>Free-form - 27% (3) |
| Which system should be supported? | **Owner-centric - 52% (9)**<br>Free-form - 35% (6)<br>Don't know - 11% (2) | **Owner-centric - 72% (8)**<br>Free-form - 27% (3) |