

Open Source Software Communities:
A Dissertation Analysis and Study Proposal

Adam Worrall

Florida State University

ISM 6109

12/7/2010

Open Source Software Communities:

A Dissertation Analysis and Study Proposal

Glen Sagers's (2007) dissertation, "Is Bigger Always Better? Toward a Resource-Based Model of Open Source Software Development Communities," asked two overall questions: first, how are open source software (OSS) communities sustained? and second, are such communities necessary for the software they develop to succeed? He adapted Butler's resource-based model of online communities to apply not just to the sustainment question, but also to the success of the software product and the community itself. Sagers believed that more communication activities would take place in larger communities, which would lead to both more successful software and a more successful community. This would, in turn, increase the resources (members) available to the community, resulting in a cyclical model. Thus, his overall hypothesis was that the size of the active community would be a critical factor in sustaining the OSS project and in making it successful. This paper summarizes and analyzes Sagers's dissertation study, then proposes a follow-up study that builds upon and improves it.

Summary

Chapter 1: Introduction

Sagers's (2007) dissertation began with an introductory chapter which briefly introduced open source software (OSS), noting that its success "suggests that collective action can be sustained by a community of volunteers" (p. 2). However, despite other research on OSS, the sustainment of these volunteer communities and their influence on the software's success had not been explored. Other, "typical communities" faced issues with social loafing and difficulties in communication, but some believed OSS communities did not face these issues to the same degree, the benefits of a large volunteer community greatly outweighing them. Sagers concluded

this brief chapter by summarizing the questions he intended to answer—discussed above—and the practical significance such answers would have for software development.

Chapter 2: Conceptual Model and Theoretical Foundation

The second chapter of Sagers's (2007) dissertation featured the majority of his literature review. He began by presenting the concept of open source software (OSS) itself along with the Open Source Initiative's (OSI) definition of OSS, "a community-accepted definition of what OSS licenses should allow" over freeware and commercial software (p. 6): (a) free redistribution and availability of the software and its source code, (b) the ability to modify the software, and (c) no discriminatory practices. Sagers also briefly compared OSS development with commercial (or "proprietary," p. 7) software development, noting major differences: OSS is developed by volunteers who "are not bound to the project by any type of formal contractual agreement," self-assign their tasks, and design and develop the software in an evolutionary manner intended to lead to continual improvement (p. 7).

OSS as public good. Sagers next conceptualized OSS as a "public good" (p. 7), given that it is "not used up in its consumption" nor is anyone excluded from its use (p. 8). He noted that "public goods are typically subject to ... free riding behavior" (p. 8), and that many users use OSS but do not contribute back to the OSS project. Governments handle this problem by levying taxes to fund production, but OSS projects do not have such power. Nevertheless, they still exist due to the phenomenon of "disjunctive production" (p. 9): a few motivated members of the community will selflessly produce the software to benefit all, including free riders. Sustaining OSS as a public good, however, is a trickier problem. Sagers argued that OSS is an "anti-rival" good (p. 10), benefiting from network externalities as the number of users increases,

they share the software with others, and they improve the software by finding and reporting bugs (p. 11). At least a few users will continue to contribute and thus sustain the community.

Communities and projects. Sagers next defined communities as “groups of individuals who work together to help each other achieve a common purpose” after Cothrel and Williams (1999, as cited in Sagers, 2007, p. 12). Individuals interacting “within a community ... form social bonds with others ... often referred to as social capital”; these form the basis of feelings of trust and a sense of belonging (p. 12). Sagers also briefly discussed the life cycle of communities, noting “attraction and retention of membership is vital” for the sustainment of the community over its life (p. 12). Sagers defined the “ideal type” of an OSS project to “consist of a software program ... under an ‘approved’ license ... developed by a community of ... (mostly volunteer) developers who meet (primarily) via the Internet” (p. 14). Common elements of such projects included a Web site (the meeting place for the community); the software itself; and codified knowledge: documentation, communication (e-mail lists, forums, etc.), and databases of bugs and proposed features. The OSS project community consists of both developers and active users, with other, free-riding users falling outside the project community as conceptualized.

Prior OSS research. Sagers (2007) next reviewed previous research that had studied open source software. He broke this into five categories and “highlight[ed] each research stream” rather than including “all literature on OSS development” (p. 17). He first found business and advocacy-related studies to be of “limited value in scholarly research, except to illustrate trends” (p. 17). Second, demographic and motivation studies had found that most contributors were volunteers and had various motivations: “recognition from peers, the desire to improve one’s skills, a better reputation within the community ... [and] belief in the OSS philosophy” (p. 18); he felt, however, that “a complete analysis” of these was beyond the scope of his dissertation (p.

18). Third, studies of the economics of OSS had found a degree of private economic benefits to both individual volunteers and firms who had open-sourced their software. Fourth, code quality studies often found OSS code to be of higher quality than proprietary code. Finally, some OSS projects were organized in hierarchies—despite this not being the “vision” of many OSS advocates—and others had “onion-like layers” or featured a “social network”-like organization (p. 19); governance relied on “social norms to induce compliance” and on recruitment of volunteers (p. 19). Sagers (2007) thus identified a gap in the previous research: no empirical investigations of whether a larger, “active ... community will improve the success of OSS projects” (p. 25). He thus formulated his two research questions (p. 26) as presented earlier.

Resource-based model. Sagers next reviewed Butler’s “Resource-Based Model of Online Social Structures” and related literature (p. 27). It proposed that a sustainable community “must provide positive benefits to its members ... [that] outweigh membership costs in order for the volunteers to remain with the community and allow growth” (p. 27). Resources such as “knowledge, time, energy, money, and material resources” (Butler, 2001, p. 347, as cited in Sagers, 2007, p. 27) must be aggregated by the community; more members leads to more resources. These resources are then converted to benefits such as social capital and other “support structure[s] for collective action” (p. 28). This conversion requires communication, so Sagers proposed that more resources and communication—hypothesized to increase with size—would result in more benefits and success for the project. As size increased, however, the “increased communication activity” and resources would require additional management (p. 28). Citing key literature on success in software development, Sagers noted the primary benefits would be the attraction and retention of participants. Given the importance of such success,

Sagers added it to Butler's model as an additional step after the creation of benefits and before attraction and retention (shown in his Figure 3, p. 31).

Chapter 3: Research Model and Hypotheses

Sagers's (2007) third chapter continued the literature review and presented additional details on his research model (p. 42) and hypotheses. He particularly examined literature on factors from his research model and their interrelations, including resources, communication activities, success, and sustainability. He noted that support and development activities were "two primary" and important "types of communication activities" in OSS projects (p. 34), hypothesizing a larger OSS community would feature more of these. Based on the literature he argued that a successful OSS project required both software and community success; he hypothesized that greater amounts of each type of communication activity would lead to greater software and community success and more social capital—operationalized as "trust and a sense of belonging" (p. 38)—within the community. Finally, he hypothesized more successful software and higher social capital would lead to greater retention and attraction of members.

Chapter 4: Methodology

Sampling. Sagers (2007) chose to select projects hosted from Sourceforge, a commonly used OSS project hosting service, restricting his population to those with active communication as indicated by inclusion on Sourceforge's "Top Forum Posts Count" list (p. 46). This list of 100 projects was reduced to 44 by purposively sampling those meeting three criteria (p. 46): (a) use of the "'Open Discussion' forum"; (b) no use of other forums or "venues for communication"; and (c) at least 10 posts made to the forum "during a six-week period." Sagers did admit a built-in bias towards active projects, but felt they showed sufficient variety to ensure generalizability.

Data collection and measures. Sagers (2007) used a multi-stage, longitudinal collection process. The initial community size was measured by counting the number of unique usernames in a collection of all non-anonymous messages posted to the project's forum between September 15 and October 31, 2003. Support communication activities were measured through a second set of messages collected between November 15 and December 31, 2003; the activity level was measured through post count, while the user name, thread title, and timestamp for each post were also collected. The Herfindahl-Hirschman index, originally used by Butler in his 2001 study, was used to analyze thread titles to determine "the diversity of topics within the forum" (Sagers, 2007, p. 49). Development communication activities were measured through counting the combined number of bug reports and feature requests made during the same time period.

Software success used four total measures, two subjective and two objective. The latter included (a) the percentage of the bug reports and feature requests collected earlier that were closed as of February 18, 2004; and (b) the number of downloads of the software during January, 2004, used as a proxy for the number of users. One subjective measure employed a short online survey, sent on January 1, 2004 to all non-anonymous users previously identified as forum posters. Hartwick and Barki's (1994) four-item attitude scale was used, labeled "software performance and utility" by Sagers (2007, p. 51). The second subjective measure was taken from Sourceforge data and measured changes in development status (e.g. from alpha to beta or beta to stable) between September 15, 2003 and February 18, 2004, when the survey was closed.

Community success was measured through further questions on the survey, focusing on trust and sense of belonging as elements of social capital (Sagers, 2007). Trust was measured via a 4-item Likert scale drawn from a 1998 study by Jarvenpaa et al., while sense of belonging used a 3-item Likert scale drawn from a 1996 study by Jones et al. Finally, attraction and retention

were assessed by comparing usernames that posted messages from June 1 to July 15, 2004 with those from the initial period, calculating the number of members attracted and retained.

Chapter 5: Analysis and Results

Sagers (2007) had to drop 5 of the 44 projects due to their moving away from Sourceforge or the forums before all data were collected, “leaving 39 projects for analysis” (p. 55). The survey link was sent to 1,724 users, with 122 “returned as undeliverable”; 355 “usable responses were returned” for a response rate of 22% (p. 55). No significant response biases were found “compared to findings of other research,” based on chi-square testing (p. 55).

PLS results. Sagers (2007) used partial least squares (PLS) structural equation modeling to test his hypotheses, after “aggregat[ing] the survey data by project” with appropriate precautions (p. 56). Due to sample size issues, he could only include two predictors of software success; Sagers chose to include the number of downloads and Hartwick and Barki’s attitude scale, since they appeared to be the best predictors in an initial analysis. The resulting PLS model featured adequate convergent, inter-item, and discriminant validity as shown by appropriate statistical testing. However, some indicators did load above a 0.50 cutoff on other scales, albeit more highly on the intended scale (see p. 60); this issue is not addressed by Sagers.

Sagers’s (2007) testing of the PLS model found some, but not all, hypotheses to be supported; the significant paths in his research model are shown in his Figure 11 (p. 61). The hypotheses relating size to increased support and development communication activity were supported (at $p < 0.05$ and $p < 0.01$ respectively). The hypotheses relating these communication activities to software and community success, however, had mixed results; only the relationships between bug report counts and number of downloads (at $p < 0.10$) and topic variation and sense of belonging (at $p < 0.05$) were found to be significant. Many of the hypotheses relating software

and community success to attraction and retention were supported, but not all. Both the number of downloads (at $p < 0.01$) and the reported sense of belonging (at $p < 0.10$) were significantly related to attraction of new members, but Hartwick and Barki's (1994) attitude scale and the reported trust level were not significantly related. Trust did highly significantly relate to retention (at $p < 0.01$), but no other success measures did.

Exploratory analysis of other paths. Sagers (2007) also used a seven-step “exploratory analysis of other possible paths between model constructs,” since a number of his hypothesized paths had not been supported (p. 65). The results (p. 69) found only one additional significant path: a direct relationship between the size of the community and retention ($p < 0.05$). The exploratory model also explained less variance in downloads, but more in trust and attraction.

Chapter 6: Discussion

In this chapter Sagers (2007) related the findings to the literature and considered their theoretical implications. Since larger OSS communities communicated more about both support and development, he noted this proved that “the time, energy, knowledge and other resources provided by members [would] be converted to benefits for the members and for the community as a whole,” as predicted by Butler and others (Sagers, 2007, p. 71). The direct relationship between size and retention found in the exploratory model, however, contradicted Butler's findings. Sagers believed this was “due to the differences in [what was] produced in each type of community” studied (p. 72); Butler's listserv communities produced knowledge, a process that is less visible and harder to observe, while OSS communities produce software, a process that is more directly observable and visible to outsiders. Sagers (2007) further found the non-significant “path between a sense of belonging and retention” to be “especially surprising” based on the “predictions of volunteerism [in the] literature” (p. 73); his suggested explanation drew on

Powers's research, arguing that most volunteers in OSS projects are "spot volunteers" who choose their own tasks and volunteer only when "affected by a bug or [needing] a new feature," not contributing as frequently as those "who remain involved with the project" (p. 73). He also noted the lack of a path between subjective performance and utility and attraction indicated community members "may not actively spread the word about their project" (p. 74); even if they do, it is possible others may not be convinced by their efforts to use the software as well.

In discussing those hypotheses which were not supported, Sagers (2007) considered the possibilities that users participating in support communication activities did not directly move from forum threads to reporting bugs as often as expected, did not need to download the software they already used, or already felt unsatisfied prior to posting, thus leading to the lack of paths found between support communication activities and software success. A significant path *was* found between topic variation and sense of belonging, but it was in the opposite direction to that hypothesized. Sagers suggested this could be due to higher topic diversity leading to a lack of "common interests within the community" (p. 75). He also argued the lack of a significant relationship between development communication activities and community success could be due to most survey respondents being users rather than developers, and thus not heavy users of the bug reporting or feature request sections. Finally, Sagers noted the significant, but negative path found between sense of belonging and attraction "could indicate that communities whose members feel a strong sense of belonging become closed ... not accept[ing] newcomers" (p. 76). Similarly, the lack of significance of the relationship between trust and attraction implied that trust was "not apparent to outsiders and therefore [could not] induce them to join" (p. 76).

Chapter 7: Conclusion

In his final chapter, Sagers focused primarily on the limitations and contributions of his study. While not all of the limitations were discussed, he did identify quite a few (pp. 78-79): (a) the small sample size leading to issues of generalizability and in detecting small effect sizes; (b) the limited time period—nine months—covered by the study and the low frequency of samples obtained, which did not show patterns and the life cycle as well as “more frequent samples ... over a longer period of time” would; (c) the possibility of non-response bias; and (d) data collection was limited to only the “Open Discussion” forum, ignoring other venues. Sagers did feel his study helped fill the gap in the literature he identified, confirm some of the findings of others in related studies, and contribute to theory by extending Butler’s model and supporting his findings (p. 80). He summarized the answers to his research questions: “OSS project communities are sustained by the interaction of members”; “member growth and retention depend on the social capital generated ... [and] upon the success of the software,” (p. 79); and “an active community does matter in producing, supporting, and sustaining OSS” (p. 80).

Sagers also briefly discussed potential follow-up research. He said “the first step ... should be” replication with larger numbers of OSS projects, both hosted on Sourceforge and hosted elsewhere (p. 81). This would overcome limitations of low statistical power and improve generalizability. He also argued a longer time period should be used, despite its “daunting” nature (p. 81). Finally, he suggested utilizing “different constructs for software success and community success” and for “support and development communication” (p. 81).

Analysis

Sagers’s dissertation most obviously crosses the topics of system success (DeLone & McLean, 1992; Seddon, 1997) and systems development and implementation (Cooper & Zmud,

1990; Curtis, Krasner, & Iscoe, 1988; Hartwick & Barki, 1994; Kemerer, 1987; Orlikowski, 1993). The relation to the former is clear from Sagers's (2007) research questions, particularly the second's relation "to the success of an OSS project" (p. 26). This falls under multiple categories of DeLone and McLean's (1992, p. 62) taxonomy of information systems success measures, as well as outside it to a degree. Sagers's antecedents measuring support and development activity by counting messages, bug reports, and feature requests could be considered proxies for system use, falling in that category of DeLone and McLean's taxonomy. His measure of the number of downloads also serves as a proxy for use. The subjective performance measures taken from Hartwick and Barki (1994) fall under user satisfaction, albeit they were originally intended to measure the related concept of attitude toward the system (see Hartwick & Barki, 1994, p. 463). The measures of community success and social capital drawn by Sagers (2007) from Jarvenpaa, Jones, and their colleagues—trust and sense of belonging—fall less directly into DeLone and McLean's system-centered categories. Since they are individual-level measures, they are probably best considered under the category of individual impact; nevertheless they have a social component that goes beyond most of the literature DeLone and McLean reviewed. In addition, the measurement of attraction and retention rates at the end of the data collection serves as a measure of organizational impact, albeit not focusing on performance as suggested by DeLone and McLean.

Comparing Sagers's (2007, p. 54) research model against the information systems success model developed by DeLone and McLean (1992, p. 87), it can be seen that Sagers left out system and information quality altogether and mixed use, user satisfaction, and individual impact together. Seddon's (1997) respecified version of DeLone and McLean's model is perhaps a better fit with Sagers's research. One critical change Seddon made was to conceptualize

benefits to individuals, organizations, and society as separate constructs that impact on perceived usefulness and user satisfaction (p. 245). This parallels Sagers's consideration of how support and development communication activities create benefits that then have impacts on software and community success (see Sagers, 2007, p. 54). In addition, under Seddon's model attraction and retention would be characterized not as organizational impact measures, but as measures of continued use of the software and participation in the community, following Seddon's feedback loop. Sagers's research model thus parallels Seddon's model much more closely than DeLone and McLean's original model.

The papers on systems development and implementation have varying degrees of relatedness to Sagers's dissertation, with at least one (Kemerer, 1987) having no direct relevance. Cooper and Zmud's (1990) study examined the adoption and infusion of MRP systems; while the diffusion of OSS software projects relates to software success and attraction of further users, the focus of Sagers's study was not on the diffusion of the software to a larger community. A follow-up study could consider the feedback loop in Sagers's and Seddon's models and the role that diffusion of innovation theory might play in the growth processes of an OSS project community; however Sagers's model and methodology is more directly applicable. Orlikowski's (1993) study shares a longitudinal approach with Sagers's, but little else; her focus on organizational change and usage of a qualitative and interpretivist approach to describe the adoption and use of computer-aided software engineering (CASE) tools has little direct bearing on an analysis of Sagers's findings, only serving as to guide a potential follow-up (e.g. to study the adoption of tools within the Sourceforge hosting system by OSS projects).

Curtis et al.'s (1988) study is somewhat more relevant; they took an ecological, behavioral, and qualitative approach to the study of software development, examining the

problems that occurred at multiple levels of analysis and determining how performance could be improved. One conclusion they drew, that “software systems must be treated, at least in part, as a learning, communication, and negotiation process,” backs up Sagers’s (2007) consideration of communication activities as a key element of his model. The latter’s focus was on a quantitative study of systems success, however, rather than a qualitative study of the processes used to develop a system. While content analysis of the messages Sagers collected could provide insight into the development process and the communication activities that took place within it, this fell outside the scope of his dissertation.

Hartwick and Barki’s (1994) study has the clearest importance to Sagers’s (2007) dissertation, not least because he used a portion of their instrument to measure “subjective performance and utility”—attitude, in reality—of OSS software projects. They proposed that “user involvement [served] as an intervening variable between user participation and system use” in information systems development (Hartwick & Barki, 1994, p. 440), developing a model drawing from previous literature and Ajzen and Fishbein’s Theory of Reasoned Action (TRA). Their testing of the model is best illustrated in their Figure 4 (p. 450); initial attitudes and user involvement were not found to influence participation in development, but “participation ... was found to influence [post-implementation] involvement and attitude,” with those who participated developing “higher levels of involvement and more positive attitudes toward the system” (p. 450). Their findings thus imply that participation in an OSS community will—regardless of initial conditions—lead to greater involvement, more positive feelings about the software, and overall continued use and participation in the community and the development of the software. However, attitude fell out of Hartwick and Barki’s model, leaving substantial question as to why Sagers used it in the first place; its falling out of *his* model is not surprising at all in this context!

The falling out of attitude parallels more recent research on the Technology Acceptance Model (TAM; see e.g. Adams, Nelson, & Todd, 1992; Davis, Bagozzi, & Warshaw, 1989). Sagers (2007) unfortunately did not include any discussion of TAM besides a brief citation of Davis's (1989) scale construction article. While the focus of his dissertation was not on predicting the acceptance and use of the OSS software projects he studied, their adoption, diffusion, and acceptance certainly should have played a role in the success of the software products—both in terms of downloads and in terms of the subjective measures of performance—and in the attraction and retention of members. The most recent extension of TAM, the Unified Theory of Acceptance and Use of Technology (Venkatesh, Morris, Davis, & Davis, 2003), does include social influences and the voluntariness of use as factors, also shown by Sagers (2007) to factor in the success of OSS projects and communities. Unfortunately, UTAUT (like TAM before it) does not explicitly consider longitudinal use behaviors. Nevertheless, the factors identified in it as significantly relating to use could—and given its high reliability and validity and the lack of significance of attitude as a construct, should—be used as measures of product success in future research based on Sagers's model.

There are also a few other issues with Sagers's (2007) dissertation to consider in the context of possible follow-up studies. Chief among these is the relatively low importance placed on the motivations of those participating in the community and thus in creating benefits for the OSS project. Sagers does touch on these a few times in his literature review (see pp. 11, 18, 20-21), but they are not a part of his model nor does he suggest future research should take them into account. These motivations clearly have a large bearing on the sustainment of the community; without motivation volunteers will not contribute benefits via communication and will leave the community. Given the importance of size and communication activities found by

Sagers, a study exploring motivations in this context would be highly insightful and useful. The life cycle of OSS communities, indirectly addressed by Sagers through his longitudinal design, could also be explored further. He focused primarily on projects in the middle stages of the cycle, rather than those just beginning or ending, stages that are perhaps most critical for the success or failure of a project and community. Further research applying Sagers's model to OSS communities and projects that are just beginning or ending would thus provide greater insight into the life cycle of OSS communities.

A few minor methodological issues exist in Sagers's (2007) dissertation: a lack of consistency in the time periods used for each measure (claimed as six weeks but at least one was over seven weeks), the collection of message header data which was never used further (in this study at least), the number of downloads being a poor proxy for the total number of users, and the presence of some indicators that loaded above the standard 0.50 cutoff value on factors other than the intended factor. The dissertation is also a little disorganized in parts; the terms "social loafing" and "free riding" are used interchangeably without evidence they are indeed the same concepts; and Sagers's understanding of p-values, alpha levels, and Type I statistical errors appears to be confused (see e.g. the end of p. 58). While these are all relatively minor issues, in sum they do add up and likely should have been addressed prior to the dissertation's defense.

Follow-Up Study

As noted, it certainly could be argued that Sagers's (2007) dissertation should have been done differently. The use of attitude-based measures that had previously been shown to fall out in Hartwick and Barki's (1994) study and in the TAM literature is probably the most prominent failing, but some important aspects of motivation and the life cycle of OSS project communities were also not addressed. In addition, the minor methodological issues could have been corrected

prior to the study's completion. However, taking into account all of these flaws would have led Sagers to a much more complicated and involved dissertation, one that would have taken substantially greater effort and time on his part. Given that it is better to have a completed dissertation that one can base follow-up research on than to have an incomplete but extremely promising dissertation, a follow-up study that expands upon Sagers's dissertation study, making changes, additions, and corrections to address these issues, is a more fruitful approach to take. Unfortunately, a literature search does not show any such follow-ups by Sagers himself; he appears to instead be focusing on mobile and wireless technology issues. Thus, this section presents a follow-up study that, it is believed, will add to Sagers's work constructively and usefully and bring substantial insights to the study of OSS communities and their sustainment.

Theoretical Framework and Research Model

The follow-up study proposed here expands upon Sagers's (2007) original theoretical framework. As in his dissertation, Butler's Resource-Based Model of Online Social Structures and Sagers's additions to it is used as a base, with particular emphasis on the concept of the community aggregating resources and converting them to benefits through communication activities, thus leading to the success or failure of both the software product and the community and to the community growing or shrinking. Sagers's addition of success as both an important step in the model and an important element to measure is retained, given the importance of both system success and organizational success in the MIS literature (see the Analysis section earlier in this paper). His conception of OSS as a public good is also retained, since it encapsulates the purpose of OSS projects and how they are seen by the broader community of OSS advocates.

However, a change to the framework that incorporates Seddon's (1997) model of system success is proposed. Seddon's model drew more directly from the wider MIS literature on

system success and technology acceptance, but also included the key factors of the net benefits of the system and its use. In particular, he included the benefits to individuals, organizations, and society, thus including multiple levels of analysis that address many of the elements present in conceiving of OSS projects as a public good. It also is clear that OSS projects and communities feature interactions and communication on all three of these levels; their explicit inclusion in a framework is essential to capturing these processes. Seddon's inclusion of measures of system success and technology acceptance drew especially from literature on TAM (see e.g. Adams et al., 1992; Davis, 1989; Davis et al., 1989), including the key measure of perceived usefulness. Perceived satisfaction, the other element in his model, has been replaced in more recent research by a wide variety of measures, many of which are included in UTAUT (Venkatesh et al., 2003). These measures are thus used to substitute for user satisfaction in this proposed study, as will be discussed later.

Seddon's (1997) model also explicitly included a feedback loop, as did Butler's model; Sagers (2007) made this implicit in the model he developed and tested in his dissertation. This proposed follow-up study will treat that feedback loop more explicitly, focusing on how each cycle through Sagers's model impacts the next cycle. In this way, some of the concerns over the life cycle of OSS communities will be addressed. Finally, the quality measures in Seddon's (1997) model (and that of DeLone and McLean, 1992, before him) will still not be explored in this proposed follow-up study; it is believed including measures of them would complicate the study too greatly. Nevertheless, they are important factors that should be considered in future research building upon Sagers's dissertation and the follow-up study proposed here.

The overall research model used in this study is shown in Figure 1 (next page). As noted, it includes factors from UTAUT for subjective measures of product performance, as well as the

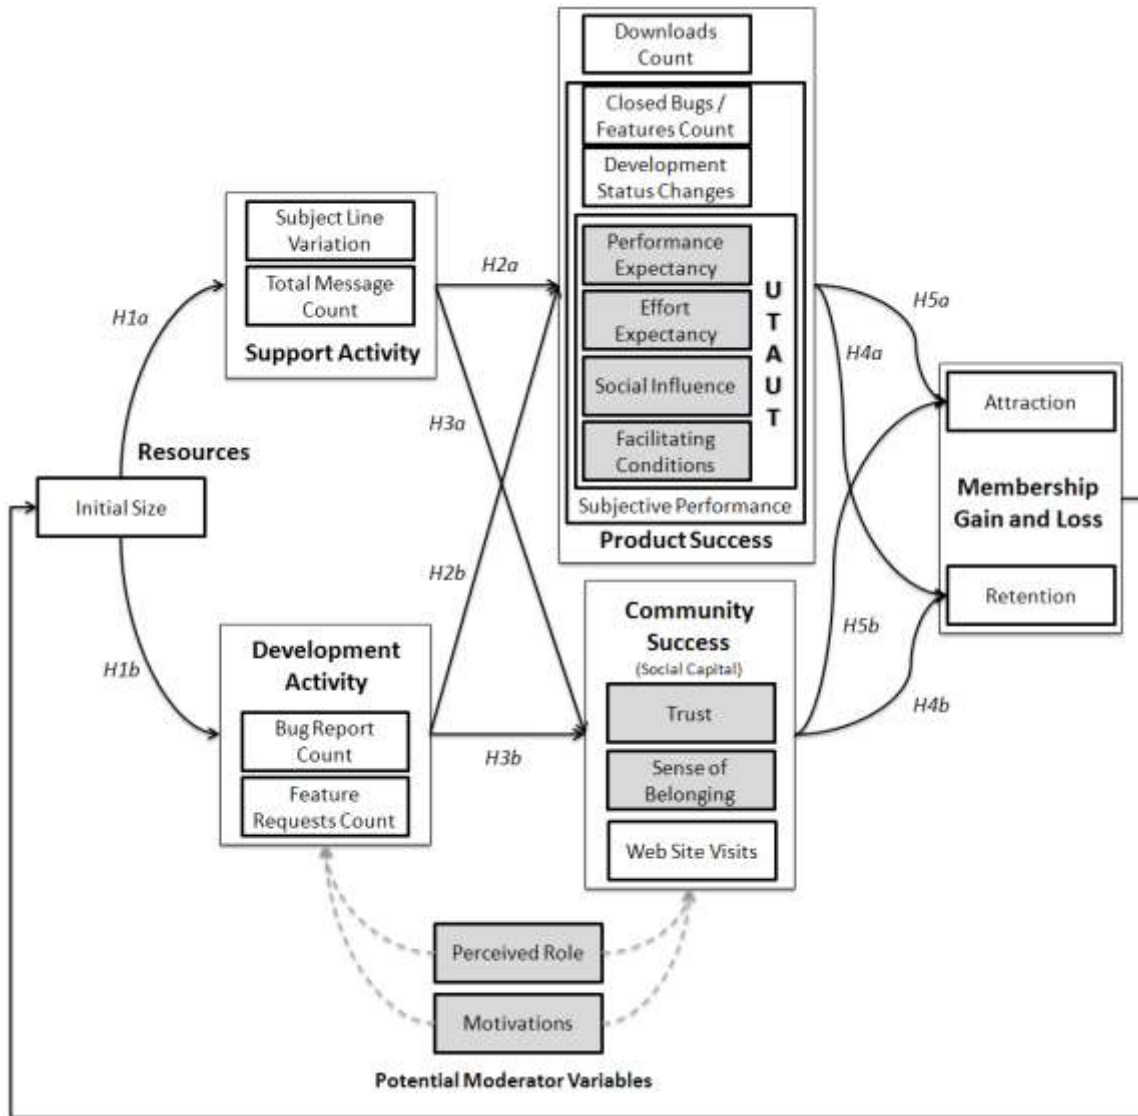


Figure 1: Research Model

explicit feedback loop present in Seddon’s (1997) model; potential moderator variables are also included. The remainder of the model draws directly from Sagers’s (2007, pp. 42, 54) model. Hypothesized relationships are shown by black arrows, and correspond—as indicated by the labels—to the hypotheses that Sagers tested but are expanded to consider all new measures introduced; arrows to each individual measure are not shown for clarity. Potential moderator effects are shown with dashed gray arrows (at the bottom of the figure). Measures drawn from

Sourceforge data are in white boxes; gray boxes are those measures which are drawn from a survey. The specifics of these measures are discussed in the Data Collection section below.

Population and Sampling

Arguably, Sagers's (2007) restriction of his population to only the 100 most active Sourceforge projects led to a relative lack of generalizability for his study. A sampling frame listing all projects hosted by Sourceforge (or project identification strings for these) can most likely be obtained by requesting such a list from the site's administrators. If this turns out to not be successful, the Sourceforge API (<http://sourceforge.net/apps/trac/sourceforge/wiki/API>) should be able to be used to generate such a sampling frame. In either case, efforts will be made to ensure the list includes the date the project was created on Sourceforge and the date of the last activity on the project, to aid sampling efforts.

Using such a frame of all projects will allow for much larger samples to be taken. However, there is clearly a limit to the amount of data that can reasonably be collected and analyzed without overburdening the researcher and assistant(s). To that end, a sample of 100 Sourceforge projects will be chosen. 50 of these projects will be purposively sampled, 25 from projects that have recently been created—and thus are deemed to be in the beginning stages of the life cycle—and 25 from projects that appear to be at the ending or dying stage of the life cycle, having had little activity over the three months prior to the selection of the sample. It is hoped selection of these will be able to be aided by the date data intended to be included in the sampling frame; projects will be randomly selected from within the sampling frame, but purposively chosen based on matching the desired criteria. In addition to these two sets of 25 projects at the beginning and ending stages of the life cycle, 50 other projects will be randomly sampled based on the sampling frame. To ensure fair representation of projects in the middle

stages of the life cycle in this random sample, projects that are less than a month old or have not been active in the last three months will be discarded from the sample if chosen and replacements drawn. It is believed these sampling procedures will balance the need for projects at all stages of the life cycle with the desire for a random sample that is representative of the population of OSS projects hosted on Sourceforge.

Data Collection

Multiple methods of data collection will be employed in this proposed follow-up study. The data collection from Sourceforge as performed by Sagers (2007) will be retained, albeit with a few changes. The initial level of resources will be measured by collecting all messages posted to all forums and listservs for the project that are hosted by Sourceforge, rather than focusing on only the “Open Discussion” forum. By doing this, a wider variety of project-based resources and communication activities will be considered, leading to a more accurate consideration of the resource level of the community. Support communication activities will also be measured across forums and listservs, to ensure accurate measure of the benefits generated through these activities. The collection of these communication activities will include the headers of the posts, as with Sagers, to allow for retention of an analysis of topic / subject line variation as an additional measure. User names and timestamps will also be collected to allow for further follow-up analysis in future research, if so desired.

Development communication activities will also retain the same measures as used by Sagers (2007), although the counts of bug reports and feature requests will be kept separate to allow for comparison of these two categories of activity. Identification numbers will be collected to aid the determination of which reports and requests were closed later in data collection; user

names, timestamps, and subject lines will also be collected to allow for further follow-up analysis in future research, if so desired.

While support and development communication activities may also occur outside Sourceforge, capture of these would be trickier; it is also believed such efforts would require too much time on the part of the researcher and assistant(s). Thus, this proposed study will continue to restrict itself to communication activity and resources present on Sourceforge; this limitation is admitted, but it is believed it is necessary at this time.

The software success measures Sagers (2007) drew from the Sourceforge data will also be retained and collected in this proposed follow-up study. The number of downloads, the percentage of bug reports and feature requests closed as of a particular point in the data collection cycle (see the discussion below regarding the study timeline), and changes in development status will be collected as per Sagers's dissertation. Attraction and retention measures will also be collected similar to his study, although as with the initial resource level and support communication activity measures all project forums and listservs hosted by Sourceforge will be examined to ensure comparability with those previous measures.

The survey methodology employed by Sagers (2007) will be used in this proposed study as well, but substantial changes will be made to its content. The measures of community success that Sagers drew from the Likert scales of Jarvenpaa, Jones, and their colleagues will be retained, but the factor loadings obtained for these items will be carefully examined to ensure they load on the correct factors and do not cross-load at values higher than the 0.50 cutoff, as observed in Sagers's study with one item from each scale. Different subjective measures of software success will be employed, however; given that Hartwick and Barki (1994) found their own attitude scale to not be a significant factor in their model, its continued use would not make sense. Instead,

measures will be drawn from UTAUT (Venkatesh et al., 2003) to measure the success of the open source software developed by the project. Survey questions to assess performance expectancy, effort expectancy, social influence, and facilitating conditions—all factors that affect behavioral intention under the UTAUT model—will be taken from the instrument used by Venkatesh et al. (2003, p. 460) to test the model, with the word “software” substituted for “system.” These measures have been proven to be highly valid and reliable in previous research, and so should serve better here than other alternatives.

In addition to measures from UTAUT, further survey questions will be included to aid analysis. Two questions asking how recently and how often the project member has visited the project’s Web site will be included, as a measure of community success paralleling the downloads measure used for measuring software success. While this measures *reported* use of the project’s Web site rather than *actual* use, and thus is admittedly somewhat limited, collecting data on actual use of the Web site by project members only—rather than all visitors to the site—would be ethically and methodologically difficult. Another survey question will be asked to obtain the perceived role the respondents see themselves as playing in the project: (a) developer, (b) support, (c) active user, (d) prospective user, or (e) other. This will aid exploratory analysis of these different subgroups of community members to see if the relationships between the constructs in the theoretical framework differ across—i.e. are moderated by—role types. Three exploratory questions on the motivations of community members will also be included, intended to begin the study of this factor and its relation as a possible moderator on software success and OSS community sustainment. The first question will ask members to check off possible motivations from a list drawn from the motivational studies reviewed by Sagers (2007, p. 18); the second question will be open-ended and will ask members to briefly discuss why they were

motivated to join the community in the first place, while the third question (also open-ended) will ask why they continue to be a member and contribute. Finally, to aid potential future analysis of the results, demographic data on the gender, age, and overall technology experience of project members will also be collected.

An additional method will be employed in addition to the survey and Sourceforge data collection employed by Sagers (2007). This proposed study will use qualitative, semi-structured interviews to follow-up with interesting cases from the surveys. Sampling for these interviews will purposively select at least eight individuals per data collection cycle (see the timeline discussed below) whose survey responses indicate an interview with them should contribute to further understanding of the applicability of the research model, particularly communication activities, software and community success, community sustainment, and their motivations for participating and contributing to the community. Interview questions and prompts will be based on the literature on these topics—much of which was cited and discussed by Sagers—and on the findings of the data collection prior to the interviews. Face-to-face interviews would be difficult to arrange given that many members of Sourceforge-hosted OSS project communities will be a significant distance from the researchers; however, visual and audio cues are often key in fully understanding interviewee responses. Individuals will thus be interviewed using audiovisual conferencing software such as Skype where at all possible, with telephone and e-mail as second and third backup choices respectively. As noted above, at least eight individuals will be interviewed per cycle, but interviewing will continue until saturation is reached and no new significant findings are obtained from each successive interview. If saturation is not reached by the end of interviews for the last cycle of data collection, then additional purposive sampling of previous survey respondents may be required.

Timeline

The proposed study will employ a longer period of study than Sagers's (2007) study, nearly doubling the total time for data collection to 480 days. While it is acknowledged that this study will thus take quite a long time to complete, the need for more extensive longitudinal data and more frequent data collection—also mentioned by Sagers—is sufficient enough to require such a time frame. The 480 days will be divided into five planned cycles of data collection, one beginning every 90 days; the last 30 days will be used to finish measuring the size of the community at the end of the last cycle. Scheduling will be consistent for each cycle, following the timeline shown in Table 1 (see next page).

This timeline remains relatively close to the overall schedule used by Sagers (2007), with a few minor changes. First, the extra two weeks or so between many data collection phases are removed to speed data collection and allow for additional cycles. Second, the determination of the final size of the community per cycle is combined with the determination of the beginning size for the next cycle, essentially providing for six measures of size throughout the project. Finally, the time span for each phase of data collection—excepting those collected on Day 90 and the interview data collected after each cycle—is kept equal at 30 days, rather than the varying six to seven weeks that Sagers ended up using in his dissertation study. Overall, it is believed that a more consistent schedule and the inclusion of multiple cycles of data collection improves greatly upon Sagers's design, allowing for additional longitudinal data to be collected that can aid analysis.

Data Analysis

As with Sagers's (2007) study, analysis of the quantitative data collected will use partial least squares (PLS) to test the relationships expected. Exploratory analysis will also be conducted

Table 1: Study Timeline

Date	Data collection activities
Days 1-30:	<ul style="list-style-type: none"> • Size measurement (level of resources, also attraction and retention)
Days 31-60:	<ul style="list-style-type: none"> • Posts to forums and listservs (support communication activities) • Bug reports (development communication activity) • Feature requests (development communication activity)
Days 61-90:	<ul style="list-style-type: none"> • Number of downloads (software success) • Subjective measures via survey: <ul style="list-style-type: none"> ○ UTAUT measures (software success) ○ Trust (community success) ○ Sense of belonging (community success) ○ Web site visits (community success) ○ Perceived role ○ Demographics ○ Motivations
Day 90:	<ul style="list-style-type: none"> • Percentage of bug reports closed (software success) • Percentage of feature requests closed (software success) • Change in development status (software success; subjective)
Days 91-180:	<ul style="list-style-type: none"> • Interview data (after each cycle, timing to vary)

to identify any other significant paths that were not theorized *a priori*. Like Sagers, this study will apply intraclass correlations (ICCs) to ensure the suitability of aggregating data across projects. Reliability and validity will also be assessed through examination of the average variance extracted (AVE), composite reliability / ICR scores, and factor loadings of the indicators. Any indicators that load above the 0.50 cutoff on a different factor—unlike Sagers’s approach—or load highest on a different factor to that expected will be removed prior to further analysis of the results. Coupled with the use of measures of software success adapted from UTAUT (Venkatesh et al., 2003), this should ensure the reliability and validity of the proposed study are high and improve upon these measures compared to Sagers’s dissertation study.

The qualitative data gathered from interviews, along with the findings of open-ended survey questions, will be analyzed through coding and analysis of their content. First, the responses to open-ended survey questions on motivations will be coded, employing a codebook partially derived from the literature and partially emergent from the results as they are analyzed. Next, interviews will be transcribed and coded using a set of codes partially derived from the literature, partially taken from the open-ended survey questions, and partially emergent from the interviews themselves. Initial stages of coding for the first cycle of data will include multiple group coding sessions, where the researcher and assistant(s) working on the project will compare codes and notes to help ensure intercoder consistency and reliability. Codes for interviews and survey responses from later cycles are expected to become less emergent and more confirmatory of the existing codebook structure. After initial coding of all survey questions and interviews is complete, codes and data will be analyzed and interpreted to determine how they relate to the research model and conclusions that can be drawn from the qualitative data. The reliability of coding will be established by having at least two individuals code each survey response and

interview transcript, with intercoder reliability to be tested after each cycle and corrective measures taken if necessary. Finally, after both quantitative and qualitative data analysis are complete the two will be considered in light of each other to ascertain what further conclusions can be drawn from the results of the study as a whole.

Rejected Methods

Finally, a brief discussion of two methods that were rejected for this follow-up study is helpful in understanding the choices made in this design. First, content analysis of the subject lines and bodies of the forum and listserv messages, bug reports, and feature requests could prove to be useful, insightful, and interesting in light of Sagers's (2007) model and findings. In particular, his discussion of how topic variation and diversity may have led to a lack of "common interests within the community" suggests further analysis of the topics discussed in these messages could help confirm or deny this proposition. Content analysis would also help to further explore communication activities and motivations, similar to the interviews included as a method in this proposed study. However, the inclusion of content analysis alongside the other methods proposed here would cause the study to take even longer and thus be too complicated to be completed in a timely fashion. Certainly, however, content analysis of messages, bug reports, and feature requests should be considered for inclusion in further follow-up studies after this one.

Second, more qualitative, ethnographic research could be included examining OSS projects and communities in light of the organizational change caused by their use of Sourceforge and other changes in the technology use of the project and community. Such methodology would draw particularly from Orlikowski's (1993, 1996; Orlikowski & Robey, 1991) work as well as from the other organizational studies cited by Sagers (2007, pp. 19, 23-25). Although these previous studies have explored organizational change in the context of

software development, further research in this area that particularly focuses on OSS projects and communities in the context of Sagers's model would be useful. However, this would also likely be better as a further follow-up study, given that it would have a slightly different purpose than that intended here. This study idea also needs further development and careful consideration to be fully successful in the context of Sagers's model and findings.

Conclusion

Sagers's dissertation research (Sagers, 2007) did contribute towards answering the research questions he posed on how OSS communities are sustained and whether they contribute to the success of the software they develop. His theoretical model, while not fully supported by the testing he performed in his study, nevertheless extended Butler's earlier model and supported his findings in a different context. Nevertheless, there were flaws in Sagers's study that should not have been present, most notably the use of an attitude scale that had been proven to fall out by its own authors, Hartwick and Barki (1994). There were also other areas that Sagers, due to his focused dissertation, did not explore or consider in his study but that are important factors in the sustainment of OSS communities and their contribution to the success of the software they develop. To that end, this paper has proposed a follow-up study to Sagers's dissertation that, while perhaps a little "daunting" (Sagers, 2007, p. 81), provides a fuller, more generalizable test of his model while incorporating more valid and reliable measures of the subjectively-measured success of the software product. This study should add constructively to Sagers's research and other studies of OSS communities, working towards a fuller, richer understanding of how they develop successful software, the key role communication activities play in this, and how and why such development is sustained over time by volunteers.

References

- Adams, D. A., Nelson, R. R., & Todd, P. A. (1992). Perceived usefulness, ease of use, and usage of information technology: A replication. *MIS Quarterly*, *16*(2), 227-247.
- Cooper, R. B., & Zmud, R. W. (1990). Information technology implementation research: A technological diffusion approach. *Management Science*, *36*(2), 123-139.
- Curtis, B., Krasner, H., & Iscoe, N. (1988). A field study of the software design process for large systems. *Communications of the ACM*, *31*(11), 1268-1287. doi:10.1145/50087.50089
- Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, *13*(3), 319-340.
- Davis, F. D., Bagozzi, R. P., & Warshaw, P. R. (1989). User acceptance of computer technology: A comparison of two theoretical models. *Management Science*, *35*(8), 982-1003.
- DeLone, W. H., & McLean, E. R. (1992). Information systems success: The quest for the dependent variable. *Information Systems Research*, *3*(1), 60-95. doi:10.1287/isre.3.1.60
- Hartwick, J., & Barki, H. (1994). Explaining the role of user participation in information system use. *Management Science*, *40*(4), 440-465.
- Kemerer, C. F. (1987). An empirical validation of software cost estimation models. *Communications of the ACM*, *30*(5), 416-429. doi:10.1145/22899.22906
- Orlikowski, W. J. (1993). CASE tools as organizational change: Investigating incremental and radical changes in systems development. *MIS Quarterly*, *17*(3), 309-340.
- Orlikowski, W. J. (1996). Improvising organizational transformation over time: A situated change perspective. *Information Systems Research*, *7*(1), 63-92. doi:10.1287/isre.7.1.63
- Orlikowski, W. J., & Robey, D. (1991). Information technology and the structuring of organizations. *Information Systems Research*, *2*(2), 143-169. doi:10.1287/isre.2.2.143

- Sagers, G. W. (2007). *Is bigger always better? Toward a resource-based model of open source software development communities* (Doctoral dissertation). Retrieved from Florida State University ETD database (etd-03302007-173851). <http://etd.lib.fsu.edu/theses/available/etd-03302007-173851/>
- Seddon, P. B. (1997). A respecification and extension of the DeLone and McLean model of IS success. *Information Systems Research*, 8(3), 240-253. doi:10.1287/isre.8.3.240
- Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly*, 27(3), 425-478.