AARHUS UNIVERSITY

Coversheet

This is the accepted manuscript (post-print version) of the article.

Contentwise, the post-print version is identical to the final published version, but there may be differences in typography and layout.

How to cite this publication

Please cite the final published version:

Haakonsson, D. D., Bach, L. A., Snow, C. C., & Obel, B. (2017). Building A Collaborative Community: An Agent-Based Simulation Study. *2017*(1), 11293. doi:10.5465/ambpp.2017.11293abstract

Publication metadata

Title:	Building A Collaborative Community: An Agent-Based Simulation Study
Author(s):	Haakonsson, D. D., Bach, L. A., Snow, C. C., & Obel, B.
Journal:	<i>Academy of Management Proceedings</i>
DOI/Link:	https://doi.org/10.5465/ambpp.2017.11293abstract
Document version:	Accepted manuscript (post-print)

General Rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

BUILDING A COLLABORATIVE COMMUNITY: AN AGENT-BASED SIMULATION STUDY

ABSTRACT

A collaborative community is an emerging organizational form that is increasingly being used in knowledge-intensive environments where success depends on the ability to collaborate. We analyze how community member composition influences the development of a collaborative community with two kinds of members: problem providers and problem solvers. In an agent-based simulation study, we examine how the dynamics and survival of a collaborative community are affected by: (a) the initial number of agents of each kind and (b) the types and levels of skills agents possess. The simulation results show an early drop in the number of community members followed by continued growth in community size. Independent of the number of members of each kind at the formation of the community, the number of members of both kinds is approximately the same during the growth phase. The community's growth rate depends on agent skills (both type and quality). The results suggest that homogeneous communities experience a small drop in the number of members during the early phase. Heterogeneous communities show a larger drop early on, but tend to converge to a higher average agent skill level than a homogeneous community. Sensitivity analyses demonstrate that the simulation results are consistently robust. Our use of the simulation method allows us to explain the behavior of a collaborative community over time, and we discuss the implications of our findings for building collaborative communities.

Keywords

Collaborative communities, collaborative problem solving, new organizational forms, organization design, self-organization, self-management, agent-based simulation

Collaborative communities increasingly are found in knowledge-intensive environments such as computer software development, drug discovery, healthcare, public-private partnerships, global professional services, and national defense. A collaborative community is a way of organizing in which actors – individuals, teams, firms – share knowledge and use commonly held resources to develop new products, markets, and technologies on a continuous basis (Adler, Kwon, & Heckscher, 2008; Bøllingtoft et al., 2012; Heckscher & Adler, 2006; Miles, Miles, & Snow, 2005). In contrast to traditional organizations, which rely on hierarchical mechanisms for control and coordination, actors in a collaborative community are able to self-organize, achieving control and coordination via direct interaction and infrastructure mechanisms (Fjeldstad et al., 2012). Successful collaborations, both within and across organizations, have been shown to reduce risk, speed products to market, decrease the cost of product development and process improvement, and facilitate access to new markets and technologies (Clark & Wheelwright, 1992; Eisenhardt & Schoonhoven, 1996; Hagedorn, 1993; Kogut, 1988).

Because community-based organizing is still relatively new, organization designers and managers do not yet have a clear recipe for building a collaborative community: how to compose the membership, how members should interact with one another, what the activity system of the community should include, and how to grow the community to a desired size (Foster-Fishman et al., 2001). In such a situation, it is appropriate to turn to computational modeling as a means of exploring future alternative states rather than observing existing cases (Burton & Obel, 2011).

In this paper, we present the results of an agent-based simulation study designed to examine the dynamics and survival of a collaborative community. We analyze how member types, the number of different member skills, and the quality level of each skill type influence the development of a collaborative community over time. Our simulation model represents a

collaborative community with two kinds of members: problem providers and problem solvers. The model simulates a collaborative community with a varying number of problems to be solved and a problem-solving capacity represented by members' skill types and levels. Our model is inspired by two real-world cases: Blade.org, a collaborative community of firms in the computer server market (Snow et al., 2011) and the Organizational Design Community (www.orgdesigncomm.com), an international community of academics and practitioners established in 2010 and dedicated to advancing the theory and practice of organization design. The Organizational Design Community was organized according to the principles and experiences of Blade.org.

Our paper proceeds as follows. First, we discuss the theoretical background of collaborative communities, focusing on their member composition and dynamics. We then specify our simulation model and present the results of different scenarios and sensitivity analyses. We conclude the paper by discussing the implications of our findings for designing and building collaborative communities.

THEORETICAL BACKGROUND

The purpose of a collaborative community is defined by the shared interests of its members. Governance of the community relies heavily on trust, norms of reciprocity, shared values, and a commitment to the common good rather than on a hierarchical authority structure (Hess & Ostrom, 2006; Ostrom, 1990). The overall effectiveness of a collaborative community hinges on the ability of its members to self-organize and self-manage – such an organization is said to be “actor oriented” rather than hierarchically oriented (Fjeldstad et al., 2012). Compared to traditional (hierarchical) organizations, we know relatively little about how to design and build

a collaborative community. In case studies of Linux and Wikipedia, Garud, Jain, and Tuertscher (2008) found those communities' organizational designs to be incomplete at the outset and continuously evolving over time. Kolbjørnsrud (2014) found a similar phenomenon in the four collaborative communities he studied. He concluded that the distinction between design as 'plan' and design as 'process' is blurred, as a complete design at the stage of community formation may not be possible or even desirable. Studies of open source software communities also have found that governance structures and processes are emergent; they are enacted over time through the actions and norms of community members (O'Mahoney, 2007). Raban, Moldovan, and Jones (2010) found that the number of members in itself does not determine community longevity and that there is a significant positive correlation between member heterogeneity and survival. Williams and Cothrel (2000) found that member heterogeneity is important to the effectiveness of an online community and, in addition, they argued that a critical mass of members was needed for community survival. Anderson (2016) concluded that member heterogeneity (measured by number of skill types) is important to understanding both the structure and functioning of collaborative communities.

Findings from these and other empirical studies suggest that the organizational designs of collaborative communities are partly planned and partly emergent; the eventual design of a given community results from the characteristics and interactions of community members. For this reason, we decided to examine the effects of number of member skill types and their quality level on community dynamics and survival by means of an agent-based simulation.

METHOD

Agent-based simulation is a method appropriate for our research purposes because it puts

emphasis on emergent processes driven by interactions among actors. Moreover, a simulation enables us to precisely measure and vary our variables of interest, such as member skill types and levels, while simultaneously holding all other variables constant. Further, our intent is to examine the long-term impact of different model parameters on the sustainability of a collaborative community. Essentially, simulation enables us to study what *might* be (e.g., alternatives and possibilities) and what *should* be (e.g., the best choice among alternatives) (Burton & Obel, 2011). Accordingly, we can use the results of our study for both theory development (Davis, Eisenhardt, & Bingham, 2007) and for making recommendations about how to build sustainable collaborative communities.

Simulation Design

As a starting point, we use the Organizational Design Community (ODC) (www.orgdesigncomm.com) to inspire our simulation model. Established in 2010, ODC is a collaborative community of academics and practitioners dedicated to advancing the theory and practice of organization design. ODC seeks to be the preeminent community where research, practice, and learning intersect to produce valuable design knowledge and applications. We chose ODC to guide the design of our simulation for two reasons. First, it is a community composed of different kinds of members with different skill types who collaborate with one another to provide and solve organizational design problems. Second, ODC itself was designed using the organizing principles and processes of Blade.org, described in the literature as a successful collaborative community of firms in a commercial context (Snow et al., 2011). Thus, our simulation model is reality-based rather than hypothetical.

ODC's purpose is to solve organizational design problems as they are put forward by its members. Some problems are advertised by the Principal Office, and members can choose to

react to the advertisement. This could be a call for papers on a thematic series in ODC's open access journal, the *Journal of Organization Design*. It could be activities related to writing a paper or inviting speakers for ODC's annual conference. Other problems are presented by the members themselves. One such example is the establishment of an accreditation program to certify organizational design courses and workshops offered by consulting firms and other course providers. Here ODC practitioner members proposed the problem, and academic members developed the minimum required course content. Administration of the accreditation program is conducted jointly by ODC academics and practitioners.

Model Specification

Similar to ODC, which has two kinds of members (academics and practitioners), our model has two kinds of agents: problem-providing members (referred to as A_p) and problem-solving members (referred to as A_s). Hence, the total population consists of these two types of members, shown formally as:

$$A = (A_p, A_s)$$

The purpose of the simulated community is to solve a series of recurrent problems. A given problem has a required number of skill types (e.g., marketing, production, and leadership skills are needed) each with a required skill level (from low to high). Each problem-solving agent may or may not have the skill types needed to solve the problem, and each problem solver may or may not have the skill level for a particular skill type required to solve the problem. If the problem-solver's skill types and quality levels both meet the problem's skill requirements, there is a match (i.e., the problem is solved). In the simulation, problems are initiated by attributing skill requirements to the A_p agents. The problem-solving agents (A_s) are similarly attributed skill

types and quality levels.. Every agent is initially specified to have a quality level for each skill type that ranges from S-min (lowest possible skill quality level) to S-max (highest possible skill quality level). The particular skill level is drawn randomly from within a range of 0-100. We attribute different S-min and S-max to investigate this parameter. In the simulation model, the number of skill types an agent can have can be altered between simulation runs.

Problems are proposed, and solution agents react to the proposal. In this way, a problem provider is coupled with a problem solver. Two procedures to couple a solver agent with a proposer agent are simulated: (1) a random assignment of problems to solvers and (2) an assignment procedure where solvers are ranked in terms of skill level and then picked from the top end of the ranking. A problem is solved when the problem solver's number of skill types and levels match or exceed the required levels of the problem for each skill type. Formally, agent members with N skill types and a quality level of the skill type n equal m_i , respectively, are evaluated and matched in the following way:

If $A_p(n_i) \leq A_s(n_i)$ for all $n=1, \dots, N$, then the problem is solved.

When a problem is solved, the A_p will generate a new problem, which is similar in difficulty level for each skill type to the problem the agent proposed in the previous round. How similar a new problem is depends on the mutation level set in the simulation.

A problem agent may also present the solution agent with a problem that is too difficult given the solution agent's skill types and levels. Then the problem is not solved. Formally:

If $A_p(n_i) > A_s(n_i)$ for just one n , $n=1, \dots, N$, then the problem is not solved.

In the case where a problem is not solved, both the A_p and the A_s will have experienced one idle round. In ODC, this has been shown to affect the likelihood of a member leaving the community. In our model, one ‘round’ (the basic time unit in the simulated community) represents one year. This is equivalent to the annual membership renewal in ODC. During each round, problem solvers and proposers interact bilaterally such that solvers (A_s) can offer their skills to proposers (A_p).

Exit Behavior of Problem Providers and Problem Solvers

In ODC, we observe that some members decide not to renew their annual membership. We suspect that members who decide to leave the community do so because they are unsuccessful in solving problems or in having their problems solved. Members who are successful in solving problems, or in having their problems solved, tend to stay in the community. Typically, members who leave the community do not do so after a single unsuccessful year. Usually, inactive or unsuccessful members will stay between 2-3 years before leaving. Therefore, we believe that each member has a ‘patience’ level that must be exceeded before he or she will leave the community. In the simulation, we model agent patience by the number of idle rounds (i.e., rounds in which a member has not solved a problem or had a problem solved) before leaving the community. Formally, we model member exit such that if, at any given time ‘i’, any agent A_j experiences a number of unsuccessful rounds t_{ij} higher than the agent’s threshold T_{rj} , then the agent leaves the community. Hence, the per round total outflow = $\sum A_j (t_{ij}, T_{rj})$, where $t_{ij} > T_{rj}$.

Recruitment of New Agents

In ODC, new members are most often recruited by existing members who recommend the community to colleagues in their professional networks. This means that new members tend to

be similar in profile and skill level to existing members. We designed the simulation model to allow us to specify, and vary, the degree of similarity of new agents to existing agents. For lack of a better term, we call the parameter controlling the similarity of newly recruited agents ‘mutation’. In other words, ‘likes’ recruit ‘likes’ -- but less so when mutation is high. Formally:

The per round in-flow of problem-providing agents = $P_p * C_{np}$

The per round in-flow of problem-solving agents = $P_s * C_{ns}$

meaning that with probability P_s a number, C_n , of clones of an existing problem-providing agent will be added to the community each round (and likewise for the problem-solving agents).

New agents will have a degree of resemblance of their characteristics (skills), S_k , to recruiting agents, depending on the mutation parameter μ_s or μ_p . Formally:

$\text{new_}A_s(\text{characteristic } k) = \text{old_}A_s(\text{characteristic } k \pm \mu_s)$

and

$\text{new_}A_p(\text{characteristic } k) = \text{old_}A_p(\text{characteristic } k \pm \mu_p).$

Provider-agent behavior is hypothesized to be similar to solver-agent behavior. That is, when a problem is solved, the provider agent stays as a member of the community and presents a new problem in the subsequent round. The new problem can be more or less similar to the old problem depending on the problem-agent mutation parameter. In condensed form, the replacement could be expressed with the mutation parameter μ_{pR} . Formally:

$\text{new_}A_p(\text{characteristic } k) = \text{old_}A_p(\text{characteristic } k \pm \mu_{pR})$

This specification allows new problems to vary (increasing or decreasing in difficulty level) as the rounds progress.

Model Parameters

The specific model parameters are listed in Table 1. The number of problem solvers shows how many problem solvers will be present at the outset of the simulation. Likewise, the number of problem providers specifies how many providers will be present in the first round of the simulation. These initial numbers can be set independently. Number of rounds indicates the number of years in a given simulation run. Complexity refers to the number of different skill types needed to solve a problem (e.g., marketing, accounting, human resources). Each problem has a number of required skills in order for it to be solved (e.g., marketing, accounting, human resources). Thus, the number of skill types is, by design, the same for both kinds of agents during a given simulation run. Bound on initial skills (min – max) sets the bounds for initializing (randomly) the skill quality levels for new solver agents (e.g., advanced accounting skills). Bound on problem difficulty level (min – max) sets the bounds for randomly initializing problem difficulty levels for each skill type required from the problem solver. Mutation level sets how much a newly recruited agent deviates from the agent that invites the new agent (measured in the units that skills and problems are defined). Mutation sign shows the sign of the deviation (up (+), down (-), or both up and down (+/-)). Problem mutation sets the deviation in skill requirements for a new problem from the solved problem that it replaces. The sign of problem mutation determines, as above, the sign of the deviation (up (+), down (-), or both up and down (+/-)). Removal threshold for problem-solving agents specifies how many idle rounds will cause a problem-solving agent to leave the community. Removal threshold for problem-providing agents specifies how many idle rounds will cause a problem provider to leave the community. Chance

of new problem-solving agent recruit is the probability a problem-solving agent recruitment event takes place in a given round. Chance of new problem-providing agent recruits shows the probability a problem-providing agent recruitment event takes place in a given round. Number of new recruits shows how many new agents will be recruited in the event of recruitment.

----- Insert Table 1 about here -----

Base Scenario

The specific numbers in the model's base scenario are shown in Table 1. The scenario starts with 100 problem-solving agents and 85 problem-providing agents, each with four types of skills each of which has randomly initialized skill levels between 0 and 100 (full range). This corresponds to a generalist community with a large spread in problem requirements and skill types/levels,, just as it corresponds to the distribution of membership types in ODC. The base scenario continues for 100 rounds in order to let the community settle into a quasi-equilibrium of agent in-flow and outflow. Every three rounds, three new recruits enter the community and hence constitute the in-flow of the community. This means that on average one new agent enters in each round, both for problem-solving agents and problem-providing agents. New members will look like their recruiting agent, but with a small deviation (+/- five skill units). Likewise, a solved problem will be replaced by a new and similar problem associated with a provider (+/- two skill units). Outflow is generated by agents leaving the community after a threshold number of five idle rounds.

RESULTS

Base Scenario

Figure 1 shows the results of the base scenario on the number of agents in the community. We observe that the dynamics of the two agent types are similar. At first, the number of agents goes up, as recruiting starts before any agent exceeds its patience threshold. Then, after a relatively stable period (five rounds in the base scenario) that depends on the patience level of the agents, the number of agents starts to drop. The initial drop in members continues until it reaches a level well below the number of members of the rarest kind (in this case, the problem-providing agents). After the initial drop, the number of both kinds of agents increases. We tested whether these results were sensitive to variation in the number of agents of the two types and found that the result does not depend on whether there were more problem solvers than problem providers. Further, we found that during the drop the number of agents converged to be the same for the two kinds of agents. Results of these runs are contained in Appendix A, and they show a consistent ‘hockey stick’ pattern, regardless of the number of agents or the ratio of the two kinds of agents.

----- Insert Figure 1 about here -----

Effects of Agent Skill Levels

Because we suspected that the hockey stick dynamic was caused by achieving a balance between the skill levels of the problem-solving agents and the problem difficulty levels, we examined the concurrent dynamics of the average skill levels and the average problem difficulty levels in the community as well as their variances. As seen in Figure 2, both the initial drop in community size and subsequent growth leave a footprint on the patterns of average skill levels and average problem difficulty levels, respectively. The average skill level of an agent is calculated as the

average over all the skill types. The average skill level of the community is calculated as the average for each of the two kinds of agents.

----- Insert Figure 2 about here -----

Figure 2 shows how the average skill level of problem-solving agents starts to increase at about the same time as the dip in Figure 1 occurs. Further, the spread around the mean (shown in lighter red) decreases. These two observations suggest that the hockey stick dynamic is caused by a ‘purging’ process in which problem-solving agents with low skill levels leave the community due to inactivity (i.e., lack of success in finding problems that match their skill level and thus in solving problems). As also shown in Figure 2, there is a slight decrease in the average problem difficulty level corresponding to the situation where very difficult problems are not being solved. Therefore, problem-providing agents with very difficult problems tend to stand idle and, consequently, leave the community. Thus, the purging process gets rid of very difficult problems, thereby lowering the average difficulty level. However, we also see that the average skill level starts to increase slightly around rounds 10-12, suggesting that more highly skilled problem-solving agents become available. This may explain why the average problem difficulty does not continue to decrease as might be expected (i.e., the outflow of difficult problems does not continue).

The spread around the average skill level of the problem-solving agents decreases in the purging phase, as seen by the decreasing standard deviation in Figure 2. This is associated with an increase in average skill level in the community. Less skilled problem-solving agents are leaving the community at a rate related to the patience parameter. As seen in Figure 2, the community shows an increase and then convergence in the average skill level together with a

reduced spread around the mean. This means that the community becomes somewhat more homogeneous with respect to problem-solving agents and problem-providing agents. We further examined and found that this effect was robust to changes in the other parameters in the base scenario. The recruitment rate had the strongest effect in determining the speed of recovery in community size (slope of the increasing community size as seen in Figure 1). (These results are contained in Appendix A.) Moreover, the finding that problem-proposing agents with difficult problems leave the community after a number of consecutive idle rounds helps to explain how the community achieves continued development. This finding led us to investigate the effects of heterogeneity in agent skill levels.

Effects of Agent Skill Heterogeneity

Based on the findings in the base scenario, we wanted to determine whether the critical phase – the initial decline in the number of agents which could lead to community extinction – was affected by the spread (initial heterogeneity) in the skill levels of the problem-solving agents and the difficulty levels of the problems, respectively. To examine this, we ran another scenario, which in its startup phase had more homogeneous problems as well as more homogeneous problem-solving agents. Instead of drawing from the range of 0-100, as in the base scenario, the draw was from the range of 40-60. This differs from the base scenario, on which Figures 1 and 2 are based, in that the baseline represents communities with heterogeneous and hence unspecialized members. The result is shown in Figure 3. In this figure, we see a weaker drop in the number of agents in the case with low specialization than in the base scenario with heterogeneous agents at initialization.

----- Insert Figure 3 about here -----

Next, we look at the dynamics of the average skill level of problem-solving agents and average level of problem difficulty, respectively, under the same parameter set-up as in Figure 3. The results shown in Figure 4 suggest that a homogeneous community at the outset undergoes few skill-level adjustments. Further, the small standard deviation is maintained throughout. Although purging is moderate, we still see the footprint of the moderate drop in Figure 3 in that the average skill level increases slightly. A comparison of Figures 2 and 4 shows that in the base scenario (Figures 1 and 2), the average agent skill level increases substantially more than in the homogeneous scenario shown in Figure 4. This difference stems from the fact that a heterogeneous skill pool leads to a higher probability of selecting relevant skilled agents as problem solvers. When you restrict the problem-solving agents' skill levels to be more homogeneous at initialization, and similarly for problems, the community will not grow much beyond the most skillful agent. In a sensitivity analysis, we found that specialized communities produce moderate hockey stick patterns (see the Appendix for all sensitivity analyses).

----- Insert Figure 4 about here -----

Thus, the number of agents at the lowest point in the hockey stick pattern depends on the degree of specialization of the agents. However, the overall hockey stick pattern remains unchanged. As in the base scenario, the number of members starts to increase after the drop. This means that a specialized community will not face the same risk of extinction as a community with high initial heterogeneity in problem-solver skill levels and problem difficulty levels. The less pronounced critical phase stems from a weaker basis for the purging of extreme agents (very low skill levels) or problems (very high difficulty). Such agents are hardly present at the outset, and hence we see only a weak effect on the numbers of both agent types.

Community Adaptation to Difficult Specialized Problems

From the above simulations, we learned that changes in the number of problem-providing and problem-solving agents are affected by the underlying distribution of skill levels and problem difficulty levels in the community. We therefore went on to examine the effects of differences in average and variance of skill levels and problem difficulty levels further.

We ran a simulation where the skill level requirements of the problems had lower variance than the skill levels of the problem-solving agents. Additionally, the average of the problem requirement levels is higher than the average skill level of the problem-solving agents. The results are shown in Figure 5. Here we see the hockey stick pattern as in the previous runs. However, as shown in Figure 6, it is the skill level requirements of the problems that are driving community growth as problem-solving agents pursue problems and not the other way around. Problem-solving agents with a low skill level will leave the community if the community has difficult problems. As seen in Figure 6, problem-providing agents seem to be tracking a specialized group of problem-solving agents with difficult problems. However, the basic hockey stick effect is also the case here, hence the critical phase has to be passed in order for the community to survive.

----- Insert Figures 5 and 6 about here -----

Effects of Number of Skill Types on Community Size

In the previous runs, we have seen that the number of agents in the community at any point in time depends on the initial conditions with respect to the number of agents and the homogeneity of skill and problem difficulty levels. Up to this point, the number of skill types was held

constant at four. We ran a number of simulations where we varied the number of skill types from one to ten. Otherwise, the parameters were set as in the base scenario. Comparing these simulations we found that the higher the number of skill types, the lower the size of the community. This relationship is the same at any point in time. Figure 7 shows the size of the community after 100 rounds.

Because number of skill types is negatively associated with long term community size, it is more difficult to ensure a minimum number of members in complex communities where members are required to have many different skill types than in a simple communities where members need only few skill types. This means that the more complex the initial set of problems to be solved in a collaborative community the higher the probability that the community drops to a level where it cannot survive.

----- Insert Figure 7 about here -----

Sensitivity of Initial Conditions

We ran simulations of the model's initial conditions with many different parameter settings and found the hockey stick dynamic in all cases (see Appendix). We also changed the random assignment of problems to solvers to an assignment procedure where solvers were ranked in terms of skill level and then solvers were picked from the top end of the ranking. Doing so only changed the results slightly, and the hockey stick dynamic was still present. Thus, the findings are robust with respect to various parameter settings.

DISCUSSION AND IMPLICATIONS

We investigated how the initial design of a collaborative community influences its dynamics. Previous literature has indicated that the size of a community (i.e., number of members) and member heterogeneity are influential factors in the behavior and survival of a collaborative community. To examine the effects of community size and heterogeneity, we included two kinds of agents, agent skill types, and the level of each skill type in an agent-based simulation model. The model enabled us to examine not only the immediate effects of these individual factors but also their effects over time.

We found that with two kinds of agents of unequal proportion, the number of agents drops after a short stable period representing the patience level of the agents. The initial drop in agents continues until it has reached a level below the rarest agent type where the number of agents is similar for each kind of agent. After the initial drop, the number of agents starts to grow with similar growth rates for each type of agent. However, it takes several cycles of recruitment to regain the number of agents lost in the first critical rounds. We referred to this dynamic as the ‘hockey stick’ effect. We tested the hockey stick dynamic on different agent compositions and found that the effect was consistent. Therefore, the crucial factor in the survival of a community appears to be that the initial drop in number of agents does not bring the size of the community to a level where it ceases to exist. The lowest number of agents necessary for a community to survive will depend on the particular community. In ODC, for example, it was related to income from member dues used to finance the Principal Office. In the community described by Kolbjørnsrud (2014) that did not survive, it was also the income to run the Principal Office that fell below a sustainable level. Because the drop is dependent on the number of the member type

of which there are the fewest, this type of member should be set as large as possible in order to avoid extinction of the community.

Next, we examined whether the hockey stick effect was caused by the problem-solving agents' skill levels for the various skill types as well as the required skill levels of the problems. Previous literature has indicated heterogeneity is important to community survival. This literature has referred to heterogeneity as relating to both the broader notion of 'specialization' (e.g., Williams & Cothrel, 2000) as well as the more specific 'skill set' of individual agents (Anderson, 2016). We extended this literature by examining both the effects of the skill level (variation around the average skill level of members' skill types) and the number of different types of skills required to perform the community's tasks. Regarding skill level variation, we examined the balance between the skill levels of the two types of agents. We did so by examining how the average skill levels of the two kinds of agents would influence community dynamics. In that sense, heterogeneity is the difference in average skill levels between the two agent types.

Our results indicated that low heterogeneity decreased the drop in number of members. With high heterogeneity in skill levels, the initial drop in numbers was severe. However, high heterogeneity also led to an increase in the average skill levels of agents. This was caused by the selective replacement of low-skilled agents with moderately and highly skilled agents (according to the recruitment protocol, 'surviving' agents tend to have higher skills, all other things being equal). These results add to previous literature by indicating that heterogeneity, defined as average skill level variation, has negative effects on the size of the community. This seems to contradict the findings by Raban, Moldovan, and Jones (2010) who found that group

heterogeneity was the most prominent predictor of community success in the long term. At the same time, heterogeneity had positive effects through the increased skill level of agents and thereby served to make the community more stable due to lower member outflow. If the community is founded with only a few members, therefore, it has to be much more homogenous with respect to skill levels. If many members establish the community, it can be much more heterogeneous, and the community will converge into a coherent focus. Overall, our results contribute to the literature by demonstrating the differential effects of that part of heterogeneity relating to skill level.

We also investigated how the number of different kinds of skill types would influence the dynamics of the community. In that sense, heterogeneity was examined as the number of different skill types required to solve problems. We examined this by studying the effects of increasing the number of skill types required. Our results indicated that the more skill types required, the higher the probability that the community would not survive. This result adds to the literature by demonstrating that when a collaborative community is established, more actors are needed at the outset if the community is dealing with complex problems requiring many skill types than if it is dealing with simpler and narrower problems requiring only a few skill types. This finding also contradicts previous findings (Raban, Moldovan, and Jones, 2010) insofar as we did not find initial heterogeneity to be beneficial to the survival of a community.

To the extent that these simulation results are generalizable, our study has implications for both theory and practice. With respect to theory, our study adds specificity to the literature on the dynamics and longevity of collaborative communities. Our use of a simulation model enabled us to study the effects of the number and heterogeneity of actors which previous data, given the limited empirical literature, have not been able to examine thoroughly. Rather than observing

existing communities, our use of an agent-based simulation allowed us to explore the future consequences of a set of factors previously demonstrated to be important and thereby provide insights into the organizing principles of collaborative communities. Future studies should continue to examine the role of member composition and link it to the activity system of various kinds of collaborative communities. Simulation is especially useful in examining the effects of actor behavior over time.

From a design perspective, member composition is the most important decision to address once the overall purpose of a collaborative community has been defined. For example, if starting with a small number of members, the community should be homogeneous to survive in the short term. However, such a community will not grow very fast, all other things being equal. By contrast, a large number of members at the outset allows the community to be more heterogeneous and possibly more adaptable to changing goals and more challenging tasks. Most governance processes can be refined after the members begin collaborating with one another, with the objective of having the community be largely self-organizing and self-managing. It is the collaboration among community members around their shared interests that drives the sustainability of the community.

CONCLUSION

The collaborative community is a growing form of organizing, especially in knowledge-intensive environments. Although the absolute number of existing collaborative communities is not large, there is a good deal of variation in the shape of those communities – the missions they pursue, the types of members they include, the activity systems created to perform their work, and the governance structures and processes used to control and coordinate this particular type of

organization. We believe that many more agent-based simulations of collaborative communities are warranted so that newly designed communities have the most likelihood of achieving early and sustainable effectiveness.

REFERENCES

- Adler, PS (2001) Market, hierarchy, and trust: The knowledge economy and the future of capitalism. *Organization science* 12(2):215-234.
- Adler, PS, Heckscher CC (2006) *The firm as a collaborative community: Reconstructing trust in the knowledge economy*, Oxford University Press Oxford, UK.
- Adler, PS, Kwon, SW, Heckscher, CC (2008) Perspective-Professional work: The emergence of collaborative community, *Organization Science* 19(2):359-376.
- Anderson, KA (2016) A model of collaboration network formation with heterogeneous skills. *Network Science* 4(02):188-215.
- Appley, DG, Winder, AE (1977) An evolving definition of collaboration and some implications for the world of work. *The Journal of Applied Behavioral Science* 13(3): 279-291.
- Benkler, Y (2002) Coase's Penguin, or, Linux and "The Nature of the Firm. *Yale Law Journal*: 112(3):369-446.
- Burton, RM, Obel, B (2011) Computational modeling for what-is, what-might-be, and what-should-be studies-And triangulation. *Organization Science* 22(5):1195-1202.
- Burton, RM., Obel, B, Håkansson, DD (2015) *Organizational design: A step-by-step approach*, Cambridge University Press, Cambridge, UK.
- Burton, RM, Børge, O (2004) *Strategic organizational diagnosis and design: The dynamics of fit*, Springer Science+Business Media, Inc. Publishers, New York, NY.
- Bøllingtoft, A, Müller, S, Ulhøi, JP, Snow, CC (2012) Collaborative communities of firms: Role of the shared services provider. In *Collaborative Communities of Firms* (89-104). Springer, New York.
- Clark, KB, Wheelwright SC (1992) Organizing and leading "heavyweight" development teams. *California Management Review* 34(3):9-28.
- Davis, JP, Eisenhardt, KM, Bingham, CB (2007) Developing theory through simulation methods. *Academy of Management Review* 32(2):480-499.
- Eisenhardt, KM, Schoonhoven, CB (1996) Resource-based view of strategic alliance formation: Strategic and social effects in entrepreneurial firms. *Organization Science* 7(2):136-150.
- Fjeldstad, ØD, Snow, CC, Miles, RE, Lettl, C (2012). The architecture of collaboration. *Strategic Management Journal* 33(6): 734-750.
- Foster-Fishman, PG, Berkowitz, SL, Lounsbury, DW, Jacobson, S, Allen, NA (2001) Building collaborative capacity in community coalitions: A review and integrative framework. *American Journal of Community Psychology* 29(2):241-261.
- Garud, R, Jain, S, Tuertscher, P (2008) Incomplete by design and designing for incompleteness. *Organization Studies*, 29(3):351-371.
- Harrison, JR, Lin, Z, Carroll, GR, Carley, KM (2007) Simulation modeling in organizational and management research. *Academy of Management Review* 32(4):1229-1245.
- Hagedoorn J. (1993) Understanding the rationale of strategic technology partnering: Interorganizational modes of cooperation and sectoral differences. *Strategic Management Journal*, 14(5):371-385.
- Hess, C, Elinor, O (2006) A framework for analysing the microbiological commons. *International Social Science Journal*, 58(188):335-349.
- Kogut, B (1988) Joint ventures: Theoretical and empirical perspectives. *Strategic Management Journal* 9(4):319-332.

- Kolbjørnsrud, V (2014) *On governance in collaborative communities*. PhD Dissertation. Series of Dissertations: 12/2014, BI Norwegian Business School.
- Powell, WW, Koput, KW, Smith-Doerr, L (1996) Interorganizational collaboration and the locus of innovation: Networks of learning in biotechnology. *Administrative Science Quarterly*, 41(1):116-145.
- Miles, RE, Miles, G, Snow, CC (2005) *Collaborative entrepreneurship: How communities of networked firms use continuous innovation to create economic wealth*. Stanford University Press, Stanford, CA.
- O'Mahony, S (2007) The governance of open source initiatives: what does it mean to be community managed? *Journal of Management & Governance*, 11(2):139-150.
- Ostrom E. (1990) *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press: Cambridge, U.K.
- Raban, DR, Moldovan, M, Jones, Q (2010) An empirical study of critical mass and online community survival. *Proceedings of the 2010 ACM conference on Computer supported cooperative work*, ACM.
- Shuman, JC, Twombly, J, Rottenberg, D (2001) *Collaborative communities: Partnering for profit in the networked economy*. Dearborn Trade Publishing, Chicago, IL.
- Snow, CC, Fjeldstad, ØD, Lettl, C, Miles, RE (2011) Organizing continuous product development and commercialization: the collaborative community of firms model. *Journal of Product Innovation Management*, 28(1):3-16.
- Williams, RL, Cothrel, J (2000) Four smart ways to run online communities. *MIT Sloan Management Review*, 41(4): 81.

TABLES

Table 1. Base scenario parameter values

Number of problem-solvers	100
Number of problem-providers	85
Number of rounds (years)	100
Complexity: number of skill types of each solver or provider	4
Bound on initial skill level (min – max)	0
Bound on problem difficulty level (min – max)	0
Mutation level (deviation of new recruits from existing members)	5
Mutation sign	+, -, +/-
Problem mutation (when problems get solved and replaced by new ones)	2
Sign of problem mutation under replacement	+, -, +/-
Removal threshold for problem-solving agents	5
Removal threshold for problem-providing agents	5
Chance of new problem-solving agent recruitment event	33%
Chance of new problem-providing agent recruitment event	33%
Number of new recruits	3

FIGURES**Figure 1. Dynamics of the two agent subpopulations**

Figure 2. Average skill level of the two kinds of agents, with corresponding standard deviations

Figure 3. Number of agents in the two subpopulations with low skill level heterogeneity

Figure 4. Average skill level of problem-solving agents and average problem difficulty level in a community with low heterogeneity

Figure 5. Number of agents in the two subpopulations.

Figure 6. Average level of skills and problem requirement levels in a community with differential initial conditions.

Figure 7. Long-term community size for different levels of skill complexity

APPENDICES

APPENDIX A. Hockey stick effect occurs regardless of the initial number of agents of different types. The following plots of the numbers of agents show examples of different initial numbers and proportions of member types while keeping the other parameters constant according to the baseline scenario. All other parameters are as in the base scenario. Note the different y-axes.

Figure A1. Problem-solver members: 400 Problem-provider members: 100

Figure A2. Problem-solver members: 200 Problem-provider members: 100

Figure A3. Problem-solver members: 100 Problem-provider members: 50 **Figure A4. Problem-solver members: 50 Problem-provider members: 100**

APPENDIX B. AVERAGE SKILL LEVEL AND SPREAD AROUND THE MEAN: ROBUSTNESS CHECKS TO DIFFERENCES IN HOMOGENEOUS AND HETEROGENEOUS COMMUNITIES

Figure B1. Number of agents when the community is initiated with intermediate heterogeneity, i.e. skill values randomly drawn in the range 30-70 (recall in the baseline scenario it was from the full range 0-100 and figures 2a and 2b the more narrow range (40-60)).

Figure B2. Average skill levels shown for same scenario as in B1.