

Starting New Online Communities

Paul Resnick, Joseph Konstan, Yan Chen, and Robert Kraut

In Kraut, R. E. & Resnick, P. (Under contract). *Evidence-based social design: Mining the social sciences to build online communities*. Cambridge, MA: MIT Press.

Carving Out a Niche	2
Opportunities Model	3
Communities with Multiple Spaces	5
Competing for a Niche.....	11
Getting to Critical Mass	14
Bootstrapping: Leveraging Early Members to Get More Members	15
Attracting Early Members.....	18
Increase Stage 1 Value of the Community	21
Early Adopter Benefits	25
Setting Expectations for Success	28
Signals of Convener Quality and Commitment	28
Signals of Positive Member Response.....	30
Summary of Design Alternatives.....	34

Build it and they will come. If only it were that simple.

In reality, most online communities never really get off the ground. On SourceForge, for example, which offers free tools to open source projects, there are thousands of projects that have been created, but only 10.3% have more than three members (Jan. 2008 data pulled by Nate Oostendorp). In an effort to test the effects of an online community for helping people quit smoking, researchers gave 684 people access to an online community in addition to the informational website smokefree.gov, but so few people used the online community features that they were not able to report on its effectiveness (Stoddard, Augustson, & Moser, 2008).

Online communities fail to take off for many reasons. For some, it's just not clear whether the community offers services or experiences that potential members want. In these cases, why anyone would want to join? In other cases, the people who would want to join never find out about it. Some lose in a competition for members with another community. Some attract a stream of potential members, one at a time, but fail to hold them because there never seems to be anyone else around.

There are three major challenges in starting a new community. The first is to carve out a useful niche. The second is to defend that niche in the ecology of competing communities and alternative ways that potential members can spend their time. Meeting these two challenges requires making strategic choices about the scope of the community and about its compatibility

and integration with other communities. The third challenge is to get to critical mass. A new community must recruit members before it has become the kind of community that they will value. There are a number of design approaches to meeting this challenge, including substituting professionally-generated content for user-generated content in the early stages, leveraging early participants to attract later ones, and setting expectations about the likely future evolution of the community.

Throughout the chapter, we rely on simple cost-benefit assessments. Because so little empirical research has studied how new communities start, the design principles in this chapter rely more upon argument from first principles than does discussion in other chapters. Participating in a community involves opportunity costs of time and effort that could be spent elsewhere. Some of the costs, such as learning the community's software and getting integrated socially, are only incurred on entry. Others are incurred on an ongoing basis. Participating also can bring benefits, such as information, social connection, or a sense of identity and purpose. The first challenge can be reframed as ensuring that net utility, benefits minus costs, must be positive for all members in steady state. The second is that net utility must be higher than that of competing communities. The third challenge, of getting to critical mass, can be reframed as assuring net positive utility for each of the members as they join, even though the community has not yet reached steady state.

Carving Out a Niche

There are three major design decisions to make in order to carve out a niche. One is the scope of the community, in terms of the breadth of topics to cover, the kinds of people to try to attract as members, the activities to support, and the purpose of the community. Sometimes the community is defined by a topic and activities, and attracts a set of people who coalesce around the topic. For example, a Minnesota Twins fan community is defined by the topic of the baseball team and the activity of discussion about the team. The purpose is to inform and entertain the members. A diabetes support group for teens will include teens with diabetes, and perhaps their caretakers, with the purpose of providing information and emotional support. Sometimes a community is defined around a pre-existing set of people, such as an alumni group, with a purpose of staying connecting and the topics and activities emerging from the actions of the members.

The second major design decision is the extent of compatibility and integration with other sites, including borrowing of features and user interface elements, sharing of user identifiers, and import and export of content and people. Many new communities are integrated in some way with existing communities or platforms. For example, Facebook applications often define new communities and new groups are formed on platforms like Yahoo! Groups, Ning, and Meetup, as well as Facebook. Building on an existing platform typically means that the new community will share user interface elements with other communities using the same platform, that members will use the same user id across communities on that platform, and possibly that some content will be shared between communities. Stand-alone communities can also have some degree of compatibility and integration with other communities. For example, if they use common software packages like MediaWiki, Drupal, or Microsoft SharePoint, then user interface elements will be shared. User identities can be shared using technologies like OpenID. Content can be imported. It can also be exported to other communities, using technologies like RSS feeds or Facebook Connect.

The third is the internal organization of content, people, and activities within the community. All but the smallest online communities let people do more than one thing. Forum-based communities offer multiple forums covering related topics and frequented by different, overlapping subsets of the membership. Open source software projects subdivide into work on different modules. Chat communities offer multiple rooms. Some communities, like Facebook, depend on a pull model, where members regularly visit the community to see what new activity there is. Others, like Meetup.com, make use of push alerts, where members are notified when there is something potentially worth attending to.

We first assess the impact of those design alternatives on costs and benefits of participating in steady state, and then add considerations of initiation and switching costs to and from competing communities. People may get some fixed psychic value from membership in a community even if they do not participate, but for steady state analysis we focus on the marginal costs and benefits of participation, that is, consuming and producing content, and interacting with people.

Opportunities Model

While online communities are complex constellation of activities, we begin with analysis of the basic building block, a single interaction space. We model each space as a collection of interaction opportunities. For example, upon visiting FaceBook, a user might look at the news feed, which has a list of status updates or other activities from the user's friends or groups, each update an opportunity to be amused, bored, or annoyed. In a forums-based site, each forum is a collection of individual messages. In a chat community, a single chat room at any one time offers the opportunity to interact with each of the other people currently present.

To analyze the costs and benefits of visiting a space, think of each opportunity as having an expected *match_value*, the expected utility of examining the opportunity and possibly engaging with it. Note that the expected *match_value* may be negative: for example, spam messages or messages about a topic that does not interest the user. We define the *collection_size* as the number of new opportunities since the last visit, such as the number of messages in a forum or the number of status updates on a page of a Facebook news feed. In the typical pull design, where a user visits a web page (or an email folder or an RSS feed page) to get a collection, there is a per-collection *navigation_cost*, the cost of getting to the space and waiting for its contents to display. Clearly, when a space is nearly empty (i.e., *collection_size* is low), the *navigation_cost* may outweigh the total value, and the net benefits may be negative. The following expression gives the expected net benefit from accessing a space¹:

$$(\text{match_value} * \text{collection_size}) - \text{navigation_cost} \quad \text{[pull model]}$$

When there are few interaction opportunities, an alternative push model may be appropriate, where users are notified each time a new interaction opportunity becomes available. For

¹ Of course, interaction design options can have a big impact on the expected value of each opportunity, even for identical content. For example, consider a dating site. If the potential matches are presented as a list of names, each one has to be clicked on, and there is a 10-second delay before the page for a person loads, then the average match value may be low. By contrast, the same set of potential dates may have higher average value if there is a page that shows photos and a few key attributes and thus it takes very little time to weed out those who are incompatible. Interaction design alternatives are largely beyond the scope of this book. For the purposes of the model, assume that the interaction design has been optimized to yield the best possible *match_value* and *navigation_cost*.

example, in forums with few posts, new posts may be forwarded by email to subscribing members. In synchronous interaction communities like multi-player games or Second Life, users may receive alerts when their friends log on or enter particular spaces. The problem with the push model is that users are interrupted once for each new interaction opportunity, whereas in the pull model users expend time once to navigate to the whole collection. The following expresses the net benefit of push notifications for a collection of interaction opportunities:

$$(\text{match_value} * \text{collection_size}) - (\text{interruption_cost} * \text{collection_size}) \quad [\text{push model}]$$

The lower the per-item interruption cost is, the more attractive push notifications are. For example, many people, depending on their email reading habits, find email notifications to be a relatively small interruption cost. Pop-up alerts, on the other hand, may have somewhat higher interruption costs, since they may interrupt an activity that is more time-sensitive than email reading (Horvitz & Apacible, 2003).

The other consideration in assessing push notifications is the time-sensitivity of the opportunities. For example, when a collection offers synchronous interactions (e.g., a chat room), the match value of an interaction opportunity exists only while the other person is present. Thus, for example, push notifications about favorite chat partners coming online or favorite chat rooms having people in them may be especially useful. On the other hand, notifications about new messages in a forum would be less useful if the messages are not time-sensitive (e.g., a site like lightbulbjokes.com dedicated to collecting old light bulb jokes).

Design Claim 1: Lower volume and higher time-sensitivity of interaction opportunities, and lower interruption costs increase the benefits of push notification.

In the remainder of the analysis, we focus on the predominant pull model, where users visit a collection. Since there is a per-visit navigation cost that is independent of the number of opportunities actually present, the challenge is to increase the collection size while still maintaining the expected match value for each opportunity. One way to increase the volume of activity is to expand the scope, adding more topics and people with more diverse interests. As a thought experiment, imagine simply merging spaces with different topics. As an extreme example, consider a forum with posts about the Minnesota Twins baseball team, Pablo Picasso's paintings, and U.S. Civil War re-enactments. Mixing the three topics may triple the number of messages, but each visitor will discover that 2/3 of the content is uninteresting, and the community will likely shrink or fail. We will refer to lumping several independent topics together as a mixed-topic scope.

Design Claim 2: A mixed-topic scope reduces expected match value.

Even when the scope for an interaction space is not deliberately mixed, ambiguity about its scope may make it de facto mixed. Different people, thinking that the "true" scope is different, will use it differently. Match values will be lower when people interact with others who have a different interpretation of what the community is about, or the content they have created.

Design Claim 3: An ambiguous scope for an interaction space reduces expected match value.

In some cases, mixing different topics need not reduce match value. Consider assembling the various Major League Baseball fan communities into a larger MLB forum. Will this help? If people are simply there to talk about their own teams, then probably not. The best we can hope for in that case is a set of separate communities sharing infrastructure. But if fans of different

teams get some value from interacting with each other, then the situation changes. There are two ways this might occur, bridging activities and transcendent identities.

Bridging activities occur when there is an intersection of interests between two topical scopes. For example, fans of two teams that are about to play each other may enjoy talking with each other. Detroit Tigers fans may be interested in interacting with New York Yankees fans to discuss the performance of a player who switched from one team to the other or upcoming games between the two teams. Or fans of other teams may join into discussion with Minnesota Twins fans about their new stadium. When a space has a mixed-topic scope, designers would do well to identify potential bridging activities and seed the community with them.

Design Claim 4: Activities that bridge interests in different topics increase match value in spaces with mixed-topic scope.

Members may also have an interest in or identify with a broader topic. To continue with the baseball fan analogy, some people are fans of Major League Baseball as a whole, instead of or in addition to any particular team. They may be happy to discuss the use of instant replay in baseball, or of the use of the All Star Game results to determine the World Series home field advantage, or the differences between professional baseball in the U.S. and Japan.

Over time, names arise for those topics and identities that transcend component topics and identities. Over the years, most major sports leagues and events have expended significant marketing resources to create widely recognized names like the NBA, the World Cup, or Wimbledon, to encourage fans to identify with the league or event as a whole rather than just individual competitors or teams. One useful rule of thumb for designers, then, is to look for topical scopes which do not require compound names to describe. If the simplest description for a community is that it covers topics A, B, and C, it is a sign that there is not yet a transcendent identity.

Sometimes, the connections between specific identities or topics can itself become a new topic around which identities and a community can form. Etienne Wenger refers to these as bridging communities of practice (E Wenger, 1999). For example, a forum devoted to reconciliation (or just argument) between Israelis and Palestinians may have high match value for people who want to make connections between the groups. In the case of bridging communities, it may be fine if the only descriptive name is a compound one that describes the component topics and the fact that connections between them are the purpose of the community. Contrast that with a forum devoted to two separate topics, Israeli politics and Palestinian politics. If it attracted two groups of people, each interested in discussing only one of the two topics, but their discussions were mixed together, there would be low match value.

Design Claim 5: A transcendent or bridging topical identity increases match value in communities with mixed-topic scope.

Communities with Multiple Spaces

Combining multiple spaces in a single online community offers several advantages. First, it amortizes the fixed costs of development and branding over more users and uses. For example, craigslist.org has a unified software infrastructure, but interaction spaces are partitioned geographically. Second, it reduces search costs for members: once someone finds a useful community or a space within it, they are likely to be interested in several of the spaces. Third,

because there is overlap in who participates in the different parts of a community, synergies are created: a personal connection that two people make through interacting in one space also increases the value they get from interacting in another space within the community. We refer to the second and third advantages, reduced search costs and benefits of overlapping memberships, as synergies among spaces within a community.

The first question facing designers when a community has multiple sub-spaces is whether these should be globally defined, looking the same to all members, or whether there should be personalized views of the community’s activity. Facebook provides groups and pages, which look the same to all users who visit them, but also provides a personalized News feed, containing content posted by the user’s friends. Since those friends may not be friends with each other, they may not see the same content in their news feeds that the user sees.

As another example, upmystreet.com is a site in Britain that provides information about local businesses and real estate. It also provides a message board for conversations with one’s neighbors. This feature may be more popular in some regions than others. Rather than defining a separate forum for each city or region, all conversation goes into a single container, with messages indexed by the location of the person posting. In the screenshot, messages from the previous two weeks that were posted by people near the city of Cambridge are displayed. In a city with more participants, such as London, it automatically selects conversations within a smaller geographic radius. When a person in a city such as Cambridge with no recent conversations explores the site, they are more likely to view the site favorably if presented with activity in nearby towns than if they are shown no activity at all.

UpMyStreet
Rate local businesses
click here

Conversations near 'London'

24 hours | Week | **Fortnight** | Month | Year

Subject	Distance	Replies	Last post	Originator
Jamie Oliver - saint or shit?	1.03 miles	215	04/10/08 22:18 by Catland3	William
garden maintenance	2.16 miles	1	02/10/08 14:49 by Hortense Culture	athina
finding a place for primary school	2.85 miles	0	26/09/08 15:27	giani
bicycle repairs	3.61 miles	0	03/10/08 00:00	achilles
Walthamstow	4.04 miles	3	25/09/08 20:12 by Seashell	debbierae
HITCHIN, Oakfield Ave?	4.21 miles	0	03/10/08 06:57	arsenalboy
gardener	4.21 miles	0	25/09/08 09:02	felineweb
WOOD FLOOR	4.22 miles	1	02/10/08 23:26 by achilles	Vero
Whats Norbury like?	4.33 miles	4	24/09/08 22:57 by Ken Stevens	duskykitten
StudioFlat for Rent Immediately	5.95 miles	0	05/10/08 20:14	studioflat

> Next page

Figure 1. Personalized views in UpMyStreet.com

The same idea could be applied to other distance metrics. For example, sites could determine what to display for a particular user through text-processing algorithms that select the most relevant content or people profiles, or recommender algorithms that select content or people most likely to be liked by a particular person. Again, this would have the effect of always showing the closest content or people, so that something would be displayed even when a perfect match does

not exist. By contrast, dividing the people and content up into spaces creates hard boundaries so that when a space is empty it looks truly empty. Computer simulations suggest that these personalized displays of information provide more benefit to participants than do non-personalized displays, which show all content or show the same subset of content for everyone (Ren & Kraut, under review).

What is lost with personalized spaces is a shared context. In a forum whose boundaries are the same for everyone, it is reasonable to write a comment that alludes to other recent threads in the same forum. Most readers will also have seen those threads. In upmystreet, responding to a post at a distance of 4 miles with an allusion to another post would risk confusing readers who may live on the other side of town and not have seen the other post. Similarly, when commenting on a Facebook status message, it is unreasonable to assume that other commenters, or even the poster of the status message, will have seen the same other posts in their newsfeeds (though this doesn't seem to stop some people from posting comments that assume such a shared context on Facebook.)

Design Claim 6: Personalized collections of “most related content” enhance match_value but reduce shared context.

Assuming designers do create rooms, groups, forums, or other globally-shared contexts, the next question is which ones and how many. A search cost perspective suggests that it is important for people to be able to easily identify those spaces that will have high utility for them, meaning high match_value for the contents and a high level of activity so that there are many interaction opportunities.

A common mistake that online community designers make is to imagine all the topics that their hoped-for members might want to discuss, and create separate forums or chatrooms for each topic. Initially, at least, most of these spaces will be empty. People who choose a space to visit based solely on the topic names will find an empty space. Since there a cost to each such visit, the net benefit of spaces with few items will be negative.

An alternative is to provide initially just one space for interaction. While the interaction in that space may cover a broader range of topics and thus not completely match a visitor's interests, it will at least limit the number of separate access costs that need to be paid to examine the community's content. As a space gets enough activity, it can be split into multiple spaces that are all active. A secondary benefit is that the community designers or decision makers will have the opportunity to observe the naturally occurring topics of conversation and creates spaces based on the topics actually of interest to the membership rather than on a priori assumptions about the hot topics. On Usenet, there is a long history of broader groups forking into more specialized groups: for example, the denizens of the rec.humor newsgroup eventually split it into subspaces, rec.humor.funny for actually funny posts and rec.humor.d for discussion of humor. Similarly, the news and discussion site Slashdot.org started with a single collection of news items, under the tagline “news for nerds”. After Slashdot grew in popularity, it subdivided into specialized sections including topics such as hardware, games, science and rights online, each of which attracted enough commenters to maintain a lively feel.

Design Claim 7: Subdividing spaces after they become active creates more net benefits for participants than having lots of inactive spaces.

One way to mitigate the negative effects of empty spaces is to decrease the chances that visitors will stumble on them. This can be done through navigation aids that filter or sort based on activity, or at least provide activity indicators. For example, the drupal.org online support community for the drupal software package lists the time of the most recent message for each forum topic (see Figure 2). In this case, all of the support forums are relatively active. In startup communities, however, sorting the forums by most recent activity could help newcomers find the most active forums.

The screenshot shows the Drupal.org forums page in a Mozilla Firefox browser. The page title is "Forums | drupal.org - Mozilla Firefox". The address bar shows "http://drupal.org/forum". The browser's tab bar includes "Most Visited", "Getting Started", "Latest Headlines", "Gmail - Inbox (1) - res...", "Gmail - Inbox (6) - resnick.paul@g...", "Google Docs - Folder - Handbook", "New Communities_Oct 2008 - Goo...", and "Introduction to CommunityLab Ha...".

The forum page has a header with "Forum" on the left and "Topics Posts Last post" on the right. Below the header is a blue "Support" banner with the text: "Try searching the forums using the advanced search option first or a specific project's bug reports. Remember all support on this site is on a volunteer basis, so please visit the forum tips for posting hints." Below the banner is a table of forum topics:

Forum	Topics	Posts	Last post
Post installation Drupal is up and running but how do I ...?	62165	223554	10 min 23 sec ago by razzel
Before you start Is Drupal a viable solution for my website? Please see the documentation Before you start before posting.	3391	14230	1 hour 28 min ago by No_Idea_Yet
Installing Drupal Installing Drupal? Please see the documentation in the handbook and the video resources for Drupal 5 and Drupal 6 for additional installation resources.	9266	39317	2 min 58 sec ago by biosv
Upgrading Drupal Questions regarding upgrading an existing Drupal site. Don't forget to read the UPGRADE.txt that comes with every Drupal download.	2713	10911	24 min 35 sec ago by vicrey!
Converting to Drupal Need help migrating your site to Drupal?	1271	5101	2 hours 54 min ago by aufumy
Module development For assistance with module development.	15030	51978	32 min 4 sec ago by brightondesigns

Figure 2. Drupal support forums with indicators of recent activity

In the group chat community PalTalk, most of the open chat rooms at any one time have very few participants in them—people who are regulars in those rooms may be satisfied with the small number of participants and occasional interactions that occur in them. New users, however, who are exploring and evaluating the service, are less likely to find such rooms attractive. The user interface to select chatrooms encourages users to select a room to visit from a display of the rooms that are currently most active, sorted in descending order by number of people in them (see Figure 3).

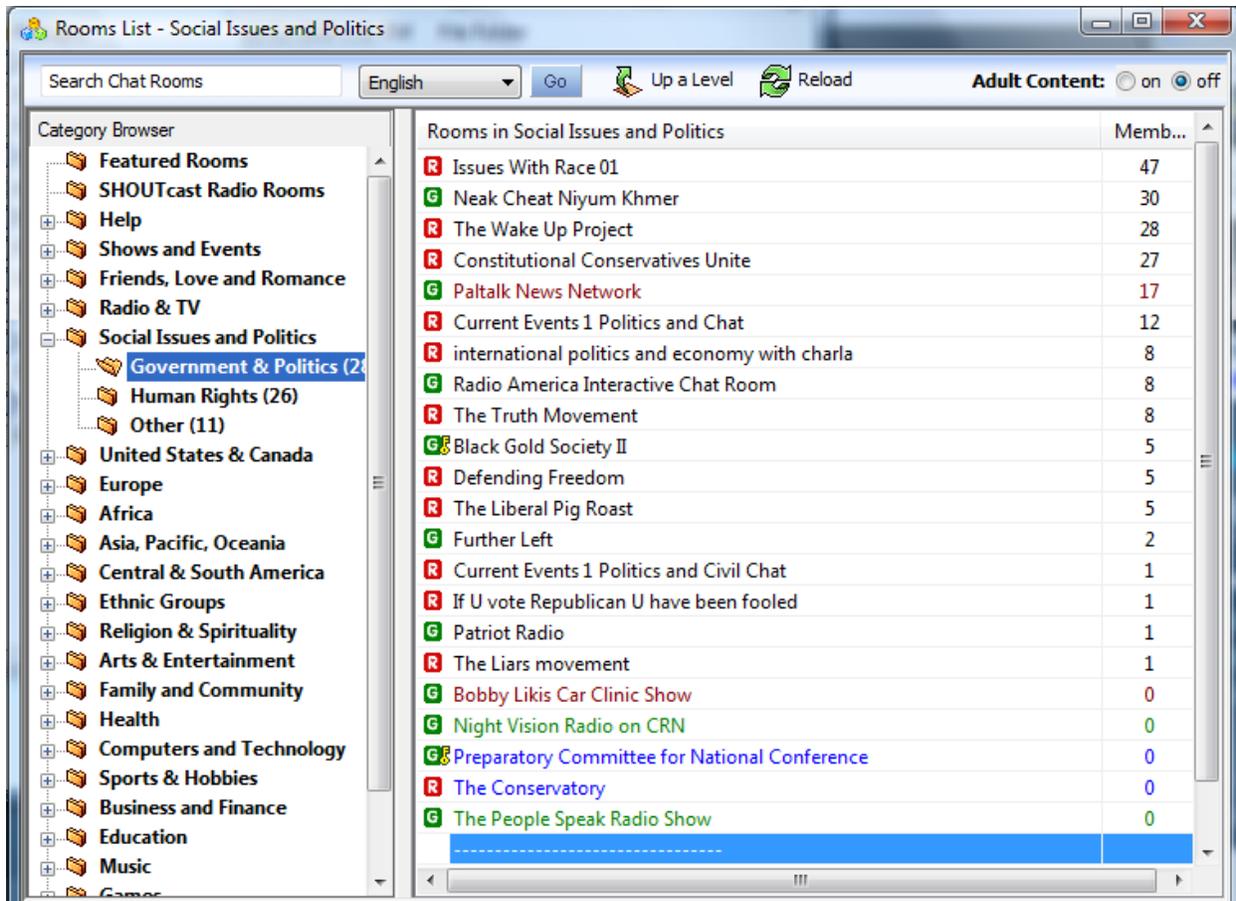


Figure 3. The PalTalk chat rooms in the category “Government & Politics”, sorted by number of people currently logged in to them.

Design Claim 8: In communities with lots of interaction spaces, navigation aids that highlight more active spaces will increase the net benefits members experience.

In communities with synchronous interaction spaces, like chat rooms or islands in Second Life, if they are not always active, it may be helpful to schedule times when they are expected to be active. That way, people can avoid visiting only to find the space empty and, if the space is compelling enough, can plan to visit when it will be open and active. The times may be announced when the space is created, based on intuitions about when people will want to participate, or just the availability of the moderator or creator of the space. For example, The Wellness Community organizes small online support groups for cancer patients. Each group has a scheduled 90-minute weekly synchronous session, in addition to a private asynchronous discussion forum. Similarly, on PalTalk, for example, someone who stumbles on the room titled “Bobby Likis Car Clinic Show” and finds it empty will be informed that it is, “Live every Saturday from 10a to 12n ET.” In the absence of fixed schedules, information may be provided based on historical data about when a space has been active, with the “expected popular hours” posted.

Design Claim 9: In synchronous spaces that are not always active, a schedule of “expected active times” can coordinate visitors and become a self-fulfilling expectation.

Navigation aids can not only reduce search costs by identifying active spaces but also by identifying spaces that are likely to have high match value. A visitor who finds one space she likes can be directed to other spaces that are “similar” in some way, where similarity may be defined by text matching (content-based filtering) or based on who participates in them (collaborative filtering). For example, in a chat community where there are hundreds or thousands of separate rooms, someone might be informed of other rooms frequented by the people who she has interacted with previously.

Design Claim 10: In communities with lots of interaction spaces, recommender systems that help people navigate to spaces that best suit them will increase the net benefits people experience.

While ambiguity about the scope of individual spaces is largely harmful, some ambiguity about the scope of the community as a whole is natural and sometimes beneficial. Organization theorists point out that uncertainty and disagreement about the purpose is a natural and unavoidable part of the startup phase of any new community of practice, including new online communities (Etienne Wenger, McDermott, & Snyder, 2002). It can be more or less painful, and cause more or fewer problems in the retention of members, depending on design choices that are made. But the negotiation of a shared purpose cannot be eliminated entirely. Even if the founder of a community announces a purpose, the members may not accept it as their purpose -- it only becomes a shared purpose through the actions of members that serve to reinforce or challenge it. One example, described by Amy Jo Kim (Kim, 2000, p. 19), is the L'eggs pantyhose web site discussion area that, against the wishes of its owner, became a site whose main purpose, shared by its active participants, was to provide an anonymous forum for men to discuss the joys of wearing pantyhose.

Some ambiguity about the eventual scope of the community has some advantages. First, it allows the community designers to learn from the members what the members want. A topic may attract an unexpected audience or the audience may be interested in different topics or activities than the designer first intended. Fighting against what the members want, by trying to stick to the original vision, can alienate them. For example, the founder of friendster.com, the first widely popular social networking site, alienated many of its members when it refused to allow them to engage in playful uses of the site, with fake profiles that did not accurately describe themselves. According to boyd and others, the active deletion of these “Fakester” accounts ruptured trust between the company and its users, causing many to leave (Anderson, 2003 ; d. boyd, 2006; Ellison, 2007). Second, the activity of negotiating the scope, especially through explicit meta-discussion about it, can itself be a rewarding activity for some members, and lead them to feel “ownership” of the community, and thus commitment to it, as discussed further in the chapters on commitment and regulation.

Design Claim 11: Ambiguity of scope for the community creates opportunities for adjustment and member ownership.

Even when everyone is interested in a narrow topic, communities with too many people may have a reduced match value. As described in more detail in chapter **TKTK**, in so-called bond-based communities, the primary source of commitment to the group derives from interpersonal bonds with other individual members. Those interpersonal bonds depend on repeated exposure to the same people. Consider, for example, a cancer support community. A person who reads 30 messages per day will likely get more emotional support if the same people write those messages each day. Indeed, the match value might be enough higher that it would be worth restricting the

community size (or subdividing it so that it effectively creates several smaller communities). For this reason, the Wellness Community, mentioned previously, deliberately launches small sub-communities of just twelve people, in addition to offering a set of forums for the community as a whole.

Design Claim 12: A larger community leads to lower match-value in bond-based communities.

Competing for a Niche

Some new communities enter a crowded landscape. If a company introduces enterprise social networking software, it needs to be cognizant of the other social networking sites its employees may already be participating in, such as Facebook and LinkedIn, because the internal site will be competing for employees' time and attention with those other sites. If a new site for cooking enthusiasts starts up, its potential members will have many other options for places to interact online with other cooking enthusiasts. Even when a new community does not have any obvious competition, it will be useful for designers to conduct a competitive analysis, for two reasons. First, potential members always have the option of muddling through with their existing communication patterns. For example, a newly diagnosed diabetic, instead of joining a diabetes support group has the option to just communicate with his existing email, IM, or Facebook networks, even if the people in those networks are not very interested in or knowledgeable about diabetes. Second, even if there is not currently a crowded landscape, there may be competitors in the future. For example, if you start a fan club for an obscure band, there are likely to be competing clubs eventually if the band becomes popular. By conducting a competitive analysis, the designer can carve out a niche that can be defended against future competitors. Shapiro and Varian analyze competitive strategies for information products more generally (Shapiro & Varian, 1999, chapters 5-8). Drawing on that analysis, we pull out implications for online communities.

Our initial interaction opportunities model for analyzing costs and benefits considered only the marginal costs of continued participation once someone had already joined. For competitive analysis, we need to consider also a new member's switching costs of leaving a community and getting started with a new one. Startup and switching costs come from a number of sources. First, a new user has to learn to operate the community's software. Second, he has to learn his way around, finding the areas of the community that have high match value for him. Third, he has to learn the social norms of appropriate behavior. Finally, to achieve maximum benefits from participation, he has to build up social connections and gain status in the community. We will model the sum of all these as a single quantity, the `switching_cost`.

These switching costs must be amortized in some way if we are to compare them against the benefits from participating over some period of time. We simplify this comparison by assuming there is some expected duration of participation. We roll up the expected net participation benefits (i.e., benefits - minus costs) from the opportunities model of the previous section, over the expected duration of participation, in to a single number, the `participation_benefits`².

² For those readers familiar with economic models, think of a discount rate δ for benefits or a per-period probability of exit δ . The `participation_benefits` quantity would then be the expected discounted net benefits, integrated over all time periods.

In a competition between an incumbent community that someone is already a member of and a competitor which the person could join, the competitor must offer participation benefits that are enough better than the incumbent to compensate for the switching cost. Incumbents gain strategic advantage when switching costs are higher: the costs serve to *lock-in* members. In the startup stage, a community is in the role of competitor and will generally want the switching-costs to be low, making it easy for people to join. Once people have joined, however, the community is in an incumbent role, and a community designer will want the costs of those same members switching out to be high.

Choices about compatibility will affect switching costs. For example, if the new community uses the same interface elements as other communities, it reduces the costs of learning to use the software, thereby reducing switching costs. One way to do that is to use a popular software platform for online communities, like phpBB, Drupal, Ning, or Yahoo! Groups. Employing innovative user interface elements will only make sense in a community's startup stage if they create significant additional participation benefits, enough to outweigh the additional switching costs. Moreover, innovative interface elements that are truly valuable may not convey a long-term edge against future competitors, as other communities can copy them.

Another compatibility decision that affects switching costs is sharing user IDs or profiles. A competitor that allows members to register and login using OpenID credentials issued by an incumbent (or by some other popular site such as Facebook) can reduce potential members' switching costs. If a user can import her profile and friend links from the incumbent site as well, switching costs are further reduced. Of course, the incumbent may not allow the sharing of IDs and profiles, in order to keep switching costs high (e.g., eBay does not allow its members' feedback profiles to be imported to other sites). When sharing IDs and profiles is an option, designers will have to judge which is more important, the initial strategic value when recruiting members, or the later strategic value when trying to retain them. In addition, in some cases differentiated IDs may be critical to the participation benefits in the community, preserving the ability for members to separate their different online social contexts rather than collapsing them as happens when people from different parts of one's life access the same online persona (d. m. boyd, 2008). For example, a health support community might advertise itself as a safe environment for sharing personal experiences, which inherently requires that the user id not be shared between the health support community and other communities.

Design Claim 13: Differentiated user interface elements in the competitor community create startup costs and thus favor the incumbent community in any competition over members.

Design Claim 14: Non-shared user IDs and profiles between incumbent and competitor communities creates startup costs and thus favors the incumbent community in any competition over members.

Design choices about import and export of content from competing communities will affect the relative benefits of the communities. Importing content from a competitor community incorporates some of the benefits of the other community into the new community. In equation **TKTK**, $\text{participation_benefits}_2$ increases without any change in the other terms, which makes the new community more attractive relative to the incumbent. In the next section, we will also see that temporarily importing data can be an effective tool to get through the initial stages, before there is sufficient participation in the community.

The opposite analysis holds for exporting data: if the new community exports content, then its competitors are able to increase their net participation_benefits and thus compete more effectively. Why, then, would any community export its data? One possibility is that it does not think it is in competition with the communities to which it exports its content. A second is that it increases visibility for the community that is the source of the content, allowing it to compete for members who would not otherwise have been aware of it. Especially if the exported content is read-only, also it may serve as a lure to those who wish to respond or join a discussion. A third is that some content contributors may be motivated to reach as wide an audience as possible: exporting the content may increase their benefits of participation in the community where they post. For example, someone who has many friends on Facebook but few followers (so far) on Twitter, might post only on Facebook if she had to choose one or the other. An application that automatically reposts her tweets as Facebook status messages might tip the balance so that she would post on Twitter.

As with many design choices, then, tradeoffs are involved. A community that is not yet well established may benefit from importing content (to enhance the value of reading there) and from exporting content (to enhance the value of writing, and to enhance awareness of the community). Sharing with less well-established competitors, however, may help the competitor more than the incumbent.

Design Claim 15: content sharing between competing communities raises awareness of the exporting community and the value of posting there, but raises the value of consuming content in the importing community.

In competitive situations, it is especially important to clearly convey the benefits of a community. Many people will not carefully investigate the community, but will instead assess the community based on short descriptions or reviews from others who do investigate. Given limited attention from evaluators, it is difficult to convey the value of a large set of small improvements. Instead, it is far more effective to identify one or two key elements. In television advertising, Rosser Reeves argued for conveying a “Unique Selling Proposition”, a benefit to the user that the competing product doesn’t offer (Reeves, 1961). In the online community setting, the unique selling proposition or core selling proposition may be a topic not covered elsewhere (e.g., “the Wikipedia of news translation”), a different set of participants (e.g., “the dating site for artists and lawyers”), a different set of activities and interaction tools (e.g., “an easier way to connect with your friends”) or a different set of social norms (“a truly *supportive* community to help you lose weight”).

Design Claim 16: Conveying a succinct unique selling proposition will attract members.

Economists describe some competitions as “winner-take-all”. For example, if there are two communities competing for exactly the same niche, if either community attracts most of the available members, their membership will make that community more attractive than its competitor. In a winner-take-all situation, it is more important to convince people that your community will succeed than to convince them that your community is inherently better. If everyone thinks that others will join community C2, they will join it, but if they all think others will join C1 they will join that one. If either C1 or C2 becomes what is called a focal point (Schelling, 1958), widely shared expectations of success will be self-fulfilling.

As Shapiro and Varian (1999, p. 181) argue, “The aura of inevitability is a powerful weapon...” Cultivating public awareness is one way to create that aura. If the public is much more aware of

C1 than C2, then even people who are aware of both know that many people are only familiar with C1. Therefore, they will expect C1 to win and will prefer to join it. Consider, for example, Angie's List, a site that provides reviews of local contractors and businesses that serve household needs, has spent much more visibly than other sites in this arena. It has advertised extensively on NPR and even television spots (<http://www.allbusiness.com/marketing-advertising/4163708-1.html>). In a city where there are competing online communities providing this serve, many more people will be aware of Angie's List than its competitors. Even people who are aware of competitors may participate in Angie's List simply because they expect others to eventually join it.

Celebrity endorsements can also help to create a focal point. For example, online poker sites have emerged and grown rapidly in the past several years as poker itself has grown and become a major televised activity. The sites themselves are organized into poker tables based on the game played and the betting limits. Since play is against other members, a site without enough traffic to have members populating nearly every type of table is likely to lose out to its competitors. Pokerstars.net heavily advertises its team of celebrity poker professionals. Its home page prominently features three recognizable "world champions", and its television advertisements (run during televised poker tournaments) highlight the professionals associated with the site. Indeed, all of the major online poker sites build their image around a set of nationally-visible poker stars (and with mottos like "play where the pros play"). Even though amateurs will probably never play with the pros (unless they are willing to play for high monetary stakes and lose a lot of money to them), the endorsements from stars help to create a focal point for where amateurs will think to go when they want to play poker online.

Design Claim 17: Advertising and celebrity endorsements can help to create awareness of a community and thus make it a focal point in a competition between communities.

Getting to Critical Mass

The third major challenge is to get past the initial growing pains to a critical mass of participation. There are two ways that a community in the startup stage can fail to provide the value to its members that it eventually could provide. First, it may not have enough members to provide the content and interaction opportunities that some prospective members want. Second, the members may not yet have a shared purpose, including rough agreement about the scope of activity and membership, along with the norms and governance needed to achieve that purpose, so that less valuable content and interactions may crowd out the valuable activity. Social science theory can help us understand the challenges, and point us in the direction of strategies to help online communities succeed through the startup stage.

The economic theory of network externalities or network effects explains situations where one person's value from using a product or system increases with the number of participants in the system (Katz and Shapiro 1985, Farrell and Saloner 1985). A good example is the fax machine. Being the only person in the world who owns a fax machine does not generate much utility. A user's value for a fax machine increases with the number of other people who use the fax machines. Metcalfe (1976?) argued (for the specific case of Ethernet) that the cost was proportional to the number of machines connected, but the value was proportional to the *square*

of that number. The term Metcalfe's Law is often used to describe such situations where, beyond some critical mass, the benefits of increasing the network size dwarf the the costs of doing so.

There are two problems, however, when individuals decide sequentially whether to join a network. First below the critical mass size, adding a member may not create enough benefits to outweigh the costs. In the extreme case, there can be an unfortunate equilibrium where people sequentially decide not to join, even if they could all benefit from a coordinated decision to join. Second, joining creates an externality, a benefit for all the other members. Thus, even when the total benefits to the community outweigh the new member's costs, the benefits to the new member may not outweigh the costs, and the member may not join.

The same concepts can easily be applied to online communities. A social networking site such as LinkedIn or Facebook has little or no value if nobody else is using it, and much greater value if many of the friends or colleagues of a prospective user are already signed up. Communities built more around content rather than connections not only face the public goods under-provision problem, but they may also have a notion of critical mass-- a level of usage at which it becomes clear that one's efforts are worthwhile. There must be both enough people editing and reading, and enough high-quality articles, at a site such Wikipedia for new users to perceive it as a venue worthy of their own participation.

Given these problems, there are two useful approaches. One is to make the community more attractive to early joiners. The other is to make more effective use of the early members, leveraging them to attract additional members. An understanding of how to leverage early joiners has implications for which types of people are most valuable as early members. Thus, we first analyze how to leverage early joiners. Then, in our discussion of how to attract them to join early, we are able to suggest design alternatives that are tailored to attracting the most valuable types.

Bootstrapping: Leveraging Early Members to Get More Members

A community may go through a series of states, each of which is attractive enough to engage a new subset of members who improve the quality of the community for the next wave of membership. When the presence or actions of early members leads other people to join the community, we describe it as bootstrapping. One approach is to organize activity so that it creates content that will be attractive to future users. A second approach is to include viral elements so that current members bring in their friends and acquaintances, who increase the value of the community to others.

The natural use of a site by early members may not always be sufficient to generate content that attracts others. But they may be more inclined if they are paid to do so. For example, epinions.com, the product review site, initially paid contributors for providing reviews. Chapter TKTK examines ways to motivate contributions more generally. Here we note some special considerations about tangible rewards for early content contributions. The fact that other members are paid to perform certain activities may demotivate volunteers from performing those same actions for free. Once offered, it is also demotivating to take away rewards for actions. When epinions.com reduced the payments to reviewers and changed the terms, it lost many contributors and had to put a lot of customer service energy into quelling dissatisfaction.

Discounts or free service may be a less problematic way to encourage early members to contribute. It is easier to discontinue these once the community reaches a critical mass. They may also be continued for the early members, without offering them to those who join later, creating a form of early-adopter benefit that we will discuss below. Discounts and free service may avoid creating the kind of envy that direct cash payments would generate.

Design Claim 18: incentives for early members to generate content can increase bootstrapping.

In order to leverage early members, it is most useful to encourage them to create primary content. In online communities, the primary content consists of blog entries, forum posts and responses, audio, video, and text documents, items for sale, and the like. Metadata consists of tags, ratings, commentary about primary content, and behavioral clickstream data about which primary content items were accessed. Metadata adds a lot more value once it is needed to help navigate through a large quantity of primary content. For example, imagine a new, specialized photo-sharing site. Suppose it had only three photos posted, but hundreds of tags for those photos and votes from users about which of the three they liked best. It would be less likely to attract new members than one with hundreds of photos and only a few tags and votes.

Design Claim 19: User-generated primary content will do more to bootstrap additional membership than will user-generated metadata, in the community startup stage.

In addition to generating attractive content, members can directly attract other members. Because people are members of multiple groups and communities, awareness of one community can spread to other communities. Indeed, a viral spread of membership is a natural phenomenon that may occur organically, without special intervention from designers. For example, Backstrom et al found examined patterns of joining communities within the blogging site LiveJournal, and of joining (participating in) academic conferences that occur annually. They find that people are more likely to join groups the more of their existing friends are members and the more those friends are friends with each other (Backstrom, Huttenlocher, Kleinberg, & Lan, 2006).

One trick for designers who want to supercharge the normal viral dissemination process is to get some marker of membership to display on web pages or in .sig lines of messages that are viewed by people who are not members. For example, within the overall Facebook site, there are lots of specific communities. When a person joins one of those communities, that membership can be displayed on the user's profile, visible to any of his/her acquaintances who view the profile. The action of joining the group may also be propagated to the feeds of other people on Facebook who are official "friends" but who may not yet be a member of the group.

Design Claim 20: services that enable displays of membership that are visible to non-members will lead to bootstrapping.

An even more effective way to leverage multi-membership is to make actions a user takes in one community visible to their acquaintances outside the community. For example, Presidential candidate Barack Obama's website allowed members to create blogs on the site, and take various actions. It gave members the option of linking with their existing Facebook account, so that some of their actions on the Obama site would generate entries in Facebook, and thus be visible in the feeds of the Facebook "friends". Numerous Facebook apps work similarly. For example, many

RECENT ACTIVITY



Figure 4. FourSquare news feeds

people have been introduced to the FarmVille and FourSquare apps by seeing status messages, in their main Facebook news feeds, describing their friends notable, or not so notable, events with those apps (see Figure 4).

Design Claim 21: services that make members' actions in the community visible to their acquaintances outside the community will lead early participants to attract later participants.

News and content sites often provide a feature that lets readers notify friends about content that they found interesting. For example, the New York Times website offers buttons that allow readers to forward content by email, or to recommend the article on news aggregator services like Digg. Because the site facilitates the forwarding, the notification about the interesting content can also highlight the existence of the site as a whole. The same mechanism can be provided in online communities for forwarding interesting content. For example, <http://www.foodnetwork.com/> has an Email button, which allows a user to email the recipe to a friend with an embedded sentence, "Checkout this awesome recipe on FoodNetwork.com!"

Design Claim 22: services that allow members to forward content from the community to their acquaintances outside the community will lead early participants to attract later participants.

The viral marketing approach is even more evident in some of the explicit social networking sites. When signing up for the professional social network site LinkedIn, members are invited to let the software scan their instant messenger contacts and their e-mail contacts both to help them connect with those already in the site (increasing value) and to invite those who are not there

(bootstrapping). Thus, each new member potentially brings in other acquaintances, and the site provides tools that reduce the effort needed for such invitations.

Design Claim 23: services that allow members to invite acquaintances outside the community to join will lead early participants to attract later participants.

As with content-generation, explicit incentives may be provided for referring new members. Physician community sermo.com has offered cash and prizes to members who refer their friends. Cloudmark's SPAMNET offered a free month of service for each customer referred. One challenge with pay-for-referral schemes, and the reason that most free sites don't use them, is ensuring that members refer actual people who are likely participants rather than their own six other e-mail accounts. Sermo handles this with an elaborate sign-up process; even though the site is free to members, those members have to demonstrate that they are licensed physicians. Cloudmark's model is similar to many pay services; you only get the credit after the referred member has been a paying subscriber for 2 months. Other commercial sites offer members the chance to give discounts with their referral (10% off the items you recommend) or give a percentage back to the referrer, a form of multi-level marketing.

Design Claim 24: pay-for-referral and revenue-sharing from referrals increase bootstrapping

Attracting Early Members

We adapt the previous cost-benefit analysis framework to consider the problem of attracting early member, who decide whether to join before the community has reached a critical mass of participation from other people. There are many possible futures for a community in its early stages, with each providing a different level of possible benefits for members. The key idea for understanding the impacts of different design choices is that they affect the likelihood of those future states, or the value that members will get from them. To present that idea in its simplest form, imagine that, after an initial time period in its current state (the first stage), there are only two possible future states, one where the community fails completely and one where it succeeds completely, having attracted a lot of participants who have developed a sense of shared purpose and social norms about how to act in the community. A potential member will have a belief about the *success_probability*, the likelihood that the community will achieve the success state.

We model the expected utility of a decision to join early and compare it to the expected utility of waiting and joining later.³ Joining early, in the first stage, requires paying the *startup_cost*, as in the model from the previous section. It also yields an expected total net benefit from the first stage, prior to either achieving critical mass or failing, which we represent as *participation_benefit_{stage1}*. If the community goes to the failure state, assume that no one will continue to use it, so there are no additional costs or benefits for anyone. If the community succeeds, participation in the second stage will yield additional benefits, *participation_benefit_{stage2}*. In addition, as we shall see, in the success state there may be

³ The decision is taken against a backdrop of some other ways the person could spend their time. All utilities in the model should be thought of as net utilities relative to the outside option, the best other way the user could spend their time. That is, we have normalized the value of the outside option to 0, allowing us to omit it from the models for simplicity.

additional benefits for the early members, such as status in the community, which we model with a variable, *early_adopter_benefit*.

In summary, we have expected utility of joining the community in the first stage:

$$\begin{aligned} \text{util}(\text{join now}) = & \\ & \text{participation_benefit}_{\text{stage1}} - \text{startup_cost} + \\ & \text{success_probability} * (\text{participation_benefit}_{\text{stage2}} + \text{early_adopter_benefit}) \end{aligned}$$

Instead of joining in the first stage, a member has the option of waiting until stage 2. If the community succeeds, she can join then. The member gets no benefits in stage 1, but has the advantage of not paying the startup costs in those cases where the community is not going to succeed. The expected utility of that option is:

$$\begin{aligned} \text{util}(\text{wait}) = & \\ & \text{success_probability} * (\text{participation_benefit}_{\text{stage2}} - \text{startup_cost}) \end{aligned}$$

For someone to prefer joining now rather than waiting, the expected utility must be higher:

$$\text{util}(\text{join now}) - \text{util}(\text{wait}) > 0$$

Expanding the two quantities and simplifying, people prefer joining now only if:

$$\begin{aligned} & \text{participation_benefit}_{\text{stage1}} \\ & - \text{startup_cost}(1 - \text{success_probability}) \\ & + \text{early_adopter_benefit} * \text{success_probability} \\ & > 0 \end{aligned}$$

From this inequality, it is easy to read off approaches that make it more attractive to join now rather than waiting. Subsections below take up each of the approaches in detail.

- 1) **Increase stage 1 value of the community.** Ideally, the net benefits of early stage participation can be made positive, despite the lack of other members. We describe ways to provide non-social value that is independent of other members' participation, in the form of access to content or services that are exogenously generated rather than provided by community members.
- 2) **Reduce startup costs of joining the community.** Reducing startup costs makes early joining more attractive relative to waiting because these costs are paid by early joiners even when the community fails. The effect of design choices such as UI compatibility that affect these joining costs were already explored in the previous section and will not be repeated here.
- 3) **Early adopter benefits.** We present three approaches for creating early adopter benefits. One focuses on the benefits of having skills and familiarity when the community becomes popular. A second focuses on the benefits of establishing a reputation as a leader early in a community's life. A third, in communities expected to generate revenues, is to promise a

share of future monetary revenues to early contributors, analogous to the way startup companies offer stock to early employees.

- 4) **Expectation-setting.** Increasing expectations about the probability of the community's success will both reduce the chances that an early joiner's entry costs will be wasted. It also increases the probability of receiving any early adopter benefits that will be available if the community succeeds. Expectations can be shaped through signaling about commitments to future resource expenditure. Other expectation-setting approaches include conveying a trajectory of growth, conveying images of what the community is expected to be like at some future time, drawing analogies to other successful communities, and communicating contingent participation commitments that other potential members have made.

All of these approaches can be targeted toward specific sub-populations. One possibility is to target the natural "lead adopters", those who need the least extra encouragement to join early. The lead adopters for many innovations are people with special needs for the good or service being offered. For example, men seeking sexually oriented entertainment were early adopters of the French Minitel system in 1982 and 1983, an early videotex service. Shortly after Minitel offered text-based communication, sexually-oriented chats known as *messagerie rose* (*pink messaging*) became "one of the main forces behind Minitel's success, especially between about 1983 and 1987, by which year 'pink sites' were clocking up staggeringly high usage figures" (Jacobs, 2003, p. 81). Doctors, factory managers and isolated farm women, all social categories with strong needs for communication, became early adopters of the telephone system in the United States (Fischer, 1992). Translated to the online community domain, early adopters would be the people who are most interested in the topic to be discussed or who most want to use the products (e.g., open source software) the community will create. One study found that the people most likely to make use of the online community features of a pedometer-based walking program were those who had the least offline social support for their efforts to become more physically active (Richardson, et al., Under review).

Those who need the community most may not always be the best to target for early adoption. Rogers (1995, chapter 7) notes other characteristics of early adopters of innovations, some of which are relatively easy to discern among potential members of a new community and some of which are less so. In terms of their demographic characteristics, early adopters tend to be richer, of higher social status and better educated than later adopters. In terms of personality, cognitive and attitudinal individual differences, early adopters tend to be more intelligent, rational, better able to deal with abstractions and have more favorable attitudes towards change and toward science than later adopters. In addition, they are better able to cope with risk and uncertainty. Taken together these characteristics fall along two dimensions: the education, intelligence, rationality, and ability to deal with abstractions that allows them to understand better the potential future benefits they may receive from adopting the innovation, and the wealth, status and ability to deal with uncertainty that allows them to cope with the risks associated with adopting an innovation. These characteristics can also partly explain what Rogers calls the 'innovation-need paradox' (Rogers, 1995, p. 275) – the observation that in many cases those who most need the benefits of a new idea (e.g., the poorer or less educated) are among the last to adopt it. In the United States, the poor, the less educated, and the elderly are among the last to use online resources, including health support groups, to deal with illness (Fox & Jone, 2009). They may not have the financial resources to absorb the risks associated with innovations nor the information or education to understand sufficiently the potential advantages. In starting an online

health support group, then, it may be better to focus first on attracting younger, more connected people, even though they may have less need than people at a more advanced illness stage who are homebound and more socially isolated.

Another useful population to target is the people who provide the greatest leverage in attracting other members. We have already argued that content production is more valuable than metadata-production for attracting additional members. In many communities, there are some people who are natural content producers and others who are more naturally content consumers, or producers of meta-data. For example, on a blogging site like LiveJournal, with many groups, there are some people who primarily benefit from having an audience for what they write, while others primarily benefit from being the audience. In the early stages of starting a blogging community, it is more important to attract those who want to write than those who want to read.

More generally, economists describe two-sided networks [(Armstrong 2002, Caillaud and Jullien 2003, Rochet and Tirole 2003, Parker and Van Alstyne 2005)]. These are situations where there are complementary types of participants, each of whom produces value for the other type. For example, in a LiveJournal group, the presence of a lot of readers makes the site more attractive for writers, and vice versa. The presence of other writers may not be so attractive for writers, as they create competition for reader attention, though their presence may be tolerated if it is the reason that there are many readers. Dating sites for heterosexuals are another type of online community with two-sided externalities. eHarmony is more attractive for women if there are more men participating, and vice-versa. Thus, a useful strategy in the startup phase may be to provide incentives for one gender to participate, in order to attract the other, much as bars and nightclubs sometimes offer free entry to women on "ladies night", in order to make the venue more attractive for men.

When there is a choice about which type to try to attract first, there are a couple of considerations to keep in mind. First, it may be that only a few of one type are needed to attract the other type while many of the other type are required to attract the first type. For example, attracting a few blog writers may be sufficient to subsequently attract many readers. By contrast, attracting a few readers will rarely be sufficient to subsequently attract many writers. Second, one type or the other may be more patient about waiting for the second type to arrive, more willing to endure the first stage before the other arrives. If either type will attract the other, it may be useful to target the early recruitment to the more patient type.

Increase Stage 1 Value of the Community

Many online sites with successful social content started by providing services of value to their customers in the absence of critical mass for an online community. Del.icio.us is an online bookmarking site. It offers the personal productivity tool of keeping and organizing one's personal tags for web content on a server on the Internet. That allows people who use multiple computers to access their tags regardless of which computer they are using. This personal productivity aid provides sufficient value to some users that they would use the service even in the absence of any other users, and thus helped attract users even before the service reached a critical mass of social usage.

Instead of offering a service that is individually valuable to one person, it is sometimes possible to offer a group service that is valuable to a small enough group that the group can collectively

decide to join. Yahoo! Groups, for example, is a community-hosting service that allows anyone to create a group (complete with e-mail list, discussion board, photo storage, calendars, and other tools). To many of the users, the value lies entirely within the group they (or their friends) create. Presumably, however, some people who were first invited by an acquaintance to join a particular group subsequently became involved in the larger constellation of Yahoo! Groups, helping to build the critical mass for the overall service.

Shopping opportunities may provide another non-social source of benefits. Shutterfly.com is an online photo printing/storage/sharing site. Unlike some photo-sharing sites (such as Flickr), Shutterfly started primarily as a site where digital photos could be uploaded and printed, both as traditional prints and as enhanced products such as calendars, mugs, prints with fancy borders, etc. People uploaded their pictures to Shutterfly for the non-social purpose of ordering prints. Social interaction, however, was a natural side effect. From the start, Shutterfly allowed photographers to share their photos with friends and family (and to thereby allow those friends and family to order prints). Today, the site has grown to build its theme around the tag line “tell your story” and to have a gallery of community-created photo books, complete with personal profiles of the photographers and scrapbookers. An initially single-user commerce activity of photo printing attracted users to what eventually became a social site for many.

Amazon.com is another example of site that attracted initial users for shopping. It has a large collection of successful social content options within its site. It uses collaborative recommender technologies (both contextual “customers who bought/shopped for this also bought/shopped for ...” and overall recommenders for products of interest to an individual). It has an extensive set of product reviews; many people have posted gift/wish lists that others can browse (and shop from); and many users publish individual profiles. Each of these “social applications” would have had trouble launching without critical mass, however, and indeed most did not exist when Amazon was launched. Instead, Amazon amassed a critical mass of people and data by offering a distinctively non-social application—book-buying.

Design Claim 25: single-user and small-group productivity, entertainment, or commerce tools can attract people to an online space before the community features are successful.

The New York Times website (www.nytimes.com) is a content site that hosts both the newspaper’s print articles and a substantial number of online only columns and blogs. While it did not start out as a social site, today it has clearly become one. A quick review of this morning’s articles finds that many of them have reader discussion postings from 10 or more different readers, and many of those discussions are truly interactive (with messages, responses, and responses to the responses). Launching a successful news discussion site from scratch has a serious critical mass problem: many have tried but only a few (notably Slashdot, Digg, and Reddit) have been successful. Instead, by using their content to attract people for non-social purposes, reading news articles and columns, the New York Times easily created a community of readers who then generated comments and discussion.

Design Claim 26: providing access to professionally generated content can help attract people to an online space before the community features are successful.

Another way to compensate for an initially small community is to import user-generated content from elsewhere. In the section on competing for a niche, we pointed out that importing content

can create some of the same positive externalities that other members would provide, since some of the benefit their presence would provide is through the content they would provide. This may be especially useful in the startup stage, when there are few members generating content locally. Consider, for example, MovieLens.org, a movie recommendation site. Users rate movies and the system uses the ratings of other people to predict how well individuals will like particular movies. When MovieLens started, it imported a database of ratings that had been gathered by another movie rating site, EachMovie, which was no longer operational. With the EachMovie data available, even the first MovieLens subscriber received useful predictions (after entering a few ratings so the site could calibrate the user's tastes). Subsequent MovieLens subscribers benefited from both the pre-existing EachMovie data and from the ratings that other MovieLens subscribers had entered. The imported data compensated for a small membership. Once MovieLens was well established, the marginal value of the imported was reduced.

A number of product review sites appear to follow the syndicated data strategy. For example epinions.com generated a large number of product reviews as a stand-alone site. It then sold its content to other sites that wanted to include product reviews but did not have enough users to generate their own. There are a large number of travel review sites and on occasion they seem to share content. For example, we found the same review for a hotel in Florence on several different sites.⁴

As described in the section on competing for a niche, one danger with syndicated data as a strategy for jump-starting an online community is that users may discover the original source and choose to join the already-active community instead. Syndicated data provide value to individuals, but don't create the stickiness or competitive advantage that having a large community would create. In combination with other useful content or services, however, or with a novel presentation, syndicated data can increase the value of a community before it becomes self-sustaining.

Design Claim 27: providing access to syndicated data can help attract people to an online space before the community features are successful, if the syndicated data is not otherwise easily accessible or if it is presented in a novel way that adds value.

Another way to compensate for an initially small community is for paid staff to participate, providing the benefits that eventually would be provided from members. For a community where the externalities come from a large number of participants, such as a movie recommending site, this would be prohibitively expensive. For smaller niches, however, such as technical support for products, this may be a viable option.

For example, when launching discussion boards for a group of patients enrolled in a research study on how to motivate increased exercise, in the startup phase several staff members posted questions and responded to all the posts by study participants (Resnick, Janney, Buis, & Caroline R Richardson, in press) This made it more attractive to post and to check the boards for new messages, and thus helped build participation by patients.

As another example, in January 2007 Microsoft embarked on an effort to grow the peer technical support that occurred in the online forums of MSDN, the Microsoft Developer Network. Managers hypothesized that a low answer rate for posted questions (~35%) was discouraging

⁴ For example, the comment, "Was disappointed in Villa de Vedetta....thought was quite stuffy" was found at one point on en.venere.com, bookingclick.com, www.alibabuy.com, and www.tvtrip.es.

people from posting questions, so they hired staff to answer some of the questions. The answer rate went up to about 80%. Not surprisingly, the volume of posts tripled in a year.⁵

Design Claim 28: participation by professional staff can help attract people to an online space before the community features are successful.

The need for staff involvement in the early stages suggests that it may be helpful to start a community a scope narrower than its eventual aspirations. For example, a book site that is hoping to eventually have reviews for all books might start with a subset (current bestsellers or book club selections) as a way of building more content in a narrower area (rather than less across a broad area). In addition to the need for staff attention to seed the contents, staff attention may be needed for marketing.

Communities based on geographic locality, with little synergy between activity in different locations, offer the strongest case for starting with a small scope and expanding gradually. A wide range of sites from Angie's List to CitySearch to Craig's List launched over time in different cities. A smaller site—LocalFiles.com, a directory for the Indian expatriate community in the U.S.—followed the same approach, launching a site in Minneapolis, MN before expanding to other cities.

Other communities may have more synergy between activities in different parts of the eventual community scope. A restricted scope in the startup stage then involves a tradeoff between the focusing of staff resources for marketing and content creation and the lost synergies. For example, in a community for baseball fans, many people be interested in discussing a particular favorite team, and also discussing the state of the sport as a whole. If the community starts with just a single team's fans, it may not have enough people to generate good discussion of baseball as a whole, and thus it may also lose some people who would have joined if there were lively discussion of baseball as a whole in addition to their favorite team. On the other hand, trying to simultaneously seed discussion spaces for all of the teams might overtax the available staff resources.

Design Claim 29: starting with a limited scope and expanding later allows focusing of staff resources toward getting to critical mass in the limited scope.

One problem with paying staff to participate is that it may be difficult to get volunteers to assume the roles initially filled by staff. People are less motivated to contribute when they think that someone else will carry the load if they don't, as explored in detail in the chapter on motivating contributions. Worse yet, people may also be less motivated to contribute on a voluntary basis when they know that someone else is getting paid for similar work.⁶ Perhaps most critically, staff may notice tasks and complete them even before volunteers notice them, and thus volunteers may never develop the habit of taking on tasks they could do. Thus, there is a real danger that the roles for professional staff will become permanent, with a need for continued payments.

⁵ Sean O'Driscoll, personal communication, April 2008.

⁶ This is not always a show stopper, as there is a counter-vailing effect that people are more likely to participate in more successful communities. Patrick Wagstrom found that getting more "community-focused" professional developers involved in open source projects within the GNOME community actually drew in more unpaid labor, rather than alienating the volunteer participants. More "product-focused" professional developers had not net impact on volunteer participation in a project. (Wagstrom, 2009. "Vertical Interaction in Open Software Engineering Communities". PhD dissertation, CMU-ISR-09-103, March 2009. Chapter 4.)

One possible solution to a community becoming dependent on staff contributors is for them to wait for volunteers to perform tasks and take on only those tasks that volunteers do not. This not only gives volunteers enough time to find tasks, but also degrades the quality of service provided by the professionals (because of the delay) and thus encourages members to take on the tasks so that they will be completed sooner. For example, when Microsoft hired staff to answer questions on MSDN forums, as a way to encourage more posting of questions, the staff only answered questions that had gone unanswered for 24 hours. That was the approximate time for items to scroll off the front page, at which point they were unlikely to be noticed by volunteer members. Over time, the percentage of questions answered by members increased and the need for staff went down.

Design Claim 30: If professionals act as contributors of last resort, they will be needed less and less as the community achieves critical mass.

In some cases, software bots can offer a partial substitute for the value that would be provided by other participants. For example, jsettlers.com is a website for playing an online version of the popular German board game "Settlers of Catan". Three or four players are needed for a game session. Visitors to the site who tire of waiting for other opponents to arrive at the site and join their games can invite software bots to play instead. For most people, the software bots aren't quite as fun to play against, but sometimes they're better than waiting for real opponents.

Similarly, the online ESP Game (<http://www.gwap.com/gwap/gamesPreview/espgame/>) pairs two people for a coordination game in which each tries to assign the same tag the partner does, without communicating with each other. Part of the fun is the sense of connection one experiences with an unknown partner, especially when a partnership selects matching tags that are unusual or quirky. The game has become popular and often one can be matched with a partner with little waiting. When there are few participants, however, the game pairs a live player with a replay of a previously recorded game session. Knowing that one can always find a partner makes it more attractive to visit the site, even if the partner is sometimes not live.

Design Claim 31: bots that simulate other participants can help attract people to an online space before the community features are successful.

Early Adopter Benefits

We now turn to early adopter benefits as a way of convincing potential members to join early. Early adopter benefits are common outside of online communities. For example, early employees in a company receive larger numbers of stock shares than later employees receive. As another example, early volunteers in a presidential political campaign, even those with little prior experience, can expect white house jobs should their candidate win. Joining the campaign after Iowa, when it's easier to pick the winner and there are more people involved, leads to less personal benefit should one's candidate win.

We describe three approaches to providing early adopter benefits in online communities. One focuses on continuing monetary advantages for the early adopters, either in terms of payments or discounts. The other two focus on benefits within the community, either the benefits of having skills before others do or the status and privileges that stick with early adopters even after the community becomes sustainable.

First, consider promises of future discounts. One example comes from the anti-spam service SpamNet by Cloudmark. An important ingredient in the service was reports from customers of particular emails that were spam. These reports enabled Cloudmark to flag and filter similar messages sent to other customers. Clearly, this service exhibits significant network externalities, since it will only be useful to join if there are many other members reporting on spam. In its early stages, Cloudmark offered the service for free and, importantly, promised when it went commercial to guarantee a special lifetime subscription rate to its early users. While intended primarily as a retention strategy, this type of offer can be an additional incentive for early users, by increasing the expected future value of the service should it achieve critical mass (i.e., many people will find it worth paying full-price for, but I won't have to).

Design Claim 32: Promising permanent discounts to early adopters can attract early adopters to the community.

There are many less tangible benefits that early joiners of an online community can expect to have. For example, early members are likely to be better known and have higher status once the community gets large. This may occur through informal processes of preferential attachment. For example, in a blogging network such as LiveJournal, new members will be more likely to link to blogs of people who are already prominent, thus making them even more prominent for people who come even later.

Joining early can lead to higher status or positions of leadership after a community succeeds. The reputation and status benefits of being an early adopter may emerge naturally in an online community, or they may be more explicit. A conversation-based community may need to do little more than make a poster's identity visible for such benefits to emerge, since people will recognize the identifier and treat the long-time members with extra respect. When a group of old-timers reminisced about Usenet news, it occurred to us that we all remembered a large number of the early active Usenet posters, moderators, and administrators--even though nearly 25 years had passed since encountering them, and even though many of them were no longer active. In communities with clear leadership roles having differential privileges, such as Wikipedia and open source projects, many people understand intuitively that there are more opportunities to assume the high status roles if one gets in "on the ground floor". For example, the process for gaining editor privileges at Wikipedia has become more codified and presumably more stringent over time (Burke & Kraut, 2008). Empirically, longevity was found to matter in promotion decision: each additional month as editor led to a 2.9% increase in the probability of approval for admin status.

Early joining benefits may also occur through explicit status markers not tied to explicit leadership roles. For example, just as American Express prints "Member since" dates on credit cards (and advertises this fact), eBay shows a "Member since" date on its user feedback profiles. Many online forums display the date a user first entered the community alongside each post. Some communities, like Slashdot, assign user IDs sequentially and people with very low IDs have status in the community, much as people with low "badge numbers" have high status in large high-tech companies.

One of the challenges of participating in communities is the time and investment needed to "get up to speed," building skills and a social network. This challenge can be reframed as a benefit of early adoption. One of the benefits of early adoption germane to online communities is the ability to build skills ahead of other members. In online gaming communities (e.g., World of

Warcraft) there is a tremendous surge of activity as new levels are introduced, in part so that those achieving them can do so before their friends. In a virtual world, an early adopter may show off her skills at moving around, her knowledge of the world, and even the special features of her avatar that early adoption provided her before her peers joined.

Turning these potential benefits of early adoption into early adoption decisions requires careful communication, so that potential members are aware of the advantages of learning the ropes early. Potential members need both to be informed about benefits and that they are available only to early adopters. One common technique is either to fix an expiration time (“this week only”) or the number of people who can claim the benefit (“the first 100 members get...”).

Design Claim 33: Promoting the status or readiness benefits of being early to an online community can attract early adopters to the community.

It may be possible to market an online community as "undiscovered", with the implication that those who adopt early will be recognized by their peers as trend-setters. In social networking communities, early adopters may see a benefit in being the one to invite all their friends (rather than being the last one to the party). Exclusivity, or the appearance thereof, is one way to promote a site as cool. Google did this particularly well with Gmail which had a limited beta release in which you could only get an account by being invited by someone who already had one--this approach linked social word of mouth with scarcity. The goal is to reach potential early adopters, and to reframe the message from one where "new" might imply uncertain and risky to one where "new" implies the chance to get in while it's still cool and undiscovered.

Design Claim 34: Promoting a site as cool but undiscovered can attract early adopters.

The default assumption in most online communities is that supply is infinite, and hence, there is no chance of a resource running out. If the supply can be limited, or differentiated in such a way so as to create a limited amount of "preferred" membership, then individuals can be enticed to join early. The namespace of usernames within a community is a naturally scarce resource. On eBay, memorable usernames like “coinguy” and “informationist” are available on a first-come first-served basis. The one hundredth member had more options than the millionth. Similarly, nearly all of the free e-mail services and many virtual world communities allow you to select your name, so long as it is not already in use. Interestingly, many online community sites have specifically chosen different paths (e.g., many social networking sites use your e-mail address, which is already unique, and Second Life limits character names to those chosen from a template unlikely to be particularly meaningful in the real world). Nonetheless, we've observed a rush of users to new e-mail sites when announced (to claim their names). Making this opportunity salient in marketing is a way to encourage earlier membership. Namespace isn't the only resource that can be made scarce. In online games, limited quantities of special objects may be made available to early users.

Design Claim 35: Creating scarce, claimable resources can induce prospective members to join earlier.

One challenge when creating scarcity is that members may "squat" on their names or other resources without contributing to the community. That is, they may join but not participate. A common technique for preventing squatting is to require a certain level of sustaining activity to maintain the claimed resource. In the early days of the Internet, people who claimed domain names understood that they had to pay a regular fee to keep them, but they were never required

to use them. Many of today's e-mail addresses and site addresses expire if they are not used at least every 3-6 months. Indeed, greater demands can likely be made since the endowment effect (Kahneman and Thaler, 1990) suggests that people will do more to keep such "property" than they would have done to acquire it.

Design Claim 36: Contribution minima for maintaining scarce status can lead to greater contribution by status-holding members.

An extra beneficial side effect for the community designer is that members who join early in order to get in on the ground floor will be especially motivated to help make the community thrive. There is nothing valuable about knowing one's way around before others do, if it gives no competitive advantage because no one else joins. There is no value in a permanent discount if the community doesn't survive. There's nothing cool about being first to a community that failed. And it is only meaningful to have high status if there are other people who join the community who have lower status.

The investments that people make through early adoption are sunk costs that in principle should not affect a rational decision-maker's future choices. People, however, often do take into account sunk costs, in part out of a desire to make the earlier decisions seem like good ones. Thus, people who join because of early adopter benefits will therefore be more likely to help the community succeed through recruitment, content creation, or other actions. The effects of sunk costs on commitment more generally is explored in Chapter TKTK.

Setting Expectations for Success

Whether a community is likely to succeed is not always obvious to potential members. Expectations of success make it more attractive to join early, which in fact fuels success, so that the expectations become self-fulfilling. A number of design choices can affect expectations of success. Some convey signals of the community convener's skill and commitment. Some convey signals that potential members are reacting favorably or are committed. Some convey external expectations. We consider each in turn.

Signals of Convener Quality and Commitment

In the early days of the Internet, users had low expectations for usability, design, and interaction. With few choices, even venues with poor appearances might hold promise as the source for valuable content. As professionalism increased, ordinary web users learned a set of heuristics to distinguish legitimate high-quality sites from low-quality or fraudulent ones. Fogg (DUX2003) studied how users determine the credibility of web sites, finding that the two most-often cited influences on credibility were: (a) the design look of the site (e.g., professional look, pleasing graphics); and (b) the information design/structure of the site (e.g., organization). Other oft-cited influences include: (c) tone of the writing (e.g., slang, poor language); (d) functionality of the site (e.g., error messages, poor search); (e) clarity of information; and (f) readability of text.

Today the importance of credibility and professionalism is even greater. Internet users are regularly inundated with unprofessional-looking spam and phishing attacks (unwanted messages

that attempt to elicit personal information). Indeed, the prevalence of poor writing, spelling, and layout in such scam attacks is one of the reasons people generally don't fall for them (a test of various phishing attacks found that nearly perfect ones caught a high percentage of even skilled users (Egelman, Cranor, & Hong, 2008).

Design Claim 37: Professional site design increases expectations about the probability of success.

In addition signaling a convener's capabilities, through professional site design, it may also be useful to demonstrate the convener's ongoing commitment. If potential members believe that the individuals or organization who started the community are committed to additional investment, even if the community does not immediately take off, these expected future investments will factor into a judgment about whether the community will take off—there will be less risk that one's own early efforts will be stranded. It is not, however, simple for the conveners of a community to credibly convey such a commitment. After all, they might be expected to “cut their losses” if there is no immediate success.

The theory of credible signaling, as elaborated in the research literature in information economics [cite], begins with one key insight: in order distinguish between “high quality” and “low quality” actors, it must be easier or less costly for the high quality actors to provide the signal. In our situation, a “high quality” actor is one that is committed to future resource expenditures on a community, even if it does immediately achieve critical mass, because of an assessment about the likelihood of eventual success and the payoffs that will result from it. A “low quality” actor is one that needs to see more immediate indicators of success in order for it to continue making investments. It is more costly for the low quality community convener to make an immediate expenditure, because it has a lower assessment about the probability of eventual success (in part because it knows about its own lack of future commitment). Thus, we should expect the high quality community convener to be more willing to undertake early expenditures. And when we observe such resource expenditures, it is a credible signal of greater commitment to future expenditures.

Sermo.com started an online community for doctors. It is relatively easy to create such a community using generic technologies such as a Yahoo! group or an email list or a forum site using PhpBB. Indeed, there were existing competitors. One way that sermo.com signaled a commitment to continued investment in its community was through expenditure on a custom software platform that it developed. It also had paid staff whose presence was very visible in the community early on, and did one-on-one demos for physicians at medical conferences. In the early days of sermo, when it did not have more members than some of its non-commercial competitors, the visible expenditures and claims of its venture capital backing contributed to a sense that sermo would eventually be a bigger, better community than its competitors, even if it wasn't yet.

When a number of new features are available, we speculate that deliberately holding some new features back so that they can be released on a more even time interval might be a better strategy than releasing all of them at once. There are three reasons for this speculation. First, user utility for new features might exhibit decreasing returns, i.e., their marginal utility for the first new feature of the day might be higher than subsequent new features. Therefore, sequential releases might increase overall user utility. Furthermore, from a usability perspective, the sequential release approach also helps the site isolate and evaluate the effects of the new feature. Finally,

from a signaling perspective, frequent releases of new features convey a trajectory of growth and continued commitment.

This approach is especially applicable to the initial public announcement of a new community. If some features that are implemented are deliberately held back, it may make it easier for people to get started. The held-back features can then be released a week or two later, conveying an image of continuing investment from the site's founders.

Design Claim 38: visible expenditures can be a credible signal of commitment to future investment in a community, and thus help to increase expectations about the probability that the community will eventually succeed.

Signals of Positive Member Response

When the quality of a public good is uncertain, announcing early contributions provides a credible signal that the public good is valuable, and thus helps to attract later contributors [Vesterlund 2002]. This is why major fundraising campaigns begin publicly only after “quiet periods” in which major donors make commitments. Similarly, indicators that other potential members are responding favorably to a new community are a form of social proof, enhancing expectations that the community will succeed. This can be accomplished by making membership and activity visible, and by showing a trajectory of growth.

One way to convey membership is to prominently post photos of members. Simply adding photos of random people to a movie ratings site did not have a noticeable effect on user behavior (unpublished experiment with MovieLens, Tassone & Kiesler, 2004). We speculate, however, that posting photos of members alongside their user-contributed content, or posting photos of the most recent new members, may convey a signal that others like the community and are joining it.

Design Claim 39: Images of members will convey the presence of other people, and thus expectations of future success.

One way to convey activity in the community is through examples. Recent user contributed content can be made prominent, perhaps even on the front entry page. Of course, making user-contributed content visible on the entry page risks conveying an image of amateurism, which is a negative signal of quality. One solution is to include user-contributed content but confine it to a small sidebar so that new visitors will first notice the professional-looking presentation of the community. Another solution is to allow only a limited number of members, those who can be trusted to produce professional-looking content, to post messages that go on the front page. For example, the drupal.org website, the online community for developers and users of the Drupal content management system software, includes a top area with static information, a wide blog-style left column with official announcements that only select people are allowed to add, and a narrow block on the right with the titles of the most recent topics added to the forum.

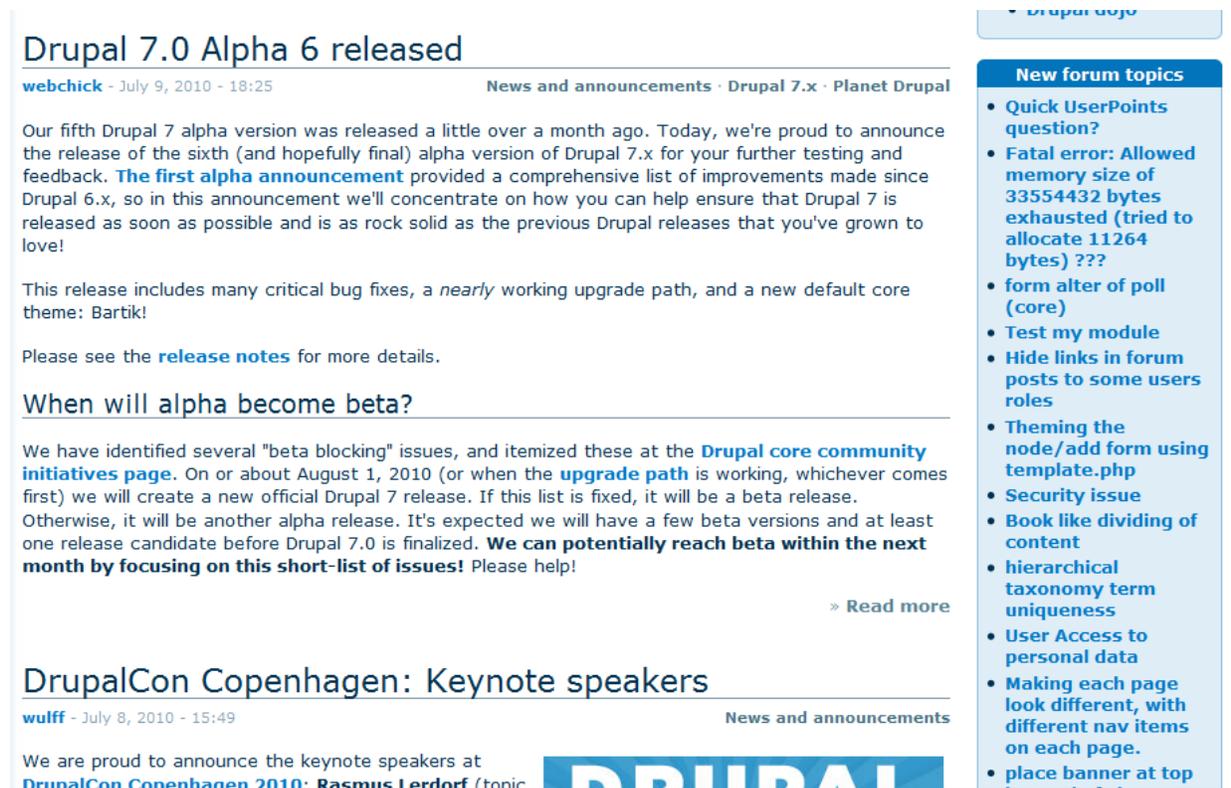


Figure 5. Making recent contributions salient in Drupal.org

Design Claim 40: Prominent display of user-contributed content will convey activity, and thus expectations of future success, as long as there is new user-contributed content.

Another way to convey activity without actually showing the contents, which may be amateurish, is to show indicators. For example, rather than showing the actual comments on blog entries, the number of comments can be shown and users can click through to see the actual comments. Or, the number of new forum topics can be shown, without showing the actual posts. As another example, many sites show an indicator of how many people are currently logged in to the site. Even if there is no synchronous interaction, many others currently using the community is an indicator that they like it. Of course, such an indicator is more effective at showing activity if, at most times when visitors would arrive, many other people would be logged in. The indicator below, conveying that there is only one current user of a site, would discourage rather than encourage visitors. In sites that are intermittently active, it may be helpful at inactive times to show indicators about times when a lot of people were active.



Figure 6. Activity indicators can highlight lack of activity

Design Claim 41: Indicators of participation levels will convey activity, and thus expectations of future success, as long as there actually is activity.

Conveying a trajectory of growth in membership and activity is especially helpful in raising expectations. There are several ways to display indicators of growth. One possibility is to publicly acknowledge each new member or each new content entry. Another possibility is to show a running tally of the current membership size or amount of content. A third possibility is to show the percentage growth. Depending on the size and growth rate of the community, one or another of these options may paint a more favorable picture of the community's current quality and long-term prospects.

For example, in the earliest days of Wikipedia, the most effective signal of growth on the front page would probably have been a list of new pages that had reached an acceptable quality level and a list of new first-time contributors. When it started to take off, the most effective signal would have been to show the percentage change in content and contributors from month to month. Now that it is wildly successful, the site creates a signal that it has already succeeded by showing the absolute number of articles in each language (2.6 million in English as of Nov. 25, 2008).

Design Claim 42: Indicators of membership and content growth signal a higher probability that the community will eventually reach critical mass, provided there really is growth.

Design Claim 43: When a community is small and slow growing, acknowledging each new member or contribution creates a more favorable signal of growth than showing total numbers or percentage change.

Design Claim 44: When a community is small and fast growing, displaying percentage growth creates a more favorable signal of growth than displaying absolute numbers.

Design Claim 45: When a community has reached critical mass, displaying absolute numbers conveys a signal that the community is already successful.

Designers can also allow for conditional commitments of membership and activity, which allow the community to convey an expectation of success. With this approach, potential members can commit to joining (or to taking certain actions), but only if enough other people also commit to membership or actions that will cause the community to succeed. Conditional commitments can reduce the risk of an early membership decision being a wasted cost when the community fails to catch on. RSVPs for meeting or parties are a good example of commitment mechanisms: if not enough people indicate a willingness to attend, an event can be cancelled without anyone suffering the cost of attending an empty event.

Meetup.com makes extensive use of conditional commitments. In the first couple of years of meetup's operation, there were no long-standing groups. Instead, each hobby or other interest group had a standard day of the month to meet at a local cafe, restaurant, or other venue. If you RSVP'ed for the next meeting of a group, you were notified if not enough others have RSVP'ed to hold the event. More recently, meetup has also reified the notion of a group, and provided a conditional quasi-commitment mechanism for memberships. If you are interested in a topic, you can register that interest with the site. Meetup.com then notifies you when other people also register an interest in the topic. When enough people have indicated an interest, and someone has sufficient energy to lead the formation of a group, meetup provides a mechanism to contact all the people who have registered an interest in the group. Since there is relatively little cost to

signaling an interest in a topic, people are willing to signal interest in topics that do not yet have a critical mass of other interested participants.

Another example in a slightly different setting is Groupon.com. Each day there is an offer for one item, typically a gift certificate to a restaurant or entertainment venue. People commit to buy an offer, knowing that the offer is only on if enough people buy it to meet the maker's threshold (e.g., we'll give out \$40 off dinner coupons, but only if 250 people buy them for \$20 each). In principle, the conditional commitment here has two sides. On the seller side, it ensures enough participants to make it worthwhile to produce the items. On the buyer side, it is a safety measure – a buyer may not want to commit to an unknown restaurant, even at a good price, unless hundreds of other people (who may have heard of it) also do.⁷

How It Works

1. Each day we feature something cool to do at an unbeatable price.
2. You only get it if enough people join that day... so invite your friends!
3. Check back the next day for another awesome Groupon!

Figure 7. Commitment contingent on others' participation at Groupon.com

The idea of conditional commitments can be applied to forming a new purely online community, or to starting up a new forum within an existing online community. Rather than simply launching, which may lead to a problem of a sequence of visitors each seeing an empty forum and giving up, people can pre-register an interest in a topic, or even tag their content as being applicable to the potential new forum. When enough people have registered an interest, or enough relevant content has been accumulated, the new forum can be launched and people who registered an interest can be notified.

Design Claim 46: Conditional participation commitments can draw people to join communities that they would not join if they had to do so without knowledge that others were also joining.

External Signals

External signals can also fuel expectations of success. One way is to draw analogies to other communities that are successful. We previously described the value of analogies in conveying the intended scope of a community. But describing something as the “Wikipedia of Numbers” or the “Wikipedia of music” not only conveys an intended scope but subtly suggests that, since Wikipedia was successful at attracting many editors and readers, so will these other sites. Mommsource.com, an online community for mothers, rebranded itself as “Mamapedia: the wisdom of moms”.

Design Claim 47: Drawing analogies to successful communities can raise expectations that a new community will be similarly successful.

⁷ In practice, the items tend to be services where there is no minimum production level required to keep costs down, and the threshold always seems to be met early in the morning, so the signaling value of others buying may be limited.

External publicity, in mainstream media and the blogosphere, can also fuel expectations of success. Online communities that are fortunate enough to get such external publicity can post links to it on their sites, for members and potential members who might have missed it.

The screenshot shows the Mamapedia website interface. At the top, there are navigation links for 'Q&A', 'Timeline', and 'Sweet Deals', along with a search bar and a 'Login with Facebook' button. Below the navigation is a horizontal menu with icons for Home, Lists, Questions & Answers, Local Businesses, Member Perks, Voices, and Timeline. The main content area is titled 'Mamapedia™ In The News' and features a list of news items with dates and headlines. A sidebar on the right contains an advertisement for 'The Greener Cleaner Carpet Cleaning & More'.

Mamapedia™ In The News

- August 24, 2009 **KENS-TV/San Antonio - Great Day SA**
Mamapedia features on Great Day SA
- August 19, 2009 **Orlando Sentinel - Back-to-school plans**
A survey of Moms by mamapedia.com asked parents about their back-to-school plans and found that most plan to drive their kids to school, pack their kids' lunches (but use cafeteria, too) and spend from \$25 to \$50 on school supplies.
- August 19, 2009 **ABC 15/KNXV-TV/Phoenix - Valley moms share notes on back-to-school concerns**
Mamapedia is a new website that lets moms get wisdom of other moms on key parenting issues. With over a million topics and answers posted, it really reveals the pulse of the mom community.
- August 18, 2009 **Examiner.com - Mamapedia mentioned in Home & Living 101: Efficient grocery shopping**
The moms at Mamapedia.com have been comparing numbers and groceries and family needs.
- August 12, 2009 **EducationNews.org - Mamapedia Offers Mom Wisdom On The Spot**
Mamapedia gives moms everywhere the power to ask any question and get 'the real scoop' from millions of other mothers. From common parenting problems such as potty training and sibling rivalry to personal mom issues like weight loss and breastfeeding, Mamapedia offers wisdom and advice without judgment.
- August 10, 2009 **KSAZ-TV/Fox 10 Phoenix - myFOXPhoenix.com interviews three Mamapedia moms.**
There's a new resource for moms. It's a place where you can talk to other moms about the best pediatrician in town, or the best way to lose the mom weight and more.

The Greener Cleaner Carpet Cleaning & More
Troy, MI
owner: Tina E
Non-toxic, environmentally friendly carpet & upholstery cleaning.
Mamapedia members~ receive 20% off of your total carpet, tile and/or upholstery cleaning order!
Redeem/Print
See my Mamapedia listing
Member Perk™

Figure 8. Highlighting external publicity in Mamapedia.com

Design Claim 48: drawing attention to external publicity and endorsements can raise expectations about future success.

Summary of Design Alternatives

This chapter has explored two critical challenges for communities that are just getting started. The first is to carve out a useful and defensible niche in the ecology of competing communities. The second is to get to a self-sustaining, critical mass of users. We conclude with a summary of the design alternatives considered throughout the chapter.

The first category of design alternatives are ones that structure the set of interaction opportunities: selecting, sorting, and highlighting them, grouping them into interaction spaces, and notifying people about them. The various alternatives are most useful in carving out a useful niche, where people will gain benefits from participating. Techniques such as subdividing spaces after they become active also help in conveying a sense of activity, which can affect expectations about the community's likely future success, and thus help the community get through its initial stage prior to achieving critical mass.

A related category of design alternatives involve the structure of the community: its size and the breadth of topics covered. Larger communities have more activity, but that may not be so desirable in niches where the value of participating comes from interpersonal bonds with a few people. Starting with a limited scope and expanding later can to focus staff and marketing attention on getting to a critical mass of participation in a more limited scope.

A number of design alternatives involve the content, tasks, and activities in the community. Topics and activities that bridge separate interests can make mixed-topic spaces work well, spaces that can be subdivided into single-topic spaces later. Offering valuable activities or content that is not generated by other members can make the community attractive in its early stages, before achieving critical mass.

Because new communities often share people and content with other communities, the chapter also analyzed several design alternatives involve external communication and integration. Decisions about the sharing of user ids and profiles and the import and export of content are critical elements of a competitive strategy. In addition, making people's membership and participation in a community visible to their friends outside the community, and enabling members to invite friends are valuable ways to leverage a community's early adopters and get to critical mass.

There are several ways to create rewards that will help a community reach critical mass. Promises of future discounts or status for joining early, or the presence of scarce, claimable resources can make it more attractive for people to join early. And those early members can be given incentives for creating content that will attract other members, or directly for attracting other members.

Finally, we find that there is remarkable power in framing things the right way. Articulating a clear rather than an ambiguous topical scope for interaction spaces will make them more useful for most participants. Conveying a unique selling proposition makes it easy for people to see why they should join. Advertising and celebrity endorsements can make one community a focal point when there is fierce competition among communities. Presenting the community as cool but not yet discovered can make early joining attractive. Drawing analogies to successful communities, highlighting external publicity, and a variety of ways of highlighting good things happening in the community can all convey an expectation that the community is well on its way to success, even if it has not yet achieved critical mass.

There is a lot to think about, and dream about, in the startup stage of an online community. Don't just build it and hope for the best. Build it, provide content and activities that people want, structure them in a way that provides maximal benefit, give people some reason to be early members, leverage their early participation to recruit others, and frame things in a way that conveys expectations of success. Then they'll come.

Type	Design Alternative	Claim #
Selection, sorting, highlighting		
	push notification	Design Claim 1

	Mixed-topic scope for an interaction space	Design Claim 2
	Personalized collections of “most related content”	Design Claim 6
	Subdividing spaces after they become active	Design Claim 7
	navigation aids that highlight more active spaces	Design Claim 8
	a schedule of “expected active times”	Design Claim 9
	recommender systems that help people navigate to spaces that best suit them	Design Claim 10
Community structure		
	Larger community	Design Claim 12
	starting with a limited scope and expanding later	Design Claim 29
Content, tasks, and activities		
	Activities that bridge interests in different topics	Design Claim 4
	Differentiated user interface elements	Design Claim 13
	User-generated primary content vs. user-generated metadata	Design Claim 19
	productivity, entertainment, or commerce tools	Design Claim 25
	access to professionally generated content	Design Claim 26
	access to syndicated data	Design Claim 27
	participation by professional staff	Design Claim 28 Design Claim 30
	bots that simulate other participants	Design Claim 31
External communication		
	Non-shared user IDs and profiles between competing communities	Design Claim 14
	Content sharing between competing communities	Design Claim 15
	Displays of membership that are visible to non-members	Design Claim 20
	Members' actions in the community visible to their acquaintances outside the community	Design Claim 21
	Allow members to forward content from the community to their acquaintances outside the community	Design Claim 22
	Allow members to invite acquaintances outside the	Design Claim 23

	community to join	
Feedback and Rewards		
	Incentives for early members to generate content	Design Claim 18
	pay-for-referral and revenue-sharing from referrals	Design Claim 24
	Promising permanent discounts to early adopters	Design Claim 32
	Promoting the status or readiness benefits of being early	Design Claim 33
	Creating scarce, claimable resources	Design Claim 35 Design Claim 36
Presentation and Framing		
	Ambiguous scope for an interaction space	Design Claim 3
	Transcendent or bridging topical identity	Design Claim 5
	Ambiguity of scope for the community	Design Claim 11
	Conveying a succinct unique selling proposition	Design Claim 16
	Advertising and celebrity endorsements	Design Claim 17
	Promoting a site as cool but undiscovered	Design Claim 34
	Professional site design	Design Claim 37
	Visible expenditures	Design Claim 38
	Images of members	Design Claim 39
	Prominent display of user-contributed content	Design Claim 40
	Indicators of participation levels	Design Claim 41
	Indicators of membership and content growth	Design Claim 42 Design Claim 43 Design Claim 44 Design Claim 45
	Conditional participation commitments	Design Claim 46
	Drawing analogies to successful communities	Design Claim 47
	Drawing attention to external publicity and endorsements	Design Claim 48

REFERENCES

- Anderson, L. (2003 August 13). Attack of the smartasses. *SF Weekly News*,
- Backstrom, L., Huttenlocher, D., Kleinberg, J., & Lan, X. (2006). Group formation in large social networks: membership, growth, and evolution. *Proceedings of the 12th ACM SIGKDD international conference on knowledge discovery and data mining*, 44-54.
- boyd, d. (2006). Friendster lost steam. Is MySpace just a fad? *Apophenia Blog*. Retrieved from <http://www.danah.org/papers/FriendsterMySpaceEssay.html>
- boyd, d. m. (2008). *Taken Out of Context: American Teen Sociality in Networked Publics*. . University of California, Berkeley, Berkeley, CA.
- Burke, M., & Kraut, R. E. (2008). Mopping up: Modeling Wikipedia promotion processes *CSCW 2008: Proceedings of the ACM Conference on Computer-Supported Cooperative Work*. New York: ACM Press.
- Egelman, S., Cranor, L., & Hong, J. (2008). You've been warned: an empirical study of the effectiveness of web browser phishing warnings *CHI 2008: Proceeding of the ACM SIGCHI Conference on Human Factors in Computing Systems* (pp. 1065-1074). New York: ACM.
- Ellison, N. (2007). Social Network Sites: Definition, History, and Scholarship. *Journal of Computer-Mediated Communication*, 13, 210-230.
- Fischer, C. S. (1992). *America Calling: A Social History of the Telephone to 1940*. Berkeley: University of California Press.
- Fox, S., & Jone, S. (2009). *The social life of health information*. Washington DC: Pew Internet & American Life Project.
- Horvitz, E., & Apacible, J. (2003). Learning and reasoning about interruption *Proceedings of the 5th international conference on multimodal interface* (pp. 20 - 27). New York: ACM Press.
- Jacobs, G. (2003). Cyberculture. In H. Dauncey (Ed.), *French Popular Culture: an Introduction* (pp. 77-87). London: Hodder Headline Group.
- Kim, A. J. (2000). *Community building on the web: Secret Strategies for Successful Online Communities*. Berkeley, CA: Peachpit Press.
- Reeves, R. (1961). *Reality in Advertising*. New York: Alfred A. Knopf: Inc.
- Ren, Y., & Kraut, R. E. (under review). A simulation for designing online community: Member motivation, contribution, and discussion moderation. *Information Systems Research*.
- Resnick, P., Janney, A., Buis, L. R., & Caroline R Richardson (in press). Starting an Online Community on Demand: A Case Study of Adding Forums to a Physical Activity Promotion Program *Journal of Medical Internet Research*.
- Richardson, C. R., Buis, L. R., Janney, A. W., Goodrich, D. E., Sen, A., Hess, M. L., et al. (Under review). An online community improves adherence in an internet mediated walking program. Part I: Results of a randomized controlled trial. .
- Rogers, E. (1995). *The diffusion of innovations* (4th ed.). New York: Free Press.
- Schelling, T. (1958). The strategy of conflict. Prospectus for a reorientation of game theory. *Journal of Conflict Resolution*, 2(3), 203.
- Shapiro, C., & Varian, H. (1999). *Information rules: A strategic guide to the network economy*. Boston, MA: Harvard Business School Press.

Stoddard, J., Augustson, E., & Moser, R. (2008). Effect of adding a virtual community (bulletin board) to smokefree. gov: randomized controlled trial. *Journal of Medical Internet Research, 10*(5).

Wenger, E. (1999). *Communities of practice: Learning, meaning, and identity*. New York, NY: Cambridge Univ Press.

Wenger, E., McDermott, R., & Snyder, W. (2002). *Cultivating Communities of Practice*. Boston, MA: Harvard Business School Publishing.