1 Introduction

Paul Resnick and Robert E. Kraut

What does social science tell us about how to make thriving online communities? Quite a lot, it turns out—but only if we listen very closely and, at times, employ a translator. Economics and various branches of psychology offer theories of individual motivation and of human behavior in social situations. The theories generalize from observations of naturally occurring behavior, from controlled experiments, and from abstract mathematical models. Properly interpreted, they can inform choices about how to get a community started, integrate newcomers, encourage commitment, regulate behavior when there are conflicts, motivate contributions, and coordinate those contributions to maximize benefits for the community. This book makes it easier for us to hear what social science has to tell us. It amplifies relevant theories and experimental evidence and then translates them into specific claims about the likely impact of particular design choices for online communities.

1 The Promise of Online Communities

By "online communities" we mean any virtual space where people come together with others to converse, exchange information or other resources, learn, play, or just be with each other. The term applies to many social configurations, from small close-knit groups to sites with millions of participants. Online communities may be supported by a wide variety of technology platforms, from email lists to forums, blogs, wikis, and networking sites. The common feature is ongoing interactions among people over time, with some of the interactions being technology mediated.

Online communities are among the most popular destinations on the Internet. The venerable Usenet had more than 160,000 active newsgroups in 2006 and Yahoo! (http://www.yahoo.com) alone claims to host more than a million online groups. Ravelry (http://www.ravelry.com), a hobby community for people who knit and

crochet, claimed more than 1.2 million members as of March 2011. The product support community for Linksys, a division of Cisco that provides consumer and small-office networking technologies, handles more than 100,000 user sessions per day (Lithium Technologies 2009). More than 35,000 people made five or more edits on Wikipedia (http://en.wikipedia.org), an open source encyclopedia, during the month of February 2011 (Wikimedia 2011b). Facebook (http://www.facebook.com), an online social networking site, celebrated 500 million members in July 2010 (Zuckerberg 2010).

Online communities serve the same range of purposes that offline groups, networks, and communities serve. They provide their members with opportunities for information sharing and learning, for companionship and social support, and for entertainment. Online communities can also produce benefits for nonmembers—either public goods that benefit society as a whole, such as open source software, product reviews, and encyclopedia pages; or private goods, such as suggestions for product improvements or new product designs that benefit the organization that convened the community.

The promise of online communities is that they break the barriers of time, space, and scale that limit offline interactions. People with unusual medical conditions can get social support from others who share their condition but live far away, and they can do so whenever they need it rather than only at a weekly or monthly scheduled meeting. On Ravelry, knitters can share patterns with thousands more people than they could stitch with in person.

2 Critical Design Challenges

Although as a class these online communities are very successful, the success of particular communities varies widely. Some communities struggle to become successful, and others fail. For every Facebook, with its millions of current users, there is a Friendster (http://www.friendster.com) that was once successful but can no longer compete and scores of smaller social networking sites that never got enough members to be viable. Of 2,872 Usenet groups with "support" in their name, some—like alt.support.diet.low-carb, alt.support.depression, and alt.support.diabetes—are successful, with more than 5,000 people posting per year, but half had fewer than thirty posters during 2004 and a quarter had fewer than six (Kraut, 2011a). Smokefree.gov (http://www.smokefree.gov), an online tobacco-cessation program, attempted to add an online community for some of its users but was unable to garner enough activity in the community during the trial period to assess whether such a community, if it

were active, would help members to quit smoking (Stoddard, Augustson, and Moser 2008). Although the English-language version of Wikipedia had more than 3.6 million articles in February 2011, the Korean version had fewer than 160,000 articles (Wikimedia 2011a). Across the more than 9,000 public information-sharing wikis, using the same Mediawiki software that Wikipedia uses, the median number of editors who have ever contributed is only seven (Kittur and Kraut 2011).

To become and remain successful, online communities must meet a number of challenges that are common to many groups and organizations, offline as well as online. The book is organized around these challenges, as described in the rest of this introduction.

Starting a New Community

Many online communities are successful because they have a rich inventory of content that attracts new members. In a conversational community, hosted at a cancer support group, the content might be the messages exchanged by cancer survivors and their caregivers. In an open source software (OSS) development community hosted at SourceForge (http://sourceforge.net), it might be a working base of computer code that provides raw material for developers to improve. On the popular entertainment site YouTube (http://youtube.com), the content consists of video clips that participants post. When creating an online community from scratch, designers and managers are faced with a critical mass problem: the fledgling site doesn't yet have enough content to attract users and there are thus too few users to create the content that might attract others.

Attracting and Socializing New Members

Even established online communities must attract a stream of new members to replace others who leave. For some online communities, a major component of this challenge is to identify and encourage potential members who have the characteristics, skills, and motivation to contribute. Thus, open-source development projects are looking for potential members who can build software. In contrast, many Facebook groups or email-based groups hosted by Yahoo! and Google (http://www.google.com) are more open and willing to accept almost anyone. Even while seeking and welcoming members, many communities also try to screen out inappropriate members. Thus, for example, health-support groups often restrict membership to people who have a particular illness or care for someone who does. Regardless of their selectivity, online groups have special problems because newcomers, who are potentially choosing from among other similar communities to join, frequently have insufficient information to

make their choices and almost always have less commitment to a community than more established members have. These factors mean that their initial observations and interactions are likely to strongly influence whether they stick around long enough to learn whether the site provides a good match to their needs. In addition, because they have not yet learned the appropriate ways to behave in the community, their actions may disrupt the activity of existing members.

Encouraging Commitment

Commitment represents members' feelings of attachment or connection to the group, organization, or community. Commitment underlies members' willingness to stay in the community and contribute to it. Both offline and online, people who are more committed to an organization tend to be more satisfied, are less likely to look for alternatives, are less likely to leave, and tend to perform better and contribute more (Mathieu and Zajac 1990). All organizations must manage the challenge of creating commitment, but because the forces keeping someone in an online group are weaker than those operating in a conventional organization, challenges of commitment are more difficult. For example, in most conventional software companies, employees have an employment contract. If they decide to leave, they lose salary, seniority, and job status. In contrast, most developers in open-source software projects participate voluntarily, with no employment contract encouraging them to stay and contribute. The physical location of a conventional organization also places constraints on members' willingness to go elsewhere. If someone wants to leave a job, church, or club, for example, only a relatively small number of alternatives are nearby and convenient to join. In contrast, if someone wants to leave a particular online community, he or she could join any other comparable community online with no constraints imposed by geographic proximity.

Encouraging Contribution

To be successful, online communities need the people who participate in them to contribute the resources on which the group's existence is built. The types of resource contributions needed differ widely across different types of groups, from the conversations in many online health- and technical-support groups to the code in open-source development projects and the music and video in media-sharing sites. Typically, online communities exhibit a power-law distribution of contribution, with a small minority contributing most of the content. For example, in the Freenet development project, only 30 people of the 369 who participated in the discussion lists ever wrote any code. Although inequality of contribution is not necessarily a problem, under-

contribution is. For example, most Usenet support groups, as mentioned earlier, were relatively inactive. Even apparently highly successful online communities suffer from problems of undercontribution. Roughly two-thirds of the articles in the English version of Wikipedia have been classified during quality-assessment drives as "stubs," articles with only a few sentences of content that are too short to provide encyclopedic coverage of a topic (Wikipedia 2010).

Regulating Behavior

The people who participate in online groups often have different and sometimes competing interests. Most large online discussion groups—especially those that deal with controversial topics—attract trolls, people who post controversial, inflammatory, irrelevant, or off-topic messages to provoke other users into an emotional response (Schwartz 2008). Commercial spammers would like to drive traffic to their external websites. In more mundane conflicts of interest, some participants in a hobby site may prefer that the discussion stay focused on the hobby, but others may want to engage in more personal conversation with other members they have become friends with. When there are conflicting interests in a group, there must be mechanisms to help participants regulate behavior. The challenges here are to deter inappropriate behavior by group members, prevent trolls and other outside attackers, and limit the damage that is caused when inappropriate behavior occurs.

Although these challenges confront almost all groups and organizations, online communities may have more difficulty overcoming them than conventional groups and organizations because of three characteristics that are typical of online communities but unusual in conventional groups and organizations. The first is anonymity. Old-timers may be less able to vet anonymous newcomers, and newcomers may feel less inhibited by social accountability. The second is ease of entry and exit, which can lead to high turnover and thus inhibit interpersonal ties or commitment to the group and affect how sanctions and other deterrence strategies work in regulating behavior. The third is textual communication, which may be prone to misinterpretation because it lacks some of the fluidity and nonverbal cues of face-to-face interaction.

On the other hand, online communities have resources for meeting these challenges that are not available to offline groups. First, the communications—and indeed almost all of the behavior exhibited by participants in an online community—are in digital form and can be archived. Second, online communities can benefit from computation. For example, computers can summarize traces of past behavior as quantifiable and viewable reputations, as eBay (http://www.ebay.com) does with its feedback profiles. Computers can execute search and matching algorithms to introduce people

and content to each other and can notify people when events of interest occur. And computers can enforce access controls so that different people are permitted to see or do different things.

3 Levers of Change: Sociotechnical Systems Design

Students in our classes sometimes challenge the notion that online communities can be designed. A product designer can specify functional and aesthetic features in order to create a desired user experience, but an online community is not so easily controlled. Even if a designer wants an online community to be larger, or more active, or more friendly in tone, he or she may not be able to make that happen. People are the key actors in online communities, and they cannot be shaped or programmed in the way physical materials or software can.

The first central argument of this book is that despite the limited direct control of individual people's actions, online communities can be designed and managed to achieve the goals that their owners, managers, or members desire. Designers are far from powerless. Throughout this book, we identify a wide variety of levers of change, features of online communities that can be deliberately and strategically chosen. Some of these levers involve technical configuration, such as whether a chat feature is enabled or whether special privileges are required to start new conversation threads. Other levers involve social configurations, such as how much externally provided content to include and whether leaders and administrators ignore, cajole, or ban people who disrupt the community.

We classify the levers of change into eight broad categories, described as follows. We will often call these levers of change "design alternatives" or "design options" to highlight the idea that their configurations can result from deliberate choices that managers, designers, or members make.

The first category of design alternatives involves the community structure. The size of the community can make a difference, as can the degree of homogeneity of member interests, whether there is a subgroup structure, and whether membership is recruited through existing social ties.

The second category of design alternatives involves the content, tasks, and activities in the community. There can be opportunities for self-disclosure (e.g., in user profiles). Content can be imported from outside or professionally generated, in addition to that which is generated by members. Welcoming activities and safe spaces for exploration can be offered to newcomers. Tasks can be independent or interdependent and can be embedded in immersive or social experiences.

Table 1.1
Types of design alternatives and the chapters in which their implications are analyzed

Type	Chapter 2: Contribution	Chapter 3: Commitment	Chapter 4: Regulation	Chapter 5: Newcomers	Chapter 6: Startup
Community structure	x	x	x		x
Content, tasks, and activities	x	x		x	x
Selection, sorting, highlighting	x	x	x		x
External communication				x	x
Feedback and rewards	x		x		x
Roles, rules, policies, and procedures			x	x	
Access controls		x	x	x	
Presentation and framing	x	x	x	x	x

Communities often have more content and opportunities than any one person will want to take advantage of. The third category of design alternatives deals with ways to select, sort, and highlight things so that people can find the ones that are best for them. These alternatives include dividing the community into separate spaces, highlighting good content, removing inappropriate content, and friend feeds or even full-blown recommendation systems that show different slices of the content to different people.

The fourth category of design levers involves external communication. Content can be imported from or exported to other communities. Identities and profiles can be shared or hidden. Facilities can be provided to allow people to invite friends or forward content to them.

The fifth category involves feedback, rewards, and sanctions. Feedback tells people how others have reacted to their participation in the community. Such feedback can be informal or structured in the form of ratings or a button to click to indicate the liking of something. Rewards and sanctions give or remove something that people value in response to the actions they take. They can be intangible, in the form of approval or disapproval or status in the community. But they can also take the more tangible form of additional privileges in the community or even money or prizes.

Sixth, communities can articulate different roles, such as welcomers for newcomers or dispute handlers. They can also have rules and guidelines about how people should behave, which can have a big impact on the nature of interactions in the community. Finally, they can establish procedures for decision making and conflict resolution.

Seventh, there are access controls, which place limits on who can join the community and what actions they can take. For example, credentials may be checked to allow only qualified people to join, or completion of a CAPTCHA (Completely Automated Public Turing test to tell Computers and Humans Apart) such as transcribing distorted text may be required to prevent computer programs from creating accounts. Moderation privileges may be extended only to members in designated roles. Alternatively, people may need to pay, using some internal currency, to perform certain actions.

Finally, in every chapter, we find that simple communication choices—ways of framing what the community is and what happens there[1]—can have a big impact on how the community functions. For example, a community can highlight bad behavior and how it is punished or can try to hide that it ever happens. A community can present itself as similar to others or highlight characteristics that distinguish it from others. It can prime norms of reciprocity. It can choose a tagline that emphasizes different aspects of the community. These and many other communication and framing choices can affect all five of the challenges, from getting a new community started through regulating behavior in an established community.

Note that with a few exceptions, we consider only design alternatives that vary how people perceive a community and what they give to or get from it. There are a variety of other alternatives in the realm of interaction design that are beyond the scope of this book. For example, though we discuss the impact of including photos of people and associating the photos with the content they contribute, this book is silent about the size, placement, or other aesthetics of the photos. And though we discuss the time cost for people of browsing through a collection of mostly irrelevant items, we do not analyze the various techniques that could be used to reduce those browsing costs, such as showing collapsed summaries with the full contents prefetched so that they can be displayed without delay if a user hovers or clicks on an item summary. Interaction design choices can have a profound effect on the user experience of an online community and can nudge people toward certain behaviors just as well as the design levers we focus on. The subtleties of interaction design, however, are beyond the scope of our expertise and beyond the scope of this book.

4 The Morality of Design

Even if convinced of the feasibility of designing online communities, some of our students question its morality. The terms "social engineering" and "paternalism" have acquired negative connotations in American political discourse. Generally, people dislike the idea of being manipulated, even if it's for their own good. Viewed in that light, designing the interaction environment of an online community in order to elicit individual behavior that benefits the community as a whole seems morally repugnant.

Weighed against this value of freedom from manipulation, however, there is also a moral imperative to create online communities that work well. People gain immense value from the education, social support, and entertainment that online communities provide to their members and from the information products that they produce for society. If different design alternatives can make the communities more attractive for their members or more productive, then forgoing those benefits may be a significant cost.

Moreover, decisions will be made anyway—through inaction if not through action—about all the design alternatives considered in this book. Any such choices, no matter how they are made, will inevitably influence members and prospective members to behave in certain ways. There is no default, morally neutral online-community design that has no manipulative effect on members.[2]

We argue that the primary moral arguments are thus not about whether to make explicit design choices in order to achieve community goals but about which community goals are the right ones. Making an online community function better may not always be a worthy goal. In some cases, an online community that functions well may produce negative effects for its members (for example, a community that encourages and supports its members to continue their bulimia) or for society as a whole (for example, a terrorist cell). In other cases, it is not so clear what it means for an online community to function better. Most goals, if achieved, involve improving the community in the eyes of some people and making it worse in the eyes of some others. For example, trolls gain enjoyment from disrupting some communities. A design that effectively deters trolling benefits most of the community members but makes things worse for the trolls.

In the remainder of this book, we leave moral judgments—about which goals are worth designing for—to our readers. Our focus is on identifying the likely effects of particular design alternatives in meeting the fundamental design challenges of online communities. We sometimes adopt shorthand like "good behavior" and "bad

behavior," but these should be taken as good or bad relative to the goals of the designer, whomever that may be.

5 The Promise of Mining Social Science

How can an online-community designer build up intuitions about the likely impacts of alternative design choices? Previous practitioner-authors have offered many helpful insights based on design decisions that were made when building online communities that they advised or observed (e.g., Kim 2000; Powazek 2002; O'Keefe 2008). Wenger, White, and Smith (2009) provide several useful frameworks for thinking about online-community design decisions, based on conceptualizing them as communities of practice. Preece (2000) summarized terminology and findings of research related to interpersonal communication and networks and groups that may provide useful background knowledge for a designer.

The second central argument of this book is that social science findings can and should inform more directly the choices that online community designers make. There is a rich research literature in psychology, economics, and the other social sciences about the individual motivations and conditions under which individuals, groups, and organizations are successful. Although most of this research has developed in the context of offline interactions, some has now been replicated in online social settings.

Social science research can inform design in several interrelated ways. First, it can be used to identify problems or challenges that will be faced by most online communities. For example, the theory of network externalities in economics, which we discuss in chapter 6, and the empirical research from which it grew, explore the impact of how the attraction of many groups for potential members grows with the number of people who already participate. This relationship between the attractiveness of a community and its size raises problems for new communities because during their start-up phases, they do not have enough members to provide the resources that will attract other members and allow growth. As another example, the theories of public goods from economics and of social loafing from psychology predict that when individual contributions are needed to produce outcomes that benefit everyone equally, voluntary contributions will be at suboptimal levels. The information contributions that people make to online communities often have this public goods character, and encouraging contributions is thus an important challenge for many online communities.

Second, social science theories provide ideas for solutions to the problems. Thus, if, as theories of network externalities predict, new online communities struggle

because they initially have too few members and too little content to attract and retain members, creating compatibility between communities can overcome this problem (Shapiro and Varian 1999). This is the solution adopted by the makers of Scrabulous, who introduced their game in Facebook, which already had a large number of members available as potential players. As another example, theories of collective effort identify several potential solutions to communities facing problems of undercontribution. Because feeling that one's contribution will be redundant is one reason that people undercontribute, a solution is to make potential contributors believe that their contributions are important. Designers have a number of ways to make potential contributors feel that a contribution will matter, such as partitioning the group so that each contributor is a member of a smaller subgroup or reminding potential contributors about the uniqueness of their contributions.

Finally, and perhaps most important, the social science research base provides predictions about likely consequences of various design decisions. For example, theories about interpersonal bond formation yield a prediction that target members will become more committed to a community to the extent that they have repeated interactions with other members and to the extent that those other members are similar to them. As another example, theories about goal setting and monitoring yield a prediction that contribution goals will be more effective at eliciting member contributions the more challenging they are. We refer to predictions of this sort as "design claims" and describe the structure and limits of such claims in more detail in the next section.

One strand of theory we draw on starts from a premise of individuals making choices that increase their own utility, that is, the difference between their benefits and costs. Thus, many design choices are geared to reducing costs, increasing benefits, or changing individuals' ability to assess the costs and benefits. Game theoretic models enable analysis of interdependent choices and predictions about equilibrium outcomes. For example, in a situation in which many people would want to join an online community only if others also joined, there are two equilibrium outcomes: one in which everyone joins and one in which none do. In such situations, one task of the designer is to shape people's expectations about what others are likely to do. Models of incomplete information permit reasoning about situations in which there is uncertainty. For example, such models can help to understand whether a seller's previous feedback on eBay can be expected to serve as a reliable signal about his or her trustworthiness.

We also draw on a variety of other theories from the fields of social psychology and organizational behavior that predict individual behavior in group and

organizational settings. We use the plural "theories" advisedly: there is no unified theory in modern social psychology with pretensions of explaining all of social behavior. Rather, the intellectual style has been to build and test a large number of midlevel theories, each attempting to account for an interesting social phenomenon in a limited domain. For example, we draw on theories of goal setting, social comparison, persuasion, conformity, and interpersonal bond and group identity formation. Despite the lack of a single overarching theoretical framework analogous to that of evolution in biology or utility maximization in microeconomics, these midlevel theories provide a rich and empirically verified understanding of some of the central phenomena of behavior in social settings.

Although social science theory is helpful in identifying problems that online communities face, suggesting potential solutions to them, and articulating claims about the likely impacts of design choices, it has its limits. First, the theories are incomplete; they offer no guidance on some important design choices. Second, they may be incorrect; like all scientific theories, they are subject to revision based on new data from new experiments. Third, creativity and care are required to map general theories to the particular context of online communities; here, we hope that this book makes a contribution by translating social science findings into useful design claims.

6 Design Claims

We use the device of design claims to translate theory to design alternatives that achieve community goals. Design claims follow a positivist scientific paradigm, seeking to state general claims—that under certain observable conditions certain outcomes can be expected. In our case, the conditions that are of particular interest are design alternatives and the outcomes are desirable features of an online community, which we refer to as "design goals." For example, we state the following design claim in the chapter on motivating contributions: *Small tangible rewards are likely to reduce contributions for intrinsically interesting tasks*. Here, the design alternative is promising small tangible rewards and the design goal is maximizing the efforts that members contribute to tasks that benefit the community. We restrict the scope of applicability of design claims by specifying a restricted set of context conditions for their applicability. Thus, the claim applies only to intrinsically interesting tasks, not to boring tasks. The context conditions will specify properties of the community (e.g., size, purpose), properties of members (e.g., newcomer or long-time member, gender or other demographic characteristics), or properties of tasks (e.g., challenging, interesting).

Table 1.2
The logical structure of design claims

Type	Logical structure
Noncomparative	Alternative X helps/hinders achievement of goal Y under conditions Z
Comparative	Alternative X1 is more effective than X2 at achieving goal Y under conditions Z

Many of our design claims are comparative. For example, the chapter on contributions also includes this design claim: *Nontransparent eligibility criteria and unpredictable schedules lead to less "gaming of the system" than do predictable rewards*. Here there are two design alternatives: predictable rewards versus rewards with nontransparent eligibility criteria and unpredictable schedules. The claim is that one is better than the other at achieving the goal of people not doing useless or destructive actions just to get the rewards (gaming the system).

Thus, we have two logical structures for design claims (see table 1.2).

Whenever we state design claims, we offer evidence in support of them. In some cases, the evidence comes from social science theories or findings that have been articulated for more general settings beyond online communities. Usually, these theories have been tested abstractly in laboratory settings. In other cases, the evidence comes from experiments specific to the online community setting. Evidence may also come from observational studies of particular online communities. Observational data may be quantitative (e.g., counts of how many posts were made) or qualitative (e.g., analysis of their content, or subjective reports from interviewing participants). In some cases, observational studies will be used merely to offer an example consistent with the design claim (i.e., here is a site that used alternative X and it achieved goal Y). This, of course, is relatively weak evidence, as the only information it provides about whether X had anything to do with the achievement of Y comes from the subjective reports of the designers or participants.

Our project of collecting and organizing design claims is akin to efforts to codify what are called "pattern languages" (Alexander, Ishikawa, and Silverstein 1977; Rising 2001). Crumlish and Malone (2009) nicely present a collection of useful design patterns that have been used in online communities. There are a couple of differences in our approach, however, from most work on design patterns. First, although we are sometimes inspired by a bottom-up approach of noticing commonly occurring features of online communities, more often we start from a design goal and some relevant theories and try to systematically explore the space of possible design choices that

could help achieve the goal. In some cases, we identify choices that *should* help achieve the goal, but have not yet, to our knowledge, been tried in existing communities. Second, although approaches to pattern languages vary, usually the design alternative itself is the central element, presented with ancillary information about when it might be best to use it, what it can be expected to accomplish, and hints and cautions about implementing it. By contrast, as we discuss shortly, we have organized our exposition around goals and challenges, presenting together all the design alternatives that have an impact on that goal. Third, we have chosen the term "claim" rather than "pattern" to emphasize that we are laying out causal claims in which design X leads to outcome Y rather than merely observing that X occurs frequently in practice. The preface to *A Pattern Language: Towns, Buildings, and Construction* indicates that Alexander (1977) intended patterns to convey causal claims—indeed, that they should convey necessary as well as sufficient conditions (if you want to achieve an outcome Y, then X is necessary). Not all of the actual design patterns, however, in Alexander's work or among others adopting the pattern language approach, seem to make such causal claims.

It is worth noting that both design alternatives and design goals can be expressed at varying levels of abstraction. For example, a design alternative at a high level of abstraction might be to provide tangible rewards for activity. At a much more specific level of abstraction, two design alternatives might be to provide a user with a $5 gift certificate or to make a $5 donation to a charity that the user chooses. Throughout the book, our design claims are made at whatever level of abstraction is most appropriate. It is also worth noting that what is expressed as a design alternative at a high level of abstraction may be expressed as a goal at a more specific layer of abstraction. For example, at a high level, we might say that the design alternative of making people feel unique helps achieve the goal of motivating effort. A more specific design claim might state that reminding people of unusual actions they have taken helps achieve the goal of making people feel unique.

It is also worth noting that a design alternative X may be compound, combining simpler alternatives. For example, a design claim might state that having a forum and a separate email list will make it harder for either of them to get to a critical mass of usage. Or a design claim might state that for technical support communities, an email list and wiki used together are more effective than either one on its own (Hansen 2007).

The design claims are not prescriptive rules that a designer can or should follow blindly, for two reasons. First, the predictive claims state only that a design alternative X helps or hinders achievement of a goal Y, not that it will always achieve or prevent the achievement of the goal. A claim that *small tangible rewards are likely to reduce*

contributions for intrinsically interesting tasks is a claim about the effect on average. In a particular situation, a designer will need to judge how intrinsically interesting a task is, whether a reward is likely to be perceived as small or large, and whether there are any extenuating circumstances.

Second, multiple design claims may suggest implications of a single design choice for more than one design goal, and the designer may have to trade off achieving one goal against interfering with another. For example, in an open source project, the design alternative of giving lots of people commit privileges (so they can easily add their contributions to the group's code base) would be likely to increase the number of contributors but decrease the amount of effort by each person and may also increase the number of bugs in the code. A designer will need to judge whether that trade-off is worthwhile in the particular situation.

Third, the theories from which our design claims derive, and thus our design claims as well, usually state the effects of manipulations holding everything else constant. On the other hand, whether designing from scratch or changing an existing system, designers typically make a number of choices at the same time. For example, at the same time that a community introduces a point system to track and acknowledge member contributions, it may also change its tagline and frequently asked questions (FAQ) section to suggest a more collaborative, less competitive atmosphere. Design claims offer guidance on the likely impact of either of these changes separately, not their joint impact. A designer will need to rely on intuition to judge whether a set of design choices are complementary or whether they interfere with each other.

7 Organization of the Book

We have organized the book around the high-level design challenges described in section 3. Thus, design claims related to a particular goal are presented together, even though they may involve quite different design elements. This organizational scheme serves several purposes. First, for a student or practitioner new to online community design, it highlights the challenges that typically arise in online communities so that some thought can be given to them before they arise. Second, it offers a systematic way to consider and compare alternative approaches to handling those challenges. For example, a designer who begins with a particular design element—perhaps because a boss has encountered the feature in another community and asks him or her to investigate its use—will naturally be led also to consider alternative ways to achieve the goals that design element normally promotes because they are presented near each other in the book.

We try to salvage some of the benefits that could be gained from alternative organizational schemes through cross-indexing. Design claims related to achieving the same design goal appear linearly near each other in the same chapter. When a particular design element or theory used in one design claim also appears in other sections or other chapters, we indicate that. Each chapter concludes with a summary of all the design levers considered in the chapter, grouped by the eight categories, to provide an alternative index into the contents of the chapter.

Chapters 2 through 6 discuss the high-level design challenges:

- Chapter 2: Encouraging contributions to online communities
- Chapter 3: Encouraging commitment to online communities
- Chapter 4: Regulating behavior in online communities
- Chapter 5: Dealing with newcomers
- Chapter 6: Starting a community

The chapter ordering reflects our pedagogical experience of presenting this material in workshops and courses. While community startup clearly precedes the other issues chronologically for any community, the material in that chapter focuses on the special problems of the startup phase, beyond those that face communities on an ongoing basis. It is hard to separate out and focus on those special problems before developing a clear idea about how the community might operate on a steady-state basis. Similarly, while people must first be newcomers before they can be long-time members of a community, we have found that it is easier for audiences to set aside the challenges of how new people will be recruited and initiated until after thinking through how veterans will interact. Readers are, of course, free to sample the chapters in any order. Although there are many cross-references, the chapter can be read independently, in any order.

No book is ever complete. There are always more topics at the periphery that could be included. Our book says much about beginnings—new communities and the entry of new members to existing communities. By contrast, it says little about endings. A future book could usefully examine when and how to gracefully handle individual departures and how to gracefully close a community that no longer serves a clear purpose. The book also says little about the challenge of keeping an online community fresh over time. As with other organizational forms, if they last a long time, there is a danger that the world will pass them by. Designers and managers can make choices that enhance a community's ability to monitor changes in the larger environment and to innovate in its practices in response to those changes. A future edition of this book could usefully include a chapter on organizing online communities in a way

that encourages innovation. Finally, while this book discusses how to motivate effortful contributions, it says little about how to coordinate those efforts, a topic that is beginning to get a lot of research attention, including by some of the coauthors of chapters in this book. A future edition could include a chapter on coordinating effort.

Despite these limitations, we think that this book will provide useful guidance to practitioners as well as an introduction to online communities suitable for advanced undergraduates and professional master's degree students. Through specific design claims, backed up with supporting examples, the book provides a wealth of design guidance. By organizing the exposition around fundamental design challenges, however, we encourage practitioners to consider alternative solutions to challenges they face, rather than simply adopting a feature that they have seen in other sites. Moreover, by grounding the design claims in theory as well as empirical examples, readers will be better able to reason about whether a particular technique is likely to work in particular online communities.

We hope to evolve the set of design claims and their justifications over time. Please send us your examples, both those that support our design claims and those that do not. Or post a public comment on our website, SuccessfulOnlineCommunities.com: join the online community of online community students, practitioners, and researchers!

Notes

1. In the field of behavioral economics, "decision frame" refers to the decision maker's conception of the acts, outcomes, and contingencies associated with a particular choice. Tversky and Kahneman (1981) use the term more broadly, beyond the context of specific choices or decisions.

2. We are indebted to Thaler and Sunstein (2008), who nicely articulated a similar argument in the context of choice environments, such as the selection of healthy or unhealthy foods from a cafeteria or whether to set aside money from each paycheck to invest for retirement. They argued that any choice environment will predictably nudge people toward making one choice or another and that there is no way to pick a default, morally neutral choice environment. Either the apples or the chocolate bars can be at eye level in the cafeteria checkout line, and whichever one is there will be consumed more often than it otherwise would.

References

Alexander, C., S. Ishikawa, and M. Silverstein. 1977. *A Pattern Language: Towns, Buildings, Construction.* New York: Oxford University Press.

Crumlish, Christian, and Erin Malone. 2009. *Designing Social Interfaces*. Sebastopol, CA: O'Reilly.

Hansen, Derek L. 2007. *Knowledge Sharing, Maintenance, and Use in Online Support Communities*. PhD diss., University of Michigan.

Kim, Amy Jo. 2000. *Community Building on the Web: Secret Strategies for Successful Online Communities*. Berkeley: Peachpit Press.

Kittur, A., and R. E. Kraut. 2011. Unpublished data.

Kraut, R. E. 2011. Unpublished data.

Lithium Technologies. 2009. The Linksys ROI Story: Support Community Delivers Significant Savings from Call Deflection. http://www.lithium.com/pdfs/casestudies/Linksys-ROI-Case-Study.pdf.

Mathieu, J. E., and D. M. Zajac. 1990. A Review and Meta-Analysis of the Antecedents, Correlates, and Consequences of Organizational Commitment. *Psychological Bulletin* 108 (2): 171–194.

O'Keefe, Patrick. 2008. *Managing Online Forums: Everything You Need to Know to Create and Run Successful Community Discussion Boards*. New York: AMACOM Books.

Powazek, D. M. 2002. *Design for Community: The Art of Connecting Real People in Virtual Places*. Indianapolis: Pearson Technology.

Preece, Jenny. 2000. *Online Communities: Designing Usability and Supporting Sociability*. New York: John Wiley & Sons.

Rising, Linda. 2001. *Design Patterns in Communications Software*. Cambridge: Cambride University Press.

Schwartz, M. 2008. The Trolls among Us. *New York Times*. http://www.nytimes.com/2008/08/03/magazine/03trolls-t.html?_r=2&oref=slo.

Shapiro, Carl, and Hal R. Varian. 1999. *Information Rules: A Strategic Guide to the Network Economy*. Cambridge, MA: Harvard Business School Press.

Stoddard, J. L., E. M. Augustson, and R. P. Moser. 2008. Effect of Adding a Virtual Community (Bulletin Board) to Smokefree.gov: Randomized Controlled Trial. *Journal of Medical Internet Research* 10 (5): e53.

Thaler, Richard H., and Cass R. Sunstein. 2008. *Nudge: Improving Decisions about Health, Wealth, and Happiness*. New Haven, NH: Yale University Press.

Tversky, Amos, and Daniel Kahneman. 1981. The Framing of Decisions and the Psychology of Choice. *Science* 211 (4481):453–458.

Wenger, Etienne, Nancy White, and John D. Smith. 2009. *Digital Habitats: Stewarding Technology for Communities*. Portland, OR: CPSquare.

Wikimedia. 2011a. Wikipedia Statistics: Article Count (Official). http://stats.wikimedia.org/EN/TablesArticlesTotal.htm.

Wikimedia. 2011b. Wikipedia Statistics: Active Wikipedians. http://stats.wikimedia.org/EN/TablesWikipediansEditsGt5.htm.

Wikipedia. 2010. Wikipedia: Version 1.0 Editorial Team/Index. http://en.wikipedia.org/wiki/Wikipedia:Version_1.0_Editorial_Team/Index.

Zuckerberg, Mark. 2010. 500 Million Stories. http://blog.facebook.com/blog.php?post=409753352130.