Catalyst: Triggering Collective Action with Thresholds

Justin Cheng, Michael S. Bernstein
Stanford HCI Group, Computer Science Department
Stanford University
{jcccf,msb}@cs.stanford.edu

ABSTRACT

The web is a catalyst for drawing people together around shared goals, but many groups never reach critical mass. It can thus be risky to commit time or effort to a goal: participants show up only to discover that nobody else did, and organizers devote significant effort to causes that never get off the ground. Crowdfunding has lessened some of this risk by only calling in donations when an effort reaches a collective monetary goal. However, it leaves unsolved the harder problem of mobilizing effort, time and participation. We generalize the concept into *activation thresholds*, commitments that are conditioned on others' participation. With activation thresholds, supporters only need to show up for an event if enough other people commit as well. *Catalyst* is a platform that introduces activation thresholds for on-demand events. For more complex coordination needs, Catalyst also provides thresholds based on time or role (e.g., a bake sale requiring commitments for bakers, decorators, and sellers). In a multi-month field deployment, Catalyst helped users organize events including food bank volunteering, on-demand study groups, and mass participation events like a human chess game. Our results suggest that activation thresholds can indeed catalyze a large class of new collective efforts.

Author Keywords

Collective action, activation thresholds, critical mass.

INTRODUCTION

By lowering barriers to participation, the web has transformed how people work together [12]. Successful collective efforts include the compilation of encyclopedic knowledge on Wikipedia [26], disaster relief for major earthquakes [48], and planning and organizational support for the Arab Spring uprisings [25]. However, for every successful cause, there are many that never reach critical mass. For example, Wikipedia was the only large-scale success out of eight independent, simultaneous attempts at creating a collaborative online encyclopedia [23]. Moreover, successful causes are typically limited to online tasks. Causes that try to mobilize offline participation, effort, and hours have far less design support.

CSCW'14, February 15–19, 2014, Baltimore, Maryland, USA.

ACM 978-1-4503-2540-0/14/02...$15.00.
http://dx.doi.org/10.1145/2531602.2531635

Many efforts struggle because participation incurs significant costs and risks, regardless of the eventual outcome [27]. Supporters run the risk of committing time to a cause that never captures enough interest to achieve its goals. Similarly, organizers risk putting time and emotional investment into a cause that never catches on. Many collective efforts need critical mass to be successful, for example study groups, flash mobs, or musical ensembles. Further, the significant upfront cost of organizing a cause, and discouragement that results from past failure [10], may prevent organizers from trying again.

Catalyst is a platform designed for crowds to take action predicated on the concept of *activation thresholds*. Activation thresholds condition commitment on the participation of others: a cause gets activated only if enough participants commit. Organizers can pitch ideas, only putting effort into those that reach success thresholds, and participants need only participate if enough others also commit to doing so. For instance, a benefit concert might get put on only if at least 100 people commit to attending. Activation thresholds are an explicit manifestation of the *threshold point* described in models of collective behavior, where the benefits of participation begin to outweigh its drawbacks [21, 38]. As each participant's utility is dependent on the number of fellow participants, they typically want to join efforts that reach this threshold and are guaranteed to succeed. As the commitment to organize or participate in an event is conditioned on the activation threshold being reached, one's time is only spent on coordinating or attending successful events. Thus, these thresholds can significantly reduce the upfront risks of organizing causes.

The funding model used in crowdfunding systems such as Kickstarter [1] is essentially a type of activation threshold: a project is funded only if supporters commit a sufficient amount of money. Motivated by the success of using activation thresholds in crowdfunding, we hypothesize that these thresholds can generalize to new classes of efforts that raise person-hours, and not just money. In doing so, we broaden the design space of such systems to include commitments that require later action (such as attendance), instead of an automatic transfer of funds.

Many collective action efforts also involve more complex forms of coordination than a threshold number of participants. In many volunteer organizations, for example, significant effort is spent in coordination and communication [51, 48]. Beyond simple participation thresholds, Catalyst minimizes back-and-forth coordination by enabling more highly-structured activation thresholds, for example by requiring that

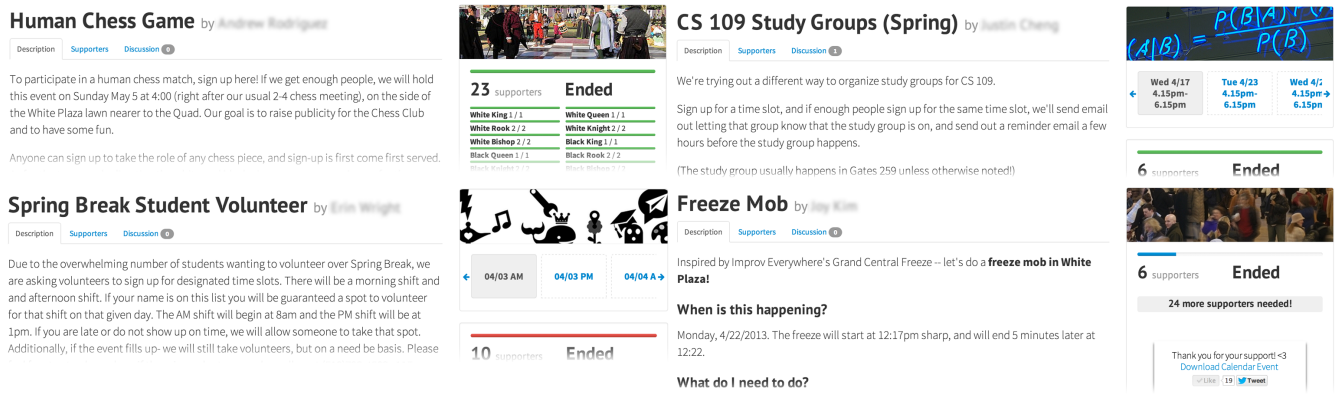

Figure 1. A variety of causes were organized on Catalyst, each using different activation thresholds, and succeeding to different degrees.

different roles or time slots all be filled. Roles end up being critical in many types of events: bake sales require both bakers and sellers, and office hours require both teaching assistants and students. This approach can also help overcome the free rider effect and promote peer commitment [15].

In this paper, we first review theories of collective behavior and existing online systems for collective action. We then describe the design of Catalyst, its support for activation thresholds and other mechanisms for coordination. Next, we detail Catalyst's evaluation through a field deployment where it was used to organize real-world causes including volunteering at a large food bank, on-demand study groups and flash mobs. We find that activation thresholds help both in group coordination and deferring commitment, with most participants following through on their initial commitments.

RELATED WORK

Activation thresholds are inspired by the threshold theories (e.g. critical mass theory) of collective action advanced by Granovetter and Oliver [21, 38], and generalize the use of similar thresholding mechanisms found in many crowdfunding platforms. Drawing on the literature on specialization and small group sizes and their impact on collective action, we extend these thresholds beyond simple minimum participation to support roles and time slots.

Collective action

The term *collective action* is used differently in fields such as sociology, political science, economics, and social computing. In particular, we refer to collective action broadly as a group of people coming together to achieve a shared goal. This definition has been applied to social computing efforts such as Wikipedia [23]. However, some literature focuses on social movements such as revolutions [25] or environmentalism, which is largely outside the scope of this work.

Triggers

There is a tipping point beyond which the benefit of participating in a cause exceeds its cost [40]. Threshold models suggest that individuals have unique tipping points, with each adopting the cause only if the proportion of people in the network who also support it exceeds this threshold [21, 38]. With the right conditions, initial support can precipitate a chain reaction that ultimately spreads to all members of a community.

These theories aim to explain several characteristics of collective action, including numeric thresholds and their tendency to either fail completely or succeed wildly [38]. A common framing is that the value of participating as a function of other number of participants may be "S"-shaped. In particular, collective action begins with a period of startup costs and low but increasing marginal returns from each additional participant, which accounts for the high failure rate of many potential causes. This then leads to a period of higher marginal returns (i.e. the tipping point), and finally satiation where returns are high, but the marginal returns diminishing. In other words, there is "some threshold of participation that is crossed before a social movement 'explodes' into being" [38]. Catalyst makes explicit this notion of a numeric threshold. Users who commit to Catalyst events are, in a sense, agreeing that its activation threshold is at or above their own personal threshold.

Efforts that fail to reach this tipping point, or critical mass, including online encyclopedias and WikiProjects, may never achieve their goals [49, 19]. Critical mass is also key to groupware acceptance [22]. Solving critical mass problems can thus benefit a large class of social computing platforms.

Collective action platforms

Where many systems exist to support specific types of group action including crowdfunding [1], crisis organization [37], creative projects [32], and public deliberation [28], Catalyst represents the first attempt at supporting the organization of arbitrary threshold-based collective activities. Motivation schemes can also be layered on top of such platforms; for example, the MIT DARPA Red Balloon Challenge team offered financial incentives to encourage effective organization [41]. In an organizational setting, crowdfunding was shown to encourage diverse proposals and promote collaboration [36].

"All or nothing"

Thresholding is ubiquitous in online crowdfunding web sites such as Kickstarter. On these sites, the funding model is predicated on the concept of *critical mass*, or that a project succeeds only if it reaches a specific funding goal. Crucially, they use the concept of "all or nothing" in canvassing support

from its users. When supporters back a project, their money is first held in escrow; depending on whether the project's funding goal is met, this money is either released to the project creators, or returned to the supporter. This "all or nothing" approach ensures supporters risk-free commitment: either the project succeeds and they obtain what the project promised, or it fails and they get their money back. Groupon also uses thresholding to activate deals: a deal is only active when enough users have claimed it. This thresholding positively correlates with the number of deals sold [14]. Sites such as Zokos [5] also require a minimum number of attendees. Still, the primary goal and commitment mechanism in many of these systems is the monetary pledge, which is optional in Catalyst. Our work generalizes these mechanisms for financial critical mass to commitments involving time and effort. Rather than raising money, Catalyst raises person-hours.

Motivating commitment

Social signaling can help people estimate an event's probability of success. A project's funding status acts as social proof, with the amount of funding and number of backers signaling the quality of the project. "People are persuaded more by the actions of others than by any proof we can offer," [15]. For example, peers' listening patterns have large effects on an individual's preferences [47].

Social signaling helps later participants make inferences based on previous participants' decisions, and can create social pressure to conform. When faced with a lack of information, people tend to mirror the decisions that others made previously [8], with the strength of this effect correlated with the size of the conforming group [34]. This signaling also affects a participant's perception of whether an event will achieve a critical mass of users. This perceived critical mass was shown to significantly influence a user's intention to use groupware technology [30]. Thus, Catalyst provides information about the current state of an event such as how close an event is to its threshold.

Many types of collective action also involve specialized roles. Wikipedia is organized hierarchically into administrators, editors, trusted users, and untrusted users, with resource and interest heterogeneity shown to be key factors contributing to its sustainability [43]. In creative systems like Pipeline, leaders delegate the creation of specific art to different individuals [32]. Further, individuals are more likely to contribute to online communities when given specific tasks, or reminded of the uniqueness of their contributions [11].

In contrast to crowdfunding, where contribution is purely monetary, in collective action, "time is the ultimate resource" [39]. This introduces several complications, most importantly the free-rider problem, as well ensuring follow-up commitment. Collective efforts also suffer from social loafing: the larger a group, the larger the loss of individual motivation and coordination [45]. Smaller groups can mitigate this effect as individuals are more likely to take action when they stand to gain a substantial proportion of the total benefit [40]. While Catalyst can support events of arbitrary size, roles and time slots help break up larger events into smaller, more manageable ones.

In contrast to crowdfunding, the commitment contracts on Catalyst are primarily social, implicitly with both the organizer and other participants. Thus, it is significantly harder to enforce followthrough on Catalyst than on many crowdfunding web sites, where money is immediately taken into escrow, or when the commitment contract is material or financial. In such cases, commitment devices can serve to discourage people from reneging on their promises to participate [20], and act to filter out the less committed.

CATALYST

Motivated by the challenges of organizing group action, our work translates threshold models into a general design mechanism of *activation thresholds*. In activation thresholds, users' commitments are only called in if the event gathers a minimum number of participants. In this section, we introduce Catalyst, a platform that demonstrates the opportunities for activation thresholds, including variants such as minimum thresholding, maximum thresholding, and role-based thresholding. We motivate Catalyst's design with an example, then the system and its use.

Alfred is trying to organize a neighborhood bake sale in support of a local charity. He creates an event on Catalyst, specifying the date and time of the sale, as well as a minimum number of people that need to participate to ensure its success (the *activation threshold*). As he needs volunteers to donate baked goods, decorate the booth for the bake sale, and sell the baked goods, he creates specific roles for these (*role-based thresholding*, Figure 2a); only when enough people fill these roles will the sale take place. Alfred links to the event on Facebook and posts flyers to recruit potential volunteers. Later, Beth sees that 3 of her neighbors have already signed up for the event in various roles. Confident in her oatmeal cookies, she signs up as a baker by entering her name and email address (*coordinating action*, Figure 2b). Several signups later, Cole, who can neither bake nor decorate, decides to help out as a seller instead (*differing levels of commitment*). Cole's support brings the number sign-ups to the requisite 10, including committed individuals for each role (*the threshold point*, Figure 2c). Having reached its threshold, the bake sale is on! Catalyst emails everyone who signed up with an email Alfred had prepared to confirm the event.

Activation thresholds

In Catalyst, activation thresholds are the minimum number of people who need to commit to an event for it to occur. To create a sense of urgency, and drawing from crowdfunding techniques, the organizer must attract the minimum number of commitments within a specific time span. Catalyst reinforces this urgency by reminding committed participants roughly 72 hours before a deadline.

Catalyst primarily focuses on lower bounds for a number of participants (minimum thresholding). However, early field studies of the system made clear that there was more complexity to collective action than Catalyst captured. Thus, organizers may also limit the number of participants (maximum thresholds) to create additional urgency, or specify multiple roles or repetitions of an event.

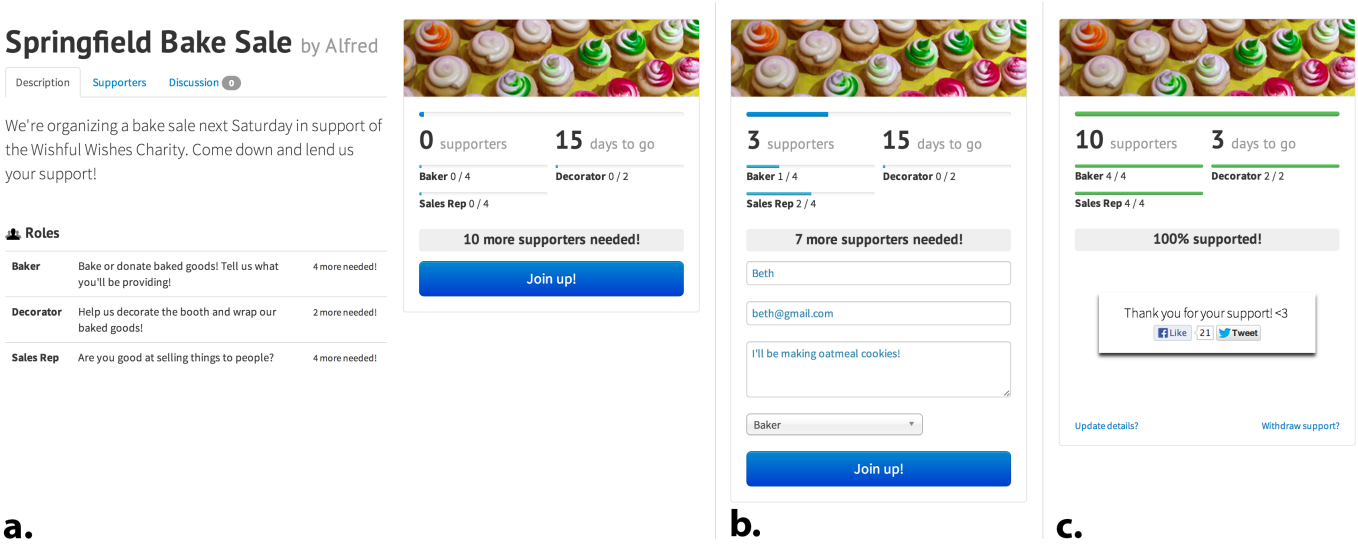

Figure 2. (a) A Catalyst event page shows the event description, the roles that participants can take on, as well as the current status of the cause. (b) A user supports a cause by typing in her name and email address, along with an optional message, picking a role if necessary. (c) A successful cause where the requisite number of supporters was achieved, with every role filled.

Minimum thresholding

In Catalyst, minimum thresholding is based on the assumption that many events require at least some number of people to be successful, and is grounded in critical mass theory [21, 38]. Additionally, minimum thresholding implicitly sets a participation goal to meet, which has been effective in persuading people to donate to charitable causes [17].

Maximum thresholding

While minimum thresholds show how "close" one is to success, maximum thresholds create artificial limits on the number of signups, thus making them appear more desirable. The time limit imposed on every cause also acts as a type of maximum threshold. Artificial scarcity causes participation to be perceived as more valuable, and can motivate participants to support and subsequently participate in an event [15]. In particular, the effect of limited supply on perceived value is strongest together with popularity, suggesting that maximum thresholds are most effective with popular events [50].

Role-based thresholding

While simple activation thresholds are intuitively appealing, many events require complex coordination — a potluck dinner might require different attendees to bring different dishes; study groups may occur at different times in a day, with each requiring one tutor and five students. It may also be necessary to condition spots on the roles that other participants pick. For example, supporters may only be able to sign up as students if there is at least one peer tutor in a study group session.

Also, many events are temporal and repetitive in nature: study groups might meet several times a week, and beach cleanups might occur over several different weekends. Thus, Catalyst supports multiple time slots for a single event, allowing participants to choose a slot that fits best in their schedule. All time slots that reach their threshold will be activated.

By introducing roles or specialization, participants can concentrate on a narrower set of tasks, improving their effectiveness [9]. In collective action systems like Wikipedia, participation begins at the periphery, with users assuming roles of increasing importance as they get more familiar with the system [13]. Similarly, roles in Catalyst allow for supporters of differing expertise, and differing commitment levels.

In many volunteer organizations, volunteer coordination, which includes scheduling and role assignment, presents an even larger issue than recruitment [51]. Role-based thresholds partially alleviate the cost of coordination by requiring participants to commit to specific roles or time slots.

Monetary commitment

Unlike crowdfunding or fundraising, where the commitment is automatically executed through a donation held in escrow, it is always possible that a Catalyst volunteer may fail to follow through on a commitment. As commitment devices can serve to discourage people from reneging [20], Catalyst also allows organizers to require a higher-friction commitment, for example donating $1 using Paypal to sign up.

Coordinating action via email

Catalyst allows organizers to contact supporters to remind them about their commitments. For instance, activation goal emails let supporters know if a cause they supported has been activated, while failure emails let supporters know a cause failed to get off the ground. Reminder emails sent a few days before can act to increase the tension and urgency in supporting the cause, and can be effective when used with negative frames ("we're not going to succeed without enough supporters!") and statistical support ("we're only 10 supporters away from our goal!") [50].

Designing for Collective Action

The organization of an event involves different costs: those of supporters *committing* or *participating*, and those of organizers *preparing* for, and *coordinating* the event. Each factor represents different tradeoffs which affect an event's scale,

impact or success. The design of collective action systems such as Catalyst can encourage events with particular cost levels, and they may even aim to raise or lower these costs.

Here, we characterize the design of Catalyst and other systems along these axes. Considering different points in this design space may lead to platforms that support very different purposes — for example, an online system for reporting sensed data require initially high commitment (e.g. setting up one's phone to sense the environment), but subsequently less effort to participate (walking around with a phone). One focused on localized, realtime events (low commitment and preparation, but high coordination cost) may focus more on short-term dynamics and recruitment strategies.

Commitment cost relates to the initial effort required to participate in a cause. Events may require high commitment costs to guarantee participation, while those hoping to gather as many people as possible may seek lower costs. Systems for volunteer organizations like VolunteerMatch err on the side of high friction signups, requiring potential volunteers to fill out several forms and pass in-person interviews. While this filters for highly committed volunteers, potential participants may be discouraged by the up-front effort [46]. On the other hand, if the initial commitment is low-cost, participants are less committed to the cause [6]. Crowdfunding, while having a low participation cost in requiring only a monetary pledge, has high commitment cost because this pledge is up-front. Catalyst aims for a balance: just enough friction to provide an honest signal about the intent to participate [18]. Since commitment devices like an initial monetary commitment can potentially improve follow-through, in Catalyst, organizers can require a monetary donation to support a cause, and can offer to return the donation when participants show up to the event.

Participation cost describes the amount of time or energy desired from participants. For instance, volunteering a morning for a beach cleanup may be a substantial commitment, but supporting a petition requires significantly less effort. Differing motivations also influence a participant's desired commitment level — activities with clear personal benefits (e.g., office hours) may be more motivating than those without (e.g., a flash mob). As Catalyst is designed primarily for offline events where participants needed to show up at a specific location and time, events tended to require greater effort. Nevertheless, roles allow for variable participation cost, where not all participants are involved to the same extent.

Preparation cost describes the up-front cost to organizers in preparing an event. For instance, large concerts may require venue reservation months in advance (high cost), and tend to use event management platforms that support ticket sales, as the event's occurrence is usually guaranteed regardless. On the other hand, social media platforms like Twitter favor low-cost, ad-hoc event organization [48]. Planned events are generally well-organized and coordinated, but lack the flexibility of ad-hoc systems which can respond quickly to crises. Catalyst requires organizers to create events, but supports low-cost event creation, so preparation can be deferred until an event's success is guaranteed.

Coordination cost describes the effort required in coordinating and organizing participants. Petitions and fundraisers require little coordination between organizers and participants, as participants need only donate or sign a petition. Thus, platforms like Kickstarter and We The People [4] are designed to primarily support this single transaction, but not more complex coordination. For instance, platforms like ConsiderIt require the crowd to collectively decide on public issues [28]; Doodle requires all participants to report availability and then agree on a specific time. Through activation thresholds (or the lack thereof) and roles, Catalyst supports events which have low to medium coordination costs, such as minimum participation, or a division of labor. In particular, Catalyst can reduce this coordination cost for smaller-scale, local causes that lack resources to develop specialized platforms. As the centralization of ties has a positive effect on collective action [33], a single organizer is responsible for a cause on Catalyst.

DEPLOYMENT AND EVALUATION

Catalyst instantiates our design hypothesis that activation thresholds support collective action. To test this hypothesis, as well as understand the types of events that are well or poorly supported by activation thresholds, we launched Catalyst and studied its use from February through May 2013. We did not pursue a controlled field experiment as early pilots suggested that events exhibited large variance in participation. We instead aimed to attract a large amount of naturalistic usage so we could probe the strengths and weaknesses of the system.

Catalyst attracted interested organizers from a local volunteer food bank, lecturers and faculty at a large American university, student interest organizations, organizers of a massive online open course, and others. The authors mined contact information for these organizations from the web, and then reached out to the organizers via email to explain how Catalyst might support their goals. As each organizer was responsible for organizing their own event, we did not perform any additional participant recruitment. Organizers used their organization's web site or email to direct participants to the Catalyst web site. Table 1 presents a sample of these events, the types of activation thresholds they used, and whether they succeeded. In the case of the volunteer food bank, a computer running Catalyst was also set up on-site.

For events that occurred locally, at least one author was present to verify the number of people that showed up to events that succeeded. To understand subjective perceptions of commitment on Catalyst versus in-person or social-network based invitations, we sent an email survey at the end of the deployment to all organizers and participants who had registered for events on Catalyst. The questions asked can be found in Appendix A.

Results

Catalyst was successful in coordinating hundreds of supporters across tens of events. From February to May 2013, Catalyst saw 2,300 visitors, 15,000 page views, and 368 supporters across all events. 16 leaders created 30 events, with 24

Event (Type)	Activation threshold	Supporters (Sparkline)	Slots	Conv. %	Description
Human chess game (Mass participation)	Min 18 2 Players, 16 Pieces	23	—	8	Act as a chess piece in an outdoor human chess game.
Peer office hours (Self-improvement)	Min 4, Max 6 1 Leader, 3-5 Learners	55	13	23	Tutor students on demand, or be tutored.
Massive online course discussion sessions (Self-improvement)	Min 10, Max 20 1 Mod, 9 Learners	177	28	28	Moderate or attend optional weekly small-group discussion sections on Google Hangouts.
Peer study groups (Self-improvement)	Min 4, Max 6	35	12	12	Attend peer-led, on-demand study groups.
On-demand labs (Self-improvement)	Min 20	61	4	N/A	Lab topics that attracted enough interest were added to the course schedule.
Spring Break volunteering (Volunteer)	Max 8-10	79	10	80	Volunteer at a food bank: there was a maximum number of volunteers each day.
Teatime meetup (Meeting)	Min 4	29	5	20	Weekly teatime meetings occurred if there was sufficient interest.
Food Bank volunteering (Volunteer)	Min 1	62	20	N/A	Volunteer at a local food bank on a given day.
Reading group (Self-improvement)	1 Discussion Leader	6 —	6	N/A	Reading group occurred each week if a discussion leader committed.
Finals study session (failed) (Self-improvement)	Min 20	4	—	0.2	Committed $1 to attend a finals week study session providing candy.
Freeze mob (failed) (Mass participation)	Min 30	6	—	3	Join a flash mob.
Document translation (failed) (Volunteer)	Min 2	1 —	3	2	Volunteer to translate documents for a food bank.

Table 1. The events organized on Catalyst ranged from meetings to study groups to volunteer opportunities to flash mobs. The sparklines show the number of supporters over the lifetime of the event.

obtaining at least 1 supporter (4 of the 16 leaders were personal contacts of the authors). Analyzing participant names using data from the U.S. Census Bureau [3], we estimate that overall, 42% of participants were female; 40% were White, 44% were Asian or Pacific Islander, 10% were Black, and the remainder Hispanic. Nonetheless, the demographic makeup of participants varied depending on the event — office hours involved undergraduate students, while food bank volunteering saw a more diverse range of ages and professions.

Eleven (46%) of the causes eventually succeeded, or reached their activation thresholds, comparable to Kickstarter's 42% success rate as of May 2013 [2]. Catalyst had a mean conversion rate of 16%, meaning that on average, 16% of an event's visitors signed up for it.

Events organized on Catalyst tended to fall into three main categories—volunteer efforts, fun mass-participation events, and self-improvement (Table 1). For a volunteer event, participants committed to packing or distributing food during a particular time slot. As an example of a mass-participation event, participants signed up to act as different chess pieces in a large outdoor human chess game (Figure 3). In the self-improvement category, study groups and office hours for computer science classes were organized on-demand if enough students expressed interest. In particular, additional office hours were only organized if there was also at least one

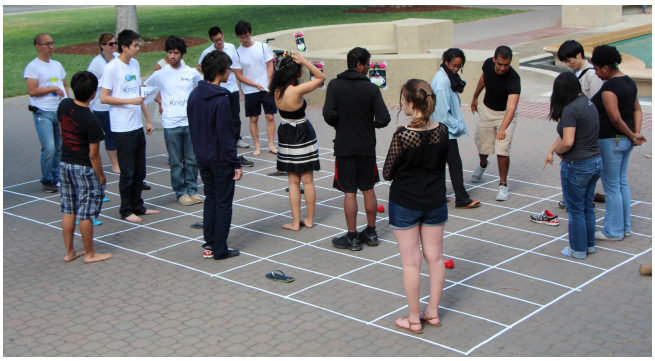

Figure 3. A local chess club used Catalyst to organize a game of human chess, where participants wearing black or white represented chess pieces on a life-size game board.

experienced peer facilitator to lead the discussion. Organizers also used Catalyst to organize meetings and social events, including informal tea gatherings and reading groups.

As predicted by prior work, the sparklines in Table 1 demonstrate empirical evidence of the hypothesized "S"-shaped production function, especially in events with substantial activity (e.g. the human chess game).

A total of 35 participants (labelled P#) and 5 organizers completed the post-deployment survey. Participants are labelled P# and organizers O#.

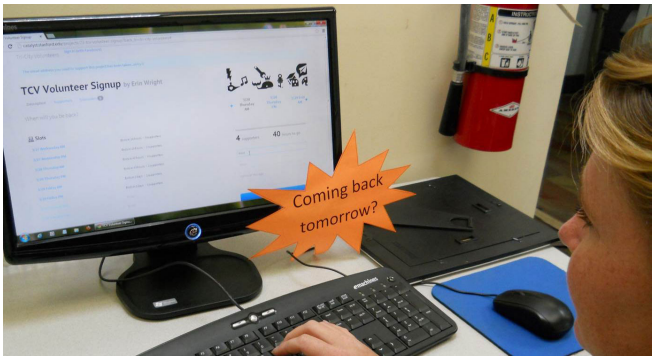

Figure 4. A computer running Catalyst at a local food bank allowed volunteers to sign up for events.

Activation thresholds supported coordination

The visibility of peer commitment served as a valuable signal of future success: participants clustered around topics or slots that others were also beginning to support. Many appreciated that "you could see who else [was] going" (P12, P13). For instance, the online class discussion sections were divided into seven different sections focusing on different topics, each studio having four different meeting times. In almost all cases, only one or two of these four time slots exceeded the minimum number of attendees. Similarly, when a class was polled to figure out which topics should be covered in labs, only two of the four suggested topics exceeded the minimum threshold of 20 supporters (obtaining 22 each), with the other two receiving far less support (6 and 11).

Through maximum thresholds, Catalyst helped generate demand and manage volunteer flow at a local food bank (Figure 4). The food bank expressed that they had previously struggled with large variance in volunteer numbers day-to-day [51]. While students would typically show up on Monday or Tuesday, with little to none on Thursday or Friday, using Catalyst, an average of 7.9 ($\sigma = 1.85$, Gini $G = 0.125$) students signed up for each of 10 slots spread out over the week.

Activation thresholds supported deferred commitment

Activation thresholds helped organizers push back preparation for an event until they had an better idea of how many people were interested. On Catalyst, mass-participation events like flash mobs had the highest failure rates out of all event types. Deferred commitment helped organizers gauge interest before committing time and resources to coordinating participants: "Using Catalyst allowed me to make the event page without committing to it happening, which was good because I was unsure if enough people would be interested" (O4). Similarly, O5 mentioned that during busier weeks, he would "not...[hold office hours] for [only] 2 people."

In other cases, Catalyst also helps organizers better estimate the number of people who participate, as was the case for volunteer food bank signups. For example, the local food bank coordinator used Catalyst to better forecast the number of volunteers to expect each day, then prepare an appropriate number of tasks for those volunteers.

Committed supporters followed through on commitments

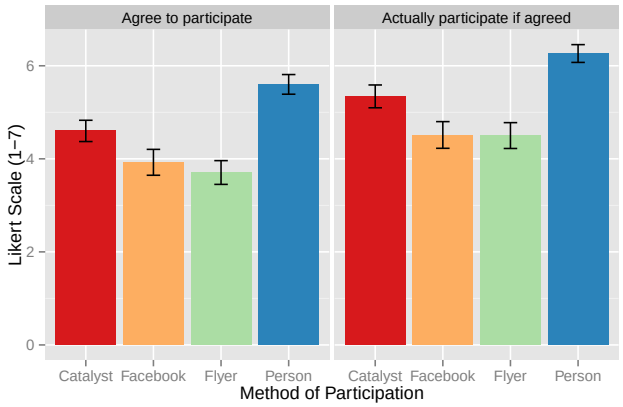

Figure 5. Participants were more likely to both agree to attending an event on Catalyst, and participate if they previously agreed to, compared to an event on Facebook or one advertised through a flyer or email.

While Catalyst successfully gathers commitments, events that might appear successful can still fail to actually mobilize supporters if they do not honor these commitments. Followthrough was thus a major concern while developing Catalyst. However, across the hundreds of committed supporters for Catalyst events, we observed that the vast majority of commitments for triggered events were truthful. For instance, almost all students who signed up for peer office hours attended them. Similarly, more than 90% of participants who signed up for volunteer food bank events followed through on their commitments.

While motivation strongly influences commitment, we observed that group size also affected the follow-through rate—fewer participants showed up when there was a larger committed group. This result corroborates previous research relating larger group sizes with decreased motivation [45]. Because each peer office hour or study group session only involved 4 to 6 students, each student was highly likely to show up. On the other hand, for the human chess game, only 70% of committed participants showed up (16/23). In this case, the event's minimum threshold was just eighteen people, and the initial gathering of sixteen attracted a few interested onlookers, so the event had a sufficient buffer to still succeed; even a small group of people engaging in some action is sufficient to induce a large proportion of passers-by to also do so [34]. In the future, Catalyst might predict the number of follow-throughs for a given minimum threshold to help event organizers ensure a sufficient buffer.

Our post-deployment survey suggests that Catalyst was also effective in motivating participants to sign up or commit, as well as follow through on their commitment. We compared participants' self-reported likelihood of commitment and followthrough between common recruiting methods such as Facebook events, flyers, an in-person request (our gold standard), and Catalyst (Figure 5). A Friedman test on the Likert-scale responses revealed a significant effect of the recruiting method on both the likelihood of committing to an event ($\chi^2(4) = 90.9,\ p < 10^{-16}$) and the likelihood of following through on that commitment ($\chi^2(4) = 82.5,\ p < 10^{-16}$). Post-hoc Wilcoxon tests demonstrated that participants were

more likely to sign up for, and participate in events on Catalyst, compared to being invited on Facebook or after seeing a flyer or mass mailing ($p < 0.05$). As expected, in-person invitations would be still the most effective at getting people to participate ($p < 0.01$), possibly because "personal requests take top priority" (P3), and "in person...[one doesn't] want to hurt that person's feelings" (P20). Thus, Catalyst is more effective than other popular means of organizing events online, but less so than personal invitations. To overcome possible response bias in our results, future work could compare signups for the same event on Catalyst, Facebook and in-person.

Roles supported complex coordination

Organizers used role-based participation for the more complex event setups. For example, in peer office hours, the organizer used multiple time slots and a role-based threshold such that students could only join a specific time if at least there was at least one peer facilitator who had already signed up for that slot. Similarly, in the online class discussion sections, students could sign up as moderators to lead a Google Hangout, or to sign up as a participant to attend one.

Role-based thresholds were especially applicable to the human chess game, which required 16 different people to play 16 different chess pieces, and to peer office hours, which required a peer facilitator and 4 students — "I couldn't see it any other way" (O5).

While every event that employed roles succeeded, some organizers had reservations about using roles: while "it makes planning automation easier...there's no way to know if...role[s] encouraged or discouraged people" (O1). Future work could study the effect roles have on commitment levels and on creating artificial scarcity.

Long-running events initially succeeded, then struggled

On Catalyst, recurring events tended to stop recurring: events like study groups and class studios were designed around weekly, repeating events, where participants had to commit to attending each. However, sustaining causes is difficult unless such they demonstrate very tangible benefits to participants [7]. While the initial instantiations of such events were successful, subsequent instances saw significantly less commitment: "we had attendance for the first [discussion session], but then it just fell apart." (P8)

Nonetheless, several individual sessions in a massive online course saw sustained participation through the end of the deployment. Though there was a drop-off in attendance, users that stuck around continued to do so for several weeks. In one session, though there were 7 participants in the first week, 4 continued to participate through the next three weeks. This suggests that utilizing small-group commitment may be useful for long-running events.

Reminders played a significant role in ensuring the success of some long-running events: "A reminder 1 hour before the event would definitely help me not forget it" (P16). For the weekly social teatime meet-ups, not only was the activation threshold low, but weekly emails were sent out to remind participants to join in. Where reminders were infrequent or not used, participant attendance decreased week after week.

Motivation is the primary factor driving participation

While activation thresholds are useful in signaling quality and deferring commitment, the actual acts of committing and attending an event are still largely dependent on a participant's motivation. Many participants echoed the sentiment that "regardless of [how] I was invited," (P15, P16), "I will go only if I want to" (P15) or "when I am interested" (P9).

Nevertheless, feedback suggested that activation thresholds may be the most persuasive when participants are on the fence, unsure whether or not to participate. In these situations, their decision could be strongly influenced by peer and social signals like the number of other people going [8].

Similarly, events that conferred larger and more specific benefits to participants (e.g., study groups, office hours and class studios) saw higher success rates and conversion rates than those where the benefit of participation was less clear (e.g., flash mobs). While participants joined class studios "to get feedback on my project from like-minded people" (P6) and peer office hours because of the "good, focused attention from the professor" (P9), flash mobs "were kind of stupid" (P17).

Deployment limitations

Participant feedback was primarily positive, and our deployment allowed us to understand how Catalyst was able to support several types of causes using different activation thresholds, Nonetheless, it was not a field experiment. Experiments that control for whether the activation threshold is shown, or compare participation in similar events across different web sites such as Facebook or Meetup.com, would help quantify the effectiveness of activation thresholds. Though our deployment lasted several months, we also cannot extrapolate to how the system would be used if it gained further popularity.

DISCUSSION

In our three-month initial field deployment, participants adopted Catalyst to mobilize hundreds of people for causes ranging from education to social gatherings and fun. Here, we reflect on open challenges and opportunities for activation thresholds and for Catalyst in particular.

Enforcing followthrough

Activation thresholds generalize two components of crowdfunding: 1) the *form* of the commitment (from a monetary pledge to participation), and 2) the *execution* of that commitment (from financial escrow to a social contract for followthrough). The challenge is that a social contract may be more difficult to enforce. Some of crowdfunding's success can be attributed to the hardness of a financial commitment, which discourages participants from reneging on their promises to fund projects. Catalyst makes an explicit design decision *not* to use hard commitments such as money, and instead to focus on social commitments to encourage followthrough. In this sense, Catalyst has a strictly harder problem to solve: it not only needs to raise interest, it also needs to encourage follow-through on that commitment.

We experimented with allowing organizers to require monetary commitment similar to Kickstarter. However, the conversion rates for events with monetary commitments was sig-

nificantly lower, with few to no people signing up for them. Such commitment devices may instead be appropriate for larger events, where demand far outstrips availability. Alternatively, reputation systems like on eBay or Stack Overflow could help build trust among organizers and participants in the long term[44].

Slacktivism or activism?

While many of the flash mobs organized never succeeded, they gained a substantial number of Facebook likes (19, 12 and 7). Many people supported the idea on Facebook, but most of them did not commit to attending the event. The result is an intermediate, "slacktivist" level of commitment on Catalyst. This suggests that when faced with a decision about how much to support a cause, most people will pick the path of least resistance. Rather than joining in a demonstration, most people would rather support the demonstrators [29].

Allowing for weaker, more tentative commitment with the ability to later escalate such commitment to actual participation may thus be more effective in encouraging slacktivists to participate: "I wish Catalyst had a way to indicate that people were 'maybe' attending...people would be more comfortable responding...see who else was coming, then change their mind" (O3). We hypothesize that this approach might decrease follow-through, but the concomitant increase in participation might be enough to offset it. Another approach might allow the less committed individuals on Facebook — who may even be remote and unable to participate — to leave messages for and motivate the committed users. Participation could also be motivated with appropriate task or role recommendations (e.g. SuggestBot on Wikipedia [16]).

Controlling for participant quality

Many volunteer organizations face the challenge of retaining volunteers in the long-term, as well as ensuring that participants who signed up are committed to carrying out the tasks they were assigned [51]. Several volunteer organizations expressed interest in Catalyst but were concerned that its framing around single events could not yet support the kind of long-term commitment they require.

Nevertheless, we believe Catalyst embodies an interesting middle ground between online signups and offline volunteerism. While organization online is highly ad-hoc, and more about accomplishing a task than about forming social bonds [52], Catalyst seems to support a combination of both ad-hoc gatherings and bond-building. While intrinsic motivation plays a large factor in participation, especially in the case of recurring events, participants reported they were more likely to commit via Catalyst than through most other mechanisms. Participants who signed up to volunteer at the local food bank generally followed through, and animated, long-lived discussions sprang up in the online discussion sessions.

Limitations

While Catalyst is useful for predetermined events, it lacks the "flexibility afforded by email" (O1) or a Doodle poll, where there is more room for negotiation and the balance of power between all stakeholders more equitable. Further, though requiring separate commitments for each event in a recurring series potentially strengthens one's resolve to attending each, it can also discourage the initial commitment and hence the success of the event. And like Kickstarter, the organizer is solely responsible for the specifics of an event including the setting of appropriate thresholds and roles. This also means a large amount of a cause's success is dependent on the organizer, like in other collective action systems. Kickstarter projects which lacked video introductions or had spelling errors were significantly less likely to be funded [35]. Despite the low barrier to entry, creators reported that organizing and marketing successful, large-scale Kickstarter campaigns was significantly harder than they believed [24]. Still, such "front loading", where the goals and details of a project are specified upfront, has been shown to be more effective at gathering contributors than projects which are more spontaneous [31].

"Let a hundred flowers bloom"

While many Catalyst events reached their activation threshold, many others did not. The organizers may have been disappointed, but the fact that Catalyst encouraged failed events is a kind of success. Likewise, YouTube is valuable not just for its popular videos, but also for its publish-then-filter model that encourages failures as well as out-of-the-box successes. In this way, Catalyst leads to many failures as a side effect of generating more collective action successes. Akin to design prototyping, the low cost of failure suggests that Catalyst potentially supports *event prototyping*, where organizers create and iterate on several potential events.

Most research has focused on models for collective action and case studies of successful causes, rather than the identification and analysis of failed efforts [42]. However, many of Grudin's exhortations for groupware developers are relevant here [22]; often, the organizer's view of an event's latent interest did not match those of the potential participants. But even then, some events that did not reach their activation thresholds still succeeded — some online discussion sessions that had only six participants instead of ten still ended up being organized, subsequently generating lively discussion.

We believe that it will be important to explore how systems can help organizers predict, understand and learn from failure. For example, often organizers were left wondering whether an event failed because it received insufficient exposure, or if it succeeded in reaching people but failed because no one was interested. Analytics like page views or historical event data could help organizers design events that appeal to larger audiences and create realistic thresholds to improve the chances of future events succeeding.

CONCLUSION AND FUTURE WORK

In this paper, we introduce *activation thresholds* for collective action. Catalyst embodies this idea by allowing participants to condition their commitment on others' participation. A three-month field deployment demonstrated that Catalyst could organize hundreds of individuals to gather for small-scale events ranging from on-demand tutoring, volunteerism and social events. Future work lies in three main areas. First, many events require more thorough consideration of activation thresholds, including events with multiple stages and

long-term commitments. Second is further analysis of how to encourage participants to commit to new events and follow through once they do commit. Third, it will be important to quantify the effect of activation thresholds in comparison to other commitment mechanisms.

Our work suggests that social computing design can help support many new classes of events. Could Catalyst have helped organize Wikipedia or Tahrir Square? In its current form, probably not. However, its vision of low-risk collective efforts will lower barriers. Online classes could be kickstarted based on the interests of students; local citizens could gather to more actively improve their communities; major movements may even take root. The web has already lowered the threshold for connecting like-minded people — we envision its potential for likewise catalyzing real-world action.

ACKNOWLEDGMENTS

This work is supported by NSF Grants XXX and a Stanford Graduate Fellowship.

REFERENCES

1. Kickstarter. `http://www.kickstarter.com`.
2. Kickstarter stats. `http://www.kickstarter.com/help/stats`.
3. U.S. census bureau genealogy data. `http://www.census.gov/genealogy/www/data/index.html`.
4. We the people. `http://petitions.whitehouse.gov`.
5. Zokos. `http://www.zokos.com`.
6. Arkes, H. R., and Blumer, C. The psychology of sunk cost. *Organizational behavior and human decision processes* (1985).
7. Backman, E. V., and Smith, S. R. Healthy organizations, unhealthy communities? *Nonprofit Management and Leadership* (2000).
8. Banerjee, A. V. A simple model of herd behavior. *The Quarterly Journal of Economics* (1992).
9. Becker, G. S., and Murphy, K. M. The division of labor, coordination costs, and knowledge. *The Quarterly Journal of Economics* (1992).
10. Bédoyan, I., Aelst, P. V., and Walgrave, S. Limitations and possibilities of transnational mobilization: the case of eu summit protesters in brussels, 2001. *Mobilization: An International Quarterly* (2004).
11. Beenen, G., Ling, K., et al. Using social psychology to motivate contributions to online communities. In *Proc. CSCW* (2004).
12. Benkler, Y. Coase's penguin, or, linux and" the nature of the firm". *Yale Law Journal* (2002), 369–446.
13. Bryant, S. L., Forte, A., and Bruckman, A. Becoming wikipedian: transformation of participation in a collaborative online encyclopedia. In *Proc. GROUP* (2005).
14. Byers, J. W., Mitzenmacher, M., Potamias, M., and Zervas, G. A month in the life of groupon. *CoRR* (2011).
15. Cialdini, R. B. *Influence: Science and practice*. 2001.
16. Cosley, D., Frankowski, D., Terveen, L., and Riedl, J. Suggestbot: using intelligent task routing to help people find work in wikipedia. In *Proc. IUI* (2007).
17. Das, E., Kerkhof, P., and Kuiper, J. Improving the effectiveness of fundraising messages: The impact of charity goal attainment, message framing, and evidence on persuasion. *Applied Communication Research* (2008).
18. Donath, J. S. Identity and deception in the virtual community. *Communities in cyberspace* (1999).
19. Forte, A., Kittur, N., et al. Coordination and beyond: social functions of groups in open content production. In *Proc. CSCW* (2012).
20. Giné, X., Karlan, D., and Zinman, J. Put your money where your butt is: a commitment contract for smoking cessation. *American Economic Journal: Applied Economics* (2010).
21. Granovetter, M. Threshold models of collective behavior. *American Journal of Sociology* (1978).
22. Grudin, J. Groupware and social dynamics: eight challenges for developers. *Communications of the ACM* (1994).
23. Hill, B. M. Almost wikipedia: What eight collaborative encyclopedia projects reveal about mechanisms of collective action. 2011.
24. Hui, J., Greenberg, M., and Gerber, E. Understanding crowdfunding work: implications for support tools. In *Proc. CHI Extended Abstracts* (2013).
25. Khondker, H. H. Role of the new media in the arab spring. *Globalizations* (2011).
26. Kittur, A., and Kraut, R. E. Harnessing the wisdom of crowds in wikipedia: quality through coordination. In *Proc. CSCW* (2008).
27. Klandermans, B. Mobilization and participation: Social-psychological expansisons of resource mobilization theory. *American Sociological Review* (1984).
28. Kriplean, T., Morgan, J., et al. Supporting reflective public thought with considerit. In *Proc. CSCW* (2012).
29. Lee, Y.-H., and Hsieh, G. Does slacktivism hurt activism?: the effects of moral balancing and consistency in online activism. In *Proc. CHI* (2013).
30. Lou, H., Luo, W., and Strong, D. Perceived critical mass effect on groupware acceptance. *European Journal of Information Systems* (2000).
31. Luther, K., and Bruckman, A. Leadership in online creative collaboration. In *Proc. CSCW* (2008).

32. Luther, K., Fiesler, C., and Bruckman, A. Redistributing leadership in online creative collaboration. In *Proc. CSCW* (2013).

33. Marwell, G., Oliver, P. E., and Prahl, R. Social networks and collective action: A theory of the critical mass. iii. *American Journal of Sociology* (1988).

34. Milgram, S., Bickman, L., and Berkowitz, L. Note on the drawing power of crowds of different size. *Journal of Personality and Social Psychology* (1969).

35. Mollick, E. The dynamics of crowdfunding: Determinants of success and failure. *Journal of Business Venturing* (2013).

36. Muller, M., Geyer, W., et al. Crowdfunding inside the enterprise: employee-initiatives for innovation and collaboration. In *Proc. CSCW* (2013).

37. Okolloh, O. Ushahidi, or'testimony': Web 2.0 tools for crowdsourcing crisis information. *Participatory learning and action* (2009).

38. Oliver, P., Marwell, G., and Teixeira, R. A theory of the critical mass. i. interdependence, group heterogeneity, and the production of collective action. *American Journal of Sociology* (1985).

39. Oliver, P. E., and Marwell, G. Mobilizing technologies for collective action. *Frontiers in social movement theory* (1992).

40. Olson, M. *The logic of collective action: public goods and the theory of groups.* 1965.

41. Pickard, G., Rahwan, I., et al. Time critical social mobilization: The darpa network challenge winning strategy. *CoRR* (2010).

42. Poteete, A. R., Janssen, M., and Ostrom, E. *Working together: collective action, the commons, and multiple methods in practice*. Princeton University Press, 2010.

43. Prasarnphanich, P., and Wagner, C. Creating critical mass in collaboration systems: Insights from wikipedia. In *DEST* (2008).

44. Resnick, P., and Zeckhauser, R. Trust among strangers in internet transactions: Empirical analysis of ebay's reputation system. *Advances in applied microeconomics* (2002).

45. Ringlemann, M. Recherches sur les moteurs animés: Travail de lhomme. In *Annales de lInstitut national Agronomique* (1913).

46. Ross, L., and Nisbett, R. *The Person and the Situation: Perspectives of Social Psychology*. 1991.

47. Salganik, M. J., Dodds, P. S., and Watts, D. J. Experimental study of inequality and unpredictability in an artificial cultural market. *Science* (2006).

48. Starbird, K., and Palen, L. "voluntweeters": self-organizing by digital volunteers in times of crisis. In *Proc. CHI* (2011).

49. Suh, B., Convertino, G., Chi, E. H., and Pirolli, P. The singularity is not near: slowing growth of wikipedia. In *Proc. WikiSym* (2009).

50. Verhallen, T. M., and Robben, H. S. Scarcity and preference: An experiment on unavailability and product evaluation. *Journal of Economic Psychology* (1994).

51. Voida, A., Harmon, E., and Al-Ani, B. Bridging between organizations and the public: volunteer coordinators' uneasy relationship with social computing. In *Proc. CHI* (2012).

52. Wang, L. S., Chen, J., Ren, Y., and Riedl, J. Searching for the goldilocks zone: trade-offs in managing online volunteer groups. In *Proc. CSCW* (2012).

APPENDIX A — POST-DEPLOYMENT SURVEY

The post-deployment survey was emailed to all event organizers, and participants who had provided their email address when signing up for an event on Catalyst. Demographic information was not collected. Questions that appeared in the survey sent to organizers are prefixed with an "O", and those sent to participants with a "P".

- General questions (open-ended)
  - (O) About how many people showed up to the events you organized? How many were you expecting?
  - (O) How did you publicize the event you organized?
  - (P) There are many kinds of events that people want you to participate in: some aren't that interesting to you, some are moderately interesting, and some are extremely interesting. Would you respond differently via Facebook (where you RSVP), via Catalyst (where you are only committed if many other people also commit), or in person (where you're telling a friend if you'll go)?
  - (O, P) What did you like about Catalyst? What didn't you like?
- (O) Understanding thresholds and roles (open-ended)
  - How did creating an event on Catalyst compare to simply announcing it (via email, Facebook, flyers)? Were there differences in how and when you made preparations for the event?
  - Catalyst also allowed you to specify minimum or maximum thresholds (i.e. this event would only happen if 5 people join, or this event can have at most 10 people). Were those useful?
  - Catalyst allowed the use of roles in your event. For example, one for a bake sale may have roles for bakers and sellers. Did you use roles in your event? Were they useful? Why or why not?
- (O, P) Comparing Catalyst to other platforms for organizing events (7-point Likert scale)
  - How likely would you be to agree to attending an event you were somewhat interested in, if you were invited to it using Facebook?
  - How likely would you be to agree to attending an event you were somewhat interested in, on Catalyst?
  - How likely would you be to agree to attending an event you were somewhat interested in, if a friend invited you in person?
  - How likely would you be to agree to attending an event you were somewhat interested in, and found out about through an email list/flyer/the local paper?
  - Supposing that you've agreed to attend an event on Facebook, how likely would you be to actually show up to it?
  - Supposing that you've agreed to attend an event on Catalyst, how likely would you be to actually show up to it?
  - Supposing that you've agreed to attend an event a friend invited you to in person, how likely would you be to actually show up to it?
  - Supposing that you've agreed to attend an event you found out about through an email list/flyer/the local paper, how likely would you be to actually show up to it?