

I study collective action in online communities and seek to understand why some attempts at collaborative production – like Wikipedia and Linux – build large volunteer communities while the vast majority never attract even a second contributor. I am particularly interested in how the design of communication and information technologies shape fundamental social outcomes with broad theoretical and practical implications – like the decision to join a community or contribute to a public good. My research is deeply interdisciplinary, consists primarily of “big data” quantitative analyses, and lies at the intersection of sociology, communication, and human-computer interaction.

Using Internet-based peer production projects as my research settings, my work seeks to understand the conditions for collective action using observational data from real communities. This work has been shaped by three complementary approaches: (1) the comparison of failures to build communities to rare successful attempts; (2) attention to the role of reputation and status in the mobilization of volunteers; and (3) analysis of design changes as “natural experiments” building a deeper, and often causal, understanding from observational data. Nearly all of my work incorporates at least two of these approaches.

### STUDYING ATTEMPTS AT COLLECTIVE ACTION

Although there have been many thousands of studies of online collective action, the vast majority have only considered projects like Wikipedia and Linux that have successfully built communities – a characterization that can be extended to observational work on collective action more generally. In this sense, most previous analyses have systematically selected on their dependent variable. Instead, most of my research treats projects as the unit of analysis and collective action as the outcome of interest – comparing the successful examples of collective action to attempts that never got off the ground.

In one study, I compare Wikipedia to seven attempts to create online collaborative encyclopedia projects that were launched previously [1]. Using an inductive, grounded-theory based analysis of founder interviews and archival data, I propose four hypotheses to explain why Wikipedia attracted many more contributors. Although the paper’s methods diverge from the quantitative, “big data” approach typical of most of my work, the research question and strategy is representative.

I have also followed this strategy in a series of quantitative studies of the Scratch online community: a public website where millions of users create, share, and remix interactive media. The community is built around the Scratch programming environment: a freely downloadable desktop application that allows amateur creators to combine media with programming code (see Figure 1). Although Scratch is a community designed to promote collaboration through content remixing, only about ten percent of Scratch projects attract a second contributor.

In one study, co-authored with Andrés Monroy-Hernández and forthcoming in *American Behavioral Scientist*, I test several of the most widely cited theories associated with “generativity” (i.e., qualities of technology or content that make some works more fertile ground for collaboration). I find some support for existing

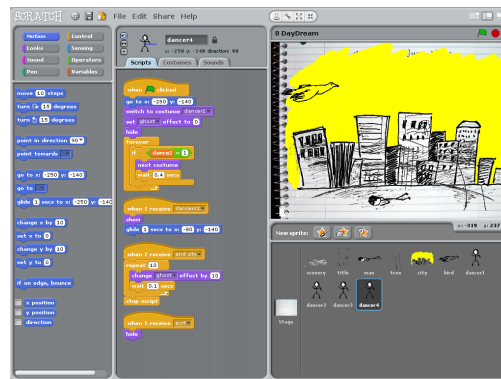


Figure 1: A screenshot of the Scratch programming environment where users create animations and interactive games.

theory but also find that, across the board, factors associated with more collaboration are also associated with less original and transformative types of joint-work [2]. In another study of Scratch written with Monroy-Hernández and Kristina Olson, I show that this type of superficial collaboration leads to negative reactions and community displeasure [3].



Figure 2: The front page of the Scratch online community where users can share and collaborate on projects.

## REPUTATION AND STATUS

Although empirical research comparing successful and unsuccessful peer production projects has been rare, theories have been widespread. No theory has been more influential than the suggestion that, in the absence of pecuniary rewards, contributions to online public goods are driven by the possibility of increased reputation and status for contributors.

In a study of status-based awards in Wikipedia called “barnstars” (see Figure 3) – a collaboration with Aaron Shaw and Yochai Benkler – I provide an empirical test of an influential status-based theory of collective action from sociology. Although the study finds support for a widely hypothesized “virtuous cycle” in which status rewards both cause and are caused by contributions, it also finds that this effect is limited to a sub-population of Wikipedia contributors – “signalers” who show off their awards [5]. This result has broad implications for both status-based theories of collective action as well the design of reputation-based rewards.

In a mixed methods study of Scratch, written with a team at Microsoft Research and nominated for best paper at the CHI 2011 conference [6], I present both a quantitative analysis of a design change and in-depth interviews of users to demonstrate how credit-giving is ineffective when it stems from an automated system because systems fail to reinforce status-ordering with credible human expressions of social deference and gratitude. These studies suggest important limits to previous theoretical work on status as a motivator for collective action, and put forward a more nuanced theoretical model.

This year, I am conducting a population-level analysis in a new dataset I have created that includes 80,000 attempts at wikis (i.e., public, editable, websites similar to Wikipedia). In my first working paper using this dataset, I consider inter-organizational competition for volunteer labor and find little support for a widely cited ecological model of collective action from sociology that treats volunteer labor as a fixed and finite resource. Instead, I show that contributions to different wikis on the same topic or theme are driven primarily by environment-level changes in interest and that projects can even benefit from complementarities and synergies [4]. By looking at failures, these studies provide tests of several of the most influential theories of the conditions for collective action, suggest important practical and theoretical limitations to existing models, and point to previously untheorized mechanisms.




Figure 3: Image of a “barnstar” social award given by Wikipedia contributors to each other to recognize positive contributions.

## DESIGN-DRIVEN NATURAL EXPERIMENTS

Although nearly all of my work has important implications for the design of socio-technical systems, I have structured much of my work around the evaluation of technological design changes. In several papers, I treat design changes as “natural experiments” that exogenously change the ways that social structure is enacted.

For example, to evaluate the impact of status-based incentives and collaboration in Scratch, I use a regression discontinuity framework to measure the causal effect of increased status for collaboration [7]. I show that highlighting collaborative projects on the Scratch web page (see the bottom of Figure 2) resulted in more collaboration but also caused a decrease in the amount of total effort exerted by contributors. Speaking to fundamental sociological work in the literature on collective action, I present evidence that this decrease is driven by both an the influx of new contributors and a decrease in the effort and contributions of established participants.

In other work with Leah Buechley, I have analyzed sales records of hobbyist microcontrollers to argue that relatively simple design changes in the *LilyPad*

*Arduino* – a electronics toolkit minimally re-designed for women and girls (see Figure 4) – led to large increases in the proportion of women contributors and drastic shifts in the type of projects created [8]. I have also explored how technical errors may be able to provide similar opportunities for analysis by interrupting normal operation of a system and revealing internal processes that are usually hidden [9]. In addition to the important theoretical findings in these studies, this type of work represents an important methodological advance in that it allows for stronger causal claims while also closing the gap between theory and design.




Figure 4: A image of the LilyPad Arduino microcontroller.

## RESEARCH AGENDA

My research agenda involves further exploration of the determinants of collection action online – especially using a series of large new datasets I have recently assembled. I plan to both continue on this research trajectory and to create new social and technical infrastructure that will allow others researchers to join me in “big data” observational research with active communities. This section outlines some future directions I plan to explore.

*Understanding the Relationship Between Collective Action and Performance* – My work has treated collective action and production as ends in themselves and has largely avoided the consideration of issues of performance, efficiency, and quality. Using my existing datasets, I plan to compare the performance of collaborative production to individually produced works to understand when successful collection action leads to increased performance. For example, in an analysis using data from Scratch which is currently under review – done in collaboration with Monroy-Hernández – I show important limitations of collaboration through remixing in regards to project quality, particularly for more artistic or media-intensive works [10].

*Integrated Theory of Design for Collective Action* – My studies of status and reputation provide a detailed understanding of the dynamics of collection action in relation to one set of important predictors. In future work, I plan to evaluate the effect of governance and different systems of authority, framing, modularity and project complexity. In the long term, I hope to offer a broad set of principles of design for online collection action.

*Toolkits for Experimental Social Design* – My research has been possible through personal relationships I have with a series of organizations with large, active, online communities (e.g., the MIT Media Lab and

the Wikimedia Foundation). These organizations, like many others, make design changes to the software that supports their communities to encourage contributions and improve their users' experiences. Most of the time, these organizations have very little idea if these changes are effective. I plan to seek funding for, and to create, a technical framework and a network of academic and practitioner collaborators to facilitate well-designed natural experiments by the hosts of large online communities and to share data that allows for academic evaluation of these experiments.

Although I study cooperation, I also practice it. In graduate school, I have collaborated with a large group of co-authors in many academic departments. I intend to continue doing so. In sum, my research uses design to contribute to social scientific theories of collective action, and uses theories of collective action to influence design. Although my research settings are online communities, I believe my work has implications for a broad range of disciplines and fields.

## REFERENCES

- [1] Benjamin Mako Hill. Almost wikipedia: What eight early online collaborative encyclopedia projects reveal about the mechanisms of collective action. Working Paper, 2012.
- [2] Benjamin Mako Hill and Andrés Monroy-Hernández. The remixing dilemma: the tradeoff between generativity and originality. *American Behavioral Scientist*, 2012. Forthcoming.
- [3] Benjamin Mako Hill, Andrés Monroy-Hernández, and Kristina Olson. Responses to remixing on a social media sharing website. In *Proceedings of the 4th AAAI Conference on Weblogs and Social Media*, pages 74–81, Washington, D.C., 2010. AAAI.
- [4] Benjamin Mako Hill. Is volunteer labor a "fixed and finite" resource? evidence from peer production. Working Paper, 2012.
- [5] Benjamin Mako Hill, Aaron Shaw, and Yochai Benkler. Status, social signalling and collective action in a peer production community. Working Paper, 2012.
- [6] Andrés Monroy-Hernández, Benjamin Mako Hill, Jazmin Gonzalez-Rivero, and danah boyd. Computers can't give credit: how automatic attribution falls short in an online remixing community. In *Proceedings of the 2011 annual conference on Human factors in computing systems*, pages 3421–3430, Vancouver, BC, Canada, 2011. ACM.
- [7] Benjamin Mako Hill. Causal effects of a reputation-based incentive in an peer production community. Working Paper, 2012.
- [8] Leah Buechley and Benjamin Mako Hill. LilyPad in the wild: how hardware's long tail is supporting new engineering and design communities. In *Proceedings of the 8th ACM Conference on Designing Interactive Systems*, pages 199–207, Aarhus, Denmark, 2010. ACM.
- [9] Benjamin Mako Hill. Revealing errors. In Mark Nunes, editor, *Error: Glitch, Noise, and Jam in New Media Cultures*. Continuum, 2010. Feature.
- [10] Benjamin Mako Hill and Andrés Monroy-Hernández. The cost of collaboration for code and art: Evidence from a remixing community. In *Proceedings of the 2013 ACM conference on Computer-supported cooperative work*, San Antonio, Texas, 2013. ACM. Forthcoming.