Let's Get Together: The Formation and Success of Online Creative Collaborations

Burr Settles
Machine Learning Department
Carnegie Mellon University
Pittsburgh PA, 15213
bsettles@cs.cmu.edu

Steven Dow
HCI Institute
Carnegie Mellon University
Pittsburgh PA, 15213
spdow@cs.cmu.edu

ABSTRACT
In online creative communities, members work together to produce music, movies, games, and other cultural products. Despite the proliferation of collaboration in these communities, we know little about how these teams form and what leads to their ultimate success. Building on theories of social identity and exchange, we present an exploratory study of an online songwriting community. We analyze four years of longitudinal behavioral data using a novel path-based regression model that accurately predicts and reveals key variables about collab formation. Combined with a large-scale survey of members, we find that communication, nuanced complementary interest and status, and a balanced effort from both parties contribute to successful collaborations. We also discuss several applications of these findings for socio-technical infrastructures that support online creative production.

Author Keywords
Collaboration; computer-supported cooperative work; music composition; online communities; social creativity.

ACM Classification Keywords
H.5.3. Information Interfaces and Presentation (e.g. HCI): Group and organization interfaces.

General Terms
Experimentation; Theory.

INTRODUCTION
Advances in information technology and low-cost production tools have led to a surge in "peer production" [9]. The creative potential of online collaboration through peer production has been realized in many forms, from epic cultural efforts such as Wikipedia, to open-source software, to smaller communities that nourish more specific artistic pursuits. For example, at Newgrounds — an online community of animators, and one of the most heavily-trafficked sites on the Internet — millions of registered users have created hundreds of thousands of animated movies and online games [2]. Likewise, nearly three million interactive media projects have been created and shared in the Scratch online community [3]. February Album Writing Month (FAWM) — a hub for songwriters who aim to write an entire album during the shortest month of the year — has helped musicians collectively write 50,000 new original works [1]. Many of these were co-created by near-strangers who are spread across the globe.

CHI 2013, April 27–May 2, 2013, Paris, France.

Working with others to achieve a shared goal can promote social, motivational, and emotional benefits. Collaboration has been shown to improve peer relationships, increase self-esteem, and develop perspective-taking skills among students in classroom settings [53]. Only recently have researchers been able to study collaboration in online peer production environments, investigating how teamwork, individual and group goals, and communication affect one another. For example, Burke and Settles [12] found that newcomers to the FAWM music community who engage in one-on-one "collabs" during their first year are not only more successful at reaching their own personal songwriting goals, but also go on to behave in more community-favorable ways (e.g., giving feedback on others' music, or donating money to the site). This suggests that collaborative efforts within these communities can lead to improved outcomes for both individuals and the group as a whole. With a better understanding of the factors that affect online creative collaboration, we can develop new social tools, technologies, and best practices that help online communities and their members to flourish. To that end, our research explores the following question:

> *What factors influence (i) the formation and (ii) the ultimate success and satisfaction of an individual's online creative collaborations?*

Perhaps the work most relevant to this question is that of Luther et al. [36, 37], who examined the role of leadership in large-scale animation projects on the Newgrounds website. They found that planning, reputation, and communication were key in actually realizing a proposed collab. While they focus on success in large distributed group productions, we consider both formation and success in collabs between pairs of individuals. This paper presents an exploratory study of the FAWM online music community, combining quantitative results (based on a machine learning analysis of longitudinal behavioral data) with qualitative insights from member surveys. We find that communication, complementary interests and status, and a balanced effort are all key ingredients in collab formation and success, and discuss how these results can be leveraged to support online creative communities.

Theory About Team Formation and Success

To help frame our investigation, we review some of the literature on social networks, online gaming, and even conventional work environments. Theory on common bond [46], social identity [55], and social exchange [18] help lay a foundation for how communication, shared interests, community status, and balance of effort affect people's desire to collaborate as well as their ultimate success.

Communication Richness

Research suggests that communication plays a key role in how online relationships form and function. In a study of online newsgroups, McKenna et al. [41] found communication frequency to be a significant factor in how relationships develop, and in whether these relationships persist years later. Frequent interaction in community forums and backchannels helps members strengthen their bond and trust [46]. Some evidence suggests that even personal information on user profiles can promote ties among members who have yet to formally interact in online communities [59].

Communication not only helps relationships to form, but it improves working relationships in existing teams as well. For example, rich communication helps to support idea generation [54], creation of shared mental models [23], and critique exchange [51] in design teams. In such creative contexts, though, communication can carry emotional undertones that derail a group. Roschuni et al. [48] studied product development groups and found that high-performing design teams frequently discuss their interpersonal relationships and spend time negotiating conflicts to achieve successful outcomes. Furthermore, discussing multiple ideas can help mitigate the emotions that arise from critique [17].

In our work, we look at how various modes of communication influence the formation and success of online creative collabs.

Shared Interests

Social identity theory predicts that members of open communities are most likely to form collaborations with people who match their view of themselves [43, 55]. This manifests itself in many ways, as people tend to be attracted to others who appear similar [13], who share the same attributes [26], and who hold common attitudes [43]. Extending social identity, Feld's "focus theory" places emphasis on how groups and partnerships form around shared activities and tasks [21]. This is evident in face-to-face contexts such as exercise partners [61], as well as online game communities such as *EverQuest II*, where group formation is highly influenced by players' shared interest in pursuing quests [27]. Similarly, the shared ambition of developing the world's best and most complete encyclopedia provides focus for Wikipedia contributors to become committed to both the group and its goal [11].

While people tend to form relationships around homogeneous traits and interests, this does not necessarily lead to better collaborations. Seminal research by Hoffman et al. [25] showed that teams of heterogeneous personalities produced higher-quality solutions to complex problems than homogeneous teams, and similar results have been observed among pairs of individuals [57]. Since this early research, performance benefits in diverse groups have been demonstrated in information sharing [24], creativity and problem-solving ability [47], and successful management teams [8], to name a few. Heterogeneous teams are believed to increase the range of knowledge and variation in viewpoints brought to bear on different problems [39]. Some point to the presence of "creative conflict" to explain increased performance of functionally diverse groups in creative settings [33].

However, research also implies that diversity can negatively affect team communication and performance [39]. In a study of product teams in technology companies, Ancona and Caldwell [6] conclude that team diversity contributes positively to creative problem solving, but impedes implementation due to poor teamwork and conflict resolution. Heterogenous team members tend not to like each other as much, and experience high turnover as a result, especially when conflict arises [44]. Van Der Vegt and Bunderson suggest that a shared group identity seems to mediate these effects; they found that teams with a strong group identity overcame the communication breakdowns and trust issues that arise from heterogeneous members [58]. Mannix and Neale [39] suggest that while surface-level social-category differences (such as race and gender) tend to negatively affect groups, the deeper underlying differences in education and skill tend to be positively associated with cooperative performance.

The models in our study consider a range of variables that represent collaborators' overlapping interests, thus shedding more light on the nuanced effects of homogenous and diverse partnerships on collab formation and success.

Status Within the Community

The formation of online relationships is also affected by the status of members in the community. Previous work has found that people are more likely to form relationships with people who share connections in their existing social network [7], in part because people feel more trust towards people who have a shared acquaintance [62]. In other words, people are drawn to friends of their friends.

People also tend to collaborate when both parties think they can gain personally in the social exchange. In this view, self-interest is not necessarily unhealthy, and can actually enhance relationships [18]. For example, in online role-playing games, players with less combat experience are more likely to join collaborative guilds with more experienced players, as this helps them to "bootstrap" their personal achievement [27]. Novices learn both implicitly and explicitly by working alongside more experienced partners [32].

In a study of Hollywood producers, directors, and cinematographers, Faulkner and Anderson [19] found that individuals tend to collaborate with people at their same level of achievement, both in terms of experience (film credits) and earnings (rental revenues). However, Cattani and Ferriani [14] found that individuals had broader creative impact (as measured in award nominations) if they formed collaborations with a heterogeneous mix of core and peripheral actors in the community. So, while people are more likely to collaborate with people of similar status, teams may be more effective if they contain a mix of "insiders" and "outsiders."

To explore this phenomenon further, our work examines the effects of community status variables on the formation and success of online creative collabs.

Balance of Effort

The amount of individual effort exerted in creative collaborations can affect team satisfaction. Equal contributions among members promotes a shared ownership of concepts, which is important for reaching a finished product [54]. However, the individual efforts of team members often fall out of balance. "Social loafing" — where an individual does not work as hard in the presence of others — was first examined in a series of experiments showing that people exerted less effort pulling a rope as a group than they did alone [29]. The mere perception of such imbalanced efforts have been shown to diminish group cohesiveness [35] and motivation [42].

In our work, we use both qualitative and quantitative methods to investigate how balance of effort (or, by contrast, social loafing) can affect collab success and satisfaction.

EXPLORATORY STUDY

Our research aims to better understand the factors that influence online creative collaboration. The present study explores how communication, shared interests, social status, and balance of effort affect collaboration in an online music community called February Album Writing Month.

February Album Writing Month (FAWM)

FAWM is an annual online music event for professional, semi-professional, and amateur songwriters [1]. The community tagline is "14 Songs in 28 Days," and it revolves around a challenge to compose at least 14 songs (roughly an album's worth) during the shortest month of the year. Since its inception in 2004, more than 7,000 participants from 30 countries have registered, collectively penning nearly 50,000 original pieces of music, including thousands of collaborations.

The main features of the site include user profiles, a jukebox of publicly-posted songs where participants can find, listen to, and comment on one another's music, and an open discussion forum. User profiles include optional musical influences and bios, links to completed songs, and a "soundboard" where others can post direct messages. Song pages include author-provided descriptive tags, primarily used to categorize songs by genre and instrument (e.g., "punk-rock" or "piano"), and community members can provide feedback at the bottom of each song page. Songs are searchable and browsable, and demo recordings are shuffled into the website jukebox. The "bulletin board" style forum contains thousands of topics on music recording, sources of inspiration, regional discussions, and "collaboration classifieds" where songwriters looking to collaborate can propose projects and team up.

Collaborative projects in the FAWM community date back to at least 2006, when three so-called "fawmers" joined forces to compose 14 songs each about different U.S. presidents (covering all 42 presidencies in history up to that point). The collection was later released as a triple-album project to much critical praise during the 2008 election season [49]. The trio then toured and performed at the esteemed South by Southwest (SXSW) Music Festival. This parallel, distributed-labor model of collaboration is reminiscent of open-source software and Newgrounds animation projects [37].

A more common mode of collaboration became popular in 2008, a leap year, when FAWM organizers jokingly upped the ante to "14½ Songs in 29 Days." Participants were encouraged (though not required) to co-write an extra half-song. This resulted in 252 documented collaborations, or 4.4% of the total musical output that year. The popularity of these pairwise collaborations have grown, comprising 7.8% of all songs posted to the website since FAWM 2009. A notable example is "Walkthrough," by fawmers *errol* and *pifie*. The song essentially outlines the steps to win the classic text-based computer game "Zork" set to ambient alternative rock music. The song went viral on the Internet and enjoys a certain notoriety among computer game enthusiasts [22].

Data and Methods

We combined four years of archival data from FAWM's web server logs with a survey of members for both quantitative and qualitative analysis. The archives include data and metadata from the years 2009–2012[1] for 6,116 users who activated their accounts, 39,103 FAWM songs posted to the site (3,047 of which are documented collaborations), and various links and interactions between them.

Through the FAWM Facebook group, Twitter feed, and official email list ($n \approx$ 5,000), we invited members to take part in a brief web survey about their collaboration experiences, and $n =$ 226 participated. The survey triaged participants into three branches, based on their answers to initial questions about their level of collaboration experience on the site. The *none* branch ($n =$ 45) included fawmers who had never collaborated on the site before, and were asked a few questions about why that was the case. The *single* branch ($n =$ 30) included fawmers who had participated in only one collaboration, and were asked more detailed questions about that experience. The *multiple* branch ($n =$ 151) included fawmers who had taken part in several collaborations. These participants were asked to think of two specific collaboration attempts — their *most* successful and *least* successful — and answer questions comparing and contrasting these two experiences. Questions in the single and multiple branches included 7-point Likert satisfaction scales about each collaboration, checkboxes to indicate the roles they played or tools they used, plus several open-ended responses, such as "Please describe how the collaboration started," and "Please describe your process of working together; how often did you communicate?" The songs and user pairs mentioned in the surveys were then mapped to the archival data for further analysis.

For quantitative exploration, we performed a series of regressions that predict collaboration formation and success outcomes as a function of variables that describe user pairs (details of these regressions are explained in subsequent subsections). We used a grounded approach to analyze the qualitative survey data, cross-referencing it with our quantitative findings to hypothesize about underlying dynamics that

[1] FAWM began official support for collaborations in 2009 (e.g., joint posting and indexing). Thus we focus on data since that time.

affect online creative collaboration. We note that survey-takers were generally more experienced and involved than the average fawmer ($p \ll 0.001$ for all): per year, they had written more songs ($\mu = 15.1$ vs. 8.2), given more comments (83.7 vs. 22.9), and participated in more collaborations (3.4 vs. 1.0). Because survey-takers were more active than average, their data is not necessarily a representative sample; therefore we mainly use it to illuminate quantitative patterns found in the more complete archival data.

HOW COLLABORATIONS FORM

To explore patterns that influence the formation of one-on-one creative collaborations within FAWM, we take an inductive, complex network analysis approach. Specifically, we adapt a machine learning model originally proposed for probabilistic reasoning over knowledge base graphs [31], and apply it to perform induction over the FAWM social network graph. The analysis presented here is essentially a *path-based logistic regression*: each data instance represents a pair of users $\langle \mathbf{A}, \mathbf{B} \rangle$, the predicted dependent variable is whether or not the pair posted a collaborative song to the website, and the independent variables are various kinds of paths that can connect users through the social network graph.

Figure 1 presents a small subset of the network to help illustrate our approach. Suppose the model is trying to predict whether user **A** will collaborate with user **B**. One way of connecting their nodes in the social graph is through the path **A** –follows→ **B**, which means that **A** has "subscribed" to **B**'s song feed (indicating that she is interested in his work). Another path is **A** ←messaged– **B**, meaning that she received a direct message from **B** on her profile. There are also longer paths, such as **A** –commented→ ♫ ←wrote– **B**, which means that **A** commented on one of **B**'s songs (song S2 in this case). We consider even longer and more complex paths, such as **A** –wrote→ ♫ –tag→ 🏷 ←tag– ♫ ←wrote– **B**, which means that both users have written songs tagged with a shared term (a good indication that they have shared interests in musical genre and style, or use similar instruments).

The network includes nodes derived from tables in the FAWM database such as users, songs[2], tags, forum topics, and the various kinds of links between them in the server logs (all timestamped). In order to discover patterns relevant to status in the community, we computed each user's eigenvector centrality [10] — a common measure of influence in social networks — using the network of communication edges. This is a real number between zero and one (inclusive), for which higher values imply greater social status. We rounded centrality values to the first significant digit and added nodes with corresponding edges to the network. In Figure 1, for example, **A** has centrality score 0.7 while **B** has 0.5. For each pair of centrality nodes, we added an edge representing the difference between them (e.g., $\Delta = 0.2$). To discover patterns relating to demographic diversity, we also added nodes and edges that compare and contrast self-reported age, gender,

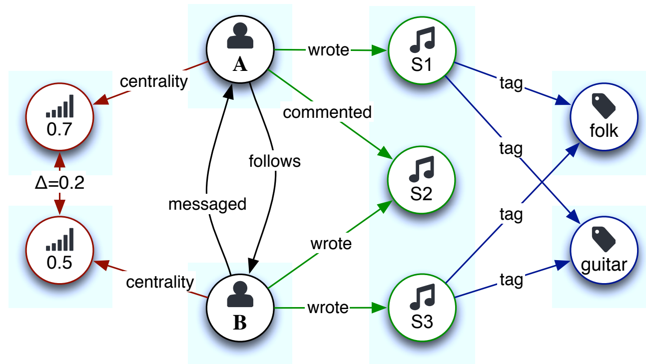

Figure 1. A small example subset of the FAWM network.

geographical location, and tenure in the FAWM community (in years, as measured by account creation date).

Our analysis treats each type of path through the network as an independent variable, whose value reflects the "strength" of that path in connecting the two users. For example, the aforementioned shared-tag path has two occurrences in Figure 1: one through the "folk" tag node, and another through "guitar." The model should recognize that this path type connects **A** and **B** more strongly than it would for two other users who share only one common tag. While it is theoretically possible to tabulate these frequencies for all possible paths connecting all possible pairs of users, that would be expensive in practice, and would not scale well to a network of our size and larger. Instead, we take a sampling approach based on "random walks," which have proven very useful for large-scale network analysis [5, 31, 45]. The algorithm begins at user node **A**, selects an edge at random to arrive at a new node, and repeats until reaching user node **B** for paths up to length four[3]. This process repeats for a finite amount of time, and the cumulative path statistics are normalized and used as the input variables that describe the pair $\langle \mathbf{A}, \mathbf{B} \rangle$.

We gather path statistics for all user pairs who posted a collaboration to the FAWM website, plus a large random sample of user pairs who did not. These pairs are assigned binary values of *collab* and *non-collab* (respectively) for the dependent variable. We then use these data to perform a logistic regression that predicts collaboration as a function of path statistics. We fit this model using the LASSO method [56], a popular model selection technique that induces sparsity in regressions with many variables. This helps to improve interpretability by eliminating statistically irrelevant variables, and also reduces over-fitting and prediction error by simplifying the model. In other words, LASSO attempts to automatically discard paths that fail to predict whether or not a collaboration will occur, leaving only path variables that have significant effects. For example, the 2012 FAWM network initially contained 2,258 unique path variables, of which only 165 were given nonzero weights by the model-fitting procedure, thus discarding 92.7% of the paths as insignificant.

[2]Caveat: we only use solo (non-collab) songs in this network. Otherwise, the model risks learning the path **A** –wrote→ ♫ ←wrote– **B**, the simple "definition" of a collaboration, as a significant predictor. Our goal is to uncover more interesting relationships than this.

[3]We have experimented with longer path lengths, but found the length-four models to be just as accurate, while remaining (i) computationally simpler and (ii) more easily interpretable.

Model Evaluation and Comparison

Before analyzing the weights in our model, we first show that the path-based regression does a good job of predicting collaboration behavior, and compare it to several established link-prediction methods. First, we compare against a baseline random walk without logistic regression [30], which scores the potential collab $\langle \mathbf{A}, \mathbf{B} \rangle$ according to how frequently a random walk from user node **A** terminates at user node **B**. Second, the Adamic & Adar heuristic [4] is a similarity measure commonly used to predict co-authorship in academic collaboration networks [34]. It scores the pair $\langle \mathbf{A}, \mathbf{B} \rangle$ according to a weighted sum over their shared neighbors, for which we use all neighboring nodes in the FAWM graph: songs, centralities, other users (via communication edges), and so on. Third, we employ a matrix factorization approach, which has been shown to be highly effective for recommender systems [28] and for predicting academic collaborations as well [34]. We represent the entire FAWM graph as an adjacency matrix, and map users to a common lower-dimensional projection using singular value decomposition (SVD). The pair $\langle \mathbf{A}, \mathbf{B} \rangle$ is scored according to the dot-product of the users' vector representations in the lower-dimensional space.

For each of the four years 2009–2012, we take a snapshot of the FAWM network after the first two weeks. We use these snapshots to train each of the methods, which then make predictions about collabs that might occur in the final two weeks of the event. To compare methods, we use two evaluation measures. *Area under the ROC curve* (AUC) — also known as the Wilcoxon rank-sum test — is a common measure of quality in link-prediction tasks: it is the probability that the method will rank a randomly-chosen collab above a randomly-chosen non-collab [20]. We also report *personalized precision at rank one* (PP@1): for users who did collaborate during the test period, this is the proportion that actually collaborated with the method's top prediction for them. This is a "personalized" variant of the precision at rank K measure common in the information retrieval literature [38]. We argue that PP@1 is a more interesting for our purposes, since it is very stringent and evaluates predictions on a per-user basis, rather than a ranking of all possible collaborations.

Table 1 reports the evaluation results averaged across all four years. Our path-based regression yields better predictions than the alternatives according both measures, and the gains are statistically significant in most cases. In particular, our model's top prediction for each fawmer was a true collaboration 34.4% of the time. After examining each year's false positives, we found several pairs who did not collaborate that year, but went on to do so in a subsequent year. A few false positives apparently did attempt a collab, but did not follow through (according to survey responses for *least* successful collabs), so their efforts are missing from the archives. In short, there is strong evidence that our path-based regression method does a good job of modeling how collabs form.

In addition to better predictions, our model can explicitly represent different path types (e.g., **A** –follows→ **B** vs. **A** ←messaged– **B**), whereas the alternatives cannot. Importantly, we can inspect the weights associated with different path types to examine how they influence collab formation.

Method	AUC	PP@1
Path-based regression [31]	**0.990**	**0.344**
Baseline random walk [30]	0.982	0.087 ***
Adamic & Adar [4]	0.953 *	0.005 ***
Matrix factorization (SVD) [28]	0.865 **	0.133 **

*** $p < 0.001$ ** $p < 0.01$ * $p < 0.05$

Table 1. Evaluation of link-prediction methods, in terms of area under the ROC curve (AUC) and personalized precision at rank one (PP@1).

Path Variable	Coeff.
A ←follows– **B**	8.433
A –follows→ **B**	7.926
A ←messaged– **B**	4.935
A –messaged→ **B**	4.183
A –wrote→ ♫ ←commented– **B**	4.160
A –commented→ ♫ ←wrote– **B**	3.879
A ←follows– 👤 ←collabed→ 👤 –messaged→ **B**	-0.434
A –follows→ 👤 ←collabed→ 👤 ←messaged– **B**	-0.484
A –liked→ ♫ ←liked– 👤 –liked→ ♫ ←liked– **B**	-0.776
A ←follows– 👤 ←collabed→ 👤 ←messaged– **B**	-1.334
A –liked→ ♫ ←liked– **B**	-1.814
(intercept)	-3.707
A –wrote→ ♫ –tag→ 🏷 ←tag– ♫ ←commented– **B**	0.868
A –commented→ ♫ –tag→ 🏷 ←tag– ♫ ←wrote– **B**	0.504
A –wrote→ ♫ –tag→ 🏷 ←tag– ♫ ←wrote– **B**	-0.388
A –centrality→ 📶 ←Δ=0.2→ 📶 ←centrality– **B**	0.818
A –centrality→ 📶 ←Δ=0.6→ 📶 ←centrality– **B**	0.614
A –centrality→ 📶 ←Δ=0.7→ 📶 ←centrality– **B**	-0.002
A –centrality→ 📶 ←Δ=0.9→ 📶 ←centrality– **B**	-0.332
A –centrality→ 📶 ←Δ=0.8→ 📶 ←centrality– **B**	-0.455

Table 2. A sample of weights from the path-based logistic regression predicting collab formation: positive weights predict collaboration, negative weights predict non-collaboration. This summarizes our analysis of the the FAWM 2012 network (results are similar for other years).

Findings

Table 2 presents a sample of the logistic weights induced from the FAWM 2012 network (results are qualitatively the same, in terms of sign and magnitude, for other years). For clarity and brevity, we do not show or discuss all 165 nonzero weights, but focus on a few of the most positive and negative predictors, as well as several specific paths that relate to the theory about communication, shared interests, status, and social exchange. Note that the intercept is highly negative, capturing the fact collaboration is unlikely by default.

1. Communication Exchanges Predict Collaboration

As theory would suggest [41], the top six predictors of collab formation have to do with communication exchanges: following a partner's song feed, direct messaging, and commenting on the partner's songs. Survey participants confirm the importance of having a rapport with your collaborator:

> *"We had both commented on the vast majority of each other's songs, and we both knew and enjoyed each other's writing styles."*

"The other person and I had both made comments like 'ohh we should totally do something together' ... "

Note that these exchanges can go in either direction (**A** ← **B** or **A** → **B**), and that the inward directions are consistently weighted a bit higher by the model. This may indicate that receiving attention from a partner makes one more willing to collaborate. As one survey respondent put it:

"I appreciated his comments on my songs throughout FAWM and thought he'd be a good person to work with."

2. Collabs Form Out of Shared Interests but Different Skills

Recall that *tags* are mainly used to categorize songs by genre or instrument, so paths that flow through tag nodes can be thought of as expressing a shared interest in musical style. The path **A** –wrote→ ♫ –tag→ 🏷 ←tag– ♫ ←commented– **B**, for example, means that **A** writes songs tagged with terms that are also used for songs on which **B** often leaves comments. In other words, **A**'s typical genre is something of shared interest to **B** according to his commenting behavior. As the theory would imply [21, 55], this path is a *positive* predictor of collaboration (as with **A** commenting on **B**'s genre). However, the path **A** –wrote→ ♫ –tag→ 🏷 ←tag– ♫ ←wrote– **B**, which means that the users often write songs in the same genre or style, turns out to be a *negative* predictor.

While this result may seem contradictory or counter-intuitive at first, it suggests a more nuanced form of homophily than is typically discussed in research on collaboration. In particular, we posit that this reflects an exchange of stylistic skill and expertise that one party has, but the other does not. Consider this post from the FAWM 2010 discussion forums:

"Sometimes I wish I had one of those screamer voices... I could do a raspy acid rage-filled rocker song. Maybe one of you rockers will take me under your AX and help me bring out the inner artistic angst??"

This member wants to stretch herself with a musical style in which she is interested but inexperienced. After looking for help in the forums, a willing collaborator volunteers his expertise in heavy metal. Survey respondents confirm that many heterogenous collaborations begin this way:

"I wrote the lyrics, wanted them to become part of a hard rock/metal soundscape, asked [him] to create one."

"The collaborator, knowing my style, pitched an idea to me that I liked. We passed ideas back and forth each doing aspects [we] could do best."

Such "interdisciplinary" collaborations are so pervasive, in fact, that there are whole forum topics dedicated to it. In FAWM 2012, for example, a discussion thread began in order to pair folk musicians with electronica artists, which resulted in at least eleven "folktronica" collaborations. Similar dynamics can manifest in other online creative communities. For example, on the Newgrounds website [2] collabs often form around animation projects of shared interest, but for which one partners has a production skill that the other does not (such as illustration or programming)[4].

[4] Kurt Luther, personal communication, December 21, 2012.

Other, mere surface forms of homophily are negatively associated with collaboration in the FAWM community (e.g., if both parties frequently "liked" the same songs). Paths involving age, gender, and tenure were discarded by the model altogether. Therefore, musical style appears to be the most important dimension of "complementary homophily" affecting collab formation in FAWM. Such nuanced effects may be similarly domain-specific in other creative communities.

3. Small Status Differences Are Associated With Collabs

Social network centrality also appears to play a nuanced role in collab formation. Theory predicts that people are more likely to work together if they are at the same status level, and less likely if further apart [19]. In our analysis, it is true that very different centrality scores among participants (e.g., $\Delta \in [0.7–0.9]$) are negatively associated with collaboration. However, the path stating that partners have the *exact same* centrality measure, **A** –centrality→ 📶 ←centrality– **B**, was discarded by the model as insignificant. Curiously, a difference of $\Delta = 0.2$ is the strongest positive predictor of collaboration with respect to centrality, suggesting that many partnerships form around a *small difference* in social status.

This result, while somewhat puzzling, is consistent from year to year and under various program parameter settings. Survey responses provide some explanation: that users of lower rank take the opportunity to reach out to their heros and other influential members of the creative community, in hopes of working together. As one fawmer put it:

"I've had a FAWM crush on [her] for ages, and I was noodling on guitar and came up with something that I thought would sound awesome with her voice."

Alternatively, members of higher status sometimes reach out to less experienced songwriters or struggling newcomers, conveying a more active mentor relationship:

"One thing I also LOVE to do is find lyrics writers that are new and put some of their work to music... there is something about seeing that new lyrics writer beam with pride with something you worked on together and for them to realize that they can write songs."

There is also evidence that experienced members use collaboration as a way to help socialize newcomers and introduce the friends that they have brought in to the fold:

"As an introduction to [FAWM, he] suggested we take a full day off work to do this."

4. General Mechanisms of Collab Formation

Survey respondents described three general mechanisms by which they came together to collaborate on projects:

(i) One person produces a "partial work" (e.g., the music or lyrics only) and recruits another to help complete the project. Alternatively, the partner stumbles upon a partial work and volunteers to help.

(ii) As the "next stage" in their online relationship, two partners decide to team up for a project before any work is done. Then they decide how to proceed together.

(iii) Partners are "paired up." This is typically done (often at random) in a forum devoted to collaboration.

Our analysis of the path-based regression model seems most relevant for understanding how collaborations form organically via the first two mechanisms. However, the model may also be useful in predicting more suitable pairings for the third mechanism (we discuss this idea further in the "Applications for Online Creative Communities" subsection).

5. Reasons for Not Forming Collaborations
Survey respondents from the *none* branch were generally still open to collaborating. When asked why they had not done so to date, most said shyness (57%) and lack of time (55%) were the primary reasons. Some also cited coordination overhead (19%), fear of quality below their standards (14%), and the awkwardness of working on music over the Internet (12%).

WHAT MAKES COLLABORATIONS SUCCESSFUL
Our analysis of success factors of collaborations in FAWM takes a slightly different approach. First, we cross-referenced responses from the *single* and *multiple* branches of our survey with website logs, and perform a set of multiple regressions. Survey-takers provided enough detail to map $n = 195$ responses to songs and user pairs in the archival data. We use two dependent variables from the survey to measure the success outcome of a collaboration: (a) a self-reported *satisfaction* rating on a 7-point Likert scale ($\mu = 5.1$, $\sigma = 2.1$), and (b) whether the respondent listed it as their *most* or *least* successful (55.3% *most* successful). We perform linear and logistic regressions against these variables (respectively), using standard model-fitting techniques.

Findings
Regression models are shown in Table 3, along with measures of predictive quality in terms of ranking (R^2, AUC) and loss (RMSE, Error), averaged over ten folds using cross-validation. Note that the sign, magnitude, and significance patterns are roughly the same for both models.

1. Balance of Effort Improves Satisfaction
To study the effect of perceived "work balance," we asked survey participants what percentage they felt they contributed to the final piece ($\mu = 49.9\%$, $\sigma = 22.4\%$). Our models include a variable computed from the Shannon entropy [52] of this response. Entropy is a measure of signal equality, which for binomial distributions is a value from zero to one, in a symmetrical "$\frown$" shape, maximal about 50%.

This notion of work balance is by far the single most significant predictor for both success outcomes. Survey-takers repeatedly praised successful collaborators for their effort and respectfulness, but expressed disappointment when they felt unsuccessful partners had dropped the ball:

> *"[She] was certainly agreeable, but had prior commitments and family responsibilities. [Throw] in a bout of the flu and our collaboration had to take a back seat."*

> *"Once I had drafted a song for us... my collaborator had apparently given up on FAWM and was uncontactable."*

	(a) SATISFACTION		(b) MOST/LEAST?	
Variable	**Coeff.**	**SE**	**Coeff.**	**SE**
work balance	3.548	0.517 ***	4.071	0.827 ***
tag similarity	0.055	0.702	0.204	0.876
$+\delta$ centrality	0.331	0.861	0.061	1.102
$-\delta$ centrality	-2.456	1.098 *	-2.614	1.406 .
$+\delta$ age	-0.023	0.016	-0.018	0.018
$-\delta$ age	0.036	0.016 *	0.037	0.021 .
(intercept)	2.315	0.540 ***	-2.997	0.806 ***
		$R^2 = 0.291$		AUC = 0.765
		RMSE = 1.760		Error = 0.291

*** $p < 0.001$ ** $p < 0.01$ * $p < 0.05$. $p < 0.1$

Table 3. Regressions predicting self-reported collaboration success outcomes. (a) OLS regression estimating a survey respondent's 7-point satisfaction rating for a collaboration, as a function of the independent variables. (b) Logistic regression predicting whether a survey respondent listed a collaboration as being *most* vs. *least* successful.

These observations are consistent with most of the work on "social loafing" [15, 35, 42], which generally cites dissatisfaction caused by the reduced efforts of one's partner. Interestingly, though, the fawmers in our study felt equally dissatisfied when *they* were the ones who dropped the ball:

> *"[I] was really flaky and disappeared for 1–2 weeks after agreeing to collab."*

> *"He recorded me as a collab — which actually made me a winner, and that was quite nice, but I did not feel I deserved much recognition for the work."*

This is a compelling and important result not often reported in the social loafing literature. It suggests that individuals are not merely trying to maximize personal gain through collabs. If that were the case, they would be more satisfied with collabs where their partner did more of the work. On the contrary, fawmers appear to feel most satisfied when they work *equally* with their partners, giving as much as they receive.

2. Stylistic Similarity Does Not Directly Affect Success
Shared interest in musical style is operationalized with a "tag similarity" variable: the cosine similarity [38] between the set of tags used by each user on songs up to that point. Stylistic similarity does not appear to have a direct relationship with success, however. Some survey-takers appreciated working with partners who shared stylistic similarities:

> *"I enjoyed working with someone who had a similar style of writing as me."*

Others took pleasure in playing off their differences:

> *"The initial lyrics were visually evocative and different from anything I'd write myself... My collaborator's different skill set ended up making me try new things, e.g., writing the composition around the bass line."*

Still others found it too difficult to bridge the gap:

> *"We were too different in what we like in songs to be successful."*

So while shared interest in style appears sufficient to influence collab formation (albeit in a subtle and complementary way), once the partnership is formed there seem to be other, less obvious factors affecting its success. This echoes some of the mixed evidence in previous work regarding the effectiveness of homogenous vs. diverse groups [6, 39].

3. Higher-Status Partners May Enjoy Collabs Less
Status in the community is again represented by differences in eigenvector centrality [10]. However, we operationalize this with two different variables: $+\delta$ and $-\delta$, which are nonzero if the partner's centrality is higher or lower (respectively) as compared to the survey-taker. Working with a more influential partner does not seem to be associated with success. However, working with a lower-status partner appears to mildly decrease satisfaction. This indicates that higher-status individuals enjoy their collaborations a bit less. We also controlled for differences in age (in years, using $+\delta$ and $-\delta$ variants), and found that individuals marginally preferred working with younger partners as well.

4. Frequent Communication Helps (Usually)
We rely only on survey responses to analyze communication and collab success, since the archives contain little interaction data during the collab process itself (most of this communication occurs off-site via email, chat, instant message, etc.). Several topics emerged from participants' descriptions of what they found rewarding or challenging about working together. One positive theme was frequent communication with an iterative, back-and-forth approach:

> *"Pretty consistent communication throughout... We collaborated on general concepts lyrically and musically, and then sort of took one part each ([her] on lyrics, me on music) and then came back together."*

Another common thread was an openness to critique and compromise throughout the process:

> *"We communicated eight times... We critiqued each other's work and our own. There was an open attitude of helping each other make the song better."*

Conversely, poor communication from the start was a common complaint in unsuccessful attempts:

> *"Lyrics sent without any context, no preparation, limited communication..."*

These are consistent with previous findings related to creating and critiquing ideas in creative teams [17, 54]. However, communication was not a requisite for all successful collabs. In fact, a few projects involved little or no interaction:

> *"Sad to say, it is simple as 1-2-3: 1. I write the lyrics 2. the musician writes the music and records the whole song 3. I say 'Thank You, that Rocks!' and we post it..."*

We asked survey-takers about the various communication methods they used. While none of these methods predicted success in general, we did find a few effects on 7-point satisfaction ratings specific to *least* successful collabs (n = 80). We observed a significant increase in satisfaction for working together in person (μ = 5.0 vs. 2.9, p = 0.006), and a decrease for using email and chat as the primary method (2.7 vs. 3.7, p = 0.004). It may be that face-to-face interaction facilitates communication and leads to a better result, or perhaps they are both mediated by a hidden variable such as friendship prior to the collaboration. The negative effects of strictly computer-mediated communication are consistent with previous findings, which argue that a lack of physical and vocal cues can hurt effective communication and result in less desirable outcomes [40, 60]. However, for both of these findings further research is warranted to be sure.

DISCUSSION
We find that communication, compatible but complementary interests, and slight differences in status are key factors in collab formation; and that balanced efforts from both parties contribute to collab success. In this section, we discuss a few applications and future directions for this work.

Applications for Online Creative Communities
One actionable idea is to summarize our findings through profiles of successful collabs in the community. By showing what these members created together, and taking care to highlight how communication, complementary interest and status, and a balanced effort played a role in their collaborative process, future members may be able to better understand how to find good partners and have more successful collabs.

Another exciting application is to integrate our models of collab formation directly into the socio-technical software that supports online communities. These models can help intelligently guide members who want to collaborate. For example, one survey respondent mentioned a way to advertise "partial work" collabs on the FAWM website:

> *"It could happen on the jukebox, where you have something tagged as an unfinished potential collaboration... and [someone] can decide to take it on..."*

The accuracy of our path-based regression makes it practical not only for analyzing behavior, but for personalizing such listings as well: potential collabs can be routed to suitable partners based on model predictions. The surprising number of "random collabs" forming in the discussion forums implies that fawmers might also use a sort of "collab-o-matic" matchmaking tool. Similar ideas have been implemented in SuggestBot [16], which helps route Wikipedia articles to suitable contributors. We suspect that collab pairings based on model predictions (or even heuristics based on our findings, if the full models do not scale well in practice) would be more effective than random pairings for such a tool.

Limitations and Future Work
Our data prevent us from drawing causal conclusions about the factors associated with collab formation and success; at this point we can only infer strong correlations. However, the "collab-o-matic" matchmaking tool described above may actually provide a framework for studying causal relationships. Users can be assigned to treatment or control groups, who are given predicted or random pairings (respectively). We can then examine how factors of interest might effect outcomes like the rate of follow-through or song quality.

Another limitation is our use of self-reported satisfaction measures for success, which keeps with the FAWM ethos of meeting a personal challenge. Other objective measures might be of more scientific interest. We have considered using five-star listener ratings for this purpose, or the number of listens and comments received from other community members. However, it is known that these are not necessarily good measures of song quality [50]. An alternative would be to hire third-party evaluators who listen to and objectively score collaborative songs for artistic merit, and study which factors are associated with these third-party scores.

Our study also only focuses on pairwise collabs, since (i) they are the dominant form of collaboration within FAWM, and (ii) our path-based regression approach is designed to predict links between *pairs* of nodes (users). However, creative collaboration often occurs among larger groups of individuals in peer production communities like Newgrounds [36, 37] and Wikipedia [11]. Further research is needed to understand how our nuanced findings about shared interest, status, and work balance generalize to larger groups and other kinds of online creative communities.

Finally, to our knowledge, we are the first to analyze complex social network behaviors using a large-scale path-based regression. The results of this novel, inductive, and quantitative approach are both informed by theory and corroborated by qualitative results from traditional surveys. We believe that this method has broader applications for analyzing and understanding complex social network phenomena, which in turn can enrich online communities in new and interesting ways. For example, Burke and Settles [12] found that receiving comments on one's songs is associated with both individual success and pro-social behavior among newcomers to FAWM. Instead of user→user *collaboration* links, our method could be used to model user→song *comment* links, in order to better explain and predict how users elect to give each other feedback. Such models could drive recommender systems that facilitate community feedback and socialization. The approach could be further modified to model *continuous values* on link edges such as success ratings, or even *hyperedges* (links among three or more nodes) such as the formation of larger collab teams. In essence, the path-based regression method could be applied to other social networks as well, and may be useful for almost any link of interest.

CONCLUSION
We have investigated the social and technological factors that influence (i) how collaborations form in online creative production communities, and (ii) what makes them successful. Through an exploratory study combining quantitative and qualitative analyses of a large online music community, we found several interesting patterns. In addition to communication, shared interests and status within the community are key predictors of whether two individuals will decide to work together — albeit in nuanced and reciprocal ways — suggesting that teams often form to complement one other's skill and influence. Once a collaboration has formed, we find evidence that an equitable division of labor is perhaps the most significant factor in its perceived success: people prefer to give as well as receive in their creative efforts.

ACKNOWLEDGMENTS
The authors acknowledge the roles of communication, shared interests, and complementary skills in this collaboration. Thanks to members of Carnegie Mellon's *Read the Web* and *Social Computing* research groups for helpful discussions (particularly Ni Lao, Bryan Kisiel, William Cohen, and Kurt Luther). High fives to all the fawmers for being our creative Guinea pigs. This work was supported in part by the National Science Foundation (grants IIS-0968487 and IIS-1208382) and Google.

REFERENCES

1. February Album Writing Month. http://fawm.org.
2. Newgrounds. http://www.newgrounds.com.
3. Scratch: Programming for All. http://scratch.mit.edu.
4. Adamic, L., and Adar, E. Friends and neighbors on the Web. *Social Networks 25*, 3 (July 2003), 211–230.
5. Agarwal, A., and Chakrabarti, S. Learning random walks to rank nodes in graphs. In *Proc. Int. Conf. Machine Learning (ICML)* (2007), 9–16.
6. Ancona, D., and Caldwell, D. Demography & design: Predictors of new product team performance. *Org. Sci. 3* (1992), 321—341.
7. Backstrom, L., Huttenlocher, D., Kleinberg, J., and Lan, X. Group formation in large social networks: Membership, growth, and evolution. In *Proc. Int. Conf. Knowledge Discovery and Data Mining (KDD)*, ACM (2006), 44—54.
8. Bantel, K., and Jackson, S. Top management and innovations in banking: Does the composition of the top team make a difference? *Strategic Mgmt. J. 10* (1989), 107–124.
9. Benkler, Y. *The Wealth of Networks: How Social Production Transforms Markets and Freedom*. Yale University Press, New Haven and London, 2006.
10. Bonacich, P. Some unique properties of eigenvector centrality. *Social Networks 29* (2007), 555—564.
11. Bryant, S., Forte, A., and Bruckman, A. Becoming Wikipedian: Transformation of participation in a collaborative online encyclopedia. In *Proc. Int. Conf. Supporting Group Work (GROUP)* (2005).
12. Burke, M., and Settles, B. Plugged in to the community: Social motivators in online goal-setting groups. In *Proc. Int. Conf. Communities and Technologies (C&T)*, ACM (2011), 1–10.
13. Byrne, D. *The Attraction Paradigm*. Academic, New York, NY, USA, 1971.
14. Cattani, G., and Ferriani, S. A core/periphery perspective on individual creative performance: Social networks and cinematic achievements in the Hollywood film industry. *Org. Sci. 19*, 6 (November 1 2008), 824–844.
15. Chidambaram, L., and Tung, L. Is out of sight, out of mind? An empirical study of social loafing in technology-supported groups. *Inform. Syst. Research 16*, 2 (2005), 149–168.
16. Cosley, D., Frankowski, D., Terveen, L., and Riedl, J. SuggestBot: Using intelligent task routing to help people find work in Wikipedia. In *Proc. Int. Conf. Intelligent User Interfaces (IUI)*, ACM (2007), 32–41.
17. Dow, S., Fortuna, J., Schwartz, D., Altringer, B., Schwartz, D., and Klemmer, S. Prototyping dynamics: Sharing multiple designs improves exploration, group rapport, and results. In *Proc. Conf. Human Factors in Comput. Syst. (CHI)*, ACM (2011), 2807–2816.
18. Emerson, R. Social exchange theory. *Annu. Review of Psych. 2* (1976), 335—362.
19. Faulkner, R., and Anderson, A. Short-term projects and emergent careers: Evidence from Hollywood. *Amer. J. Sociology 92*, 4 (January 1987), 879–909.

20. Fawcett, T. An introduction to ROC analysis. *Pattern Recognition Letters 27* (2006), 861–874.
21. Feld, S. The focused organization of social ties. *Amer. J. Sociology 86*, 5 (March 1981), 1015–1035.
22. Frushtick, R. Spoilers! Zork 'Walkthrough' rocks out. *MTV Multiplayer* (September 30, 2009). New York, NY: MTV. http://on.mtv.com/1rQiu.
23. Fuller, D., and Magerko, B. Shared mental models in improvisational theatre. In *Proc. Conf. Creativity and Cognition (C&C)*, ACM (2011), 269–278.
24. Gruenfeld, D., Mannix, E., Williams, K., and Neale, M. Group composition and decision-making: How member familiarity and information distribution affect process and performance. *Org. Behavior and Human Decision Processes 67* (1996), 1–15.
25. Hoffman, L., and Maier, N. Quality and acceptance of problem solutions by members of homogeneous and heterogeneous groups. *J. Abnormal and Social Psych. 62* (1961), 401–407.
26. Hogg, M. A., and Turner, J. C. Interpersonal attraction, social identification and psychological group formation. *European J. Social Psych. 15* (1985), 51–66.
27. Huang, Y., Zhu, M., Wang, J., Pathak, N., Shen, C., Keegan, B., Williams, D., and Contractor, N. The formation of task-oriented groups: Exploring combat activities in online games. In *Proc. Int. Conf. Computational Sci. and Eng. (CSE)*, vol. 4, IEEE Press (2009), 122–127.
28. Koren, Y., Bell, R., and Volinsky, C. Matrix factorization techniques for recommender systems. *Computer 42*, 8 (2009), 30–37.
29. Kravitz, D., and Martin, B. Ringelmann rediscovered: The original article. *J. Personality and Social Psych. 50*, 5 (1986), 936–941.
30. Lao, N., and Cohen, W. Relational retrieval using a combination of path-constrained random walks. *Machine Learning 81* (2010), 53–67.
31. Lao, N., Mitchell, T., and Cohen, W. Random walk inference and learning in a large scale knowledge base. In *Proc. Conf. Empirical Methods in Natural Language Processing (EMNLP)*, ACL (2011), 529–539.
32. Lave, J., and Wenger, E. *Situated Learning: Legitimate Peripheral Participation*. Cambridge University Press, 1991.
33. Levine, J., and Resnick, L. Social foundations of cognition. *Annu. Review of Psych. 44* (1993), 585–612.
34. Liben-Nowell, D., and Kleinberg, J. The link prediction problem for social networks. In *Proc. Conf. Inform. and Knowledge Mgmt. (CIKM)*, ACM (2003), 556—559.
35. Liden, R. C., Wayne, S. J., Jaworski, R. A., and Bennett, N. Social loafing: A field investigation. *J. Mgmt. 30*, 2 (2004), 285–285.
36. Luther, K., and Bruckman, A. Leadership in online creative collaboration. In *Proc. Conf. Comput.-Supported Cooperative Work (CSCW)*, ACM (2008), 343–352.
37. Luther, K., Caine, K., Ziegler, K., and Bruckman, A. Why it works (when it works): Success factors in online creative collaboration. In *Proc. Int. Conf. Supporting Group Work (GROUP)*, ACM (2010), 1–10.
38. Manning, C. D., and Schütze, H. *Foundations of Statistical Natural Language Processing*. MIT Press, 1999.
39. Mannix, E., and Neale, M. What differences make a difference? The promise and reality of diverse teams in organizations. *Psych. Sci. in the Public Interest 6*, 2 (October 2005), 31–55.
40. Markus, M. Finding a happy medium: Explaining the negative effects of electronic communication on social life at work. *ACM Trans. on Inform. Syst. 12*, 2 (April 1994), 119–149.
41. McKenna, K., Green, A., and Gleason, M. Relationship formation on the Internet: What's the big attraction? *J. Social Issues 58*, 1 (2002), 9.
42. Mulvey, P. W., and Klein, H. J. The impact of perceived loafing and collective efficacy on group goal processes and group performance. *Org. Behavior and Human Decision Processes 74*, 1 (1998), 62–87.
43. Newcomb, T. Varieties of interpersonal attraction. In *Group Dynamics: Research and Theory*, D. Cartwright and A. Zander, Eds. Row, Peterson, Evanston, IL, 1960, 104–119.
44. O'Reilly, C., Caldwell, D., and Barnett, W. Work group demography, social integration and turnover. *Administ. Sci. Quart. 34* (1989), 21–37.
45. Page, L., Brin, S., Motwani, R., and Winograd, T. The PageRank citation ranking: Bringing order to the Web. Technical Report 1999-66, Stanford InfoLab, November 1999.
46. Ren, Y., Kraut, R., and Kiesler, S. Applying communication identity and bond theory to design of online communities. *Org. Studies 28*, 3 (March 2007), 377–408.
47. Robinson, G., and Dechant, K. Building a business case for diversity. *Acad. Mgmt. Executive 11*, 3 (1997), 21–31.
48. Roschuni, C., Oehlberg, L., Beckman, S., and Agogino, A. Relationship conflict and feeling communication in design teams. In *Proc. Int. Design Eng. Tech. Conf. and Comput. and Inform. in Eng. Conf. (IDETC/CIE)* (2009), 87626.
49. Rose, J. Presidents' songs: Of legends and mortal men. *All Things Considered* (October 8, 2008). Washington D.C.: National Public Radio. http://n.pr/QyWjPz.
50. Salganik, M., Dodds, P., and Watts, D. Experimental study of inequality and unpredictability in an artificial cultural market. *Science 311* (2006), 854–856.
51. Schon, D. *The Reflective Practitioner: How Professionals Think in Action*. Ashgate Publishing, 1995.
52. Shannon, C. A mathematical theory of communication. *Bell Syst. Technical J. 27* (1948), 379–423,623–656.
53. Slavin, R. Cooperative learning. *Review of Educational Research 50*, 2 (Summer 1980), 315–342.
54. Sutton, R., and Hargadon, A. Brainstorming groups in context: Effectiveness in a product design firm. *Administ. Sci. Quart. 41* (1996), 685–718.
55. Tajfel, H. *Social Identity and Intergroup Relations*. Cambridge University Press, London, 1982.
56. Tibshirani, R. Regression shrinkage and selection via the LASSO. *J. Roy. Stat. Soc. Series B (Methodological) 58*, 1 (1996), 267–288.
57. Triandis, H., Hall, E., and Ewen, R. Member heterogeneity and dyadic creativity. *Human Relations 18* (1965), 33–55.
58. Van Der Vegt, G., and Bunderson, J. Learning and performance in multidisciplinary teams: The importance of collective team identification. *Acad. Mgmt. J. 48*, 3 (June 2005), 532–547.
59. Walther, J. B. Time effects in computer-mediated, groups: Past, present, and future. In *Distributed Work*, P. Hinds and S. Kiesler, Eds. MIT Press, Cambridge, MA, 2002, 235–257.
60. Wideström, J., Axelsson, A., Schroeder, R., Nilsson, A., Heldal, I., and Abelin, A. The collaborative cube puzzle: A comparison of virtual and real environments. In *Proc. Int. Conf. Collaborative Virtual Environments*, ACM (2000), 165–171.
61. Wu, M., Ranjan, A., and Truong, K. An exploration of social requirements for exercise group formation. In *Proc. Conf. Human Factors in Comput. Syst. (CHI)*, ACM (2009), 79–82.
62. Yuki, M., Maddux, W., Brewer, M., and Takemura, K. Cross-cultural differences in relationship- and group-based trust. *Personality and Social Psych. Bulletin 31*, 1 (2005), 48–62.