



Issues in Community Moderation

Managing Anonymous Participation

Should online communities permit people to participate *from behind a digital disguise*?

This is a question as old as social media. Even so august a publication as the *New York Times*—which supports its own online community—recently wrestled with the issue of anonymous participation (“News Sites Rethink Anonymous Online Comments”).

There are no simple answers to the anonymity question, but no one should attempt *any* answer without understanding two indisputable facts:

Lots of people prefer to remain anonymous on the Web. Many of them feel so strongly about it that they will avoid destinations where they must establish a true identity.

Web destinations that admit anonymous visitors while applying no governance on how people conduct themselves inevitably drown in abusive, offensive content.

These are immutable facts that, together, create great tension for any organization wanting to maintain a healthy community. These organizations must find a way to nurture an environment that suits *all* of its audience, or at least every *well-intentioned* audience member.

We at Demand Media have helped launch—and continue to moderate and manage—hundreds of healthy, dynamic communities that respect anonymous participation. Our experience has taught us that, in many communities, identity has to be earned. We’ve also learned that when it comes to controlling inappropriate or unacceptable behavior, community members themselves are often far more effective than rules of conduct.

This paper reflects our years of experience, and summarizes the insights we’ve gained concerning anonymous participation.

Demand Media 

“Anonymity is just the way things are done. It’s an accepted part of the Internet, but there’s no question that people hide behind anonymity to make vile or controversial comments.”

Arianna Huffington

Living with Digital Disguises

It is inevitable that many people want to—and will find ways to—hide behind a digital disguise. Of course, you could go to the extreme of forcing everyone to provide a credit card in order to contribute to your site, but you would lose 95% of the audience without fail.

Rather than focus on stamping out anonymity, concentrate more on living *with* digital disguises. Think about ways to *motivate* your visitors to use their real identity and, where necessary, help your visitors to create a *meaningful* digital disguise.

Encourage users to use their real identity

Some users perceive that there is a cost associated with using their real identity online. If real identities are important to you, you must compensate your visitors for that cost, or lower the cost so far that it is no longer an issue:

1. *Make registration easy.* The more onerous your registration process the less likely visitors will be to share valid information or even participate. Fortunately, there are technology solutions that simplify registration.

Janrain lets visitors register on your site using an existing identity (e.g., Yahoo, Google, Facebook). *Facebook Connect* lets visitors register on your site using their existing Facebook credentials (although there are some restrictions on how you can use Facebook Connect data that you need to be aware of).

Anonymity cannot be avoided
Identity has to be earned

2. *Reward visitors for sharing their real identity.* This seems obvious, but it is seldom done. One easy way to reward candor is to award visitors a “Verified” badge if they register using a valid identity (e.g., their Facebook account). We actually call these *assigned badges* in Pluck Rewards, our Social Reputation system.



You can also let visitors who provide their real identity do more things in your community. For example, if a person registers on your site with Facebook Connect, let them invite their Facebook friends to join them on the site, and connect locally with their Facebook friends already participating on the site.

3. *Force anonymous members to undergo a trial period.* Consider pre-moderating—that is, reviewing prior to publication—the first few posts of new visitors who choose to remain anonymous. You’ll quickly identify the folks that have something meaningful to share and you’ll easily weed out the spammers and low value contributors.

How can Pluck help?

The Pluck platform reflects our view that friends are an essential element of online community. Pluck-driven communities encourage members to import their friends, make new friends, and leverage the value of their friends as they behave socially on your site.

- Visitors who use Pluck’s Social Bridging features to establish their connections with Facebook, Twitter, or LinkedIn can easily invite their friends and followers on those networks to join them on the your site.
- When you deploy the Pluck Comments application your registered visitors can filter the entire conversation stream, limiting it to just their friends’ contributions and interactions.
- The Pluck Persona application can present your registered visitors with a feed of all their friends’ activities on your site. This and the aforementioned Pluck Comments support for friends give your visitors convenient and valuable ways to cut through the noise.
- Pluck allows community managers to set moderation policies by user tier. For example, you can assign all new visitors to a default *untrusted* tier, all members of which have their comments pre-moderated before publication to the site. Visitors with multiple approved contributions can be promoted to a *trusted* tier where their contributions publish immediately upon posting.



Help users create a meaningful digital disguise

Well-designed communities reflect community members' online activities back into their identities and give would-be anonymous users a reason to offer more than just another digital disguise:

1. *Recognize and reward valuable digital disguises.* Consider implementing a badging system that rewards visitors who invest in their digital disguises. Label highly regarded commenters with an *earned* badge that reflects the community's appreciation for their contribution. And labeling chronically abusive members with a "dunce" badge of some kind may discourage their behavior and will help others to avoid them.
2. *Encourage competition in support of desired community outcomes.* Consider implementing a leaderboard that encourages community members to make frequent, valuable contributions.
3. *Raise switching costs.* Drive value into your visitors' digital disguises and they will be less likely to do something to damage them. Make a visitor's participation about more than just the immediate conversation. Provide additional reasons to return to your site and invest in their digital disguise: let them build friendships on your site, track their contributions in one place, recognize them for their contributions, give them a place of status in your community.

How can Pluck help?

When we first launched the Pluck social media platform in 2007, we introduced the concept of a *persona*. Visitors to Pluck-driven communities use Pluck Persona to create their digital representation at that particular Web destination. By design, you can configure Pluck Persona to give visitors latitude regarding the personal information they disclose—even so far as making display names optional—as they create this digital representation. In other words, Pluck-driven communities can be set up to support digital disguises and contributor anonymity.



MANAGING ANONYMOUS
PARTICIPATION

Pluck Rewards provides a wealth of features for detecting and surfacing participant conduct, positive and negative.

Pluck Rewards supports the concept of both *earned* and *assigned* badges. Earned badges can be awarded according to programmatically determined criteria, and badge visibility is at the community manager's discretion.

And Pluck Rewards' *leaderboards* can surface—and optionally make public—members that have achieved positive or negative superlatives that the community and its managers care about and want to encourage or discourage. Microsoft, in its Windows Phone 7 campaign site, gives prizes to members who reach leaderboard status in an effort to encourage community participation.



“He [an anonymous contributor] recalls the time his comment on a column titled ‘My Lazy American Students’ drew hundreds of ‘Recommends’. Leaning back in his chair, with his left arm resting on his thick belly and his right arm jingling change in his pocket, he smiles. ‘There were more than 600 comments, and at least half of them agreed with me!’”

*“Inside the mind of the anonymous online poster”,
Boston Globe Sunday Magazine*

Create a Community-driven Social Contract

Organizations that choose to permit anonymity need mechanisms for separating anonymous winners from losers. In our experience, empowering the larger community is key. You don’t always need to know the real identity of a visitor as long as your community is empowered to exert its own social controls. Think of this as a *social contract*, a pact that the community members make, and enforce, with one another.

Your community can help you
sort out winners and losers

In most communities there is an army of contributors who voluntarily will help ensure community norms are enforced. Research has shown well-intentioned anonymous participants really care what the community thinks of them and their contributions. So how do you enlist the army of volunteers to enforce a social contract?

1. *Allow community members to score contributions.* Simply allow the community to vote a contribution up or down. When one person takes issue with what you say you are ready to do battle; when 50 people take issue with what you say you shut up and run the other way.
2. *Automatically remove contributions that exceed a threshold of abuse.* Most abuse report systems flag a contribution for review by a moderator. This is fine, but consider going a step further by automatically removing contributions that exceed a pre-determined number of abuse reports.
3. *Use abuse reports to identify more than just bad contributions.* Abuse reports, if used the right way, point out repeat offenders. It’s important that you respect the sentiments of your community by establishing policies for banishing repeat offenders from the community in addition to removing their abusive contributions.
4. *Spotlight valued contributors.* Just as you make it really hard for ill-intentioned contributors to be heard; it is important that you recognize valued contributors, anonymous or otherwise.



How can Pluck help?

Pluck offers a variety of ways—including features of Pluck Comments, Ratings and Reviews—for community members to offer feedback on member-generated content, and for community moderators to track community sentiment.

Pluck assigns demerits to visitors each time a contribution by them is flagged as abusive and deleted. This allows moderators to quickly identify a list of abusive users in the system and take action independent of abuse reports.

Moderators can even take this a step further by applying some intelligence to the sorting of abuse reports. Pluck applies a sorting algorithm that takes into account the tier of the scrutinized visitor and the number of abuse reports against a contribution. This shows moderators the most sensitive items first in the list of abuse reports.

Pluck's new Moderation Manager tool provides a number of visitor management options, such as labeling as abusive, blocking, contacting, and moving the visitor to a different, potentially more regulated, user tier. The tool's moderation action reporting function delivers comprehensive daily documentation of moderation and abuse actions.

Moderation Manager also has a number of automated features that offload significant moderation workload. Moderation Manager lets moderators deal with abusive users who reappear with new user names by applying IP and meta-data blocking. Finally, Moderation Manager helps moderators and community managers deal with spam content.



“I’m very happy about these changes. There were interesting articles in the NY Times and the Washington Post this week on trying to keep the comments area more civil. I particularly like the new Most Popular approach. It’s fine to attack ideas. It’s not fine to attack people. Can’t we all just get along? :)”

Reader reaction to new community policing features, *San Francisco Chronicle’s sfgate.com*

Give the Community Enforcement Power

It’s not uncommon for anonymous troublemakers to overwhelm community moderators. Not too long ago anonymous trolls forced comments vendor engadget to completely disable comments on their site.

Beyond establishing a community-wide social contract, Pluck customers have succeeded in directly involving their community members in the moderation process. The following four tactics, among others, have helped members of those communities to actuate the crowd’s opinion and shut down chronically abusive anonymous members:

1. *Sorting.* Most of us are familiar with basic chronological sorting of contributions (e.g., *oldest to newest, newest to oldest*). Consider empowering the crowd by implementing two additional sorting options, *highest rated* and *most active*. *Highest rated* promotes the consensus most valuable contributions. *Most active* recognizes content that has been most often reacted to by the community. Most active contributions may not have the overall highest score, but they most likely are intriguing.

Give each user the power to cut through the noise

2. *Filter by Friends.* The influence of friends on social networks is well established. For example, a recent study by research and consulting firm Morpace Inc. showed that Facebook users were 68% more likely to visit or buy from a retailer if it received a positive referral from a friend. Savvy community managers take advantage of this phenomenon by letting community participants filter conversations to see just what their friends are saying.
3. *Ignore.* In the end, everyone has an opinion and it is impossible to please everyone. The *ignore* feature lets you stop trying. Letting community members create a personalized list of ignored contributors relieves them from having to interact with those contributors. If the feature is well designed it displays a placeholder in place of each contribution which the member can expand if the contribution has received a high score or has a lot of activity.



MANAGING ANONYMOUS
PARTICIPATION

4. *Replies*. This is a simple feature that can hide a great deal of noise in the conversation. By allowing contributors to reply directly to what another contributor has said, and linking the reply to the original contribution, a well-designed *replies* feature hides replies by default and only expands them if the reader finds value in the first contribution.

How can Pluck help?

All of these tactics are standard, out-of-the-box features of Pluck Comments and Reviews.

Pluck Persona is the primary vehicle for creating social networks—with friends, messaging and the other features commonly associated with social networks—on community Web sites. Pluck's Social Bridging service makes it easy for community members to invite their Facebook or Twitter friends to join them.

Pluck Comments supports linking replies to the sorting values of the first contribution. This keeps the conversation around a contribution together and filters out the noise of replies that lack context outside of the original contribution.



“We decided instead to redouble our efforts at engaging readers in the comments, policing and blocking those that didn’t meet standards and working with our software vendor to improve the tools. With the work of reporters, editors, a full-time comment monitor and the occasional muscle of an attorney – as well as help from thousands of readers (thanks) – I’ll say the cleanup is a successful work in progress.”

Ken Brusic, Editor and Senior
Vice President
Orange County Register

Be a Part of Your Own Community

Ultimately, all of these best practices are made more successful by your presence in the community. You must inject yourself into the community from the very beginning, not only when issues arise. Your community needs to know you and trust you if you expect them to respond when you take action. Fortunately, making your presence known from the beginning is not hard:

1. *Establish a clear set of community guidelines.* Every community has different norms. You should clearly define and publish your norms in the form of written guidelines. These guidelines inform your membership, and provide a framework for your moderators to evaluate contributions.
2. *Establish a profile and actively participate.* As a community manager, make sure you have your own Persona in the community so your community members can recognize you. Use your Persona to start discussions, address questions, ask for feedback and respond to complaints. The more involved you are the more you will be recognized in the community.

You share responsibility for the
conversation on your site

3. *Assign privileges to trusted community members.* Enlist trusted users in the community to help you enforce the guidelines. The best way to identify trusted users is to monitor conversations and look for those individuals with thoughtful and meaningful contributions. You can think of these individuals as your community advocates. In the end, the more eyes you have looking for violations and hands you have reacting to them, the quicker you will be able to respond and show your community guidelines should be taken seriously.
4. *Label trusted community members.* Early on establish a badge or a format for contributions by trusted community members. These users and their conversations should be instantly recognizable in the conversation. This provides a show of force and highlights the contributions community members should be aware of.



MANAGING ANONYMOUS PARTICIPATION

5. *Get your editors, journalists and content creators involved.* No one likes talking to a brick wall. Community members interact with content on your site in hopes that others will respond. Having the original creator of the content participate in the conversation raises the value of participation and interest. It is also likely to raise the level of the conversation or remove negative contributions quickly.

How can Pluck help?

Pluck provides Moderation and Community Management Services that can be leveraged to help you define a set of community guidelines, moderate your community and recommend and apply best practices.

Pluck provides a set of user tiers that allow Community Managers to assign privileges to trusted community members.

Community Managers can use Pluck Rewards to assign badges to staff members and privileged community members. These badges are typically prominently displayed, in the interest of promoting the visibility of these individuals.



Living with Anonymity

Online anonymity is not a danger to civil online conversations, provided appropriate moderation is in place. Every community has different characteristics and different requirements, but in general responsible behavior depends on moderation that features both carrots and sticks. The carrot is the cohesion and self-interest of the larger group. The stick is the discomfort associated with being publicly associated with embarrassing behavior. Good moderation requires technology and techniques that, combined, wield the carrots and sticks required to preserve an environment that the community demands.

About Demand Media

Demand Media, Inc., is a leading online media company that informs, entertains and connects millions of people every day. Through a portfolio of vertical Web properties reaching more than 80 million monthly visitors, a global network of digital partners, a breakthrough content studio, and an enterprise social media platform, Demand Media publishes what the world wants to know and share. Founded in 2006, Demand Media is headquartered in Santa Monica, CA with offices in Bellevue, WA, Austin, TX, New York, NY and London, UK.