

The Virtues of Moderation

Remarks on a work in progress
IP Scholars Conference
DePaul University College of Law
9 August 2007

James Grimmelmann

New York Law School
57 Worth St.
New York, NY 10013
jgrimmelmann@nyls.edu

This paper is the first (itself incomplete) part of a larger project on collaboration and conflict in creative communities, especially online ones. The goal is to be precise about what really changes we go online, so as to inform Internet and IP policy. The underlying intuition is that that the “community” is the most useful unit of analysis; to see how the Internet affects creativity, we should try to understand both the diversity of creative communities and their commonalities. This paper mashes up two reasonably well-understood ideas—commons theory and layering—to lay some groundwork for seeing how communities work.

My plan of action is first to motivate the problem by giving some divergent views on online communities organized around sharing, to show that there’s an ambiguity when we talk about a “commons.” Focusing on that ambiguity and applying what we know about layering leads to the conclusion that most online communities (and the Internet itself, as a community writ large) are actually semicommons: partly private and partly held in common. I’ll finish by sketching some of the kinds of claims that that analysis makes possible.

To begin: the motivation. There is a specter haunting intellectual property; the specter of the commons. The idea is that a commons in information goods could promote massive collaboration, harnessing efficiencies of distributed production without incurring the costs associated with exclusive rights. The viability of this argument hinges on two related problems. First, who in their right mind would contribute to an information commons without the financial incentives provided by exclusive rights?

The force of this question has been blunted by abundant evidence that millions of people (all apparently in their right minds) do contribute to such commons and by more reasonably nuanced explanations of why they do so. A mix of indirect economic benefits with altruism, reciprocity, and other social motivations motivate people to contribute to everything from open source software to lolcats.

Still, even well-intentioned and well-motivated collaborators must nonetheless face the organizational problems that bedevil firms, governments, and other institutions of coordination. Here, the concern is that a mix of free riding, disruptive behavior, differences of opinion, and good old fashioned cacophony will eventually undermine whatever mechanisms a sharing-based community uses to integrate the perspectives and contributions of its multiple participants. On this view, as these processes break down—and they WILL break down at scale—the result is abandonment and collapse.

To make this issue more concrete, consider Wikipedia. Fans of the commons think that it represents a remarkable demonstration of the sustainability of voluntary information

production. As compared with traditional encyclopedias—and competing online ones—Wikipedia’s abandonment of restrictions on contribution is the key to its success.

On the other hand, skeptics have argued that the Wikipedia model is inherently unstable, and that its openness means that it will be overwhelmed by vandals, sock puppets, and script kiddies. Here are two apparently diametrically opposed views of the commons. One thinks that common access will make Wikipedia work; the other that common access will make Wikipedia fail.

I’d like to argue that they’re both right, after a fashion. There’s a subtle equivocation taking place over the meaning of “commons” in these debates. What, exactly, is the resource to which access is either open or restricted? The theoretical basis for commons production is the non-rivalry of information. The theoretical basis for the attack on commons production I’ve just described is that common access is only rarely an effective way to manage rival resources. That divergence requires more detailed elaboration, which means that it’s time . . .

. . . for a two-by-two matrix. Conventionally, goods are public or private. Public goods are nonrival and nonexcludable; private goods are rival and excludable. Since rivalry and excludability are at least partially independent, club goods and common-pool goods round out the matrix. On this conventional view, it’s non-excludability that causes problems. What we can keep others from using, we can create pricing regimes in and therefore use efficiently. The last few decades of research on the commons have challenged this view in two ways.

First, for rival goods, the conventional theory tells us that an exclusion-based system leads to efficient allocation. Without exclusion, a tragedy of the commons results from wasteful overuse. Property rights and prices—or perhaps, government regulation—are needed to keep use to appropriate levels.

Thanks particularly to Elinor Ostrom and research drawing on her work, we know now that that story is incomplete. Top-down rights and regulations are not the only way to create the necessary excludability. Bottom-up self-created and self-enforced community systems of common ownership and management can also prevent wasteful overuse. The common-pool resource literature tells us that groups with well-defined boundaries, graduated sanctions, and good fora for communication can produce stable institutions that regulate use to sustainable levels. I call this the Tragic story; it explains how common ownership can avoid the tragedy of the commons.

On the non-rival side, things are somewhat different. Here in the realm of ideas and intellectual property, conventional theory claims that a pricing system (or direct government provision) has its problems, but still often beats the alternative. Intellectual property can lock up information goods so that too few have access to them, but this sacrifice is a necessary one, since otherwise no one would have an incentive to create those goods in the first place.

Here as well, those who wave the banner of the commons argue that there is once again a better way. The seemingly intractable production problem is in fact tractable. People are natural information producers, demand creates its own supply, and combining the creativity of huge numbers of individuals can produce all the information we ever need, and more. Once the information exists, the best thing to do is share it as widely as possible—we get more creativity from connecting authors to audiences and to previous authors than we do by offering them exclusive rights. I call this story the Comedic one; it explains how a commons can catalyze collaboration on a vast scale.

The Tragic and Comedic stories of the commons both counsel against the private property rights strategy. But they have interestingly different things to say about excludability. The Tragic story says to embrace excludability to prevent waste; the “commons” is just another institution for generating excludability. This is the point that the Wikipedia skeptics are making. They see competition for server resources and for the limited attention of readers, with a wasteful cycle of overuse and abuse spiraling out of control unless there are strictly enforced limits on what one does with the resource.

The key to absolutely everything I’m saying here is that both stories are right, simultaneously. When we talk about particular online resources and their associated communities (rather than talking about a commons in information in general), these resources have some aspects that are rival and subject to Tragic effects and some aspects that are non-rival and subject to Comedic effects. They’re semicommons.

The semicommons framework comes from Henry Smith’s work on semicommons resources. His theory starts from a study of the open-field system. The underlying land was held privately in strips by farmers, but during some seasons held common for grazing by sheep. This mixture of regimes had symbiotic benefits: simultaneous use, fertilization of crops, and fodder for the sheep. But it also had costs, including risks that shepherds would selectively graze on particular farmers’ land. Smith argues that open-field communities dealt with this risk by scattering landholdings into thin strips, making it hard for shepherds to focus on any particular plot, and that this redivision of boundaries was preferable to closer monitoring of shepherds’ activities.

This framework is useful for describing online communities. It leads us to ask which aspects of the resource are private, which are common, what forms of strategic behavior this combination faces, and what institutions respond to those threats.

Specifically, the technical concept of layering provides a useful way of understanding the resource-set used by an online community. The physical infrastructure that supports an online community--the servers and bandwidth--is rival and privately held. But there’s no logical necessity that the resource be held the same way at all layers; many communities throw that infrastructure open at the content layer, and thus effectively hold it in common. Only the private owners of the chattel computers can reconfigure them or move them, but anyone is free to participate and contribute by sending content using well-defined higher-level technical protocols.

Smith's semicommons directs our attention to the various kinds of strategic behavior possible in such a system; thus, common users could act to place costs on the private owner, but the private owner could act to capture the benefits of common use. I give a more detailed model of how people use an online semicommons in the paper, built around the idea of "roles" users play in a creative community: as authors, moderators, readers, and infrastructure owners. I expect that the microeconomics involved will need a paper of their own; for example, it turns out to be ambiguous whether having more authors increases or decreases the incentives to read; it depends on how effective the moderators are.

Thus—and this is either the lacuna in the current paper or the basis of yet another one—everything depends on the institutions that the community uses to moderate between uncontrolled overuse and overly-controlled underuse. I'm working on taxonomizing these moderation patterns, but in my time remaining I'll go beyond the well-established to suggest a few of the points this form of analysis may enable, given the initial motivation.

Semicommons theory gives to my mind a satisfying explanation of why USENET newsgroups stalled out but email continues to work. USENET got the property boundaries wrong: the typical community was defined by a newsgroup, but had no link to infrastructure ownership. When the spammers came, there was no one in a position to moderate their use. Email puts much more of the responsibility for delivering messages on those who benefit from the delivery: recipients. Result: it can tolerate a general freedom to send messages at a much larger scale.

Semicommons theory also emphasizes that there's no one-size-fits-all solution. As a community gets larger and more diverse, the quality-control and credentialing problems have change dramatically. Scale is a central variable in both Tragic and Comedic stories, and the interplay between them plays out differently at different scales. The real miracle of Wikipedia is that it's changed its own institutions repeatedly and evolved new moderation patterns as it's grown.

Importantly, these new moderating institutions can't simply be described as making Wikipedia more "controlled." Nor could one say that USENET or email is "more open" across the board. There are incomparable virtues at stake. A community with strong social norms, like Metafilter, discussed in the paper, may need to take steps to place limits on its userbase growth. Slashdot has more contributors but relies a complicated and crafty rating system to discourage abuse. Community diversity should make one skeptical of claims that any particular form of ownership or moderation is "best."

And finally, there are many ways to fail. The Internet is littered with Slashdot clones—many using the same exact software powering Slashdot—that crashed and burned. But there are also many ways to succeed. The Internet is filled with sites whose communities of users are healthy and thriving, that manage to be substantially open while nonetheless deterring strategic behavior. Thinking about online communities as semicommons helps explain why such stable compromises are possible.