

# *The Wealth of Networks*

A Crooked Timber Seminar (<http://www.crookedtimber.org>)

**This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 2.5 License. For further information see <http://creativecommons.org/licenses/by-nc-sa/2.5/> .**

## *Contents*

- p. 3**            *Introduction*
- pp. 4-10**       *Norms and Networks* by **Henry Farrell**
- pp. 10-14**      *A General Theory of Information Policy* by **Dan Hunter**
- pp. 15-17**      *Why Do Social Networks Work?* by **John Quiggin**
- pp. 18-20**      *Whose networks? Whose wealth?* by **Eszter Hargittai**
- pp. 21-24**      *Mediating the Social Contradiction of the Digital Age* by **Jack Balkin**
- pp. 25-26**      *The Dialectic of Technology* by **Siva Vaidhyanathan**
- pp. 27-37**      *Response* by **Yochai Benkler**

## ***Introduction***

Yochai Benkler's *The Wealth of Networks: How Social Production Transforms Markets and Freedom* is a very exciting book. It captures an important set of developments – how new information technologies make it easier for individuals to collaborate in producing cultural content, knowledge, and other information goods. It draws links across apparently disparate subject areas to present a theory of how these technologies are reshaping opportunities for social action. Finally, it presents a highly attractive vision of what society might be like if we allow these technologies to flourish, as well as the political obstacles which may prevent these technologies from reaching their full potential. If you're interested in debates on Creative Commons, on Wikipedia, on net neutrality, or any of a whole host of other issues, this is an essential starting point.

We've put together a seminar on the book, which we hope will help spur discussion around it in the blogosphere. This is an important debate. In a (long overdue) departure from previous seminars that I and others have organized at CT, we hope to include other blogs more directly in the discussion than in the past. We'll do this by borrowing an idea from Will Wilkinson, and using this post to link to blogs which we think make substantial contributions to this set of arguments (nb that my definition of substantial is necessarily an idiosyncratic one). The material from this seminar is also available under a Creative Commons license (the Attribution-NonCommercial-ShareAlike 2.5 License) for others to re-use and add to in creative ways. The seminar is available both as a PDF and as an .rtf file for easier reading and re-use.

The contributions are in the order that they are mentioned in Benkler's response. Henry Farrell argues that not only formal institutions but also informal norms are necessary for these technologies to enable proper collaboration. Dan Hunter celebrates the book, but worries that it covers too many topics, and that it's written in language that non-academic readers may have difficulty in understanding. John Quiggin examines the underlying motivations behind the production of common resources, and suggests that Benkler's arguments point to major flaws in innovation policy. Eszter Hargittai suggests that inequalities in the ability to participate may mean that these new technologies won't do as much to flatten social hierarchies as they might seem to. Jack Balkin claims that Benkler's book isn't so much about new modes of cooperation replacing market mechanisms, as existing side-by-side with them. Siva Vaidhyanathan argues that Benkler's book is guilty of a soft form of technological determinism, which overemphasizes the positive consequences of new technologies and implicitly discounts the less positive. Finally, Yochai Benkler responds to all of the above.

As with previous seminars, please don't comment on this introductory post, except to point out formatting glitches etc that need to be taken care of. If you have general responses, you should leave them in the comments section of Benkler's post. If you have specific responses to individual posts, of course leave them in the comments sections for those posts. Finally, if you wish to link to this seminar from your own blog, please link to the introductory post (as this links in turn to all the contributions).

## Norms and Networks

*By Henry Farrell*

*The Wealth of Networks* is a very important book, not only for people involved in debates on information and technology policy, but for the left as a whole. While it clearly builds on work by Larry Lessig, James Boyle, Pamela Samuelson and others on intellectual property and the public space, the real contribution (as with some of these other writers) is to a broader tradition of thought; that of people like Jane Jacobs, James Scott, Richard Sennett and Iris Marion Young. Benkler's vision of the good society is one in which people have a high degree of autonomy, so that they have the practical capabilities to pursue their own interests and their own forms of cultural expression. He argues that new communications technologies, if they're left unhampered, might radically increase this practical autonomy. They can do this by making it far easier for individuals to become producers as well as consumers of culture, and to share their cultural products with each other, so that they can build collectively on each other's work. On the one hand, technologies such as the Internet allow us to engage with each other in new ways, and to form networks of collaboration and of conversation, creating possibility conditions for the kinds of diversity and critical thinking that democratic theorists prize. On the other, these technologies do so in a relatively non-constraining way.

This is an important set of arguments for the left. It suggests that many of the values of the left can be achieved without either grand planning initiatives by the state, or the often stifling structures of community. Benkler is far from being a libertarian – he argues that the state should support education, health and other basic social goods as it has done in the past. More to the point, he argues that the new spaces of collaboration that he's interested in are important precisely because they open up new possibilities for organizing production than free markets. But even if he's not a libertarian, he's interested in freedom; he offers a set of practical prescriptions that are intended to create spaces for the free play of diverse interests and sensibilities. His arguments about politics seem to me to have more in common with Elinor and Vincent Ostrom's work on polyarchies than the claims of either liberal or communitarian political theorists. Benkler sees the state as ideally being an enabler of actions by individuals and loosely grouped communities of interest, rather than as a hands-on regulator. This is the source of his criticisms of the (US) state, which he sees as too eager to regulate these spaces, and as representing a set of interests (those of copyright holders) that are inimical to the spread of autonomy. The transition from an economy of cultural production that is based around a sharp distinction between producers and consumers, to one in which everyone is potentially a producer, will have clear distributional consequences. There'll be winners and losers. If the potential losers have sufficient political clout, they can hamper, and perhaps even block change. There's nothing inevitable about the emergence of the new economy of collaboration; depending on institutional configurations, it may be still-born.

If it's not clear already, I'm broadly in agreement with Benkler's normative claims and diagnosis. Still, I believe that there's an important piece missing from his argument. Benkler is a lawyer, and while he's concerned with institutions, his primary concern, naturally enough, is with formal institutions – that is government issued and government enforced laws and regulations, which can be interpreted by the courts. But the kinds of exchange that he's interested in rely less on laws, than on informal institutions and norms are what anthropologists call “gift exchange” and “generalized reciprocity.” If it's not clear already, I'm broadly in agreement with Benkler's normative claims and diagnosis. Still, I think that there's an important piece missing from his argument. Benkler is a lawyer, and while he's concerned with institutions, his primary concern, naturally enough, is with formal institutions – that is government issued and government enforced laws and regulations, which can be interpreted by the courts. But the kinds of exchange that he's interested in are what anthropologists call “gift exchange” and “generalized reciprocity.” They don't rely on formal rules so much as informal institutions and norms. Benkler doesn't completely neglect this - he discusses it at several points in the narrative, but doesn't say as much as he might about the internal dynamics of these informal rules and norms, and how they affect exchange. While we (by which I mean social scientists including anthropologists, sociologists, economists and political scientists) still don't understand as much as we'd like to about the ways in which these informal institutions can change over time, exactly such change is an important problem for the kinds of decentralized order that Benkler is interested in. We can't just look at changes in formal institutions, and how these may make technology-based cooperation easier or less easy. We also need to pay attention to the factors that might precipitate changes in informal norms. If informal norms change, then so too may we expect the kinds of cooperation that they support to change. And some of these changes may have negative repercussions for the kind and level of cooperation that we might expect.

We can see this if we look at the example of the blogosphere, which occupies a prominent place in Benkler's book (and, indeed, in my own thoughts). The blogosphere, as I see it, works as well as it does because of three key norms – linking, attribution and authenticity. The first norm – linking – is that when one wants to criticize, agree with, or otherwise comment on a particular item of source material, one should link to the source material if it's available online. The second norm – attribution - suggests that when one has come across a particular piece of source material thanks to another blogger, one should credit the blogger in question by providing a link to him/her (perhaps sending some traffic back in his/her direction). Sometimes both of these norms overlap – when the original source material is a post by another blogger, then linking to that post will satisfy the norms of both linking and attribution. But sometimes they don't. To take a topical example: when I posted recently on the Wealth of Networks, Matthew Yglesias wrote a [response post](#) which both linked to the Wealth of Networks home-page (satisfying the norm of linking) and to my original post (satisfying the norm of attribution).

The final norm (or, more precisely, set of norms) is a bit more difficult to pin down – it's a norm of authenticity. Roughly speaking, I take this norm to say that individual bloggers should represent their own points of view in an honest and

straightforward fashion. The comparative advantage of bloggers vis-a-vis other kinds of pundits is that they have (or should have) a strong personal voice based on their internal beliefs. This distinguishes the blogosphere from many other spheres of publication, where individuals are expected to represent the positions of their institution, or their political party rather than their own personal position on the issue at hand. It also distinguishes blogging from genres of writing (op-eds, speeches, political autobiographies) where authorship is blurred and ghost-writing by others than the official author are considered to be perfectly acceptable. Bloggers who are perceived as not representing their own position on the issues, or as having their material written for them by others, are likely to have a hard time getting their writing accepted by other bloggers.

None of this is to say that these three norms are universally adhered to in practice; few norms, if any, are. But they are adhered to frequently enough to shape behaviour and expectations. The norms of linking and attribution provide a basis for both generalized reciprocity and for more specific patterns of gift exchange. Bloggers who don't have many readers, but who feel that they have a post that might attract a large audience have an incentive to inform well known bloggers of their post, because this norm gives them a reasonable expectation that the well known blogger won't simply steal their idea. Instead, it is likely that the well known blogger will provide a link back to the original post, perhaps providing a massive (albeit perhaps temporary) influx of readers. Over time, this may allow less well known bloggers to become well known themselves. For example, Glenn Greenwald, after a series of posts that were well received among prominent left wing bloggers, has himself attracted considerable readership and attention for his blog, and for a book that builds upon arguments first presented in his blog, which is #1 in sales at Amazon.com as I write. Furthermore, as Jack Balkin has argued, norms of linking provide at least a minimum of exchange of views across ideological boundaries. In order to criticize people who say things that you harshly disagree with, you have to link to these people in the first place. There are limits to the efficacy of this norm in promoting exchange of views. It's not clear how many people actually follow cross-ideological links (I recall, although I can't find it, a post by Mark Kleiman where he said that approx. 1% of Glenn Reynolds' readers followed the link when Reynolds said something rude about him). Furthermore, many prominent bloggers on left and right link to other, less-well known bloggers critiques, rather than directly to the subject of the critiques. Even so, Balkin's point stands – the structure and norms of linking allow a degree of cross-ideological exchange that are nearly impossible for other media (such as print) to emulate. Finally, the norm of authenticity is what gives the blogosphere its flavour. I may think that some of the bloggers whom I argue with and criticize are political hacks. I don't usually believe that they're bought-and-paid-for political hacks (and if I did, I'd probably criticize them in very different ways).

If norms of attribution and linking didn't exist, then, many of the attractive features of the blogosphere would disappear. There would be much less bubbling up of ideas from small bloggers to big bloggers, and hence to a wider audience. There would also be less exchange of views across ideological boundaries. If the norm of authenticity didn't exist, then the blogosphere would be (for me at least, but also, I suspect, for many others) a much less interesting place, something much closer in spirit to the empty and

sterile back-and-forth of flacks and politicians on op-ed pages. Now it may well be, as Balkin argues, that “the customs of linking” (and perhaps, although Balkin doesn't claim this, the norms of authenticity) are relatively robust, because they “make sense, given the way that the technologies work.” But it seems reasonable to me to at least entertain the possibility that they're not robust in this sense, and that they could change if actors' underlying incentives change, or if all actors don't have the right incentives.

Two possible factors that might precipitate this change are money and invasion. First, money. Via [Michael Fromkin](#), this [argument](#) by Edward Leamer seems worth exploring.

We [academics] are part of a "Self-Organizing Collaborative Community" called the research universities of the United States and increasingly the rest of the world. Unlike contributors to Wikipedia and Linux, we get paid for our work, not by those who consume the fruits of our labor, but by taxpayers and by donors and by our students, all of whom we have convinced are better off by virtue of the research that we do. When it got started fifty years ago, this system worked great, but it isn't working as well anymore. While we are doing plenty of worthwhile research we are also doing plenty that isn't worthwhile, and the competition for research talent defined by the fads of the moment is driving up the cost of education to unaffordable levels. Adam Smith would have understood what's wrong here. It takes sales for the invisible hand to do its magic. Begging in your work clothes when you aren't working isn't enough, even though the pastime may be lucrative. On the contrary, the more lucrative is the begging, the more likely is the conclusion that the work is worthwhile. But it takes accurate market prices to tell us what's valuable and what's not. Of course, good will and good intentions can carry a collaborative community productively for a while, but financial rewards relentlessly bend the system to their will, slowly perhaps, but inevitably. That's the invisible hand at work. Thus, open-sourcing has the same problems and the same probable longevity as the communes of the 1960s -- they worked great for a while, but the participants chose other ways to live once they got to know the people in the community.

There are two parts to Leamer's argument, a normative claim and an empirical claim, which blend together in a confusing way in this paragraph. The normative claim is that the price mechanism is superior to other metrics in telling us “what's valuable and what's not.” This can be swiftly disposed of – it's a common trope among economists, but one that is, at the very least, controversial as a general argument. The second claim is much more troubling: his suggestions that “good will and good intentions” are inevitably going to fold in the face of financial rewards. We don't have to accept his normative claim, or indeed his suggestion that this is an ineluctable social process to agree that this may be a real problem. Even if we want to celebrate the possibilities of collaborative communities, we can realistically worry that they may be vulnerable to disruption by market forces, given historical examples of how decentralized systems of collaboration have been destabilized by markets in the past. To take one example, which Benkler mentions in passing, amateur sports, in the nineteenth century sense, have given way to

professionalization in very nearly every field of sporting endeavour as star athletes have gone after the money. The point here is not to celebrate amateurism uncritically (there was a lot that was shoddy about it – distinctions between players and gentlemen etc), but to use it as an example of how collaborative fields of endeavour may be disrupted by financial pulls from outside.

How could the norms structuring the blogosphere be disrupted by money? One plausible way might be as follows. Some bloggers earn quite substantial amounts from advertising, and have tried to capitalize on this through creating miniature publishing empires, such as the Denton publishing stable, the soft porn Suicide Girls blogsite and so on. These commercial enterprises have very different incentive structures than not-for-profit blogs (or indeed blogs that only earn pocket money revenues from blogging). Specifically they have much weaker motivations to provide external links. They want their readers to stay with them (or to go to other blogs with the same owners, or to websites that are prepared to pay for external links). They're in it for the money. This suggests that they have far less incentive to follow the norms of linking and attribution than other blogs do. All other things being equal, as time passes, we might expect that there will be less and less external links from these blogs to other parts of the blogosphere. We might also plausibly expect a weakening of the ethic of amateurism, and something much closer to the ethos of the traditional print media emerging. In short, we might expect an unwillingness on the part of for-profit blogs to acknowledge other blogs by linking back to them, and a consequent degradation of the dense networks of communication that allow interesting ideas to percolate rapidly through the blogosphere. They would become more like newspapers, which are notoriously unwilling to acknowledge the scoops of competitors.

Now the crucial clause in the argument above is “all other things being equal,” because other things aren't equal, at least at the moment. We haven't seen this become a major problem among political blogs, in part because nobody has really figured out a successful way to consistently make large amounts of money from blogging. So far, efforts to create blogging conglomerates such as Pajamas Media have been dismal failures. But this isn't a structural necessity, and it may change in the future. More generally, the key point is that the success of the blogosphere as a network doesn't just rest on the technological possibilities offered by hyperlinking. It also rests on a set of norms about when hyperlinks are appropriate, and there is no *ex ante* reason to believe that these norms are invulnerable to change in actors' incentives, or that actors' incentives mightn't change in the future.

The second vulnerability that the blogosphere faces is that of invasion. Here, the norms most at risk are those surrounding authenticity. Blogs prize authenticity – but they don't have any very good means of authentication. The blogosphere is highly open to newcomers, but this also renders it vulnerable to invasion by actors who don't share the social commitments that most bloggers have. Indeed, we've seen evidence of something like this already, in the degradation of two aspects of the networked blogosphere. First has been the invasion of comments sections by spammers. I'm not one of the first generation of bloggers, or even the second, but I have been around long enough to

remember when comments sections were open to all comers. Unfortunately, exactly this openness made blog comments sections vulnerable to spammers who wanted to increase websites' Googlejuice by inserting bogus linkbacks into comments. As a result, open comments sections are relatively rare in the blogosphere today. Most blogs have at least a minimal set of hoops that commenters need to jump through before they're able to leave comments; some have strong registration requirements. While there isn't any ideal technical solution for comment spam, at least there are acceptable workarounds. The same, unfortunately isn't true for the second aspect of blogging which has been degraded - the Trackback system. Trackback, which when it worked provided a real reflexivity to conversations, by allowing the readers of a blogpost to see which other blogs had linked to that post, and follow links back to them. Trackback is almost completely unusable thanks to spam links these days, and alternative workarounds (such as using Technorati or other services) are considerably inferior.

How might invasions of this sort threaten authenticity? Through a specific kind of invasion - astroturf blogs - which I suspect are going to pop up like mushrooms everywhere during the coming electoral cycle. The example of the Thune v. Daschle race, where individuals paid by the Thune campaign ran purportedly independent blogs, which apparently helped sway local media coverage, suggests that there's a lot of scope for astroturf blogs – i.e. blogs that claim to represent authentic grassroots opinion, but are in fact written to order by paid operatives working on behalf of particular interests or political parties. As blogging becomes more closely integrated into campaign activity and public controversies, we're likely to see more and more dubious use of blogs to inject slurs and falsehoods into political debate. It seems to me that prominent bloggers like Kos and Jerome Armstrong have negotiated the tricky boundaries between political consultancy and blogging with admirable honesty – but it's not at all clear that they're going to be the rule rather than the exception. None of this is to say that these problems are at all unique to blogs – but they're likely to plague blogs in particular, because of the very low barriers to entry. I'd like to think that astroturf blogs will be less likely to succeed in winning attention than the genuine variety, but I wouldn't want to bet my life on it.

None of these problems is necessarily insuperable. There's no ineluctable historical reason why systems based on gift exchange and generalized reciprocity have to give way to markets. The two co-exist, and systems of informal trust and exchange are sometimes reproduced at the heart of capitalism (cf. the sociology of stock markets and of diamond merchants). As Benkler points out when he talks about Slashdot's system of moderation, technology can help at least to some extent. Furthermore, more elaborate forms of self-regulation of the kind found on Wikipedia can help too. But in order to understand the problems, and how they might be solved, I think that we have to pay attention not only to the formal rules that might stymie free exchange of ideas and cultural products, but also to how best to support the informal norms and institutions that structure this exchange. These informal rules too are subject to change over time, and may, in the absence of considered responses, change in ways that make conversation and interplay less likely or less attractive. It may well be (and I don't want to overstate my case here) that the forms of exchange that Benkler rightly celebrates are more vulnerable

than they appear to be at first, even in settings where there isn't formal regulatory interference.

Furthermore, to the extent that problems do arise, the best available solutions (which I take to be self-regulatory ones) also carry costs. Self regulatory solutions are likely to help reduce the risks of bad behaviour, but they also, by the same token, reduce spontaneity and reintroduce a certain level of hierarchy. The structures underlying Wikipedia seem to me to be becoming more rigid over time. I suspect that the same is going to be true of the blogosphere too over the next few years. Less freewheeling conversation; more codes of conduct, implicit guild-style self-regulation and development of technologies to inhibit bad behaviour. This seems to me likely to happen – but it's hard not to lament what we would lose in the process.

## A General Theory of Information Policy

By Dan Hunter

Yochai Benkler has been working on questions of intellectual property policy and telecommunications policy for some time now, but his work has always seemed to operate slightly orthogonally to the *canon du jour*. He has always been respected as probably the smartest guy in the room at any given conference, but when the hot topic was, say, the problems with the political economy of copyright, Benkler was [talking about spectrum commons](#); or when everyone wanted to talk about digital rights management or search engine bias, Benkler was talking about [autonomy as a normative basis for information policy](#), or the [social production of knowledge](#). It wasn't immediately apparent what were the connections between his interests, although it was clear that there was a theme that bound them all. More generally, those who thought about the myriad issues within internet policy, telecommunications regulation and intellectual property could see that all of them might be parts of a grander whole, and wondered what that whole might look like.

In [The Wealth of Networks](#), Benkler demonstrates how his interests are connected and how they have always been part of a grander vision. But he also provides something close to a General Theory of Information Policy for the networked age that begins to explain how we should think about topics as different as spectrum policy, copyright, user-generated content, network neutrality...well, the list pretty much encompasses all questions within internet law and policy. For someone used to reading legal scholarship in the cyberlaw arena, this book is remarkable. Academic legal writing tends to be driven by current policy problems—in the cyberlaw arena obvious recent examples include peer-to-peer filesharing and the music industry, the threat of digital rights management, the economics of copyright term extensions, problems with the anti-circumvention provisions of the Digital Millennium Copyright Act, and so on. This is not a criticism of the legal scholarship; indeed it is one of the great strengths of legal writing that it responds to the immediate concerns of society in a way that can guide policy and law making.

Benkler's approach is different, much more in the vein of, say, Manuel Castells. He tries to articulate the changes in society generated by access to a network that allows about a billion people (and growing) to connect with each other virtually free, virtually instantaneously, and without typical expected barriers to entry and transaction costs that have driven our approach to policy to date. The social changes rendered by the internet is by now a venerable topic, and we've had ten years of cyberutopianism (the net will change the way people think, work, sleep, play, eat, drink, etc, etc) followed by the inevitable response of cyberskeptics patiently explaining (as though to impaired children) that the net changes little, the economics remains the same, governments will still regulate it, big business will still dominate it, etc etc. Benkler manages to revive the initial utopianism of the early days of the net, but does so in a way that captures the

revolutionary opportunity of socially-produced information goods, without lapsing into the kind of boosterism that characterized early cyberutopianism and the ensuing dot com bubble.

The central argument of *The Wealth of Networks* is that the “digital networked economy” allows for a new modality for the production of information goods—a modality that Benkler usually terms “peer production”—that has hitherto been unrecognized by economists, policy-makers and scholars. Benkler begins by providing a range of examples of socially produced information goods, such as the Wikipedia, open source software, blogs, and so on. He explains how the modality that underlies the production of these goods is opaque to most theories of economic development, and makes the case for more careful attention to the significance of this modality. He also takes enormous effort to defend this modality and his vision of the internet from the myriad arguments that incumbents and skeptics might make against them. Thus, he defends “peer production” from concerns generated by economics and liberal theory. If you have ever wondered how open source might be economically sustainable, then you will find the answer here. You will also find arguments about the role of distributed networks in justice, international development, and freedom. It is here that his work moves from the descriptive to the normative, and so it’s possible to make the case that this is a general theory of how the internet is and should be.

It’s not hard to agree with [Larry Lessig’s assessment](#) that *The Wealth of Networks* is the most important recent book in the field. It generates a series of new questions about the nature of information and the policies that we have to regulate and produce it. But that said, it’s not without problems. Aside from the usual quibbles that reviewers always have—Was it really a good idea to draw an implicit comparison between the significance of Smith’s *The Wealth of Nations* and this work?—two particular problems stand out. First, the book suffers from a lack of focus. One could take at face value that the book is about social production as an alternate modality of production; but this view has some problems. Notably, if read like this then one starts to wonder what isn’t produced by social production processes. Everything gets viewed in this light: so a discussion of NGO contributions to agricultural practice in the developing world is presented as being primarily about peer production when it seems to me to be a fairly pedestrian example of underproduction of certain types of goods as a result of the limits of private markets, and hence the need for public subsidy for this sort of good. This is not to say that any of the examples presented are wrong; in some form they all illuminate the potential of social production of various types of goods. But I don’t think it helps the argument to view so many types of innovation as being about social production when the link is tenuous.

There are other examples (the scale question of networks, the history of radio ownership, a critique of Sunstein’s ludicrous internet polarization thesis, etc etc) which made me wonder whether an alternate view of the book might not be more accurate: that it is a series of observations about the character of the internet, how it has a distributed nature that generates particular opportunities and dangers for policy, and what our policy responses should be. This is not inconsistent with the “peer production” thesis; it simply

means that the observation about peer production of information goods is but one of the cyberspace policy questions that Benkler wants to address. This reading of the book indicates the scale of Benkler's ambition, but does demonstrate that the scope of the book is broader than indicated in the introduction (and subtitle).

The second serious concern is not with the thesis but rather with the presentation, and specifically the strange view that the book seems to have of its audience. There are only a small number of endnotes for each chapter, and at best these mention (briefly) a couple of the authors cited in the text. And technical acronyms like "http" or "html" are defined in brackets, and PayPal is patiently explained to be "a widely used low-cost Internet-based payment service". This approach signals that this is an accessible, populist work, like others in this field by people such as Larry Lessig (*Code, The Future of Ideas, Free Culture*), Siva Vaidhyanathan (*The Anarchist in the Library*) or Cass Sunstein (*Republic.com*). But readers without a solid grounding in economics, liberal theory, political science and jurisprudence (and possibly network theory and internet architecture) are going to struggle through the book's 500-or-so pages. It's not an easy read, and Benkler doesn't pander to the audience by dumbing down the arguments. Many sections are elaborately worked through, addressing all manner of likely concerns about the thesis presented that, for example, the internet demonstrates an intriguing scale of granularity for the purposes of democratic engagement, or why a number of the typical accounts of autonomy are consistent with Benkler's view of social production. But this means that the book is really an academic work got up to look like *Wired* magazine. This poorly serves the scholarly readers because, although Benkler references the authors and theories on which he relies, many of his positions are controversial and reliant on arguments and positions that are eternally contested. Benkler is exceedingly fair in his account of, say, the principles of autonomy, redistributive justice or welfare economics. But each of these concepts admits of a range of views that simply cannot be dealt with in a scholarly fashion in a few pages and a couple of footnotes.

Unfortunately this approach also poorly serves the interested non-academic reader who—thinking this work is exactly what she needs to understand why [Ohmynews](#) is so important, or why [Wikinews](#) isn't working—picks up the book expecting it to be both smart and accessible. It's definitely smart, but it's not built for the general reader. The danger is not so much that the reader won't understand what is going on, so much as that she will wonder why she doesn't get it. Benkler's prose is routinely described as "dense", which I once took to mean "if you don't understand this then you're dumb and it's your fault". I fear that too many readers have the same fear that the problem is with them, that they're too dull or lazy or busy to take in the lessons here. It's not impossible to follow what Benkler is saying, but I wished that he would say it directly. It's clear that he can. The text fairly comes alive when it talks about the recursive/iterative moderation system behind Slashdot, or when discussing the role of blogs in the Trent Lott, Diebold, and "BoycottSBG" cases. But there are large swathes that are a challenge for those who, like me, are lazy.

It's churlish, I know, to criticize a *tour de force* on the basis of difficult expression. (And I have deleted and reinstated this paragraph about ten times as I try to

decide whether to say anything about it; I still wonder whether it's worth commenting on this). If the work signaled that it was only scholarly, and was only talking to the limited range of scholars who address these issues professionally, then this would not be a failing and I wouldn't say anything. But I'm worried that too many of the peer-producers—the blog writers, the open source software gurus, the amateurs who create for the love of it; in short the people who this book is written about—will pick up this work in the hope of understanding how their creativity fits into the grand scheme of innovation, and what their role will be in the amateur production sphere that promises to change the way that we view information goods within society. And they may not get past the introduction.

I hope that they will try. This is a revolutionary book, in all sorts of ways. The future of innovation and creativity is different from the past. Amateur hour looms: we are entering the era when commercial producers of content will no longer dominate. We should be grateful that we have Benkler's exhaustive account of the significance of this radical change.

## Why Do Social Networks Work?

*By John Quiggin*

Yochai Benkler's *The Wealth of Networks: How Social Production transforms Markets and Freedom* is full of interesting things to discuss, but the point that interests me most is the question of why people contribute to social production and what economic and political implications it has, for states as well as for markets and freedom. Benkler previously discussed the same question in *Sharing Nicely*, and I'll talk a bit about this as well.

To the extent that there is a conventional wisdom about these things, it's Eric Raymond's idea of a gift economy, derived from his participation in the open source software movement. Raymond focuses exclusively on reputation as a motive for contributing to social production, seeking to assimilate all other motives (such as craft values) to reputation. It's precisely this kind of totalising logic, I'll argue, that is absent from social production.

On the contrary, there are a wide variety of motives which might lead people to contribute to networked social capital, for example by participating in various aspects of blogging (make posts and comments, linking and blogrolling, improving software, various kinds of metablogging). Possible motives include altruism, self-expression, advocacy of particular political or social views, display of technical expertise, and a desire for social interaction. Particularly in relation to a collective, and largely anonymous, product like Wikipedia, Raymond's central focus on reputation does not fit the facts.

As against these various motives, there are the two standard methods that have been relied on to deliver most information production and innovation for the past 150 years: markets and bureaucracies.

Benkler discusses markets, and the associated profit motives, at some length, making the point that narrowly economic motives tend to crowd out alternative forms of motivation. He mentions the classic work of Titmuss on blood donations and some other examples to show that monetary and social-psychological motives are likely to conflict, rather than reinforcing each other. By contrast, different social-psychological motives are usually complementary or at least mutually consistent.

Why is this? At a superficial level, it's obvious that people act differently, and are expected to act differently, in the context of relationships mediated by money than in other contexts. Behavior that would be regarded favorably in a non-monetary context is regarded as foolish or even reprehensible in a monetary context.

One of the most important general differences relates to rationality and reciprocity. In a non-market context, careful calculation of costs and benefits and an insistence on exact reciprocity is generally deprecated. By contrast, in market contexts, the first rule is never to give more than you get.

This rule applies in market contexts but not in social contexts, where such careful calculation is, as Benkler notes, generally deprecated, because markets create opportunities for systematic arbitrage that do not apply in other contexts. In an environment where exchanges are not carefully calculated, a trader who consistently gives slightly short weight can amass substantial profits. If trading partners assume honourable behavior, none will suffer enough to notice, but eventually arbitragers will drive out their less calculating trading partners.

Similar points can be made about other motives. There are a whole range of sales tricks designed to exploit altruism, friendship, desire for self-expression and so on. Hence, to prosper, or even survive, in a market context, it is necessary to adopt a view that 'business is business', and to (consciously or otherwise) play a role as a participant in the market economy that is quite distinct from what might be conceived as one's 'real self'. This is a prime example of what Goffman calls an obligatory role.

The crucial feature of economic motives in a money economy is not that they are less noble or desirable than alternatives such as desire for fame, but that a money economy provides a total system of rationality, from which most of the motives associated with social production are excluded.

Markets are not the only total system of rationality that operate in a modern society. bureaucracy and the state have a logic of their own. For most of the 20th century, the central issue of politics and economics was the question of where the boundary between markets and bureaucracies (public and private) should be drawn.

Benkler largely ignores the state. As he says (p. 16) 'In much of [my discussion], the state plays no role or is perceived as playing a primarily negative role, in a way that is alien to the progressive branches of liberal political thought'. But this position overlooks the critical fact that both the Internet and the World Wide Web were developed primarily by state agencies or state-funded institutions (DARPA, NSF, the university sector, CERN, NCSA and so on). Yet this outcome was not the product of rational bureaucratic planning. Rather, like Topsy, the Net and the Web 'just grew'.

Like market rationality, bureaucratic rationality implies a complete specification of behavior. When dealing with a representative of a bureaucracy, we (mostly) expect consistent application of rules, rather than an adherence to standard social norms such as 'look after your family/mates before others'. Similarly, and more crucially, bureaucracies are supposed to allocate their resources to the achievement of specified goals, not to do things because they would be fun, or even socially beneficial. All of this seems to leave little room for social production. So how did the state come to give us the Internet? More

significantly for our present purposes, what kinds of public policy will facilitate the further growth of social production and the wider distribution of its benefits?

I don't have a complete answer to either question. However, some obvious implications run counter to current developments in policy.

First, if monetary returns are weakly, or even negatively correlated with the value of social production, there's no reason to expect capital markets to do a good job in allocating resources to supporting innovation. It follows that the policy orientation of the past thirty years, in which increasing reliance has been placed on capital markets, is going in the wrong direction. The examples considered by Benkler are illustrative. Both blogs and wikis have their roots in the late 1990s, a time when capital markets were splashing hundreds of billions of dollars around on Internet-related projects. Most of these projects came to nothing, while blogs and wikis developed with little or no venture capital money to help them along.

A second implication is that the policies of New Public Management, which attempt to tighten bureaucratic accountability and focus on competitive disciplines and measurable outcomes may be misguided. Rather than seeking to drive people harder in the search for increased productivity, government macro-economic policy should be oriented towards making room for creativity and facilitating its expression. Similarly, while competition has its place, public policy should be at least as much concerned with promoting co-operation. The assumptions we have had about the competitive nature of innovation are, therefore, undersupported in the new environment in which we find ourselves. If governments want to encourage the maximum amount of innovation in social production then they need to de-emphasize competition and emphasize creativity and cooperation.

## Whose networks? Whose wealth?

*by Eszter Hargittai*

A good book raises at least as many questions as it answers. By focusing on an issue that does not really get addressed by Yochai Benkler in the *Wealth of Networks*, I do not mean to suggest that the author should have written the book I would like to see written (in fact, I am working on it so it is just as well:). Rather, I would like to suggest some issues that are worth considering in the context of this discussion.

Before I present my comments, I would like to note that I very much enjoyed reading this book and think it will appeal to and be of value to a wide range of people. I like the mix of tackling large important questions while getting into the nitty-gritty of what is happening on the ground and illustrating points with very detailed and careful examples. The book serves as a helpful reference to various online discussions and events that have occurred in the past few years and it is useful to have them documented and linked together suggesting that they were more than isolated occurrences of how people are using digital technologies for political and social purposes. The other contributions of this seminar tackle the big questions addressed directly in the book. I have decided to address an aspect of the issue that is not addressed in detail yet I believe needs to be part of the conversation.

The point I want to focus on here is the unequal distribution of opportunities discussed in the book and what potential consequences this may have. That is, to what extent are the benefits of the possibilities raised by recent innovations distributed equally among different segments of the population and to the extent that they are not, what implications might that have?

To be sure, the author addresses this point in Chapter 7 as one of the expected critiques of the enthusiastic claims. Under the heading “Digital divide” (pp.236), the author acknowledges that distribution of digital media is not equal in society. He counters this concern by noting that (1) access differences in the US are much less significant than in the 1990s; (2) “growth rates among underrepresented groups are higher than the growth rate among the highly represented groups” (p.237); (3) “the democratizing effects of the Internet must be compared to democracy in the context of mass media, not in the context of an idealized utopia” (p.237). While these claims may all be true, I will address them here point-by-point, because I think they still leave room for discussion and concern.

First, while it is true that differences in access to the Internet by various population segments in the U.S. (and many other countries) have declined since the 1990s, we do not have much evidence showing continued strong growth. In fact, [some data suggest](#) that active home use is in decline in several nations, including the United States. While home use is not the only type of access to the Internet, it is not hard to

argue that in most cases it is likely the most beneficial type of access as in many cases it will mean the highest level of autonomy in using the medium. That is, having to rely on connectivity at a library or school that is only accessible during limited hours, is located miles away, does not offer much privacy and may be limited by filters is a far cry from 24-hour access in the privacy of one's home. I do not mean to exaggerate this point and do realize that for some people using the Internet at home may afford less privacy than otherwise, but those are likely the less frequent cases. Some work I have done suggests that home use is a statistically significant predictor of how knowledgeable people are about the Internet, which then predicts the diversity of uses to which they put the medium. All this implies that differences in access and access type continue to exist despite some changes over ten years.

Second, it is not saying too much that underrepresented groups are getting connected at higher rates than those who are already online in high numbers. After all, groups whose members are already online at 80+% only have so much further to go while groups whose members are connected at 15-20% have the potential to grow considerably. In many ways this is just a numbers game, and numbers can be interpreted in lots of ways. So while the rate of increase may be higher for an underrepresented group, that does not automatically result in high penetration levels or even a much better relative position. (For a detailed discussion of this point, take a look at Figures 5 & 6 and the accompanying discussion starting on p.15. of [this paper](#).)

The third point regarding the relevant baseline comparison of the traditional media landscape versus some utopia seems less problematic, and it is certainly a point I like to call attention to as well when people challenge the limits of how digital media are being used today. That said, in the context of discussing inequalities, I think this point deserves a pause as well. The assumption seems to be that thanks to new opportunities, *everybody* will have more of a voice and input. However, is it not possible that inequalities may actually increase if some people are much more likely to take advantage of these opportunities than others? That is, if those who are already in more advantageous positions (e.g. higher socio-economic status) are more likely to use new tools and if they do so in ways that do not benefit everybody, then the disengaged may be left behind even more.

Instead of focusing simply on access figures, I have argued for years that we need to pay just as much attention to differences in uses and abilities. To this, we can add the notion of differences in participation. There are very different ways of being an Internet user and being counted in baseline statistics of the connected. However, large chunks of people who are online never use their connectivity for certain types of activities such as political engagement.

My work focuses on empirical studies of people's Internet uses so the above comments are not hypotheticals about differences in abilities and types of uses. I have published [several papers](#) that discuss skill differences among users. I have also recently collected some data – not yet published – about college students' Internet uses. This data set is based on surveys administered on 1,300 college students' (mostly first-years) in

February and March of this year. ([click here for some sample descriptives](#)). Students are a representative sample of those attending the University of Illinois Chicago campus.

It is hard, if not impossible, to come up with just one or two survey questions that would capture the extent to which young people – the most connected age group – are taking advantage of the medium in democratizing ways. Nonetheless, there are some revealing and helpful measures. For example, when asked how often they visit “blogs, discussion forums or other sites that allow you to interact with other participants”, the topic category of “politics, economics, law, policy” got the lowest rating among types of interactive sites students visit (the other options of sports, technology, arts & culture, and personal journal all ranked higher in popularity). Only five percent visit such sites daily and 63.5 percent never visit interactive sites on the topics of politics, economics, law or policy. When asked whether they have *ever* visited certain sites - undoubtedly the most conservative measure of any level of engagement with these blogs - only one percent say that they have ever visited Instapundit or Daily Kos. These participation rates are so low that it is not even possible to try to determine what types of students are most likely to engage in such online activities.

When looking at use of the Internet for different types of activities, it turns out that certain college students use the medium for statistically significantly fewer activities than others. For example, women, African-Americans, students whose parents have lower levels of education and those with lower levels of knowledge about the Internet do less online than others. So even if diffusion figures increase and access gaps are closing, use gaps remain and are not randomly distributed.

While Benkler is right to be enthusiastic about some of the opportunities made possible by new tools and services in the digital landscape and the possible implications of their implementations, it is important to put all this in the context of actual usage figures. There is no evidence to suggest that equal diffusion of all services to a broader public is simply a matter of time. So while it may be the case that a larger number of people get to participate in political and other social action, how much flattening can we expect if the already privileged continue to be the ones most likely to embrace and thus potentially benefit from the new opportunities?

## **Mediating the Social Contradiction of the Digital Age**

*by Jack Balkin*

There is a social contradiction to the digital age, a fancy way of saying that two forces are pulling in opposite directions. The lowered cost of creating, collating, copying, sharing, and transmitting knowledge and information goods pulls toward democratic participation in information production and free access to information. These features of the digital revolution mean that more and more individuals can participate in the benefits of the information economy and in the production (and co-production) of information and knowledge goods.

But the very same features and forces that make the digital revolution possible mean that knowledge and information goods become an increasingly large component of global wealth and power. Moreover, the digital revolution and associated telecommunications technologies allow greater investments in knowledge goods and expanded markets to recoup those costs.

These features of the digital revolution give some businesses strong incentives to reap as much profit as they can by propertizing knowledge and information goods and extracting rents. These incentives are made all the more urgent because knowledge and information goods have relatively high first copy costs and low to zero marginal costs. Hence, for many businesses, increasing monopolization and propertization of knowledge and information is a particularly attractive strategy, a strategy that they have pursued through successful lobbying efforts at both the national and international level.

Thus, at the very moment when the digital revolution holds out the promise of genuine democratic participation, businesses driven by the twin needs to maximize profits and protect themselves from competition have tried to assert control over the knowledge economy through expanding intellectual property rights and securing legal protection for proprietary architectures, undermining the Internet's democratic promise. This collision of interests is not accidental: Industrial, closed and proprietary models of information production and democratic, open, and commons-based models are made possible by the same technology; the struggle between these two models of information production is the social contradiction of the digital age.

How can these opposing trends be reconciled? Yochai Benkler's argument in *The Wealth of Networks* is that the contradiction can be resolved by two features of the digital revolution. The first is that not all successful business models in the knowledge economy have rested or will rest on maximizing the exploitation of intellectual property or closed and proprietary architectures. The second is that the digital networked environment makes possible and gives increased salience to commons-based peer production methods for information production. In both cases, but especially in the

second, democratic participation in information production is wholly consistent with efficient economic production and the growth of the knowledge economy. Indeed, preserving a space for democratic participation in the means of production is the best way for the knowledge economy to flourish.

The terms "social contradiction" and "participation in the means of production" might make you think of Marx and the socialist ideal. But Benkler is not talking about centralized or state-planned information production; rather, if anything, he is urging further decentralization and privatization of the methods of information production. It is important to remember that in Benkler's commons based production, people still own their own property. (For example, people access spectrum commons using their own computers). Individual property and liberty rights are simply allocated differently than they are in closed or proprietary architectures. Benkler's key claim is that allocating property rights differently leads to both greater participation \*and\* greater economic productivity.

In Chapter Five Benkler offers the parable of three different methods of producing public entertainment: the worlds of Reds, Blues, and Greens. Reds are hereditary storytellers who have the duty and the privilege to tell stories to the rest of the population. Blues elect a new storyteller every night by majority vote, and this person tells stories before the whole community. Finally, in the Green community, people tell stories to each other all the time whenever they feel like it.

In this parable, the Reds represent pre-modern hierarchical methods of information production, the Blues represent modern mass communication driven by market forces, and the Greens represent commons based peer production. Benkler thinks that the Green world is a better world. It offers more autonomy and more opportunities for more people to participate in cultural and information. Moreover, by using the example of storytelling-- a form of entertainment that is now fully commodified in our era-- Benkler is also adverting to the fact that there is no clear boundary between democratization of economic production of information and democratization of the realms of culture and democracy.

The parable of Reds, Blues and Greens exemplifies the hopeful and optimistic message of the book. It also allows us to see the wager that Benkler is making on the future. The question is one of tradeoffs. Will moving from a Blue world of proprietary and industrial information production to a Green world of open architectures and commons based peer information production trade greater participation for less quality? Will commons based peer production of information goods gain equality at the cost of efficiency? Will a Green world be a world of mediocrity-- whether in stories, or software, or telecommunications access-- that is more equal but less efficient and contains less that is excellent?

Benkler considers and rejects the argument that the Green world will produce substantially less quality, arguing that commons based peer production can produce its own methods of accreditation and quality control. He also rejects the assertion that

efficiency will be sacrificed. Although the types of stories produced may be different in the Green world than the Blue world, they will be just as good for the social purposes that these stories serve; and they will gain the added advantage of producing more stories told by more people, who will become more than mere passive spectators (as in the Blue regime) but also active producers of their informational world.

This is, in fact the book's great leap of faith-- the argument that peer production is an equally efficient (or sometimes even more efficient) method of producing information goods that is also superior in terms of autonomy and participation.

The burden of the book is to demonstrate this optimistic claim-- that democracy and efficiency are served by commons based production methods without significant tradeoffs-- in a wide range of areas, ranging from software to culture to journalism to agriculture. But the best argument for Benkler's position is that the choice is not as stark as his presentation sometimes suggests. In fact, the choice we face is not between the Blue world of proprietary, industrial information production and the Green world of commons based peer production. For the industrial model is not going away anytime soon in any of the areas that Benkler identifies. For example, as I have pointed out in the area of free expression, the blogosphere does not displace traditional media outlets; rather it routes around them and "gloms on" to them; it engages in nonexclusive appropriation of information and knowledge produced through industrial methods and uses elements of them to make a wide variety of new things with a wide degree of cultural participation. Routing around and glomming on is simultaneously a form of independence and dependence, of subversion and homage, of breaking away and bricolage.

We can make similar arguments about many of the other forms of information production that Benkler identifies; they do not displace industrial methods of information production but rather build on top of them, sometimes leading to new synergies and syntheses between older and newer forms. This is already happening in the realms of journalism and mass culture; I would not be surprised if we could not tell similar stories in each of the other areas that Benkler discusses.

Thus, the real choice we face is not between a Blue world of industrial information production and a Green world of commons based peer production. It is rather a choice between a largely Blue world dominated by a relatively small handful of players and a Blue-Green or Teal world with substantial contributions from a wide range of actors using commons based forms of information production layered on top of industrial forms that respond to and adjust to the commons based forms. How that layering occurs will be different in the realms of software production, telecommunications, scientific research, agriculture, journalism, education and entertainment. But in each case we will have multiple methods of information production, that, one hopes, will bring us the best of both worlds. This is perhaps the strongest reason to believe that Benkler's optimistic vision is also a realistic one.

But we are by no means assured of this happy outcome. I have already noted that social contradiction of the digital age features a strong pull toward proprietary and

industrial models of information production backed up by law. As a result, Benkler's book wavers between an optimistic description of what the digitally networked economy has produced and will produce and a warning that these bounties will be squandered if the legal regime goes in the wrong direction-- the direction in which it is currently headed. This alternation between prediction and warning is not accidental, for much as Benkler might wish that features of the digitally networked environment and information economics will lead us inevitably toward a blessed world of peer production, he well understands that the political economy of information production has repeatedly pushed the law along a different path. Benkler shows us a vibrant world that we are moving toward and might yet achieve; but it is up to us to realize it.

## The Dialectic of Technology

*By Siva Vaidhyanathan*

Yochai Benkler has given us a comprehensive, complex, and persuasive work in *The Wealth of Networks*. There is no better place to turn for an account of the processes of creativity and commerce relating to digital networks and the work that people do with them.

Benkler is the most original thinker on matters digital writing today. If there are observations in this book that seem familiar, it is only because others (including most clumsily, myself) have poached Benkler's shorter works and enthusiastically rephrased his ideas for common consumption. But this book is certain to last and displace all pretenders (including, sadly, my own work). It should be the first place scholars and students turn to understand the radical changes we have encountered in the culture and economics of information in recent years.

But what of chickens and eggs? Have we generated and proliferated these powerful technologies because we desired them to "extend" (as McLuhan might say) our corporeal and social capabilities? Or have these technologies pushed us to new states of consciousness and new relationships that we could not even imagine before Netscape's IPO?

These are questions I have been pondering for years. I don't have clear answers myself. Perhaps I am digging myself a philosophical hole in my attempt to grasp the dynamic relationships among developers, users, regulators, and vendors in the distributed and delicate digital environment. So please forgive me if I seem to have hunted hungrily through Benkler's book seeking wisdom and guidance in my search for answers. He has granted me much clarity to me before. So I might be guilty of holding him and his book to unfair standards. This one issue remains underwritten in the text: the story of the technology itself. Throughout the text, there seems to be an almost givenness about the technology. TCP/IP is just there. Even Cisco's notorious discriminating servers, the source of so much tension over [the end of network neutrality](#), just appear (p. 146-161). We get no sense that particular technologies are malleable, adaptable, contingent, and [socially shaped](#). We get no account of developer's wishes or users' adaptations. We only get cursory accounts of the conflicts over the future of these technologies that have unleashed (to choose a loaded term) so much creativity.

Benkler does, however, offer a general manifesto on his softly deterministic view of the relationship between technology and human affairs. I find this account (pp. 16-18) generally inadequate.

Benkler opens the section by declaring technological determinists like [Lewis Mumford](#) and [Marshall McLuhan](#) out of fashion in the academy. While he is correct, and he invokes the more solid account of technological influence that we get from [Elizabeth](#)

[Eisenstein's](#) *The Printing Press as an Agent of Change* as a model, he declares Paul Starr's [The Creation of the Media](#) as the model of a "political-choice-based" explanation of technological change. But he does not take us there.

While Eisenstein's book is foundational, we have moved beyond it. We have supplemented (if not displaced) her account of the changes wrought by the printing press in Europe. Subsequent works such as [Adrian Johns'](#) *The Nature of the Book* and [Roger Chartier's](#) *The Order of Books* have heralded a sociological account of the relationship among writers, readers, printers, and vendors. Such relationships, in the work of these two historians, are dynamic and contingent. So are the technologies they consider.

The work of technology in *The Wealth of Networks*, I am afraid, echoes Eisenstein and McLuhan more than it does Johns or Starr. Not that there is anything wrong with that. But the next account of the digital revolution (and I agree with Benkler about the revolutionary implications of these technologies and our uses of them) must engage with a more philosophical and historical account of the specific technologies.

Consider the central claim that Eisenstein raises in *The Printing Press as an Agent of Change*: The very fixity of the printed word altered hierarchies, values, and habits of Europeans. It changed minds. Thus popes and kings fell. This is fine, as far as it goes.

But then consider that the central technological claim that Benkler invokes in *The Wealth of Networks*: the digital revolution has ignited a powerful global culture of flexibility. Texts, codes, machines, and expressions are all in flux in this book. So are the minds of those who engage with them. That's the source of so much wealth and creativity. Again, it's a powerful case.

But that's not the whole story, is it? What of efforts to retard the flux? What of challenging technologies like [digital rights management](#) (DRM) and [trusted computing](#)?. These issues get play (and dismissal) in *The Wealth of Networks* (pp. 409-410), but we don't get a sense of the dialectic relationship among the groups of people pushing controlling technologies (efforts to re-install fixation as the dominant model of distribution) and those working to break the bonds (and thus enable flux). Benkler seems to be confident that flux has won and fixity must fall. I am not so sure.

Five years ago, Benkler predicted at a copyright conference that the major commercial music industry would be no more within five years. It would be destroyed by the powerful pull of file-sharing technology, he declared. His technological optimism remains. So does the music industry.

## Response

*By Yochai Benkler*

First, I would like to thank all the participants in the seminar for their generosity in time, effort, and spirit. It is a rare treat to have such a collection of intelligent and knowledgeable individuals comment on one's work; more rare yet is to have such fair minded and thoughtful remarks. I hope to be able to reciprocate with an equally fairminded response to the main claims each of the participants have made. Of the readers I beg patience, then, as each comment is substantial and each deserves, in turn, a response. These are mostly designed so that each section is read in companion to the commentary to which it responds.

*Farrell*

Henry Farrell emphasizes the role of social norms, rather than legal institutions, in defining the future of the networked public sphere. But before turning to that main point, I want to emphasize that I agree with Farrell's characterization of the angle that the book provides for the question: "what's left?" I suggest a strong emphasis on the practical capabilities and freedoms of individuals and a heavier dose of skepticism about the need and advisability of affirmative government action, as opposed to government action in a facilitating mode, than have traditionally been the "left" positions. I think the networked information environment requires and enables liberals (in the American colloquial sense) to refocus on the abilities of individuals acting alone and in voluntary association with others. But it is also important to recognize that the same set of phenomena offers an opportunity for libertarians to adopt a more sophisticated model of what constrains liberty. In particular, the relatively new vintage and the clearly regulatory history of what passes for "property" and "markets," such as wireless spectrum "property rights," copyrights, or patents, makes clearer the fact that these are nothing more than detailed and highly imperfect and coercive regulation. The common roots of these currently quite distinct political orientations are exposed in this technical-economic-social context in a way that they were not in the industrial system. These observations requires libertarians, no less than progressives, to examine the potential of actors supposedly acting in markets and enforcing legal regulations they dub property to constrain liberty.

But to the substance of Farrell's argument, which focuses on the importance of prosocial norms in the construction of the peer production processes in general, and of the networked public sphere in particular. This is an important claim, and I want to take it up in characterizing an important research agenda. But first, a bit of self-defense. Farrell claims that I, as a lawyer, am more concerned with formal institutions than social scientists are, who are focused on norms, to which I am insufficiently attentive. I thought that I had done just that in several chapters. In chapter 4, when I discuss the economics of peer production, I specifically focus on the interaction of law, technology, and norms as mechanisms to stabilize commons-based production. When I speak in chapter 7 of what characterizes the networked public sphere, the norms of "here, see for

yourself,” linking to others, and transparency, which are of a similar kind, if not identical in content, to those Farrell characterizes--authenticity, linking, and attribution--are central to my claim about how the networked public sphere differs from the “trust me, I have authority” model of mass media. Finally, in chapter 11, the main chapter on institutional ecology and the legal battles, I talk about institutional ecology precisely in terms of the interaction of law, markets and business practices, social practices, and technology; when I characterize the various pressures for and against commons-based production, I give the cultural glorification of hackers the same status as laws prohibiting circumvention-device hacking. Perhaps because I am a lawyer I am at a loss to see how one could be more attentive to the role of norms, as compared to law, without falling into the fallacy that formal law doesn't matter at all. The critical point is that law is an appropriate focus when it is aimed at shaping the underlying economics and social dynamics. If a law will pass that will force computer manufacturers to build glorified TV sets optimized for a 5000 channel broadcast world, as the MPAA has sought since the introduction of the Fritz Chip, that will have a strong and overriding effect on the set of feasible actions open to people—irrespective of their social norms. These are the kinds of interventions that I seek to diagnose, and to whose negative effect I most directly turn my attention when I move from diagnosis in the first part, and normative analysis in the second, to prescriptive policy commentary in the third.

But it is much more interesting and fruitful to engage Farrell's critique as a research agenda, rather than for me to defend my own work from the criticism that it lacks sensitivity to this important issue. Clearly, from what I have said about the role of law up to this point, I agree with Farrell that norms are susceptible to being screwed up through altered incentives. I don't think what he calls “the norm of authenticity” is so embedded in the technology that it cannot be subverted. But like Jack Balkin, I think that (1) the technology invites this kind of quotation and (2) the fact that a norm has been established, and has worked for those who have followed it rather well, will tend to entrench it. This is as true of good norms as it is for evil. If what it means to be a good blogger today is, among other things, to give attribution, link, and be authentic, then that is what new bloggers will do, and that is what readers will expect. And while none of this is determined, it is self-reinforcing for precisely the reasons Farrell himself describes—that is, it does reinforce precisely what makes the blogosphere attractive to its practitioners. So while I think norms can be dislodged, they are not so fragile once entrenched that we cannot begin to rely on their presence for practical institutional design.

More specifically, in a very important move Farrell emphasizes the potential corrupting effect of money. As more blogs will find ways of making money, they may move away from the practices of linking, authenticity, and attribution. I think this is an important concern; there is no question that sites motivated by keeping eyeballs internally for profit will have different motivations, which may lead to different practices, than “amateur” blogs. It is important to remember, however, that the blogosphere emerged as a cultural practice when traditional commercial media online were already present online, as were commercial “destination sites” like Yahoo. People were at that time even more culturally attuned to seek those out, as opposed to seeking out blog-like amateur media.

And yet the blogosphere and its norms emerged. What people flocked to in the blogosphere was precisely the attribute that they made a very different information environment possible, an environment that emerged out of either nonmarket action or small market action.

A central claim of the book is not that markets or market actors will disappear, but that, because of the low cost of action in the networked environment, a substantial new form of production can take root and provide a steady and sustainable flow of information, knowledge, and culture, alongside and around the market-based flows. The question then becomes whether the adoption, by some bloggers and some corporations, of the same technical form and a related cultural form will crowd out the cultural form we now know as blogging. My intuition is very much toward a market-based metaphor for the answer: these are very different “products,” occupying quite different market niches. As long as people are free to choose, they are likely to choose some of each, because each sector, the mass-media-like market, the noncommercial, and the small-commercial (like TPM) will fill different needs, and allow different conversational affordances. That is why the net neutrality debate is so important. The plan to charge authors for the privilege of being read efficiently threatens to dampen the supply of noncommercial content willing and able to pay its marginal cost of carriage, but not willing or able to pay the duopoly rents sought by the incumbent telephone and cable carriers who sit on the last mile of the broadband pipes.

I am extremely sympathetic with the tenor of Farrell's conclusion. We cannot be sanguine about the sustainability of the practices we today celebrate. There are internal pressures—like what he describes as “invasion” from actors such as paid political astroturf bloggers or spammers—that put pressure on the genuinely free environment, and require technological or norms-based changes from a more open norm. All this is true. The ways in which these are developing, and the responses to them represent a rich and important area of research. Going into the basic science of cooperation to try to get some of the answers is an important project, and my next major focus. But I would defend the claim that understanding all these pressures in terms of political morality is important, and that understanding the particular effect of formal institutional interventions through law is urgent. We are living through a sustained assault on the integrity and freedom of the networked environment, carried out through rent-seeking legislation. We need to understand the stakes and we need to be able to recognize where the battle is being fought.

### *Hunter*

Dan Hunter has two main complaints, one small and stylistic, the other large and focused on what real work “social production” is doing in the book. I shall briefly respond to the former first, and then turn to the second.

Hunter criticizes me for writing a book that masquerades as a popular book but then cuts no corners in the text, and is a hard read for readers that it is their own

intelligence is to blame for the difficulty, rather than that I have tricked them into thinking that this is an easy book, when in fact it is not, and they lack the basic means necessary to comprehend it (to wit, “a solid grounding in economics, liberal theory, political science and jurisprudence (and possibly network theory and internet architecture).”) I must say that I tried to be as upfront as possible. If anyone reads Chapter 1 and still thinks this is a popular text that won't have hard sections, I'm not sure what else to do other than to engage in purposeful obscurantism. It is my goal to offer my work to the intelligent, diligent reader, without compromising the intellectual integrity and completeness that I think the project requires. I have tried to signal this as clearly as possible in the introduction. I have enormous faith in the good will and intellectual capability of many people who are not specialists and professional academics. I wrote this book as I did on that belief.

Hunter's main substantive criticism is that the book lacks focus: in particular, that my claim that social production is doing a lot of work here is overstated, and that this is more of a loosely connected set of observations about the character of the internet, connected to policy questions that need to be addressed. As this goes straight to the heart of my descriptive claim, it deserves something of an answer. First, it is important to recognize that social production covers more than peer production. I spend some time in Chapter 2 explaining that peer production, large scale collaboration without price signals or hierarchical control, is only a subset of social production. The nub of the claim is that with the decentralization of capital investment in the core inputs into information knowledge, and culture—that is, computation, communications, and storage—behaviors that were once central to human well being but peripheral to economic production have moved to the core of the most advanced economies. Where once we could not decide to make an automobile for mass consumption on our spare weekends, either alone or in collaboration with friends, we can collaborate now on the production of information, knowledge and culture. This does not replace markets, but it does offer a new source of a steady flow of information and knowledge into society, and these in turn are what shape autonomy, democracy, and justice.

Hunter illustrates his concern by criticising the discussion in Chapter 9 of whether some form of social production can alleviate problems of biological and agronomic research into agriculture for food security in the developing world. He questions whether all I am doing is to take a “fairly pedestrian” “example of underproduction of certain types of goods as a result of the limits of private markets, and hence the need for public subsidy for this sort of good,” and simply renaming it “peer production.” But what I actually do there is take this “pedestrian” problem (which causes death and underdevelopment in millions of children and adults around the globe), and suggest how the networked information economy could offer some avenues for alleviation that do not depend on the political will to provide public subsidy. In particular, (1) NGO provisioning can now be more effective because of the ways in which the cost barriers to information production have been lowered and more importantly (2) that there may be new ways of organizing production, on the model of the BIOS initiative that I describe in Chapter 9, that have the potential to provide significant improvement by harnessing peer production to deliver the desired goods without appropriation. The critical point to see is

that my focus is precisely a form of “social production”--that is, production done outside of either the market or the state--and that it characterizes a new solution space to “the limits of private markets.” The solutions I explore emerges out of social behavior that is not “public subsidy.” The problem is indeed a well known one. The traditional response was indeed public subsidy, and it has fallen short for decades. What I explore is the possibility of supplementing this inadequate response with a newly possible approach.

A core claim I make about social production throughout the book is that if, when capital costs were high, we had only two modes of production, and solving production problems required either market-based solutions or state-subsidy or regulation-based solutions, the rise of social production opens a new solution space to these problems. This new solution space has different advantages and constraints relative to both the market and the state. It is in the diversity of sources of solution that new forms of solutions to the problems of justice and freedom emerge.

### *Quiggin*

John Quiggin is interested in the dynamics of social production, in the motivations, and in the relations to the market and the state. Needless to say, these are central to my concerns. As Quiggin discusses, the broad understanding of social production has moved from disbelief to some fairly crude cuts into why and how it works—like Raymond's very influential, but also very early, work. Much of my own work, and that of so many others in this area, focused primarily on free and open source software, has largely been affected by the environment of disbelief. We operate on such a strong background assumption that markets and states are the only two modalities of production that simply establishing the plausibility of nonmarket, nonproprietary production in social relations was a major task of research. The next step, I think, is to begin from the baseline understanding that social production is real, and here to stay, and to begin to characterize its dynamics in greater detail and with greater rigor. One form of work would entail close analysis of particular online efforts, such as Joseph Reagle's work on Wikipedia. There also seems to be very promising opportunities for applying the work of experimental economists, like Ernst Fehr and his collaborators at Zurich, Samuel Bowles and Herbert Gintis at the Santa Fe Institute, and anthropologists like Robert Boyd and Peter Richerson, game theorists like Matthew Rabin, and one step removed, evolutionary biologists like Eliot Sober and David Sloan Wilson. A third and related approach would involve application and synthesis of the large literature on trust and social norms. Along these dimensions we can continue in earnest to characterize practical design considerations--for both technical and institutional design.

These lines of work document and probe carefully the diversity and range of human motivations, and how they manifest in different contexts. They suggest that human populations are characterized by a fairly stable mixture of people with different motivational profiles. Some in fact seem to behave like homo economicus. Most, however, do not. Few, are true unconditional altruists. Many are reciprocators—returning good for good and ill for ill. The dynamic combination of such types causes

cooperation and public goods provisioning to operate very differently than they are predicted to under the traditional, selfish rational actor model's predictions.

One of the things we begin to learn from this literature is that people do not even think of their normal market interactions as a series of opportunities to make a quick buck by screwing a stranger, but instead as relations of trust and forms of cooperation. “Leaving money on the table” seems to play an important role in producing efficiencies even within the market context. Understanding these dynamics is going to be central to developing better interfaces between markets and social production systems. Understanding the ways in which stable cooperative communities police themselves, and build on the different advantages of selfish actors and cooperative players will provide insight into the importance of identity, as opposed to anonymity, of transparency, and of the structures and relative costs of participation.

### *Hargittai*

Eszter Hargittai takes this opportunity to emphasize the importance of considering the distribution of access, skills, and use as a central concern that should significantly moderate our enthusiasm about the networked environment. Hargittai criticizes the way in which I excuse myself from spending large chunks of the book on distribution, or the “universal access” concern that was so prominent in the 1990s, and still is in many conversations. First, she claims that ongoing studies suggest that there is a levelling off of the relative growth in underserved populations, so that my optimism that we will gradually see attenuation of the difference we are seeing in access across race or educational attainment lines is unwarranted. This may be true, and it is important. Nonetheless, the particular point she makes: that growth is not as strong as it used to be, and therefore that my hope that growth, and in particular the differentially higher growth among relatively underserved populations (which, she points out, is not surprising given their lower baseline), is somewhat at odds with the particular study she links to. There, Nielsen ratings show a decline in active home use in some countries. But this evidence raises two objections: first, the strongest growth is in Brazil, stronger than in any of the more advanced countries. This supports, rather than undermines, my point that those sectors (surveyed in that study) who were underserved before are increasing service at a higher rate. Second, home use is important, as Hargittai explains, because of privacy and convenience. Yet, as her work shows so well, skills differences are an important component of lack of access. For those who lack skills, access to a reasonably well-staffed telecenter or library may well be more valuable and more enabling than access in the privacy, comfort, and relative ignorance of the home. Certainly in developing nations, public telecenters as well as small-business telecenter kiosks and cybercafes are central pillars of access diffusion policy. Finally, she raises the concern that the young users, the teenagers, those for whom skills gaps may be closing, seem to be showing an utter lack of interest in politics. Here, since the study is not yet published, I would wish to wait before I comment. The basic question that would determine the valence of these data is whether or not the survey allows one to generate a baseline rate of interest in politics of this population when offline. If 5% of them read the New York Times op-ed

page in print, but only 1% go to Kos or Instapundit, then there is cause for concern. If the finding is simply that teenagers are uninterested in politics at that age, then it may simply be that they are honing the skills they will later use for adult pursuits by playing with teenager pursuits today. If that proves to be the case, it does not seem to be cause for concern.

At a broader level, of course, I share Hargittai's concern for justice, as well as for freedom. I think her work in measuring and describing the state of distribution of access is important, and we should continue paying attention to it. My own focus in this book, in Chapter 9 on Justice and Development, has not been on digital divide issues, but rather on what I considered to be both more important and less explored: how the networked information economy, through connectivity available to about one billion people around the planet, can improve the lives and quality of development of both the connected and the five billion who are not connected. I tried to show how the emergence of social production and the networked information economy could alter, not radically, perhaps, but still to a potentially significant degree, the organization of information, knowledge, and information embedded tools and goods. My claim was that one major problem created by the integration of intellectual property into the international trade system has created significant barriers to technology transfer, to access by developing nations to opportunities for production in the global information economy, and to competitive prices for the outputs of innovation: like textbooks, drugs, or agricultural innovations. By circumventing the intellectual property industries, social production can offer alternative avenues to obtain these factors of human development. It is around these questions that the global Access to Knowledge movement is growing (for an overview of the movement at this moment, you can browse around the website of a conference we recently organized at Yale).

### *Balkin*

It is difficult for me to offer a response to Jack Balkin's post, because he so well captures precisely what I was trying to say in so many portions of the book. Rather than restate his analysis, I would only wish to endorse it as a well reasoned and generous exploration of central themes of the book.

I will, however, use this response to take a closer look at one question where I may differ from Balkin, and certainly from others who have written on Internet culture. This is the question of the relative autonomy from, or dependence of nonproprietary social production on, market production. Balkin puts it here as the idea that the networked information economy, and in particular social production in both its individual and collaborative forms, is “layered on top of industrial forms.” Balkin elaborated this relationship in *Digital Speech and Democratic Culture*, where he characterized the relationship as one of “glomming on” to the outputs of industrial information economy as material for bricolage, and “routing around” the blockages of that system. As Mimi Ito recently discussed, this also characterizes the difference between my own treatment of

distributed cultural production and Henry Jenkins', as well as her own, work on fan culture.

The basic question is one of core creativity. Do individuals and nonmarket groups need the platform or scaffolding of industrial cultural outputs to coalesce around? The basic kinds of claims in favor of the continued necessity of industrial culture as first order culture are two. First, there is the question of basic quality, ideas, themes, on which the variations can then be riffed. Second, there is the question of whether people need some authoritative cultural reference point in order to care about it as authoritative or salient, and around which they can organize their own identity and expression. In other words, must there be industrially manufactured stars as a precondition to fan culture, or can there be distributed culture that is not dependent on these icons of identity and reference?

Present cultural production always exists in conversation with yesterday's culture, and in looser conversation with last year's, decade's, and century's culture. In the twentieth century, widely shared culture was mass culture, and mass culture was dominated by proprietary and commercial culture. The definitions of what counts as quality; what are appropriate and exciting themes; what is the cultural locus of identity, were mostly generated by the industrial system, and therefore the orientation of fans, or nonmarket users and cultural bricoleurs, was toward the outputs of the industrial system. My prediction is that as the proportion of culture that is generated in the network and through social production increases in quantity, quality, and salience, social production will in turn be able to increasingly rely on its own resources for generating new subjects, themes, judgments of worth, and identity attachments and enactments. In software, we already see self-sustaining creativity; in the blogosphere too we begin to see blogs becoming their own sources of commentary and conversation, rather than "glomming on" at any point to what the industrial media produce. Indeed, we see early signs of the inverse. But I cannot claim to have hard evidence that this is in fact the trend. I can only point to indications, and to my own skepticism that the industrial producers have some special abilities or role in generating first order creative outputs, leaving to the distributed millions only second order, dependent creativity.

That said, I want to reemphasize that I do agree with Balkin that the best case scenario is not one of displacement of market activities, but of emergence of a salient new set of social practices alongside the market. It is the diversity of modalities of production, not the dominance of one or another, that offers greater freedom and greater efficiencies. We are at no risk of entirely losing the market-based and industrial models of information production. In the practical battles caused by the social contradiction Balkin characterizes, the risk is one sided. The political power of the incumbents is such, that if either system is to end up the sole system, it is only the industrial system that could do so by undermining the basic economic structure of social production. My optimism, if such it is, is based on the fact that a combination of technology, social practice, emerging political consciousness, and the business world's adjustments to the networked information environment present a formidable obstacle to those industrial players seeking to mold the network in their own image. But a successful long term stabilization of

networked information economy is not the end of industrial, proprietary cultural production, but its relative contraction.

### *Vaidhyathan*

Siva Vaidhyathan starts off with warm and glowing words, and for these I thank him. His complaint, though, is that I wrote a book about what interests me, not about what interests him. That is, that I wrote a book about how the dynamics of how technology, society, economy, and law intersect to fundamentally alter how information, knowledge, and culture are produced, rather than a book about the dynamics of how the technology component itself got to be as it is, and how it may or may not change given present pressures.

I plead only partly guilty, and that part excused by the fact that not every book can be about everything. Perhaps Vaidhyathan is correct that a book that offers as broad a canvass as this on the networked information environment needs a chapter on the technology itself: where it originates and what are the dynamics and pressures, historically and today, that led to its past and that affect its future. Instead, I did offer, as he mentions, a brief outline of an account early in the book. I also offer the case of peer-to-peer applications and their battle with law and economic actors as a case study in the relations of law, the economic pressure, and technology, (pp. 418-425), and the brief discussion of social software (372-375) and of the emergence of platforms for human connection (369-372) as instances where one can trace the kinds of relationship and flux associated with the development of technology.

Clearly, these references were sufficient to allow Vaidhyathan to map me, correctly, as sitting somewhere between Elizabeth Eisenstein's whiggish, but eminently well-defended, history of print, and Adrian Johns's highly detailed attack on that conception of print, as well as somewhere between her and Paul Starr in my sense of the relative autonomy of the social subsystem of technological development from politics. If that is soft determinism, so be it.

What matters is, of course, how soft is "soft." Jack Balkin better captured the mixture of optimism about what the technology can enable and concern for how it could be undermined that characterizes the book. The entire third part of *Wealth of Networks* is dedicated to mapping the set of contemporary battles over the shape of the future of this cluster of technologies and their uses, and it is argued from the perspective that the institutional and political choices will affect the shape of the technology and how it will structure social relations. When I say that the primary negative effect of software patents is that they would undermine the structure of free and open source software development given the social-economic dynamics of that system of software development (438-439), I am not entirely sure what it means to say that I am inattentive to who is developing the technology and how the battles I describe affect them. When I say that the p2p wars are ambiguous in their outcome, and that the technology has developed, as have its uses, in constant interplay with the lawsuits and the rules, and that the current architecture of p2p

networks and patterns of their use are a result of the constant effort, both legal and illegal, institutional and technical, to disrupt these systems, again, I am not sure how this is determinism, soft or otherwise.

My basic goal in this book is to characterize the technological-social-economic state in which we find ourselves, and the stakes of current battles. I take the technology as it is much of the time when I try to diagnose what is important, and what is not. I focus on those attributes of what we observe that are basic and structural: the decentralized capital structure, the connectivity, the flexibility. The interfaces with society, economy, and law are described when and as I deem them important for that diagnosis and prediction. It is that ability of my approach to diagnose and predict that is the subject of Vaidhyanathan's somewhat out-of-character ad hominem conclusion. Vaidhyanathan says: "Five years ago, Benkler predicted at a copyright conference that the major commercial music industry would be no more within five years. It would be destroyed by the powerful pull of file-sharing technology, he declared. His technological optimism remains. So does the music industry."

I couldn't remember having said such a thing, so I went back to the archives of that conference. It was called "A Free Information Ecology in the Digital Environment." You can browse the archive [here](#). The "prediction" he refers to was actually Eben Moglen's opening joke to his characteristically insightful presentation, in which Moglen analyzed why effective DRM would have to eliminate free software from the playing devices, why these efforts would fail, and how the free software movement needed to gear up to meet this long and difficult challenge. It was, as is so often the case with Moglen, not only prescient but also in-your-face. It hardly counts as evidence of unwarranted optimism and a lack of attentiveness to the contingencies and the necessity of battles to establish and maintaining a free information environment against the pressures of proprietary producers and the politicians they buy. Indeed, the DMCA is here; DRM are here, but the leak-proof pipe that the industry has sought for so long is as far away as it was six years ago. What I say in the book about the music industry, in any event, is that it, among all the other copyright industries, is in the position where it offers the least value, has the fewest avenues of alternative revenue streams, and therefore is most vulnerable among these industries and most likely in fact to die, without much loss to culture or the economy. I consider its future as more tentative by comparison to what appears to be the much more stable, but still stressful, condition of Hollywood. (pp. 425-28).

### *Conclusion*

I want to conclude by again emphasizing how grateful I am for all the work and attention that the contributors have put into their comments. As a totality, they focus on some of the core methodological questions that the book raises or has to work on: the internal dynamics of social production: motivations, the role of norms, and their natural history; the relations of social production to market production; the role of the state; and

the dynamics of technology politics, and society. I have done my best to illuminate these questions in the book, but much work remains.