

## The Future of the Internet – And How to Stop It

Jonathan Zittrain

v2.1 – 16 January 2006

### Chapter 5: Wikipedia and the Generative Internet

#### Verkeersbordvrij

Drachten has undertaken an unusual experiment in traffic management. The Dutch city of 45,000 people is “verkeersbordvrij,” free of nearly all road signs. It’s one of several European test beds of a new traffic planning approach called “unsafe is safe.”<sup>244</sup> There are no traffic signs, parking meters, or even parking spaces. The only rules are that drivers should yield to those on their right at an intersection, and those who get in others’ way when they leave their cars somewhere will be towed. The result is, counterintuitively, safer than the status quo. Without signs to mechanically obey (or, as studies have shown, disobey 70% of the time), people must drive more mindfully – operating cars with more care and attention to the actual circumstances that surround them. They also must communicate more with pedestrians, bicyclists, and other drivers using hand signals and eye contact. They see other drivers rather than other cars. In an article describing the expansion of the experiment to a number of other European cities – including London’s Kensington neighborhood – traffic expert Hans Monderman told Germany’s *Der Spiegel*:

The many rules strip us of the most important thing: the ability to be considerate. We’re losing our capacity for socially responsible behavior. The greater the number of prescriptions, the more people’s sense of personal responsibility dwindles.<sup>245</sup>

Law has long recognized the difference between rules and standards – between very precise boundaries like a speed limit and the much vaguer admonishment in negligence law to “act reasonably.” There are trade-offs between each, of course. Rules are less subject to ambiguity, and if crafted well, people know exactly what they can or can’t do, even if individual situations may render the rule silly or even dangerous. (For example, we need to be able to exceed the speed limit to get someone to the hospital or to get out of the way of an accident in progress.) Standards let people act in a way tailored to the situation, but leave much more to the good judgment of the often self-interested actors – or to the second-guessing of a jury or judge who later decrees what the actor did to be unreasonable.

A crude message of the verkeersbordvrij experiment is that standards can work better than rules in unexpected contexts. A subtler one has less to do

with the specificity of a rule than its source, and how closely people affected by regulation identify with its purpose and believe in its fairness, both in origin and in application. Perhaps most precisely, the question is how closely people see the benefit of a regulation not only to themselves but to their fellow citizens – about whom they care. Under the right circumstances people can rise to an occasion, behaving charitably to one another in the comparative absence of rules (or enforcement of rules) that would otherwise compel that charity.

Asking people to behave humanely can be a powerful force. In modern cyberspace – where an absence of rules and enforcement has led both to the generative blossoming of Chapter One and to a new round of challenges of the sort described in Chapter Two – it is especially worthwhile to explore how to prevent destructive behavior without a generativity-damaging rewiring of networks and endpoints to facilitate control that will otherwise be demanded by its denizens. Circumstance is important: try removing the traffic signs and rules in Manhattan, and the result may sooner be gridlock than utopia. When a disaster befalls a city, what determines how much people will help each other or turn on each other in the chaos that follows?

The answer of “unsafe is safe” is that if people can take each others’ welfare seriously, even a well-enforced precise rule can be less effective than un-compelled goodwill. The problem of Chapter Two is that of PCs run wild, spewing spam and viruses and infected by spyware because their users either don’t know what they should be asking to install on their computers or don’t care. The ubiquity of the PC among mainstream Internet users, and its flexibility to be reprogrammed at any instant, are both its signal benefits and major flaws. To figure out how to cure the flaw without eliminating the benefit, we can look to other layers of the generative Internet where the very same type of problem arises, and where solutions are in progress – some of which incorporate verkeersbordvrij precisely because traditional enforcement is not so easy.

### **Generative problems across the Internet’s layers**

The Internet was built according to a “layers” principle. Definitions vary according to purpose, but a common portrayal has a physical layer (wires and radio spectrum), a protocol layer (Internet Protocol and the routing of data packets), an application layer (Web browsers and instant messaging clients), and a content layer (videos, Web pages, music). Thanks to the modularity of the Internet’s design, network and software developers can become expert in one layer without having to know much about the others. Some legal academics have proposed that regulation might be most efficiently tailored to respect the borders of those layers.<sup>246</sup>

For our purposes, we might look at the layers to borrow from the solutions of one for insight into the problems of another. Recall from Chapter One that

the Internet developed in a backwater, and much of its design and operation presumed that people would behave reasonably. That assumption is now sorely tested, with the result being a computing environment so open to change that its users will either flee it or insist that it be locked down as against their own poor choices about what code to run or not run. Chapter Four argued that the move to tethered appliances and Web 2.0’s streaming software is part of a migration away from generative platforms. While there are benefits to running software remotely – just as there are benefits to software that is updated near-instantly by its vendor long after it’s in users’ hands – software as service is so tethered to its maker that third parties either can’t change or build upon it, or if invited to through open APIs or interfaces for application programming that allow for third party uses, the invitation is often contingent: the generative service exists until the original vendor chooses to pull the plug or announce new conditions for maintaining linkage with third-party add-ons.

Interestingly, the pattern of generative success and vulnerability present in the PC and Internet is also present in one of its more recent and high-profile content-level applications: Wikipedia, the free online encyclopedia that anyone can edit. The story of Wikipedia may provide insights on solutions for other layers.

### **Towards a free encyclopedia**

If the inexplicably-touted (and never very popular) first use of the PC was to “keep your recipes,” the first touted use of the Internet was access to knowledge and ideas. People spoke of digital “libraries of Alexandria,” with the world’s information a few clicks away for all. As a network that began with no particular content, this was at first an empty promise, especially since most “knowledge” was understood to reside in forms that were sold bit by bit, profitable because of a scarcity made possible by the restrictions of copyright. Existing producers of educational materials, including dictionaries and encyclopedias, were slow to put their wares on the Net. They worried about cannibalizing their existing paper sales – for *Encyclopedia Britannica*, \$650 million in 1990<sup>247</sup> – because there was no good way of charging for the small transactions that a lookup of a single word or encyclopedia entry would require, and there were few ways to avoid users’ copying, pasting, and sharing what they found. Eventually Microsoft released its *Encarta* encyclopedia on CD-ROM in 1993 for just under \$1,000, drawn from *Funk and Wagnall’s New World Encyclopedia*, pressuring *Britannica* to experiment both with a CD-ROM and a subscription-only Web site in 1994.<sup>248</sup>

As the Web exploded, the slow-to-change and walled content of formal encyclopedias was bypassed by a generative proliferation of topical Web pages and search engines that could pinpoint them. There was no gestalt, though: the top ten results for “Hitler” on Google could include a biography written by

amateur historian Philip Gavin as part of his History Place web site,<sup>249</sup> a variety of texts from Holocaust remembrance organizations, and a site about “killers,” or cats bearing uncanny resemblances to the tyrant.<sup>250</sup> This scenario exhibits generativity along the classic Libertarian model: allow individuals the freedom to express themselves, and they will as they choose; we are then free to read the results. This is also the spirit of most blogging. If any of the material is objectionable, people can either ignore it, request that the author take it down, or find a theory on which to sue over it, perhaps imploring gatekeepers like site hosting companies to remove material that individual authors insist on keeping.

Other models came nearly simultaneously from two sources, one the founder of the dot-org Free Software Foundation, and the other who had achieved dot-com success in part thanks to the market for pornography.<sup>251</sup>

Richard Stallman is the first. He believes in an environment where software is shared – with its benefits freely available to all, and its code able to be understood by those with technical aptitude and modified as they see fit for their own purposes, and then shared further. This was the natural environment for Stallman in the 1980’s as he worked at the Massachusetts Institute of Technology, and it tracks the environment in which the Internet and Web were invented. It is not the natural environment for the corporate development of software, whereby employees are hired to produce work according to particular standards, the work is owned by the company, and the resulting software is typically licensed to customers who are neither permitted to further copy it nor change its code. (Though recall that proprietary operating systems like Windows are meant to allow others to build “on top” of them via the application layer, so access to the Windows source code may be less crucial for generative purposes than the free vs. proprietary software debate may suggest.) Stallman holds the same views about other forms of intellectual expression – applying his philosophy across all of the Internet’s layers – and in 1999 put forth the idea of a free encyclopedia with content coming from whoever wanted to submit it, one article at a time. By 2001 some people were ready to try out the idea. Just as he’d sought to replace the proprietary Unix operating system with a similarly functioning but free alternative called GNU (“GNU’s Not Unix”), they called the project first “GNUpedia,” then GNE (“GNE’s Not an Encyclopedia”). There would be few restrictions on what those submissions would look like, lest bias be introduced:

Articles are submitted on the following provisions:

- The article contains no previously copyrighted material (and if an article is consequently found to have offending material, it will then be removed).
- The article contains no code that will damage the GNE systems or the systems from which users view GNE.

- The article is not an advert, and has some informative content (persoengl [*sic.*] information pages are not informative!).
- The article is comprehensible (can be read and understood).<sup>252</sup>

These provisions made GNE little more than a collective blog; people would submit articles, and that would be that. Any attempt to enforce quality standards – beyond a skim to see if the article was “informative” – was eschewed. The GNE FAQ explained:

### **Why don't you have editors?**

There should be no level of “acceptable thought”. This means you have to tolerate being confronted by ideas and opinions different to your own, and for this we offer no apologies. GNE is a resource for spe [*sic.*] speech, and we will strive to keep it that way. Unless some insane country with crazy libel laws tries to stop something, we will always try and fight for your spe [*sic.*] speech, even if we perhaps don't agree with your article. As such we will not allow any individuals to “edit” articles, thus opening GNE to the possibility of bias.<sup>253</sup>

As you might predict from its philosophy, at best GNE would be an accumulation of views rather than an encyclopedia – that might account for the “not” part of the “not an encyclopedia” label, even though GNU was meant to act a lot like Unix even though it would have different – free – code underneath. GNE was a generative experiment that failed, a place free of all digital traffic signs that never attracted any such traffic. It was eclipsed by another project that unequivocally aimed to be an encyclopedia, from an unusual source.

Jimbo Wales founded the Bomis search engine and Web site in 1996, when the dot-com boom was just beginning. Bomis focused on helping people find pornography and erotica, and earned money through advertising as well as subscription fees for premium content. In 2000, Wales took some of the money from Bomis to support a new idea: a quality encyclopedia free for everyone to access, copy, and alter for other purposes. It was called Nupedia, and was to be built as other encyclopedias were, through the commissioning of articles by experts. Wales hired philosopher Larry Sanger as editor-in-chief, and about 25 articles were completed over the course of three years.

As the dot-com bubble burst and Bomis's revenues dropped, Wales sought a way to produce the encyclopedia that neither involved paying people nor a lengthy review process before articles were released to the public. He and his team had been intrigued at the prospect of involving the public at large, at first to draft some articles which could then be subject to Nupedia's formal editing process, and then to offer “open review” comments to parallel a more elite peer review.<sup>254</sup> Recollections are conflicted at this point, but at some point Ward Cunningham's wiki software was introduced to allow for a simple plat-

form for contributing and edits others' contributions, and in January 2001, Wikipedia was announced to run alongside Nupedia and perhaps feed articles into it after review. Nupedia was quickly eclipsed by its easily modifiable counterpart – fragments of its Web site exist as of this writing, a fascinating time capsule<sup>255</sup> – and Wikipedia became an entity unto itself.<sup>256</sup>

Wikipedia began with three critical attributes. The first was being verkeersbordvrij. Not only were there few rules at first – the earliest ones merely emphasized the idea of maintaining a “neutral point of view” in Wikipedia’s contents, along with a commitment to eliminate materials that infringe copyright and an injunction to ignore any rules if they got in the way of building a great encyclopedia – but there were no gatekeepers. The way the wiki software worked, anyone, registered or unregistered, could author or edit a page at any time, and those edits appeared *instantaneously*. This of course means that disaster could strike at any moment – someone could mistakenly or maliciously edit a page to say something wrong, offensive, or nonsensical. However, the wiki software made the price of a mistake low, because it automatically kept track of every single edit made to a page in sequence, and one could look back at the page in time-lapse to see how it appeared before each successive edit. If someone should take a carefully crafted article about Hitler and replace it with “Kilroy was here,” anyone else could come along moments later and revert the page to the way it was before the vandalism with a few clicks, reinstating the previous version.

The second critical attribute of Wikipedia was the provision of a discussion page alongside every main page. This provided an avenue for people to explain and justify their changes, and anyone disagreeing and changing something back could explain as well. Controversial changes made without any corresponding explanation on the discussion page could find themselves reverted without having to rely on a judgment on the merits – instead, the absence of explanation for something non-self explanatory could be reason enough to be skeptical of it. Debate was sure to arise on a system that accumulated everyone’s ideas on a subject in one article (rather than, say, having multiple articles written on the same subject, each from a different point of view, as GNE would have done). The discussion page provided a channel for such debate, and helped new users of Wikipedia make a transition from simply reading its entries, to making changes, to understanding that there was a group of people interested in the page on which changes were made that could be engaged in conversation before, during, and after editing the page.

The third attribute of Wikipedia was a core of initial editors who shared a common ethos and some substantive expertise, many drawn from Nupedia. In these early days, Wikipedia was a backwater; few knew of it, and rarely would a Wikipedia entry be among the top hits of a Google search.

Like the development of the Internet’s architecture, then, Wikipedia was simultaneously ambitious in scope but modest in execution, devoted to making something work without worrying about every problem that could come up if its extraordinary flexibility were abused. It echoed “Postel’s Law,” a rule of thumb written by one of the Internet’s founders to describe a philosophy of Internet protocol development: “Be conservative in what you do; be liberal in what you accept from others.”<sup>257</sup>

Wikipedia’s initial developers shared the same goals and attitudes about the project, and they focused on getting articles written and developed instead of deciding who was or wasn’t qualified or authorized to build on the wiki. These norms of behavior were learned by new users from the old ones through informal apprenticeships as they edited articles together.

The absence of rules was not non-negotiable; this wasn’t GNE. Instead, the procrastination principle was at work. There would be maximum openness until there was a problem, and then the problem would be tackled. The rules would be developed on the wiki, publicly accessible and editable, though in a separate area from that of the substantive encyclopedia.<sup>258</sup>

From these beginnings there have been some tweaks to the wiki software on which Wikipedia runs and a number of new rules as the enterprise has grown and diversified and problems have arisen. Try to suddenly edit an existing rule or add a new one and it will be reverted to its original state, unless enough people are convinced that a change is called for. Most of the rules are substance-independent: they can be appealed to and argued about wholly apart from whatever argument might be going on about, say, how to characterize Hitler’s childhood in an article about him.

For example, as Wikipedia grew it began to attract editors who had never crossed paths before, and who disagreed on the articles to which they were contributing at the same time. One person said that Scientology was a “cult,” the other would change that to “religion,” and the first would change it back again. Should such an edit war be settled by whoever has the stamina to make the last edit? Wikipedia’s culture says no, and has developed the “three-revert rule.”<sup>259</sup> An editor may not undo someone else’s edits to an article more than three times in one day. Disagreements can then be put to mediation, where another Wikipedian, or other editors also working on the article, can offer their views as to which portrayal is more accurate – or whether the article, in the interest of maintaining a neutral point of view, should say that there is controversy about whether Scientology is indeed a religion or a cult.

For articles prone to vandalism – the entry for President George W. Bush, for example, or the front page to Wikipedia – administrators can make it so that unregistered or recently registered users may not make changes. Such

locks are seen as necessary and temporary evils, and any administrator can choose to lift a lock at his or her discretion. How does an editor become an administrator with such powers? By making lots of edits and then applying for an administratorship. “Bureaucrats” are Wikipedians able to promote editors to administrator status – or demote them. And to whom do the bureaucrats answer? Ultimately, an elected arbitration committee, the board of Wikipedia’s parent Wikimedia Foundation, or Jimbo Wales himself.

Administrators can also block particular users’ access to Wikipedia’s editing functions, registered or unregistered. Such blocks are rare and usually temporary. Persistent vandals usually get four warnings before any action is taken, and the warnings are couched in a way that presumes – often against the weight of the evidence – that the vandals are acting in good faith, experimenting with editing capabilities on live pages when they should be practicing on test articles expressing for that purpose. Other transgressions include deleting others’ comments on the discussion page – after all, it’s a wiki page itself, so can be edited in free form, making it possible to eliminate rather than answer someone else’s argument – and threatening legal action against another Wikipedian.

Along with sticks there are carrots, arrived at bottom-up rather than top-down. Each registered Wikipedia user automatically has a user page about him or herself, and a corresponding page for discussion with other Wikipedia users, a free form drop box for comments or questions from others. The habit has developed of awarding “barnstars” – literally an image of a star – to a user that someone has deemed helpful. To award a barnstar is simply to edit that user’s page to include a picture of a star and a note of thanks. Couldn’t a user simply award herself a pile of barnstars the way a megalomaniacal dictator can present himself with military ribbons? Yes, but that would defeat the point – and would require a bit of prohibited “sock puppetry” as the user would need to create alter-identities so that those looking at the page’s edit history would see that the stars came from someone other than the user herself.

Wikipedia has followed a path from crazy idea to stunning worldwide success. There are versions of Wikipedia in every major language, even one in simplified English for those who don’t speak English fluently, and Wikipedia articles are now often among the first search engine hits for the topics they cover. The English language version topped over a million articles in March of 2006, and in November it passed the 1.5 million mark.<sup>260</sup>

Quality varies greatly. Articles on familiar topics can be highly informative, while more obscure ones are uneven. Controversial topics like abortion and the Arab-Israeli conflict often boast thorough and highly developed articles, more objectively thought-out and less prone to including inflammatory comments than one might expect. As with the development of free software by many par-



ticipants, “[w]ith many eyes, all bugs become shallow.”<sup>261</sup> A controversial study by *Nature* in 2005 systematically compared a set of scientific entries from Wikipedia and *Britannica* (including some from the *Britannica* Web edition), and found a similar rate of error between them.<sup>262</sup> For timeliness, Wikipedia wins hands-down: breaking events of note have articles near-instantly appear about them. And for any given error that is pointed out, it can be corrected on Wikipedia in a heartbeat. Indeed, some of Wikipedia’s toughest critics can become Wikipedians by simply correcting errors as they find them.

### The price of success

When the Internet and PC moved from backwater to mainstream, their success set the stage for the problems discussed in Chapter Two. Email is no longer a curiosity but a necessity, and the prospect of cheaply reaching so many recipients has led to the scourge of spam, now said to account for over 90% of all email.<sup>263</sup> Further, the value of the idle processing power of millions of Internet-connected makes it worthwhile to hijack them, providing a rationale for the creation of viruses and worms.

Wikipedia’s generativity at the content level – soliciting uncoordinated contribution from tens of thousands of people – provides the basis for a similar vulnerability now that it is so successful. Vandals might be annoying, but they have been easily dealt with, so long as a critical mass of Wikipedians exists to keep an eye on articles and quickly revert any that are mangled. Some Wikipedians even appear to enjoy this duty, declaring membership in the informal Counter-Vandalism Unit in homage to the Counter-Terrorism Unit of television’s hit show *24*. Still others have written scripts that detect the most obvious cases of vandalism and automatically take care of them. And finally there is the option of locking those pages that consistently attract trouble.

But just as there is a clearer means of dealing with the threat of outright malicious viruses to PCs than there is to more grey-zone “badware,” vandals are the easy case for Wikipedia. The well-known case of John Seigenthaler, Sr., a retired newspaper publisher and aide to Robert F. Kennedy, only scratches the surface of the problem. There, a prankster had created a Wikipedia article about Seigenthaler suggesting that it had once been thought that he had been involved in JFK’s and RFK’s assassinations.<sup>1</sup> The article sat for four months until Seigenthaler noticed it. He then wrote an op-ed in *USA Today* decrying the libelous nature of his Wikipedia article and the idea that the light gate-keeping responsibilities described in Chapter Three could make it difficult to

---

<sup>1</sup> The actual Wikipedia entry stated: “For a brief time, [Seigenthaler] was thought to have been directly involved in the Kennedy assassinations of both John, and his brother, Bobby. Nothing was ever proven.” Seigenthaler, A false Wikipedia ‘biography,’ *USA Today*, Nov. 29, 2005, available at [http://www.usatoday.com/news/opinion/editorials/2005-11-29-wikipedia-edit\\_x.htm](http://www.usatoday.com/news/opinion/editorials/2005-11-29-wikipedia-edit_x.htm).

successfully sue Wikipedia itself for what an anonymous editor wrote. Wikipedians have since agreed among themselves that biographies of living persons are especially sensitive, and they are encouraged to highlight unsourced or potentially libelous statements for quick review – by other Wikipedians. Jimbo and a handful of other Wikipedia officials reserve the right not only to have an article edited – something anyone can do, of course – but to actually change its edit history so the fact that it *ever* said a particular thing about someone will no longer be known to the general public, as was eventually done with the libelous portion of the Siegenthaler article. Such practice is apparently done not under legal pressure, but as an ethical commitment.

Still, the reason Siegenthaler’s entry went uncorrected for so long is likely that few people noticed it one way or the other. Until his op-ed he was not a national public figure, and Jimbo himself attributed the oversight to the quick rate of new article creation – overwhelming the Wikipedians who have made a habit of keeping an eye on new entries immediately after they are created. Wikipedia since altered its wiki software so that unregistered or just registered Wikipedians cannot create new articles, but only can edit existing ones.

Because of its success, there is a new category of Wikipedian between committed community member and momentarily vandalizing teenager, one that creates tougher problems. It is someone who cares nothing for the social act of working with others to create an encyclopedia, but who only cares what a particular Wikipedia entry says about something. Now that many people merely consult Wikipedia as a resource, many of whom find its entries via search engines, Wikipedia has surpassed its origins as a quaint experiment.

Wikipedia has a rule that individuals are not to create or edit articles about themselves, nor prompt friends to do so. Instead they are to lobby on the article’s discussion page for other editors to make corrections or amplifications. What about companies, or the political aides? When a number of edits were made to politicians’ Wikipedia entries by Internet Protocol addresses traceable to Capitol Hill, Wikipedians publicized the incidents and tried to shame the politicians in question into denouncing the grooming of their entries.<sup>264</sup> In some cases it has worked. After Congressman Marty Meehan’s Wikipedia entry was edited by his chief of staff to omit mention of a broken campaign promise to serve a limited number of terms and then to replace the entire article with his official biography, Meehan repudiated the changes. He published a statement saying that it was a waste of time and energy for his staff to have made the edits, “though the actual time spent on this issue amounted to 11 minutes,” because “part of being an elected official is to be regularly commented on, praised, and criticized on the Web.”<sup>265</sup> Meehan’s response sidestepped the issue of whether and how politicians ought to respond to material about them that they believe to be false or misleading – surely if the *New York Times* published

a story that he thought was damaging, he'd want to write a letter to the editor to set the record straight.

If the Wikipedia entry on Wal-Mart is one of the first hits in a search for the store, it will be important to Wal-Mart to make sure the entry is fair – or even more than fair, omitting true and relevant facts that nonetheless reflect poorly on the company. What can a group of volunteers do if a company or politician is implacably committed to editing an entry? The answer so far has been to muddle along, assuming the best of all editors and hoping that there is epistemic strength in numbers. If disinterested but competent editors outnumber shills, the shills will find their edits reverted or honed, and if the shills persist they can run afoul of the three-revert rule.

In August 2006, a company called MyWikiBiz appeared to help people and companies promote themselves and shape their reputations on Wikipedia:

If your company or organization already has a well-designed, accurately-written article on Wikipedia, then congratulations – our services are not for you.

However, if your business is lacking a well-written article on Wikipedia, read on – we're here to help you!<sup>266</sup>

MyWikiBiz offers to create a basic Wikipedia stub of three to five sentences about a company, with some links, for \$49. A “standard article” fetches \$79, with a premium service for \$99 where MyWikiBiz will check up on the client's Wikipedia article after a year to see “if further changes are needed.”<sup>267</sup>

Wikipedia's reaction to MyWikiBiz was swift. Jimbo himself blocked the firm's Wikipedia account on the basis of “paid editing on behalf of customers.”<sup>268</sup> The indefinite block was only one of a handful recorded by Jimbo in Wikipedia's history. Wales talked to the firm on the phone the same day and reported that they'd come to an accommodation. Identifying the problem as a conflict of interest and appearance of impropriety arising from editors being paid to write by the subjects of the articles, Wales said that MyWikiBiz had agreed to post well-sourced “neutral point of view” articles about their clients on its own Web site, which regular Wikipedians could then choose to incorporate or not as they pleased into Wikipedia.<sup>269</sup> Other Wikipedians disagreed with the outcome, believing that good content was good content, regardless of source, and that it should be judged on its merits, without a *per se* rule prohibiting direct entry by a firm like MyWikiBiz.

The accommodation was short-lived. Articles submitted or sourced by MyWikiBiz were nominated for deletion – itself a process that entails a discussion among any interested Wikipedians and then a judgment by any administrator

about whether that discussion reached consensus on a deletion. MyWikiBiz participated wholeheartedly in those discussions and appealed to the “Jimbo Concordat” to persuade some Wikipedians to remove their *per se* objections to an article given its sources. Wales himself participated in one of the discussions that October, saying that his earlier agreement had been misrepresented, and after telling MyWikiBiz that it was on thin ice, once again banned it for spamming Wikipedia with corporate advertisements rather than “neutral point of view” articles.

As a result, MyWikiBiz has gone into hibernation according to its founder, who strongly believes that all sources, even commercial ones, should be able to play a role in contributing to Wikipedia.

A constitutional lawyer might consider these tales of Wikipedia and see a mess of process that leads to a mess of substance: anonymous and ever-shifting users; a God-king who can choose to act unilaterally; a set of rules now large enough to be confusing and ambiguous but small enough to fail to reach a new set of challenges. Much the same could be said about the Internet Protocol development process, which is equally porous – the Internet Engineering Task Force has no “members”; anyone can participate – but also has a proliferation of standards and norms designed to channel arguments to productive resolution. These standards and norms chafed as corporate interests, who generally respond to incentives rather than argument, became more keenly interested in protocol development as the Internet succeeded. The IETF avoided the brunt of these problems because its standards are not self-enforcing; firms that build network hardware, or for-profit Internet Service Providers, ultimately decide how to make their routers behave – so IETF endorsement of one standard or another, while helpful, is no longer crucial. With Wikipedia, decisions made by editors and administrators can affect real-world reputations since the articles are live and searchable; firms do not individually choose to “adopt” Wikipedia the way they adopt Internet standards.

### **Towards a theory of netizenship**

The “rule of law” is understood to be a system in which people are treated equally, without regard to their power or station; where the rules that apply to them arise legitimately from the consent of the governed; where those rules are clearly stated; and in which there is a source of dispassionate, independent application of those rules.

Despite the apparent mess of process and users, by these standards, Wikipedia has charted a remarkable course. Jimbo Wales may have extraordinary influence, but his power in Wikipedia depends in large measure on the consent of the governed – on the individual decisions of tens or hundreds of administrators, any of whom can gainsay each other or him, but who tend to work to-

gether because of a shared vision for Wikipedia. In one extraordinary chat room conversation of Wikipedians recorded online, Wales himself laments that Larry Sanger is billed in several Wikipedia articles about Wikipedia as a “co-founder” of the encyclopedia. But Wales does not edit the articles himself, nor directly instruct others to change them with specific text, since that would violate the rule against editing articles about oneself. Instead, he makes the case that an unremarked use of the label is inaccurate, and implores people to consider how to improve it.<sup>270</sup> From time to time Wikipedia’s articles about Wikipedia note that there is controversy over the “co-founder” label for Sanger. In another example of the limits of direct power, Wikimedia Foundation board member Angela Beesley has fought repeatedly to have the Wikipedia entry about her deleted. She’s been rebuffed, with administrators concluding that she is newsworthy enough to warrant one. (She tried again after resigning from the Foundation board, but to no avail.)

Wikipedia’s ideals are not libertarian. While outside regulation is not courted, Wikipedia’s policy on copyright infringement exhibits a desire to integrate with the law rather than reject it. Indeed, its copyright policy is much more strict than the law might require; the light gatekeeping described in Chapter Four would mean that Wikipedia could wait for formal notifications of specific infringement before taking action to remove copyrighted material. And despite the fact that Wales himself is apparently a fan of Ayn Rand – whose objectivism pretty closely aligns with libertarian ideals, a triumph of the individual over the group – Wikipedia is a consummately communitarian enterprise. The activity of building and editing the encyclopedia is done in groups, though the structure of the wiki allows for the groups to naturally break up into manageable units most of the time: a nano-community coalesces around each article, often about ten to twenty people at a time, augmented by non-subject-specific roving editors who enjoy generic tasks like line editing or categorizing articles. (Sometimes articles on roughly the same subject can develop independently, at which point there’s a negotiation between the two sets of editors on whether and how to merge them.) This structure is a natural form of what constitutionalists would call subsidiarity: centralized, “higher” forms of dispute resolution are reserved for special cases, while day-to-day work and decisions are undertaken in small, “local” groups. This subsidiarity also exists within groups drawn according to language. The different language versions of Wikipedia form their own policies, enforcement schemes, and norms. Sometimes these can track national cultural standards – the Polish Wikipedia is naturally edited primarily by people from Poland – but at other times it crosses such boundaries. The Chinese language Wikipedia serves mainland China (when it’s not being blocked by the government, which it frequently is), Hong Kong, and Taiwan.

When disputes come up, consensus is sought before formality, and the lines between subject and regulator are thin. While it’s true that not everyone can

be an administrator, the use of such special powers is reserved for abuse rather than daily enforcement. It is the editors – namely, anyone who chooses to participate – whose decisions and work together collectively add up to an encyclopedia – or not. And most – at least prior to an invasion of political aides, PR firms, and other true foreigners – subscribe to the notion that there is a divide between substance and process, and that there can be an appeal to content-independent rules on which meta-agreement can be reached, even as editors continue to dispute a fact or portrayal in an article.

This is the essence of law: something larger than an arbitrary exercise of force, and something with meaning apart from a pretext for that force, couched in neutral terms only for the purpose of social acceptability. It's been rediscovered amidst a society that largely disrespects its own "real" law, following it not out of civic agreement or pride but because of a cynical balance of the Skinneresque penalties for being caught against the benefits of breaking it. Indeed, the idea that there exists such a thing as a "neutral point of view" that can be determined among people who disagree is an amazingly quaint, perhaps even naïve notion, yet it seems to have amazingly worked in the case of Wikipedia. Recall the traffic engineer's observation about road signs and human behavior: "The greater the number of prescriptions, the more people's sense of personal responsibility dwindles." Wikipedia shows, if perhaps only for a fleeting moment in history, that the converse is also true: the fewer the number of prescriptions, the more people's sense of personal responsibility escalates.

Wikipedia shows us that the naïveté of the Internet's engineers in building generative network technology can be justified not only at the technical layer of the Internet, but also at the content layer. A system that can produce running code among talented (and not-so-talented) engineers – the free software movement, but also the openness of even proprietary-but-generative operating systems and PCs to outside contribution – can be replicated among writers and artists. Larry Lessig, Eric Von Hippel, and others are working on theories of how for-profit and non-profit enterprises can integrate their labors – what Lessig calls a "hybrid" economy – and Wikipedia will be an important piece of that puzzle.

For there is one last ingredient to Wikipedia that encourages good faith contribution and serves as a check on abuses of power as it accretes among administrators and bureaucrats there: Wikipedia's content is licensed so that anyone may copy and edit it, so long as attribution of its source is given and it is further shared without restriction. This applies regardless of whether the copy is made for-profit or not-for-profit. Thus dot-com Web sites like answers.com can simply mirror all of Wikipedia's content and show banner ads to make money, something Jimbo Wales has vowed never to do with Wikipedia. (A list main-

tained on Wikipedia shows dozens of such mirrors.<sup>271</sup>) This can lead to problems for people like John Seigenthaler, who not only have to strive to correct misrepresentations in the original article on Wikipedia but in any mirrors as well, but it has the benefit of allowing members of the Wikipedia community the option of exit. If they don't like how Wikipedia is running, they can leave – and take a copy of the encyclopedia with them. It also allows for generative experimentation and growth; for example, third parties can come up with ways of identifying accurate articles on Wikipedia and then compile those together as a more authoritative or vetted subset of the constant work-in-progress that the wiki represents.

Larry Sanger, the original editor of Nupedia and organizer (and, according to some, co-founder) of Wikipedia, has done just that. He is starting the “Citizendium,” an attempt to combine some of Nupedia's original use of experts with Wikipedia's appeal to the public at large. His plan is to fork Wikipedia, and then solicit volunteers who agree not to be anonymous, so that their edits may be credited more readily, and their behavior made more accountable. If Citizendium draws people and content, links to it from other Web sites will follow, and given enough links, its entries will appear as highly ranked Google search results. Wikipedia's dominance has a certain measure of inertia to it, but the generative possibilities of its content, guaranteed by its choice of a permissive license, allow a further check on its prominence.

Wikipedia shows us a model for interpersonal interaction that goes beyond the scripts of customer and business. The discussions that take place adjunct to editing can be brusque, but the behavior that earns the most barnstars is directness, intelligence, and good faith. An owner of a company can be completely bemused that in order to correct (and have stay corrected) what he sees as inaccuracies in an article about his firm, he'll have to discuss the issues with random members of the public. Steve Scherf, co-founder of dot-com Gracenote ended up engaged in an earnest, lengthy exchange with someone known as “Fatandhappy” about the way Gracenote's history was portrayed.<sup>272</sup> The exchange was heated and clearly frustrating for Scherf, but after the intervention of an editor not previously involved in the article, Scherf pronounced himself happy if not thrilled with the revised text. These conversations are possible, even common, and they are still the norm at Wikipedia.

The elements of Wikipedia that have led to its success can help us come to solutions for problems besetting generative successes at other layers of the Internet. They are verkeersbordvrij, a light regulatory touch coupled with an openness to flexible public involvement, including a way for members of the public to make changes, good or bad, with immediate effect; a focus on earnest discussion, including reference to neutral dispute resolution policies, as a means of being strengthened rather than riven by disagreements; and a core of people prepared to model an ethos that others can follow. It's likely that with any of

these pieces missing Wikipedia would not have worked. Dot-coms that have rushed in to adopt wikis as the latest cool technology have found mixed results. Microsoft’s Encarta Web site, in a naked concession to the popularity of Wikipedia, now has an empty box at the bottom of each article where users are asked to enter comments or corrections, which will be forwarded to the Encarta staff for review – but no further feedback to the user.

Makers of cars and soap have run contests<sup>273</sup> for the public to make advertisements based on stock footage found in their respective commercials, complete with online editing tools so that amateurs can easily put their commercials together. Dove ran the winner of its contest in the Super Bowl. There is a realization that “user-generated” content can be important – part of a new hybrid economy that Lessig, Benkler, Von Hippel, and others are studying – and these solicitations of the public to manipulate corporate and cultural symbols may prove to be further building blocks of “semiotic democracy,” where we can participate in the making and remaking of cultural meaning instead of having it foisted upon us. The rise of an ever-increasing number of informational resources, including amateur and citizen journalism, has caused debates such as who *Time* will name its Person of the Year (ironically, in the end, “You”) or what will land on the cover of *Newsweek* to lose much of their relevance. The anchor’s chair of the evening news for one of the three main American broadcast networks no longer carries the power of an oracle.

But Wikipedia stands for more. It stands for the idea that people of diverse backgrounds can work together on a common project, in this case a noble one – bringing knowledge to the world. Jimbo Wales has said that the open development model of Wikipedia is only a means to that end – recall that this is the person who started with the far more restrictive Nupedia development model. Wikipedia has since come to stand for the idea that involvement of people in the information they read – whether to fix a typographical error or to join a debate over its veracity or completeness – is an important end itself, one made possible by the recursive generativity of a network that welcomes new outposts without gatekeepers; of software that can be created and deployed at those outposts; and of an ethos that welcomes new ideas without gatekeepers, but that asks the people bearing those ideas to argue for and substantiate them to those who question.

There are plenty of online services that can affect our lives. For example, Google’s choices about how to rank its search results can determine which ideas have prominence and which do not. That’s one reason why Google’s agreement to censor its own search results for the google.cn version intended for users in China has attracted so much disapprobation. But even those most critical of Google’s actions appear to wish to pressure Google as any firm might be pressured: through moral suasion, shareholder resolutions, government regulation compelling non-censorship, or a boycott inflicting financial



pressure. No one thinks that Google ought to be “governed” by its users in some democratic or communitarian way, even as it draws upon the wisdom of the crowds in deciding upon its rankings, basing them in part on the ways in which millions of individual Web sites have decided about to whom to link. Amazon.com welcomes user reviews (and reviews of those reviews), but the public at large does not “govern” Amazon.

Many people expect more of Wikipedia. They see it as a shared resource and a public one, even though it is not an arm of any territorial sovereign. The same could be said of the Internet Engineering Task Force, but Wikipedia appears to have further found a way to involve non-technical people in its governance. Every time someone reads a Wikipedia article but knowingly chooses not to vandalize it, she has an opportunity to identify with its ethos.

If Wikipedia didn’t exist there’d still be plenty of reason to cheer the generative possibilities of the Internet, its capacity to bring people together in meaningful conversations, commerce, or action. There are leading examples of each – blogs, e-Bay, Meetup, or Pledgebank, that draw on the power of individuals contributing to community-driven goals. These examples will grow, transform, or fade over time, and their futures may depend not just on the public’s appetites and attention, but on the technical substrate that holds them all: the powerful but delicate generative Internet and PC. In the next chapter, we will see how the lessons of Wikipedia, both its successes and shortcomings, shed light on how to maintain our technologies’ generativity in the face of the excesses arising from their widespread adoption.