THE FUTURE OF THE CONSTITUTION

January 27, 2011

Reza Estakhrian

A Mutual Aid Treaty for the Internet

Jonathan Zittrain

Governance Studies at BROOKINGS

INTRODUCTION

Jonathan Zittrain is Professor of Law at Harvard Law School and Harvard Kennedy School of Government, and Professor of Computer Science at the Harvard School of Engineering and Applied Sciences. He co-founded the Berkman Center for Internet & Society.

By 2030 we will have all of humanity's books online. Google's ambitious book scanning project – or something like it – will by then have generated high quality, searchable scans of nearly every book available in the world. These scans will be available online to library partners and individual users with certain constraints—what those will be, we do not know yet. It will be a library in the cloud, one that is far larger than any real world library could hope to be. It will make no sense for a library to store thousands of physical books in its basement. Rather, under a Google Books plan, there will be one master copy of the book in Google's possession.[1] The library partners display it and access it according to particular privileges. A user can access it from anywhere. One master book shared among many drastically lowers the costs of updates – or censorship. For example, if one book in the system contains copyright-infringing material, the rights-holder can get a court order requiring the infringing pages of the book to be deleted from the Google server. Google has no choice but to comply, at least as long as it continues to have tangible interests within the country demanding a change. This vulnerability affects every text distributed through the Google platform. Anyone who does not own a physical copy of the book—and a means to search it to verify its integrity—will now lack access to that material. Add in orders arising from perceived defamation or any other cause of action, and holes begin to appear in the historical record in a way they did not before.

Some people – and I am in this camp – are alarmed by this prospect; others regard it as important but not urgent. Still others see this as a feature, not a bug. What's the constitutional problem, after all? Court orders in the U.S. are subject to judicial review (indeed, they issue from judges), so can't they be made to harmonize with the First Amendment? Not so easily. Current constitutional doctrine has little to say about redactions or impoundment of material after it's had its day in court. What has protected such material from thoroughgoing and permanent erasure is the inherent leakiness of a distributed system where books are found everywhere: in libraries, bookstores, and people's homes. By centralizing (and, to be sure, making more efficient) the storage of content, we are creating a world in which all copies of once-censored books like *Candide*, *The Call of the Wild*, and *Ulysses* could have been permanently destroyed at the time of the censoring and could not be studied or enjoyed after subsequent decision-makers lifted the ban.[2] Worse, content that may be every bit as important—but not as famous—can be quietly redacted or removed without anyone's even noticing. Orders need only be served on a centralized provider, rather than on one bookstore or library at a time.

The systems to make this happen are being designed and implemented right

[1] *See generally* Google Books Settlement Agreement, http://books.google.com/googlebooks/agreement/ (last visited Apr. 7, 2010).

[2] See John M. Ockerbloom, Books Banned Online, http://onlinebooks.library.upenn.edu/banned-books.html; Jonathan Zittrain, *The Future of the Internet – And How to Stop It* (Yale: 2008), p. 116.

now, and can be fully dominant over the decades this volume asks us to chart. One helpful thought experiment flows from an incident that could not have been invented better than it actually happened. Somebody offers, through Amazon, a Kindle version of *1984* by George Orwell.[3] People buy it. Later, Amazon has reason to think there is a copyright issue that was not cleared by the source who put it on Amazon. Amazon panics and sends a signal that actually deletes *1984* off of all the Kindles. It is as if the user never bought *1984*. It is current, not future, technology that makes it possible. The only reason this isn't a major issue is because other copies of *1984* are so readily available – precisely because digitally centralized copies have yet to fully take root. This is not literally cloud computing; for the period of time the user possessed *1984*, it technically resided physically on his or her Kindle. But because it is not the user's to copy or to process, and it is Amazon's power to reach in and revise or manipulate, it is as good as a Google Books configuration—or, in this case, as bad.

By 2030, a majority of global communications, commerce, and information storage will take place online. Much of this activity will be routed through a small set of corporate, governmental and institutional actors. For much but not all of our online history, a limited number of corporate actors have framed the way people interact online. In the 1990s, it was the online service providers such as Prodigy and AOL that regulated our nascent digital interactions. Today, most people have direct access to the Web, but now their online lives are described by consolidating corporate search engines, content providers, and social networking sites.

With greater online centralization comes greater vulnerability, whether the centralization is public or private. Corporations are discrete entities, subject to pressures from repressive governments and criminal or terrorist threats. If Google's services were to go offline tomorrow, the lives of millions of people would be disrupted.

This risk grows more acute as both the importance and centralization of online services increase. The Internet already occupies a vital space in public and private life. By 2030, that place will only be more vital. Threats to cybersecurity will thus present threats to human rights and civil liberties. Disruptions in access to cloud-hosted services will cut off the primary and perhaps the only socially safe mode of communication for journalists, political activists, and ordinary citizens in countries around the world. Corrupt governments need not bother producing their own propaganda if selective Internet filtering can provide just as sure a technique of controlling citizens' access to and perception of news and other information. This is the Fort Knox problem: a single bottleneck in the path to data, or a single logical trove where we put all our eggs in one basket.

This scenario has implications here for both free speech and cybersecurity. The Fort Knox mentality exposes vulnerable speech to unilateral and obliterating censorship: losing the inherent leakiness of the present model means we lose the

[3] *See* Brad Stone, *Amazon Erases Two Classics from Kindle. (One Is '1984.')*, N.Y. TIMES, July 18, 2009, at B1.

benefits of the redundancies it creates. These redundancies protect our civil liberty and security in ways as important as a constitutional right scheme. Indeed, the Constitution is interpreted with such reality in mind. The ease with which an order can be upheld to impound copyright infringing materials, or to destroy defamatory texts, can only be understood in the context of how difficult such actions are to undertake in the world of 2010. Those difficulties make such actions rare and expensive. Should the difficulties in censorship diminish or evaporate, there is no guarantee that compensating protections would be enacted by Congress or fashioned by judges.

Moreover, threats to free speech online come not only from governments wishing to censor through the mechanisms of law but from *anyone* wishing to censor through the mechanisms of cyberattack, such as denial of service. If a site has unpopular or sensitive content it can find itself brought down – forced to either abandon its message or seek shelter under the umbrella of a well-protected corporate information hosting apparatus. Such companies may charge accordingly for their services – or, fearing that they will be swamped by a retargeted attack, refuse to host at all. This is why the more traditional government censorship configurations are best understood with a cybersecurity counterpart.

That which appears safer in the short term for cybersecurity – putting all our bits in the hands of a few centralized corporations – makes traditional censorship easier.

The key to solving the Fort Knox problem is to make the current decentralized Web a more robust one. This can be done by reforging the technological relationships sites and services have with each other on the Web, drawing conceptually from mutual aid treaties among states in the real world., Mutual aid lets us envision a new socially- and technologically-based system of redundancy and security.

The Problem

The threats that are creating market forces for this kind of consolidation are real, and their destructive potential has already been amply demonstrated. For example, a venerable form of cyberattack involves hitting a Web site with so many requests for information that it has, in essence, a nervous breakdown. A coordinated attack in 2007 paralyzed a huge swath of Estonia's Internet infrastructure[4]—and there's no reason it couldn't happen again tomorrow. In 2008, a single Internet service provider in Pakistan sought to prevent its subscribers from getting to YouTube, on orders from the Pakistani government. It misconfigured its Internet routers to make that happen, and within minutes YouTube was unavailable not just to the ISP's subscribers, but to nearly everyone

[4] *See* Steven Lee Myers, Estonia Computers Blitzed, Possibly by the Russians, N.Y. TIMES, May 19, 2007, *available at* http://www.nytimes.com/2007/05/19/world/europe/19russia.html .

in the world.[5] The harm from such disruptions grows with our level of reliance on the Internet. And when natural disaster strikes—whether a hurricane in New Orleans or an earthquake in Haiti—the loss of vital communications lines can eliminate Internet access when it is needed most. The current vulnerabilities in the Internet's structure, combined with the increasing sophistication of would-be cyberterrorists and restrictive government censors, suggest a frightening image for the future of the World Wide Web: an unstable and constantly besieged space where only the biggest and baddest sites are able to stay online.[6] We have already seen a preview of this world in the hacktivist fallout to the Wikileaks diplomatic cable leak in November of 2010. Large sites, such as Paypal, Amazon, Mastercard and Visa sudden found their websites targeted as thousands of hacktivists engaged in organized, voluntary DDOS attacks.[7] Large corporate entities weren't the only targets. Smaller, less well-defended sites were also targeted, with the hacktivists responding quickly to any perceived slight against the Wikileaks organization. Easy-to-use DDOS tools, such as the "LOIC" program favored in the "Operation Payback" attacks, are freely available for download across the net, and LOIC was downloaded from SourceForge over 88,000 times in a week's time.[8] It seems unlikely that the clock can be turned back. Rather, it seems more and more likely that hacktivist actions – whether launched by individuals, groups, or governments – will become a reality of existing on the Internet, driving sites to obtain their own private security forces, something that is both inefficient and deeply unfair, as sites without resources would remain exposed—or compelled to take shelter in configurations that amount to corporate consolidation of Web hosting.

A centralized deliberative solution to the problem may not be possible; attempts to gather major public and private stakeholders in one room—literally or metaphorically—are necessarily limited. Furthermore, it is far from clear that a centralized response is the most desirable. Responses implemented by governments are necessarily subject to government control. The tools to facilitate such regulation can, especially when ported to regimes that don't embrace the rule of law, threaten free expression and even the safety of political dissidents. It was originally thought that the decentralized nature of the Internet would cause problems for restrictive governments—that they would face the attenuating problem of regulating their own people because of the difficulties of regulating the

---

[5] *See* Bonnie Malkin, Pakistan ban to blame for YouTube blackout, THE TELEGRAPH, February 25, 2008, available at http://www.telegraph.co.uk/news/uknews/3356520/Pakistan-ban-to-blame-for-YouTube-blackout.html.

[6] *See, e.g.*, Ellen Nakashima, *FBI director warns of "rapidly expanding" cyberterrorism threat*, WASHINGTON POST, March 4, 2010, available at http://www.washingtonpost.com/wp-dyn/content/article/2010/03/04/AR2010030405066.html; Posting of Kim Zetter to Wired Threat Level Blog, Report: Critical Infrastructures Under Constant Cyberattack Globally, http://www.wired.com/threatlevel/2010/01/csis-report-on-cybersecurity/ (January 28, 2010).

[7] *See* Christopher Walker, A brief history of Operation Payback, SALON December 9 2010 available at http://www.salon.com/news/feature/2010/12/09/0

[8] *See* James Finke, Wikileaks hacktivists look to improve attack software, REUTERS December 15 2010, available at http://www.reuters.com/article/idUSTRE6BE5LB20101215

internet.[9] However, any government can still lean on plenty of intermediaries to control its citizens' access to the Internet, and a centralized regulatory scheme intended for the salutary purpose of increasing security would make such censorship easier.

Traditional corporate interventions are similarly unsatisfying in the cybersecurity context. On first glance, it may seem tempting to centralize web hosting in a few major, stable players; doing so would improve the overall robustness of smaller individual sites sheltered under a larger umbrella. However, concentrating control over vast amounts of Web content in only a few hosting providers introduces a host of undesirable security and political control problems—starting with the Fort Knox problem.

A more promising approach to the cybersecurity problem would capture the essence of cooperation and concerted action that made the Internet possible to begin with, and has roots in the physical world as well: mutual aid. By reframing the issues most commonly mentioned as problems for the fabric of the Internet to encourage mutually beneficial, reciprocal actions among many Internet participants, it becomes clear that a spirit of mutual aid, backed up by corresponding technologies and practices, can make a difference, while avoiding some of the gravest hurdles and unintended consequences that arise when governments alone, or formally chartered multi-stakeholder groups, attempt to regulate or intervene.

The Internet's structure is often conceived as a series of layers. In its most basic form, the network includes a physical layer, wired or wireless, by which signals representing data are transmitted and received; a logical layer, which comprises the protocols—then implemented in Internet-compatible hardware and software—that allow data to be routed and understood properly between sender and recipient; and an application layer, where services and software visible to the general public are placed. Modularization means that expertise in one layer need not implicate much about another, and the firms and other parties involved in providing connectivity for one layer need not exercise control over another.[10]

The vision for mutual aid can be implemented at each layer of the Internet. These implementations are meant to show the elasticity of the mutual aid principle as a means of dealing with very real problems; any specific suggestion, such as the proposal offered below, is just one contender in a parched field of ideas. The particular details can always be refined and varied. A consideration of mutual aid also identifies those parties best positioned to take a lead in bringing such implementations to life, showing how the most helpful parties may vary from problem to problem and layer to layer. This is just another reason to emphasize an underappreciated, collaborative framework for problem-solving rather than the

[9] *See* James Boyle, Foucault in Cyberspace: Surveillance, Sovereignty, and Hardwired Censors, 66 U. CIN. L. REV. 177 (1997).

[10] *See generally* DAVID G. POST, IN SEARCH OF JEFFERSON'S MOOSE: NOTES ON THE STATE OF CYBERSPACE (2009).

process-oriented institutional framework that is effective in so many other areas of interest to the global agenda.

Mutual Aid for a More Robust Web

The World Wide Web is perhaps the most successful application to which the Internet has been put. Indeed, the Web is now so fundamental to being online that it is itself rightly viewed as a form of infrastructure. Many applications, whether interacting with your bank, buying and selling things, or browsing and conveying news, rely on functioning Web servers and browsers.

As Figure 1 below indicates, when you access a site (Step #1) and read a Web page there (Step #2), links generally look the same whether they point to another page on that site or to another site entirely. By clicking on the link, you ask your browser to visit a new destination and see what's there (Step #3). Too often, links don't work: you click and nothing happens (Step #4). This can be because a server has crashed or eliminated the page in question; because it is experiencing a cyberattack in the form of "denial of service," where a stream of requests overloads its ability to serve most visitors at all; or because there is some network interruption between you and your desired destination.

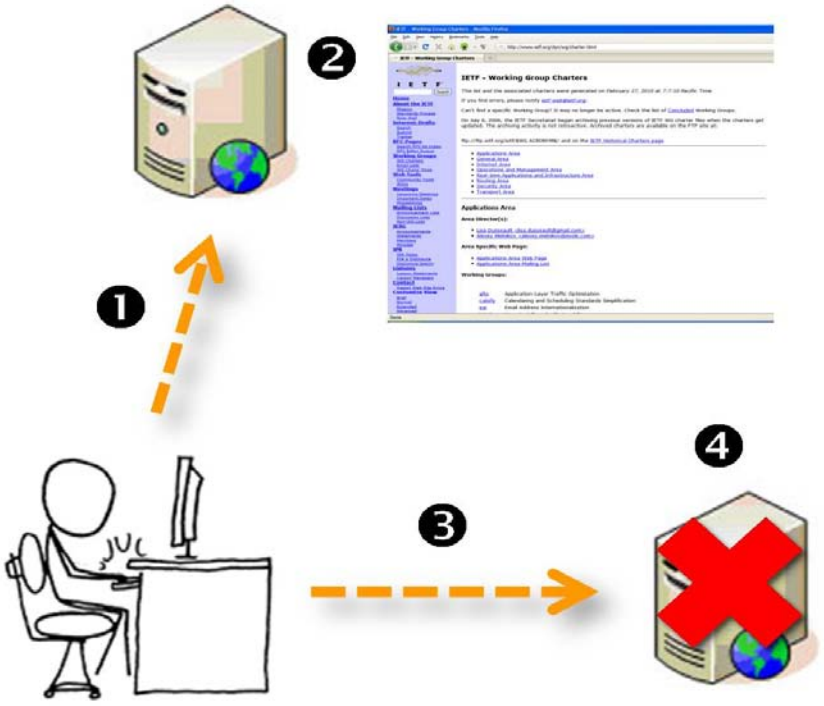

[11]

Figure 1: Anatomy of a Failed Link

[11] Internet user icons in Figures 1 and 2 courtesy of XKCD, licensed under a Creative Commons Attribution-NonCommercial 2.5 License, http://xkcd.com/license.html.

Cybersecurity is a multifaceted problem, but the risk of any site's becoming unavailable at any moment as a consequence of a sudden denial of service is a prominent threat. Currently, only those sites with significant financial resources can hope to weather such attacks, and then only through nontrivial tactics.

Current state-of-the-art thinking on cybersecurity relies on traditional defense notions honed during the Cold War, such as effective deterrence. To achieve effective deterrence one must know how to reach (and usually to identify) an attacker, in order to exact a price for the attack. This in turn has led to some calls for rearchitecting the Internet to solve the "attribution" problem: to enable effective identification of the source of any particular set of data sent over the Internet. Many implementations would require a fundamental rewriting and re-implementing of Internet Protocol—a practical challenge—and, if successful, could carry dire consequences for civil liberties. Not only could repressive governments demand access to the identities of political enemies within their countries, but the mere existence of data linking online activity with identity would deter the free and open discourse for which the Internet is currently uniquely suited.[12]

So how might the philosophy of mutual aid help? Imagine a very small tweak to the way Web servers work. At the choice of a site operator—and that of the external sites referenced on its Web pages—a site could implement a practice that might be called "Mirror as You Link," shown in Figure 2 below.

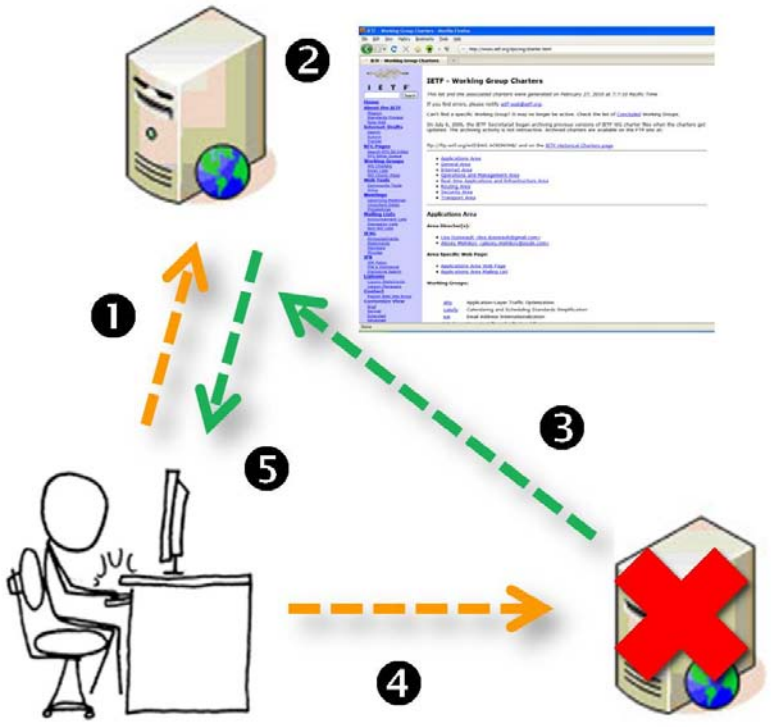

Figure 2: Mirror as You Link

[12] Anonymous speech has long been recognized as a fundamental right protected in America by the First Amendment. *E.g.* McIntyre v. Ohio Elections Comn'n, 514 U.S. 334 (1995).

Here, you visit a Web site (Step #1) and are shown a page there (Step #2). As the page is displayed, the site follows (or has previously followed) the links that it's showing to you and, if the destination server is amenable, makes a copy of what's there for safekeeping (Step #3). In the meantime, you might click on a link on the page that points that other site, try to get there, and find that you can't (Step #4). You then have the option of asking the *linking* site to show you its best rendition of what you're missing at the *linked* site (Step #5).[13]

Such an approach could have an impact on this particular dimension of the cybersecurity problem in a way that helps everyone, not just those with substantial money to spend on private defenses, and without the downsides of more traditional interventions on the table. The decentralized nature of the solution renders it more affordable at the same time as it reduces the problems associated with the centralization of Internet content and regulation. An inexpensive solution that can be adopted by individual users draws upon the principles of self-help that marked the Internet's beginnings, while perpetuating the values of free expression. The practical implementation details are not nearly as fundamental as reworking Internet Protocol or establishing a broad-based identity scheme—an elusive if ongoing goal. Rather, because nearly 80 percent of Web servers worldwide are accounted for by two vendors, Apache and Microsoft, the option to join such a mutual aid scheme could be added through updates to two pieces of software.[14]

In addition, as with most forms of mutual aid, many actors can be motivated to participate through self-interest, rather than charity. Consider longstanding practices on the high seas, where private parties will respond to an SOS without any obligation to do so—indeed, even those of business competitors to the firm owning the ship in distress. They do so not only because they may consider it the right thing to do, but because they adhere to a larger scheme of reciprocity: if they should be the ones in trouble next time, others will help. This is also the animating logic behind a standard military alliance: an attack upon one is an attack upon all. A state may see the downsides of joining—being called upon to defend another state in an unasked-for fight—balanced against the benefit of knowing that others will come to its aid should it be attacked. So long as no one knows ahead of time who will be the next target, it can make sense to band together. Moreover, such a scheme creates a natural overall deterrent: cyberattacks on participating sites will be naturally less effective, since inbound-linking sites will have copies of the data the attacker is attempting to take offline.

Participation in a mutual aid scheme of this sort can be entirely voluntary, since it is naturally incented. Those sites not desiring to have their data mirrored for any reason can decline to participate, and see that respected by inbound-linking sites. And those who fail to participate by actively mirroring other sites

[13] A similar system has been proposed by a group of academics in the Netherlands. http://www.globule.org/.

[14] *See* Netcraft, February 2010 Web Server Survey, available at http://news.netcraft.com/archives/2010/02/22/february_2010_web_server_survey.html.

may find that they are not themselves mirrored—eliminating free riding as a reliable option.

Mirror as You Link is only one concrete example of the larger principle of distributed reciprocity aided by appropriate technologies. Of course, there are many further details and issues to be worked out for such a tangible proposal. For example, one might want to safeguard interactive database-stored content and transactional services where successful connection to a Web site is crucial for purposes other than simply viewing information on a page ideally accessible to all. And one might also want to see a principle of mutual aid applied to link persistence, not just accessibility, leading to the ability to naturally see what used to be at a link—even if something different is there now. (Currently such a service is provided by the private, non-profit, but administratively-centralized Internet Archive's Wayback Machine, started by an Internet user with a vision for preserving the contents of the Web at large.[15]) The idea here is to sketch a proposal that demonstrates the power of looking at an Internet problem through the lens of mutual aid. Existing standards organizations such as the World Wide Web Consortium might be impelled to flesh out implementing protocols that in turn could be made available by the makers of Web servers and implemented by Web site owners one server at a time.

Possibilities for mutual aid are evident at other layers of the Internet. For example, at the physical layer, should traditional Internet service be disrupted in times of natural disaster or unrest, so-called ad hoc mesh networking can allow each Internet user's device to also serve as a router, allowing data to hop from one device to another until it finds its way out of the troubled zone—with no Internet Service Provider needed in the middle. The key value is of reciprocity: participating Internet users would see their devices expending power and bandwidth carrying others' data—in exchange for having their own data carried similarly.

At the application layer, another problem of cybersecurity might be addressed: that of users' personal computers and other machines becoming compromised by malware. Often a Web site visited by that user will be able to detect that his or her machine is under attack, but will fail to issue a warning lest the messenger be blamed. One could imagine a network of Web sites who agree to warn simultaneously: a compromised user will see persistent warnings from one site to the next, and thus alerted to a problem regardless of his or her configuration of extra security software on the infected machine, and in a way that does not implicate any particular Web site.

[15] *See* Internet Archive: Wayback Machine, available at http://www.archive.org/web/web.php.

Conclusion

By emphasizing and applying a particular principle—mutual aid—rather than a particular institutional structure in a domain containing many overlapping organizations and forums, it may be possible to generate solutions to cybersecurity problems that arise from the same collaborative principles that gave rise to the Internet itself in the first place. Cooperation can arise from a recognition of mutual interests and the implementation of technologies and practices designed to further those interests—technologies and practices that can come about in a variety of ways, from any number of existing stakeholders. Historically, the Internet domain has seen infrastructural advances not through carefully planned interventions forged self-consciously at any given moment among stakeholders participating in a worldwide summit, but rather through an open architecture that allows ideas to be floated for general adoption: applications that in turn can become infrastructure. On the communications level, it is this same pattern of evolution that renders the Internet such an effective forum for speech, discussion, and the percolation of new and potentially risky ideas. Society relies too heavily on the Internet to allow the cybersecurity problem to persist unaddressed. But the form of the solution must reflect a commitment to preserving online speech as an avenue for dissidents, activists and the politically unpopular—and for ensuring that speech deemed contraband cannot be put down the memory hole.

The critical lesson is that multi-stakeholder cooperation can take many forms, and the Internet can be mediated minute-to-minute through technology and praxis as much as through formal hierarchy. Mirror as You Link is only one concrete example of this promising approach. The original structure of the Internet, which allows any node to join on equal terms as sender or receiver without gatekeeping or negotiation, itself suggests a realistic frame for taking on problems that, left unchecked, will otherwise call for much more costly and complex forms of intervention.

The author wishes to acknowledge the contributions of Tim Berners-Lee in the development of this paper, in particular in originating the mirror-as-you-link idea. This paper grew out of a report the author prepared for the Global Agenda Council for the Future of the Internet at the World Economic Forum. Thanks also to Heather Casteel and Molly Sauter for impeccable research assistance.

Jonathan L. Zittrain is a professor of Internet law at Harvard Law School and the Harvard Kennedy School, a professor of computer science at the Harvard School of Engineering and Applied Sciences, and a faculty co-director of Harvard's Berkman Center for Internet & Society. Previously, Zittrain was Professor of Internet Governance and Regulation at the Oxford Internet Institute of the University of Oxford and visiting professor at the New York University School of Law and Stanford Law School. He is the author, most recently, of *The Future of the Internet and How to Stop It*; and co-editor of the books *Access Denied* (MIT Press, 2008) and *Access Controlled* (MIT Press, 2010).

Zittrain works in several intersections of the Internet with law and policy including intellectual property, censorship and filtering for content control and computer security. He founded H2O, a project at the Berkman Center for Internet and Society that develops classroom tools.

Governance Studies
The Brookings Institution
1775 Massachusetts Ave., NW
Washington, DC 20036
Tel: 202.797.6090
Fax: 202.797.6144
www.brookings.edu/governance.aspx

Editors
Jeffrey Rosen
Benjamin Wittes

Production & Layout
John S Seo

E-mail your comments to gscomments@brookings.edu