Anarchy, State, or Utopia?
Checks and Balances of Power in Internet Governance

M. Christopher Riley*

Abstract

In the beginning, the Internet was managed primarily through a social contract. Good behavior from all parties involved produced a ripe environment for invention and innovation and generated tremendous benefits for the entire world. But over time, the influx of money and power began to reward selfish behavior more and more, breaking open the Internet's utopia and leading to crime, censorship, and fights over control. As a result, many are questioning whether national or international governmental bodies should play a more active role in Internet governance. As it is frequently framed, this question of "more or less government" on the Internet is overly simplistic. Today, Internet governance is a complex system of checks and balances among users, businesses, and governments – and too great a disparity of power, for any of these parties, could create imbalance and undermine some policy goals in favor of others, to the detriment of the Internet as a whole.

* Adjunct Professor, School of International Service, American University; Internet Freedom Program Analyst, ATSG LLC, with the U.S. Department of State, Bureau of Democracy, Human Rights, and Labor. The opinions, arguments, assumptions, and assertions in this paper are the author's alone, not those of his employer(s).

I. Introduction

Internet governance is a widely discussed topic in the realm of technology policy. Some frame the issue through a binary question: whether or not the Internet needs more government oversight. In practice, this question appears in many contexts, including high profile domestic debates over issues such as net neutrality, paid peering and interconnection, and even higher profile global conversations around Internet freedom and censorship, cybercrime, and intellectual property.[1] Some suggest that local, national, or international governmental bodies ought to play a larger and more direct role in managing activity on the Internet, while others say such a change would effectively end the Internet as we know it. But in any policy context, the "more governance or not" question is overly simplistic. It posits a choice between government control or a lack thereof, between empowering governments through greater regulation or empowering people by removing as much government presence as possible. This illusory choice focuses on only one small piece of a complex, dynamic system; the reality is much more nuanced.

The purpose of this article is to provide context for the Internet governance question, not to answer it. It challenges the premise of the question "more governance or not" by exploring in greater detail the complex factors that contribute to Internet governance. This is not an article evaluating existing international governance institutions or proposals to change them, nor is it about domestic policy questions in the United States or other countries over cybersecurity, net neutrality, privacy, or other issues. Instead, this piece abstracts these questions out to identify the normative goals at stake in Internet policy, particularly those that relate to information censorship and control, and it describes at a high level the potential for conflicts among those goals. It also examines the parties that play a role in shaping the Internet – a question that, like the Internet, is not neatly defined by national borders – and thereby influence its ability to function as an effective engine for commerce and society around the world: governments, businesses, and users. These categories of entities strive to achieve their individual, distinct objectives, which can effectively be framed in terms of the normative goals at stake, and more specifically to a preferred balance among those goals.

In other words, the right question is not "should governments play a role" in Internet governance, because the Internet is "governed" by the actions and interactions of domestic and international government institutions, businesses from a range of sectors, and Internet users from around the globe. No global or domestic regulation can completely remove any of these parties from relevance, or eliminate entirely their impact on the Internet. Instead, laws – like technologies, market forces, and a host of different actions and contexts – work to adjust the balance of power dynamics among the parties. This is not to say laws cannot have a substantial impact, or that the question of "more or less law" is not salient in both domestic and global contexts. Rather, the impact of law in practice is heavily contextualized by other powers and limitations, so the question of "more or less law" does not always have a single answer, and cannot be fully understood in isolation.

[1] For example, both the Republican and Democratic parties included Internet policy issues including censorship and governance as major elements of their party platforms in the late days of the 2012 election cycle.

At the core of practical Internet "governance" is a set of complicated checks and balances among the active parties – governments regulating or promoting users and businesses, and users and businesses jousting with each other in markets while attempting to avoid (or employ to their advantage) laws and other government actions. Change to this system takes the form of strengthening or weakening various actors in various contexts. Thus, in a utilitarian sense, the study of Internet governance translates into an examination of the checks and balances that impact Internet governance in practice, with the goal of determining whether incremental change could improve one normative goal without doing more harm than good for other goals.

The objective for those who seek to preserve the Internet as an engine for social and economic welfare is not, and should never be, to identify or pick one source of power above others and always promote that one at the expense of others; such an approach would produce more harm than good, whether the preferred power is governments, businesses, or even Internet users. Users, businesses, and governments are far from monolithic, of course, and some (most, even) of each category do not engage in harmful behavior; but empowering a group often enables the bad actors as well as the good. There are circumstances where greater government oversight, if properly cabined, is very helpful on balance, just as there are circumstances where reduced government oversight would be a great improvement overall. But too much or too little oversight – like too many or no checks on business activity, or too many or no checks on Internet users – will create significant problems for one or more key normative goals of Internet policy. To preserve the Internet's unique value as an engine for a range of socioeconomic benefits, the goal must be defending a careful balance among governments, businesses, and users.

II. The Internet's Broken Social Contract

As the Internet began to grow and questions of governance online were first raised, many leading thinkers believed that the technology would take care of itself. Code itself had become the new law, or so the conversations went, and traditional social ordering mechanisms such as policy and law and governments were irrelevant in the land of the Internet.[2] Governments could do what they chose to do, but their actions wouldn't matter, because "[t]he Net interprets censorship as damage and routes around it."[3]

Their vision was compelling, and for a time, the Internet flourished without active government involvement, either positive or negative. Businesses with online operations rarely engaged in harmful technical behavior, such as deliberately blocking or slowing down other businesses' traffic; and the community self-managed those few who acted out of turn. And users were (largely) well behaved; the first computer virus was many years away, and "hackers" were still using cereal box whistles to get free long distance telephone calls. Even in the first days of active censorship, workarounds were trivial – website content was mirrored or hidden, and governments and service providers could never catch up to the changing hosts. In small

[2] The most iconic reference to this point is John Perry Barlow's "A Declaration of the Independence of Cyberspace," available at https://projects.eff.org/~barlow/Declaration-Final.html. A far more nuanced take was given in Lawrence Lessig's book "Code and Other Laws of Cyberspace."

[3] Phillip Elmer-Dewitt, "First Nation in Cyberspace," *Time International* no. 49 (Dec. 6, 1993) (quoting John Gilmore, founder of the Electronic Frontier Foundation).

communities, the most effective governance systems can emerge from the communities themselves, without a formalized order; the early Internet demonstrated this phenomenon well.

But the Internet outgrew the small community of its founders and early adopters, and problems continued to grow. Two changes took place, one anthropological and one technological. The Internet became the most important social and political communications medium in the "real world," as well as its most vibrant, global commercial marketplace. Seeking a selfish advantage – which means, sometimes, not playing by the same rules as everyone else – began to offer the possibility of substantial returns, whether for users or businesses or governments. At the same time, technologies of Internet censorship and control improved dramatically, hallmarked by the increased power and decreased cost of deep packet inspection technologies that enable governments and ISPs to monitor and analyze traffic in real time,[4] as well as inexpensive yet powerful viruses and infection vector attacks that enable "script kiddies" to wreak havoc on unprotected systems around the world.

Today, Internet policy problems exist at every level, arising from many different sources and impacting a broad range of interests. Individual Internet users engage in cybercrime, sometimes for fun but often for profit. Internet businesses act anti-competitively to the detriment of other businesses and users by blocking or manipulating Internet traffic for parochial purposes. Some governments censor speech of individual Internet users, and restrict innovation and growth by businesses in a variety of ways. Across these issues, the early assumptions of the continued self-governance of the Internet have broken down.[5]

In the face of new incentives, opportunities, and means to engage in selfish behavior on the Internet without significant risk of retribution, cracks in the "social contract" that once governed effectively have appeared at every level. With few or no fail-proof opportunities for self-defense, those whose interests are not adequately served now turn to policy and politics as one means to fight for an advantage – at times, against government actions contrary to their interests.[6] Whether for good or for ill, governments became relevant again.

This is true for international policy contexts as well as for domestic policy contexts in many countries around the world. In the United States, disputes have arisen over privacy, net neutrality, competition policy, and a host of other issues as technologies and market structures evolve, and as once-relevant legal structures strain to match – a pattern witnessed in parallel, though with variations, in Europe and many countries in Asia. Globally, Internet freedom, cybersecurity, and Internet governance have become major issues of international tension and disagreement, as the borderless Internet and its composite entities struggle with conflicting legal and normative environments in major population centers and markets. The international context

[4] *See, e.g.*, M. Chris Riley and Ben Scott, "Deep Packet Inspection: The End of the Internet as We Know It?", White Paper, *at* http://www.freepress.net/files/Deep_Packet_Inspection_The_End_of_the_Internet_As_We_Know_It.pdf.

[5] On government activity and blocking in particular, *see, e.g.*, TJ McIntyre, "Child Abuse Images and Cleanfeeds: Assessing Internet Blocking Systems," to appear in Ian Brown, ed., *Research Handbook on Governance of the Internet* (forthcoming), available at http://ucd-ie.academia.edu/TJMcIntyre/Papers/794093/Child_Abuse_Images_and_Cleanfeeds_Assessing_Internet_Blocking_Systems.

[6] *See, e.g.*, Ian Katz, "Web freedom faces greatest threat ever, warns Google's Sergey Brin," *The Guardian* (Apr. 15, 2012), at http://www.guardian.co.uk/technology/2012/apr/15/web-freedom-threat-google-brin (referencing increased government censorship and control as one of the most significant threats to the Internet).

is, arguably, made far more complex by the presence of multiple, powerful governments with very different normative goals for the Internet's proper use and management. But even though domestic and international contexts are quite different, the core patterns, players, and processes have many similarities: The social contract that once maintained order is dissolving, to be replaced by conflicts among governments, businesses, and users in pursuit of inconsistent normative views of the "right" way to manage the Internet.

In particular, the breakdown of the social contract has left in its wake an ever-increasing volume of calls for legal reform, both domestically and globally. Users, businesses, and governments share a perception that there are opportunities for improvement to align the Internet to better serve their parochial interests. At the same time, those interests vary widely, even among different users, businesses, and governments. To understand the resulting tensions that arise, as well as the other factors that impact the effectiveness or risks of any changes in government oversight, it helps to explore the normative goals of Internet policy in greater detail, and the range of varying preferences among those goals.

III. Goals of Internet Policy

The goals of Internet policy take many forms. For purposes of this paper, I will identify five high-order goals: 1) freedom of speech/association; 2) privacy; 3) security, including cybersecurity; 4) economic growth; and 5) social order. Depending on the interpretation, these goals may have significant tensions between them, arising from circumstances where advances for one goal come at the expense of another. Yet each is part of the "uber-goal" of preserving an open and societally beneficial Internet – so a deficiency in any one, in any context, will lead some parties to argue for change in the current structures of Internet control and governance.

The first, freedom of speech and association, lies at the heart of many active Internet policy debates, particularly internationally. Its diametric opposite is censorship, actions that are designed to restrict speech or association, including for political or religious purposes. At the technological level, censorship includes targeted censorship, through blocking of specific web servers or search terms, as well as broader censorship, such as a "kill switch" designed to cut off communications with the outside world. Correspondingly, Internet policies that oppose or circumvent censorship by promoting speech and access to information fall within this goal. Although all of the five goals are connected, freedom of speech relates perhaps most strongly to the second goal, privacy. Privacy stands as a separate goal from free expression in part because intrusions to privacy are distinct. The intentions behind violations of privacy can differ greatly from those for free speech – they may be political as with censorship, to identify dissidents, or they may be commercial, to sell targeted products or services. Privacy issues arise in different fora than free speech issues as well, including within nations where speech is generally considered to have strong legal and normative protections. The technologies behind privacy violations include deep packet inspection, as with censorship, but they also include tracking data, information retention and analysis, and a host of other practices. Practical freedoms of speech and association often require reasonable protections of privacy – particularly privacy from those who have a desire to censor. The goals are also connected pragmatically, in that many technologies to promote privacy are also effective at promoting freedom of speech.

Whereas freedom of speech and privacy apply primarily to individual Internet users, the third goal, security, impacts each of the three interest groups evenly. Individual users and businesses highly value the security of their computing systems and personal data. Governments have a powerful economic interest in preserving security of their citizens and businesses, but they also have their own valuable and sensitive data that can be at risk. The antithesis of cybersecurity is cybercrime, including in particular actions that violate laws and take or damage the information and/or property of others.

The fourth goal of Internet policy, economic growth, is impacted not only by cybercrime, but also by the architecture of the Internet itself. The massive economic benefits associated with the Internet derive in large part from its extraordinary prowess as an engine for invention and innovation. Many attribute this capacity to the Internet's extraordinary openness or "generativity,"[7] which in turn derives from its uniquely open, end-to-end controlled architecture.[8] Others contend that targeted optimizations or other traffic manipulations can improve services and enable new functionality, despite being contrary to end-to-end design principles, and can even improve the Internet's effectiveness as an economic engine.[9] For either camp, the desire to promote global economic growth constitutes a significant goal of Internet policy.

Related to both free expression and security, protection of social order is itself a significant goal of Internet policy for many individuals and governments.[10] All countries and cultures share *some* concerns in this space, particularly over extreme harm to values, such as the use of the Internet to transmit child pornography (which is a form of cybercrime, but not germane to cybersecurity in principle). But some countries go much further, though, and strive to develop governance systems to prevent or punish Internet communications that violate strict moral norms, including against rude language, adultery, or alternative religious beliefs.[11] Other countries engage in rampant control of political expression or attempts at assembly with the goal of stifling any threats to "public order," including any challenges to the authority or legitimacy of the current government.

These broad normative goals do not fully articulate the complexity of objectives at play in the Internet governance ecosystem. But this level of abstraction provides perspective into how the actions of users, businesses, and governments impact the practical governance of the Internet.

[7] *See* Jonathan Zittrain, *The Future of the Internet – And How to Stop It.*

[8] *See, e.g.*, Barbara van Schewick, *Internet Architecture and Innovation.*

[9] The scope of economic growth as a factor for Internet policy goes far beyond architecture, of course; this is merely one set of Internet economic policy issues.

[10] A strong case could be made that this goal is a subset of free expression. However, social values and free expression are distinctly positive goals, and countries can and do seek to advance both – even though tensions can arise between the two, just as tensions arise between privacy and cybersecurity, and other pairs of goals. Because they can be somewhat independent – depending on whether the limitations on free expression are driven by an intention to protect social values, or by political motivations or some other goal – they are best treated separately.

[11] *See, e.g.*, Choe Sang-Hun, "Korea Policing the Net. Twist? It's South Korea," *New York Times* (Aug. 12, 2012), available at http://www.nytimes.com/2012/08/13/world/asia/critics-see-south-korea-internet-curbs-as-censorship.html.

IV. Checks and Balances

A baseline assumption of this article is that the Internet should be preserved largely as it was created – as a general-purpose, powerful engine for expression, innovation, communication, and global economic and social welfare. In the Internet's original state, under its social contract, all of the five categories of goals were preserved. Social norms discouraged users from engaging in cybersecurity or values violations; economic growth and free expression flourished organically without barriers or gatekeepers; and, by and large, the technologies and business models that lead to privacy harms did not exist. But now that the social contract is cracking, what is taking its place? The answer is a dynamic, complex, evolving system of checks and balances among parties with different goals that reacts constantly to changing external circumstances such as cultural and technological shifts. Rather than trying to articulate all the good or bad aspects of the current state of the system, a question that would require volumes examining every possible policy environment, this section will attempt to explain the mechanisms of checks and balances in general, and the possibilities for change within the system.

The most high-profile (though perhaps not the most effective) mechanism is that wielded by governments: law. Governments can directly restrict or shape the activities of users and of businesses on the Internet through laws. These take a variety of forms, from criminal laws, to property laws that are enforced through civil litigation, to complex and specialized regulatory structures. Laws aren't perfect – they often include interpretive or enforcement gaps – but they are powerful tools nevertheless. Governments have other tools at their disposal as well, such as their purchasing power as major consumers of Internet services, but most of their activity in Internet governance involves laws, whether through adoption and enforcement or even merely creating the possibility of imminent or future legal action. Governments often act to influence and promote cybersecurity and social order (though, as noted above, with greatly varying interpretations and levels of emphasis), which are most germane to internal government interests. Some governments also play an active role in supporting the Internet policy goals of businesses and individual Internet users (including freedom of expression and privacy), through multilateral diplomacy, grant awards, and other means.[12]

Businesses, particularly global businesses, have several tools at their disposal to check the power of governments that are behaving in ways contrary to their interests. One method, not specific to Internet or Internet-enabled businesses, is through investment. Business investment generates significant benefits for countries, including jobs and livelihood for citizens as well as revenue through various taxes. If a country's government is overly restrictive of economic growth, businesses have the option to set up headquarters in another country, and grow and expand operations in new countries to diversify business and create new growth opportunities. Effectively, this creates a competition for investment that rewards governments for engaging in good behavior and penalizes those that engage in bad behavior. Also, geography limits legal jurisdictions, which for Internet businesses has substantial practical effects. When Microsoft or Research in Motion maintains a physical server presence in a country, they are subject to that country's laws, and may be forced to allow that country access to data contained on that server, subject to due process. In some companies, this may rise even to the extent of sharing decryption

[12] *See,e.g.*, Internet Freedom Fact Sheet, U.S. Department of State (Feb. 15, 2011), *at* http://www.state.gov/r/pa/prs/ps/2011/02/156623.htm.

keys for encrypted data, where feasible.[13] Similarly, a local government can (at least attempt to) force an Internet company to censor communications through its services in that country.[14] Moving servers out of countries represents a way for businesses to avoid such forms of censorship.[15]

Users have a substantial range of tools to respond to governments that act in ways contrary to their interests, and in particular to government actions that restrict freedom of expression or violate privacy rights, their primary Internet policy goals. Some of these are soft responses, such as advocacy by civil society and other efforts to persuade governments to improve their legal systems. In the proper contexts, these efforts can be quite effective, particularly internationally where governments take up the cause and engage in collective political pushes for change, or domestically where democratic voting systems can oust politicians who aren't sufficiently responsive to public pressure. In other circumstances, though, advocacy for political change is not only ineffective, but also creates risk of arrest or other harm. Internet users seeking to preserve their free speech and assembly rights and protect their privacy in the face of aggressive censorship frequently adopt more direct measures: building and using Internet tools that can defeat the technological mechanisms of censorship, and can preserve privacy and anonymity in the face of active surveillance. These range from simple single-hop proxy systems and VPNs, to advanced onion routing systems that can preserve anonymity even in the face of multiple compromised network nodes, to ad-hoc wireless networks that can avoid commercial and government-controlled systems altogether.[16]

Users and businesses both exert checks on each others' behavior as well, primarily through the market system. Businesses establish contractual terms to limit what can be done with their products and services, and surround themselves with property and intellectual property protections to defend against some forms of abuse of individual power.[17] Users, in turn, can punish businesses by refusing to purchase their goods – though this power is greatly diluted in a market environment that lacks robust competition. These powers are generalized and not specific to the Internet context, though they take on additional complexity with Internet-based services where businesses can employ a broad range of technological controls over user activity to

[13] This is far from a hypothetical issue. The government of India has demanded that RIM hand over master decryption keys for its Blackberry servers, which would allow the government to read all emails sent into, within, or out of the country. RIM maintains that such functionality does not exist. *See, e.g.*, Bill Ray, "India: We DO have the Blackberry encryption keys / RIM: Er, I think you'll find you don't," *The Register* (Aug. 2, 2012), *at* http://www.theregister.co.uk/2012/08/02/rim_keys_india/.

[14] Google and China had several rounds of changing policies and computing systems over this issue; in the final resolution, Google effectively stopped operating servers in mainland China, and instead redirected users looking for Google.cn to its Hong Kong portal, Google.com.hk, so that the company could offer uncensored search without facing Chinese legal liability. *See* Blog post of David Drummond, Google (Mar. 22, 2010), *at* http://googleblog.blogspot.com/2010/03/new-approach-to-china-update.html.

[15] Such a move increases latency, of course, and it cannot prevent a country from censoring communications passing through border routers – but, users themselves have ways to circumvent these restrictions.

[16] In recent years, Western governments, led by the United States, have begun actively supporting a broad range of internet user efforts to advance free expression and assembly around the world, as one component of international human rights support. *See, e.g.*, Nicole Gaouette and Brendan Greeley, "U.S. Funds Help Democracy Activists Evade Internet Crackdowns," *Bloomberg* (Apr. 20, 2011), *at* http://www.bloomberg.com/news/2011-04-20/u-s-funds-help-democracy-activists-evade-internet-crackdowns.html.

[17] Of course, this mechanism, like many others, depends on some amount of rule of law to enforce legal rights.

enforce their terms of service – and where users may have technological means to circumvent those controls.

This discussion has only scratched the surface of the checks and balances of Internet governance; however, it provides some context for the assertion that most Internet policy debates can be framed in terms of norms, and then analyzed from the point of view of powers and limitations on users, individuals, and governments to advance their valued norms. To attempt to preserve and promote the Internet's general-purpose socioeconomic value by respecting all of the goals, a careful balance of power must be preserved across users, businesses, and governments. Imbalance can result in some goals disproportionately advanced over others. For example, authoritarian governments given too much power might be very effective at promoting their parochial interpretation of security and social order through complete control over all Internet activity. But such governments have less reason to promote economic growth, and they may have few or no incentives to promote free expression or privacy on the Internet, and consequently will curtail these values whenever advantageous. (Other governments do, of course, promote these goals, quite strongly, reflecting and projecting the wills of their constituents, while still pursuing their own security and social order agendas.) In contrast, users have powerful and personal incentives to push for their free expression and privacy rights, and may be very effective at protecting these if given substantial control over the behavior of businesses and governments as they affect the Internet. But if unchecked, some users would also violate security (and privacy); others would hamper global economic growth to advance their own; and many would frustrate states' and other users' interest in protecting cultural and social values. And overly powerful and unchecked businesses would be able to focus exclusively on their own individual, parochial economic advancement, undermining each of the other Internet policy goals as much as the relevant market allows (and possibly damaging global economic growth in the process).

But if gaps and imbalances of power across governments, businesses, and users can be perceived and corrected early on, perhaps an unstable yet effective state of equilibrium can be maintained. How the gaps can be perceived and then corrected, using solutions that don't create more problems than they solve, is a question that cannot be resolved in the abstract, or with a single or simple answer. It must be tackled on an issue-by-issue basis, separating out individual domestic and international contexts, and looking at specific policy concerns that rise to significance, and possible solutions to those concerns. The challenge is engaging in specific public policy debates about the trees without losing sight of the complex, multistakeholder forest.

V. Conclusion

Stepping back, one of the most active debates today concerns the single question of whether international governance bodies ought to take a more active role in policing the Internet. Some proposals would substantially increase the role of governments in shaping Internet governance, and could thus generate a seismic shift in the balance of normative goals.[18] There is, of course, plenty of room for debate concerning Internet governance structures. But the success

[18] This risk is particularly high when discussing changes to international Internet governance, as some of the responses outlined above (including investment and democratic pressures) are diluted or even entirely ineffective.

or failure of these proposals will not end the questions, nor establish a permanent, static system of Internet governance; regardless of the evolution of international law, actions and reactions of individual governments, Internet users, and businesses will continue to produce an ever-changing fight among divergent interests with varying goals, and in these balances and fights practical Internet governance will continue to be determined.

Similarly, domestic policy questions abound over the proper role of governments in managing the behavior of Internet businesses and users, whether in the contexts of privacy policy, net neutrality, competition (at all layers of the network stack and in many distinct Internet-based services markets), cybersecurity, cybercrime, or others. And, similar to the international context, there is plenty of need to debate the best roles for Congress, the FTC, the FCC, and other bodies – and to debate the best principles and processes for each. But legal changes (whether in the form of more or less law) will not affix a static system in perpetuity, nor remove the impact of businesses or users in the Internet's day-to-day operation.

And all Internet policy debates will be more productive, and likely to lead to better outcomes, if the role of governmental bodies in the complex system of Internet governance is understood in context, not in isolation – and if the diverse normative goals of governments, businesses, and users are identified and articulated as part of the analytical process. For those whose goal is to preserve the Internet as an engine for technology invention and innovation, empowering any one party too strongly over any other creates risks that some of the normative goals will be promoted to the detriment of others, shifting the Internet's balance significantly, and for the worse.