2015

Framing the Question, "Who Governs the Internet?"

Robert J. Domanski
CUNY Graduate Center

Chapter 1

Framing the Question, "Who Governs the Internet?"

As the Internet continues to become further integrated into all aspects of the global culture and economy, society has an increasing stake in pursuing socially beneficial and collective goals. Most people would agree, for instance, that society has a definite interest in preventing the dissemination of illicit child pornography or in mitigating the effects of widespread computer virus outbreaks. Some type of governance is vitally necessary to serve the interests of the public community, and indeed, such governance of the Internet has already emerged—although how these systems have emerged remains something of a puzzle. How have government institutions, private commercial firms, and the scientific academic community been able to create and implement rules and procedures for both the functional operation of the Internet and the behavior that takes place on it? To what extent and in what ways have these governance policies and arrangements emerged as a result of institutional decision-making and public policy processes at the federal level in the United States?

This book's main objectives will be, first, to develop a new model that deconstructs the Internet into four conceptual layers with the aim of helping scholars and policymakers better understand various Internet policy issues, and, second, to use this model in formulating a new political architecture that accurately maps out and depicts authority on the Internet by identifying who has decision-making authority and, therefore, a clear ability to shape behavior. We will then assess this four-layer model and its resulting map of political architecture by performing a detailed case study of U.S. national cybersecurity policy, post-9/11.

This study will examine the Internet from a public policy perspective, with a particular focus on policymaking processes and institutional arrangements. Specific institutions of various types have played a crucial historical role in

shaping the direction of both how the Internet has evolved technologically as well as in setting the rules for how people use it. The Internet did not emerge spontaneously, nor did its present incarnation develop by accident. Rather, the Internet and all of its characteristics were consciously shaped as a direct result of explicit policy decisions.

The central question, then, is who governs the Internet? Which institutions, individuals, or other actors are shaping both the substance and direction of Internet governance policies? As the Internet continues to become more culturally and economically significant, it is important to investigate what type of governance is emerging and why it is emerging in that way.

WHAT DO WE MEAN BY "GOVERNANCE"?

Let's begin by dispelling a common myth. If you ask a random person off the street today, "Who governs the Internet?," the reply is typically, "Nobody does." Such is the prevailing wisdom—originating out of early Internet folklore that mistakenly equated decentralization with ungovernability. However, to the individual replying that nobody governs the Internet, a follow-up question along the lines of, "So can you do anything you want on the Internet without fear of consequence?," the reply is also just as frequently, "Of course not." And there's the disconnect. Many people still stubbornly hold onto the notion that the Internet cannot be governed, and yet nearly all of those same people freely admit that others have authority over their actions.

My contentions are that the Internet is, in fact, being governed; that it is being governed by specific and identifiable networks of policy actors; and that an argument can be made as to how it is being governed.

For one undeniable example that the Internet is being governed, take the existence of its DNS, or domain name system. The reason why when a user types the Uniform Resource Locator (URL) "www.google.com" into their web browser they can reliably expect to reach the website of the Google Search engine is that a public-private hybrid institution named ICANN (the Internet Corporation for Assigned Names and Numbers) has been developed over time to create a system for administering Internet domain names, ensuring their uniqueness, linking them to IP addresses, creating requirements for registration, and implementing formal dispute resolution mechanisms. ICANN is a nonprofit institution that was originally created by private-sector actors in response to a mandate issued by the U.S. Department of Commerce under the Clinton administration, which sought to cede control over the management of the Internet's system of centrally coordinated identifiers, for reasons which will be explored in Chapter 2. ICANN manages the DNS system, maintaining the Internet's operational functionality, using a multi-stakeholder

model that incorporates businesses, governments, civil society organizations, and academic and scientific organizations, and is international in scope.[1] It is responsible for overseeing the Internet's core root name servers and all of its 639 generic top-level domains (gTLDs), including the core six—.com, .org, .net, .edu, .gov, and .mil,—and 248 country code top-level domains (ccTLDs),[2] as well as for making decisions over the adoption of future TLDs, which have sometimes proven to be controversial.[3] The very fact that the DNS exists and keeps the Internet operational is direct evidence of governance policy, and certainly refutes notions of the Internet being "ungovernable."

So what do we mean, then, in asking, "Who governs"? In academia, the phrase "Internet Governance" is fairly well established. Milton Mueller, a scholar of political economy, was an early researcher in the field and adopted a narrow definition, referring to the phrase only in terms of the Internet's functional operation with research focusing on the ICANN and DNS systems.[4] To this day, the phrase "Internet Governance" largely is used only in reference to the administration of the DNS system. Subsequent scholars like Laura DeNardis have similarly adopted this narrow definition in examining the adoption of the Internet's technical protocols like IPv6 and the roles played by a small handful of semipublic international consortium groups like the Internet Engineering Task Force (IETF) in standards-setting processes.[5]

But why should use of the phrase "Internet Governance" be so limited? I argue that a far broader, more comprehensive definition is called for. It ought to include the study not only of who makes decisions over the Internet's technical functionality, but also of who holds authority over much of the activity that takes place on it. In other words, we also want to be able to study the follow-up question of, "Can you do anything you want on the Internet without fear of consequence?" and "Why or why not?"

So, when asking, "who governs," the definition which will be used here adopts a broad policymaking approach and views governance as having three criteria: (1) the ability to constrain behavior; (2) the ability to enable behavior; and (3) the ability to produce intentional effects. Actors are said to govern when they have clear decision-making authority to create and implement policies with intentional effects that meet all three of these criteria.

To be clear, the issue here is one of governance, not government.[6] From Robert Dahl to C. Wright Mills, scholars have long sought to determine who has power, why they have it, and how they use it. In pluralist theory, power has many dimensions and is held in varying degrees by numerous actors—from individual people to large corporations to formal governmental institutions. Indeed, Dahl's approach in famously asking "Who Governs?" was to question how various interest groups compete in the political sphere, and that governance is ultimately determined by the relative capacities of different

actors to influence governmental decision-making.[7] The questions at hand, in the context of the Internet, remain how all of those different actors are organized in creating and exercising their relative levels of authority. However, what sets the Internet apart from Dahl's analysis, as will be demonstrated time and again, is that on the Internet, it is not merely a matter of government having final decision-making authority, but also, to a considerable degree, numerous private actors as well. The Internet governance dynamic is characterized by various competing interest groups not only trying to influence government, but also competing to influence each other, and sometimes government trying to influence them. Identifying who holds authority versus who is trying to wield influence, perhaps more clear in Dahl's day, is an increasingly difficult task. Thus, not only do we need to ask who has power, but also who has more power than whom?

Understanding who has the power to govern is a political question. Governance is inextricably linked with concepts of power, and in that context, both must be defined for the purpose of this project. This is not to say that we plan on comprehensively defining these two ideas at the heart of Political Science—governance and power—once-and-for-all. Rather, it is necessary to clearly state which definitions will be used to carry out our specific research.

The literature on governance has markedly shifted in recent years from focusing on hierarchical governmental structures toward greater reliance on horizontal, hybridized, and associational forms of governance.[8] In the field of public administration, for instance, scholars such as Frederickson and Smith have observed this re-focus from the bureaucratic state and direct government to the "hollow state" and "third-party government."[9] Governance theories that incorporate ideas about the role of "conjunctions" or "associations" among organizational entities have become increasingly widespread.[10]

This "governance fever"[11] focusing on horizontal relationships between public- and private-sector actors has seen a deconstruction of the governance concept into several categorical types. *Network governance*, most frequently used for characterizing the Internet, is commonly associated with ideas of "self-governance" or "self-regulation." It refers to loosely structured coordination among numerous actors that function like an "organic or informal social system."[12] Network governance, as Taylor has argued, arises because of modern societies' complexities and their consequent requirement for distributed knowledge acquisition and decentralized problem-solving.[13] In contrast, *hierarchical governance* embraces the activities of government, law, and statutory regulation.[14] It describes processes that are characterized by vertical integration and managerial control within a set of lead institutions, and is the traditional method of analysis for studying top-down bureaucratic organizations. Meanwhile, *market governance* is equated with the forces of effective free-market competition with the invisible hand governing

behavior.[15] There have also recently been new additional theories developed as scholars have sought to meaningfully depict what's occurring on the Internet specifically. *Adhocratic governance*, for example, is based on the idea of policy being made "ad-hoc," meaning in an improvised, on-the-fly type of manner, and that decision-making is guided by simply dealing with problems as they arise.[16] According to scholars like Mintzberg, "adhocracy" is a system superior to bureaucracy and one that will even eventually replace it. It is "any form of organization that cuts across normal bureaucratic lines to capture opportunities, solve problems, and get results."[17]

These are some of the theories about governance, but when it comes to actually defining the broader concept of the term, specifically from a policy-making perspective, the approach undertaken by Lawrence Lessig and others is most helpful.[18] This is the Foucauldian conception of power that involves both *constraint* and *enablement.*[19] Actors are said to hold power if they have demonstrated the ability to (1) constrain certain forms of behavior as well as to (2) enable other forms of behavior. This is echoed by Mills who defined the power elite as being "in positions to make decisions having major consequences" and that "whether they do or do not make such decisions is less important than the fact that they do occupy such pivotal positions."[20] It must be stated, however, that also central to our understanding of governance is the importance of *intentionality*. Bertrand Russell is famous for arguing that power is "the production of intended effects,"[21] and considering the level of intentionality of potential governing actors is extremely important for our discussion insofar as intentionality signals causality. We want to be able to distinguish between those actors who are structurally positioned to make decisions and create policies with intentional effects versus those who can be weeded out from the governance discussion because their role in causality is hazy, at best.

If such a definition for power is utilized then identifying who holds power on the Internet can be answered more scientifically. The goal of this research is to identify those actors who simply have influence in the policy process versus those who have repeatedly provided evidence of their decision-making authority through policymaking. Who has *influence* versus who has *authority* is a critical distinction.

With regard to the Internet, it follows that *governance can be defined as the practical exercise of decision-making authority through a demonstrated ability to create policies that constrain or enable behavior with intentional effects*. Recurring throughout the existing Internet governance literature is the idea that governance is the persistent shaping of the environment through explicit decision-making.[22] We will build on this notion to show that the Internet's policies—inclusive of policies not only made by governments, but by various private actors as well—are authoritative insofar as they meet the

criteria of our definition above, and that empirical evidence comes in the form of existing statements of policy intent that correlate with evidence of policy actions. Actors who have the decision-making authority to create policies that effectively constrain or enable Internet behavior can reasonably be said to govern.

Determining who has this ability to govern through policymaking can further be analyzed by examining what Marcus Franda has called "single controlling points."[23] We will examine the numerous "single controlling points" on the Internet where behavior is constrained or enabled—examples include the Web hosts that operate servers, the Internet Service Providers (ISPs) who deliver Internet access, the websites that control user accounts through Terms of Service (TOS) agreements, and the local and national governments who can still assert their territorial jurisdiction. By analyzing exactly where Internet policies are being created that intentionally constrain or enable behavior, it is here where our inquiries for determining governance will focus.

This interpretation of governance, then, refers to coordinated efforts among various types of actors operating at multiple levels in their efforts to achieve desired ends. Because of the complexity involved, what we will refer to as the Internet's "political architecture" is a mapping of power and authority that includes the relationships among various institutions and other influential actors and policymakers who are best positioned to directly affect change in their environment. This is why our discussion encompasses the full governance spectrum, and not merely the public policies that are made and enforced by formal governmental institutions. Governments and the public sector are limited in their policymaking capabilities as a result of, first, the global dimension and "borderlessness" of the Internet, second, the decentralized architecture of the environment, and third, the limits of technological capabilities. These, along with a unique developmental history characterized at least as much by grassroots movements as by governmental agencies, are the reasons Internet policymaking is differentiated from, by comparison, other traditional policy venues occurring in real-space.

Defining governance in this way helps to place the title question at the heart of this study in context. For years, legislators of governments around the world have often grown frustrated when trying to transpose their authority in their attempts to regulate Internet content and behavior. Problems inevitably arise involving territorial jurisdiction and frequent anonymity achieved through technical measures, and, as a result, many such governmental policymaking processes and implementation strategies have been rendered largely ineffectual. Attempts by U.S. national, state, or local governments to generate policies using a strictly vertical governmental approach have largely been ineffective at achieving desired ends—thus relegating such policies to the status of being merely symbolic actions. Rather, policies of governance,

emphasizing coordination among various public, private, and hybrid institutions at every stage throughout the policy process, have become the primary mechanisms for constraining and enabling different aspects of Internet behavior. To be clear, governments are still extremely relevant and essential in the policy process. However, the role of formal governmental institutions has often been fundamentally transformed in the Internet sphere to that of leading coordination-based strategies, acting as a policy catalyst for private-sector actions, or formalizing and legitimating previously made policy decisions after other actors had already propelled the policymaking process forward.

WHAT DO WE MEAN BY "THE INTERNET"?

The Internet is a rather generic term that often means very different things to different people. So in asking the question, "Who governs the Internet?" we need to clarify exactly what it is we are referring to.

In terms of a functional definition, the Internet is a global decentralized network of computer networks, each of which is independently managed in whichever ways its administrator deems fit. Decisions, particularly over technical protocols, are often made by "rough consensus," and their implementation relies completely on voluntary measures being adopted in order to facilitate reliable interconnection and communication. Moreover, the term refers to both the hardware and software components that connect the various networks and computing devices to each other.

In conceptual terms for our discussion of governance, the various entities and ideas that together form the basis of the Internet must be deconstructed into their constituent parts in order to analyze what specifically is occurring with regard to governing the Internet as a whole.

The model we propose in addressing this problem for explaining governance of the Internet is based on the conceptual scheme first put forth by economist and legal scholar Yochai Benkler.[24] This framework conceptualizes communications systems into three layers: the physical architecture, the logical infrastructure (or the code), and the content layers. Benkler originally devised this scheme to understand structural media regulation, arguing that emerging modern network technologies make a decentralized and democratized information environment possible—"enabling small groups of constituents and individuals to become 'users' (or participants), rather than simply 'passive consumers'." Benkler's three layers were conceived as a means of presenting "a new set of regulatory choices" that governments have in decentralized networked environments, and though pertaining primarily to media regulation, I argue that they are valuable for conceptualizing entire modern information communications systems, including the full reach of the Internet itself.

Benkler's framework was later applied by Lawrence Lessig, who used the three-layer model to argue that the Internet "mixes freedom and control at different layers." In his attempt to assess notions of property rights and "the commons" in cyberspace, Lessig extended Benkler's model in two fundamental ways. He utilized the three layers as a way of conceptualizing the Internet specifically, and he used them as a lens for analyzing systems of control—what is free, what is shared, and what is owned in cyberspace.[25] This is particularly important for our purposes in determining governance.

My proposal is to build upon this framework, yet also modify Benkler and Lessig's code layer to create a new distinction *within* the code layer. This study will demonstrate that, when identifying the various actors and institutions involved in Internet governance, two fundamentally different types of actors emerge within the code layer, and therefore it is important to draw this distinction in order to formulate a better understanding of governance arrangements. This will be done by emphasizing the difference between code, understood as technical protocols, versus code as the software developer's tool for creating applications which the end-user encounters. The result is the emergence of what may ultimately be deemed a fourth layer, separating the code layer of Benkler into a protocols layer and an applications layer. This will highlight not only the differences between institutional actors who either create technical protocols or create private, proprietary web applications, but also the different *types* of actors involved in decision-making.

Thus, in contrast to Benkler's three-layer model of (1) the physical architecture, (2) the logical infrastructure (the code), and (3) the content, I propose a new model be introduced that aims to conceptualize the Internet into four layers:

1. The Infrastructure
2. The Technical Protocols
3. The Software Applications
4. The Content

The purpose of this four-layer model is to create a lens for policymakers who seek to produce intentional effects, and this is accomplished by breaking down the different political dynamics at each layer so that policymakers' goals can be better aligned with implementation strategies. These various political dynamics will then be analyzed by addressing three questions within each layer: (1) Why is it important? (2) Who governs it? (3) How are policies being made within it?

By building upon this framework, the task of determining who governs the Internet becomes far more manageable. Public policies and governing efforts at each individual layer, examined independently and separate from

one another, can be more clearly ascertained as coherent strategies and tangible entities. Actors at each layer are readily identifiable, and their roles in the policymaking process provide a greater capacity for reasonable analysis. In other words, my approach to answering, "Who governs the Internet?" will be broken down into "Who governs at each layer?" and "Who is governing across layers?"

LITERATURE REVIEW

The field of Internet governance is relatively new by academic standards having only just emerged in the past two decades, and has been developed by a strikingly multidisciplinary cross-section of scholars originating from the fields of law, economics, public administration, international relations, and more. Books on the subject loosely use the terms "govern," "rule," "regulate," and "control" almost interchangeably, which belies the point that governing is based on a more complex political architecture of authority. Understanding both the technical and political dimensions across disciplines is vital. Too often policymakers draft regulatory laws applying to Internet technologies with little understanding of the technologies themselves. Likewise, far too few programmers of such pervading technologies have any involvement in, or knowledge of, the legal systems or the political systems which they are so deeply affecting. The intention here, therefore, is to help bridge this gap by building upon the existing literature across disciplines to develop a new framework central to understanding the governance on the Internet.

There are two general approaches that scholars have used to study Internet governance and public policy: (1) How the Internet is reshaping government and politics, and (2) How government and politics are reshaping the Internet. Our focus shall be on the latter.

The academic literature exposes several distinct arguments in answering how the Internet is being governed. There is an evolution of ideas in answering the question of who governs the Internet—and the wide range of answers include code, national and local governments, international regimes, self-regulation, private engineering consortium groups, and more. Each of these not only serve as a potential counterargument to what will be presented in the chapters that follow, but they also help frame the scholarly evolution of the debate, placing the discussion to follow in better context.

In his path-breaking scholarship isolating architecture as a constraint on behavior online, Harvard Law Professor Lawrence Lessig famously argued that *code* governs cyberspace, meaning that software is programmed to set the rules for behavior, and therefore code and its designers are the central authority.[26] This "code governs" argument is extremely insightful in

emphasizing how, in digital environments, technical decision-making has inherently political consequences. Because code itself is an agent of authoritative power—constraining and enabling behavior by defining what actions are even possible in a given space—programmers have a disproportionate amount of authority at several of the single controlling points already mentioned. For example, whether it is controlling the operations of web servers, setting the TOS on social media websites, or establishing network bandwidth caps, programmers make binding decisions over the private virtual spaces that all users of their service must adhere to. They may not completely have free rein—again, the dynamic is too complex than to say any one group of actors controls everything and can do whatever they choose—but in their ability to write code to shape the environment, these programmers definitely prove themselves to be a large part of the governing equation.

However, the great limitation of this argument is that Lessig—to his credit—only claims that code governs cyberspace; not the Internet as a whole. This is a crucial distinction often misunderstood. Though commonly used as a synonym for the Internet, the term "cyberspace" actually refers to only one aspect within the Internet—the virtual environment where people interact with one another and where content, such as websites, images, ideas, and experiences, proliferate. As scholar David Bell has explained it, cyberspace is a cultural artifact—a "product of and producer of culture simultaneously." It is the part of the Internet that "is lived."[27] By contrast, the Internet itself is a communications network defined by its physical infrastructure comprised of wires and cables connecting devices. Its hardware can be found at specific geographical locations; it can be touched. To briefly put this in context, it can be said that someone may post a digital video to a website in cyberspace so long as their computer remains connected to the Internet. Lessig is correct in asserting that code governs cyberspace because, in this context, that is where code is deployed. However, code and its programmers play a far smaller role in the governing dynamic when examining different aspects of the Internet—namely, for instance, the regulation of the physical infrastructure. This is the great limitation of the "code governs" argument—it is immensely valuable for understanding how policies get made regulating cyberspace, but not comprehensive enough to apply it to the Internet as a whole.

A contrasting argument was put forth by legal scholars Jack Goldsmith and Tim Wu who countered with the proposition that *local and national governments* increasingly govern cyberspace, as such governments have begun taking more proactive roles in formulating vertically designed public policies affecting cyberspatial content.[28] Again the focus is on cyberspace, however their narrative suggests that national and local governments derive their power from an already-existing and clear ability to regulate the physical aspects of the Internet—notably, through a re-assertion of their territorial

jurisdiction. By leveraging their authority over the physical world—and, hence, the physical infrastructure of the Internet within their sovereign borders—but applying it to regulating content in cyberspace, Goldsmith and Wu are significantly taking a "cross-layer" approach, albeit in a limited fashion, by seeking to explain how authority over one aspect of the Internet can translate into powerful consequences in another. This is an important point that will be revisited shortly.

Meanwhile, international relations scholar Marcus Franda argues that the Internet is governed by an *international regime* consisting of both the public and private sectors, and formalized through international agreements between governments.[29] This is certainly a more comprehensive view of the Internet in its totality, and it utilizes a similar definition of governance based on coordination among multiple actors at multiple levels. However, Franda explores Internet governance from a strict international relations perspective, and as a result his conclusions focus almost exclusively on formal institutions and organizations at that level. Ultimately, his approach is a comprehensive model that can be applicable to the Internet as a whole, but by under-emphasizing the role of individuals and grassroots efforts that have historically played a vital role in driving the Internet's evolution forward, his argument doesn't portray the full picture of power arrangements and policymaking efforts that are occurring in venues other than at the international level.

Then there is the aforementioned Milton Mueller and the cadre of scholars who define "Internet Governance" in the more narrow terms of its technical functionality, arguing that ICANN, the IETF, and a small handful of semi-public *international consortium groups* comprised mostly of academics and engineers govern the Internet, focusing on the administration of the DNS system and standards-setting processes. This narrow definition certainly lends itself to solving the problem of who creates policies regarding the functional, day-to-day operation of the Internet—and international consortium groups like ICANN clearly demonstrate a decision-making authority in that realm. Unfortunately for our purposes, defining Internet governance in this way is generally unhelpful for understanding any Internet issue area other than those focused on technical functionality.

Countering these different notions of code or some type of public or private institutions governing the Internet are proponents who argue that the Web is increasingly *self-regulated* by the masses of users, or netizens, who actively engage in cyberspatial activities and social networks. There are several problems with this argument. First, again, the argument is only intended to apply to cyberspace; not the Internet as a whole. Second, and more importantly, even just as the argument applies to cyberspace, there is a seemingly endless list of examples that contradict the notion that self-regulation is what is currently taking place. The anarchic vision of cyberspatial behavior having a complete

lack of oversight is more a part of Internet mythology than it is reality. When people visit websites, they are subject to several single controlling points such as the rules of the website, the web server, the ISP, the telecommunications carrier, and the government or governments who can claim territorial jurisdiction. As will be demonstrated repeatedly throughout this project, self-regulation is a normative, not empirical, depiction of Internet governance today.

Furthermore, much of the debate identifying who governs the Internet has centered on such normative issues of alternative cyber ideologies regarding systems of control. The libertarian model for Internet governance was famously crystallized in John Perry Barlow's classic *Declaration of the Independence of Cyberspace* in 1996, calling for governments of the world to completely stay out of cyber affairs, and that "self-governance" by users will inevitably arise.[30] However, scholars like Barbrook and Cameron have offered direct challenges to the cyber libertarian model, dissecting the principle components of the "Californian school" by seeking to expose it as little more than "an incursion of capitalist values."[31]

There is a longer literary history concerned with the political nature of technologies. As it relates specifically to an Internet context, this is embodied by the debate over how technological systems institute control and order in people's online activities.[32] The architecture of the Internet enables and constrains certain forms of political behavior, and therefore that technical architecture and the policies which sustain it must be viewed as inherently political as well.[33] This is a major point that ought not to be undervalued. In the context of Internet policymaking, *technical decisions often have very political consequences, and thus are often political actions in and of themselves*. As will be explored throughout this study, decisions over which technical protocols to adopt or what type of software code to create have a direct effect on setting the rules for what types of behaviors are even feasible in different cyber spaces, and such decision-making, therefore, inevitably embodies certain political values at the expense of others.

It is in this vein of the technical becoming political that Lessig's "code is law" argument gains so much credence. He purports that just as laws regulate behavior in real-space, code regulates behavior in cyberspace, as "the software and hardware that make cyberspace what it is *regulate* cyberspace as it is." Technology is powerful but not uncontrollable, Lessig notes; it can be designed by human intervention to embody certain values. In the final analysis, cyberspace is made of code, created by people. How people write that code—the type of architecture they set up to protect certain values—will determine if cyberspace will become "free" in the libertarian sense, or "regulable." Indeed, he claims, the invisible hand of cyberspace, guided by commerce, has already constructed an architecture based on control and highly efficient regulation.[34]

As to some examples of when code is law, Lessig cites (1) how in some places you must enter a password before you gain access, while in others you can gain access whether identified or not; (2) how in some places the transactions you engage in produce traces that link those transactions back to you, while in others this link is achieved only if you want it to be; or (3) how in some places you can encrypt your communications, while in others encryption is not an option.

However, Tim Wu formulated a direct counterargument to Lessig's "code is law" argument, publishing an article in the Virginia Law Review actually titled, "When Code Isn't Law."[35] He asks, if the goal is to understand the net effect of code's regulatory forces, how can we not examine the reaction to those forces? In other words, code only has the effect of law if it is largely being complied with, and in cyberspace compliance is certainly not always a given. Rather, he argues, code is more a mechanism for avoidance of the law than it is for change, or even a form of law itself. As he states, "Nothing the code designer does rewrites laws. Instead, code design defines behavior to avoid legal sanctions." The examples he cites to illustrate how code is actually used for avoidance of the law include (1) virtual child pornography, (2) overseas gambling, (3) junk e-mail, and, (4) P2P file-sharing. Thus, according to Wu, code isn't law because, although it can influence the success or failure of a law's effects, it is more accurately viewed as a tool that interest groups use to avoid legal sanctions or use for legal advantage.

Aside from the debate over to what extent code is, or isn't, a type of law, there is also a challenge to the Mueller-led techno-centric approach that focuses on the DNS system and standards-setting processes. Scholars such as Richard Collins have emphasized how, despite the Internet being a global medium, most of the scholarship takes the United States' experience as its focus. While conceding the value of much of this work, Collins writes that, "the idiosyncrasies of the U.S. has, misleadingly, constructed a world of for-profit domain name registries, fretting about network neutrality and the like as a global experience. It is not." He further goes on to highlight three myths of Internet governance that are commonly made in the academic literature: (1) that Internet governance works best when the market decides; (2) that self-regulation is both pervasive and effective (national policies are only marginally important); and (3) that the Internet regulatory environment is distinct from legacy media.[36]

The main problem with the body of Internet governance literature to date is that each of these approaches ultimately leads to a far too narrow understanding of the governance of the entire Internet. Internet governance, particularly viewed through a policy lens, is far too complex to suggest that there is just one answer to the question—akin to one single individual or conspiracy of organizations behind the magic curtain pulling all the levers.

The aforementioned literature either focuses on only one particular aspect of the Internet or oversimplifies a very complicated topic in order to arrive at a single coherent answer. In the former case, it leaves the reader unsatisfied; in the latter, unconvinced.

THE POLITICAL ARCHITECTURE OF THE INTERNET

There exist different sets of primary actors and political arrangements at each Internet layer. As a result, the policies that govern at each layer often have fundamentally different motivations underlying them and seek to achieve different, and often conflicting, objectives. The consequence of this dynamic has been the emergence of policy processes which often address issues and formulate policy alternatives too narrowly, failing to incorporate all four Internet layers. In my conclusion, I will argue that a more comprehensive policy process involving all of the layers is needed for effective governance of the Internet, and that such a process ought to be open and transparent.

The Internet is, in fact, being governed. Staking out a historical-institutionalist approach, it will be demonstrated that policies have been intentionally developed which have shaped and continue to reshape the Internet itself. Again, direct evidence that governance is indeed taking place can be found in the case of ICANN and the governance of the DNS system.

Not only is the Internet, in fact, being governed, but it is being governed by specific and identifiable networks of policy actors at each of the conceptual layers. Governments and public institutions, private commercial firms, public-private hybrid institutions, international agencies, and various NGOs—including specific interest groups and engineering consortium groups—are all actively involved in coordination-based governance policies.

What this study will seek to accomplish in Part I is an identification of which types of policy actors have decision-making authority—an ability to govern, by our previously stated definition—at each conceptual layer. To be certain, there exist different politics, relevant actors, institutional arrangements, and types of public policies at all four of the Internet's layers. It is their identification that is the primary task at hand.

In Chapter 2, we will develop a brief narrative of the Internet's history from a governance perspective. After reviewing its evolution from being a Defense Department project to being transferred under National Science Foundation (NSF) control to, finally, being largely privatized and commercialized, we will see how all four of our conceptual layers came about chronologically and evolved through very different processes. We will argue that this historical development, including the parallel roles of both the public and

private sectors, still has tremendous ramifications for understanding Internet governance in the four conceptual layers today.

In Chapter 3, we will examine governance of the Infrastructure layer of the Internet, consisting of the wires, cables, and airwaves that make up the physical network itself. We will determine that the Internet's wired network is governed by a small handful of private telecommunications firms and cable companies who own and operate the infrastructure, and the national governments around the world that, to varying extents, regulate them, and we will explain the political dynamic using an advocacy coalitions framework. Meanwhile, when it comes to governing the Internet's wireless spectrum, we will assert that the Communications Act of 1934 and the spectrum-allocation auctions of recent years serve to demonstrate how and why the federal government—primarily the F.C.C—is the central governing authority, along with an epistemic community of engineers that is paramount in guiding its decision-making.

In Chapter 4, we will examine governance of the Protocol layer, referring to the technical standards and protocols that facilitate digital communication over the network. We will argue that decision-making authority is held by a small handful of international engineering consortium groups—primarily, the Internet Society (ISOC), its IETF, and the World Wide Web Consortium (W3C)—and we will then analyze the constitutional makeup of these organizations and assert that policymaking is best characterized by the "rough consensus" principle. Finally, we will assert that the decisions over which technical protocols to adopt, and how they are to be designed, are, in themselves, an important form of policy which constrain and enable behavior on the Internet.

In Chapter 5, we will examine governance of the Applications layer, referring to the software applications that enable people to use the Internet. We will illustrate how the code underlying both desktop and web applications is a form of policy itself. These software applications enable and constrain the actions of every Internet user on a technical basis, and thus we will demonstrate how code constitutes a unique type of policy, one in which the environment itself is designed to deny the user even a *capability* to act in defiance. We will then argue that a relatively small handful of the most well-capitalized private commercial software firms govern the Internet's applications the most—and this will be demonstrated based on several usability metrics. Ultimately, we will assert that Lawrence Lessig's "code is law" argument best explains how code constrains and enables Internet behavior, only, we will argue, that the code written by private commercial firms often indicates an implicit recognition of the sovereign authority that traditional governmental institutions retain over them.

Finally, in Chapter 6, we will examine governance of the Internet's Content layer, the most highly visible and controversial layer of them all. By highlighting several prominent issue areas such as the regulation of pornographic material online, efforts to mitigate spam, and the regulation of file-sharing over peer-to-peer (P2P) networks, we will argue that while national governments certainly have governing authority over Internet content to an extent, ISPs and private website operators (through their TOS Agreements) also have demonstrated their authority to make policies that directly constrain or enable behavior with intentional effects, particularly in the transnational context.

Fundamentally, these layers are not sequential, nor are they necessarily mutually exclusive. Policies made at one layer typically have significant consequences for shaping the policy environment at the other layers. For example, at the Protocol layer, the decision to adopt the Transmission Control Protocol/Internet Protocol (TCP/IP) protocol, which is open and universally accessible, rather than alternatives that may have allowed for far more centralized control, directly led to the development of the open and decentralized Internet that currently exists. If decisions had been made to adopt more closed, rather than open, standards and protocols, the policy environment affecting cyberspace at, say, the Content layer would be fundamentally different, allowing for greater government-imposed systems of control—as centralized authority would be built into the technology itself.

As a result, from a prescriptive point of view, I will argue that *policy objectives can be best achieved by either identifying which layer is most appropriate to a specific problem and designing narrowly targeted policies with the context of that specific layer's political dynamics in mind, or by targeting one layer with the direct intent of causing cascading effects at another layer entirely*. In other words, whether policymakers choose to work within the political architecture of one specific layer, or whether they choose to take a cross-layer approach seeking cascading effects, either way it is the conceiving of Internet-based problems in terms of our conceptual layers that will ultimately prove to be a valuable tool for policymakers. Doing so will enable the development of better Internet policies that can more reliably achieve desired outcomes.

Policymakers ought to utilize this conceptual model because it accounts for the Internet's complexities, both in technical and political terms. The four-layer model, and its resulting map of political architecture creates four distinctive policy arenas, each with its own set of criteria for determining what policy designs are most appropriate, and each with its own political dynamics that will ultimately influence to what extent a policy will be effective in achieving the desired outcomes. The question is as old as Political Science itself: If something needs to get done, who has the power to do it? The four-layer model and its resulting map of political architecture provide the answer.

Table 1.1 Summary of the Internet's Political Architecture

		Why is it Important?	*Who Governs?*	*How Are Policies Being Made?*
LAYERS	Infrastructure	Enables the actual connection between network devices	National governments, private telecom firms	Wired: Advocacy Coalitions; Wireless: Epistemic Communities
	Protocols	The languages by which devices communicate over the network	International engineering consortium groups	"Rough Consensus" principle
	Applications	The tools which allow people to make use of the network	Private commercial software firms	"Code is law" principle
	Content	The actual material that people see, read, listen to, download, watch, and interact with while online	Private ISPs, hosting companies, website operators, and national and local governments	TOS Agreements, Issue Networks

THE CASE OF U.S. NATIONAL CYBERSECURITY POLICY

In Part I of this study we will explore each of the four conceptual Internet layers—the Infrastructure, the Technical Protocols, the Software Applications, and the Content. At each, it will be ascertained why that layer is important, who governs it, and how are policies being made that affect it. Viewed in its totality, this will define the current political architecture of the Internet.

In Part II we will apply this new four-layer model and resulting political architecture by performing a detailed case study on U.S. national cybersecurity policy, post-9/11. As will be demonstrated, this case is a prime example both of what works and what doesn't when policies are designed to coordinate actions among governments, private commercial firms, hybrid institutions, and the software and engineering communities—in other words, within the context of the political architecture that will be laid out.

The story of U.S. cybersecurity policy can be thought of in two parts. First, in the initial years following the terrorist attacks of September 11th, 2001, the story is about the policymaking process that ultimately led to the *National Strategy to Secure Cyberspace (NSSC)* policy document. Second, in the years since, the story is about the formation of a new bureaucratic regime headed by the U.S. Department of Homeland Security (DHS).

Our objective will be to utilize our four-layer model and its resulting map of political architecture by analyzing the issue of national cybersecurity from a broad public policy perspective in order to test the hypothesis, and

commonly held perception, that cybersecurity policy's failures are the result of a flawed policy design that focuses almost exclusively on voluntary public-private partnerships.

First, we will conduct a descriptive analysis of the problem definition underlying the issue. The generalized problem which U.S. national cybersecurity policy is designed to address—namely, digital threats to the nation's critical cyber assets—can be made more specific by deconstructing the problem using a layer-based approach. At the Infrastructure layer, the threats include outright destruction of the Internet's physical components, such as critical telecommunications lines or operating centers, and the hijacking of industrial control systems, such as regional power grids. At the Applications layer, the threat is comprised of malicious code infiltrating vulnerable software applications to steal data or hijack network devices. At the Content layer, the threat comes in the form of defacement of websites or websites being taken offline completely.

The problem definition will be further analyzed by highlighting the categorical and specific mechanisms by which threat agents pursue their goals at each of the aforementioned layers. We will introduce a new typology that draws important distinctions between cyberterrorism, hacktivism, cracktivism, and cyberwarfare, and place specific deployment mechanisms like viruses, worms, botnets, and distributed denial-of-service (DDoS) attacks in this context. Again, our objective is to clarify the problem that cybersecurity policy is designed to address, and conceptualizing this complex, often vague, problem in terms of layers will prove useful in understanding the subsequent policy analysis.

Second, we will perform a detailed analysis of the primary document currently guiding U.S. national cybersecurity policy—the Bush administration's *NSSC*.[37] The policy design of this document is important in how it implicitly addresses all four layers in our conceptual framework. It calls for enhancing the protection of the nation's critical cyber assets by bolstering the defenses of the physical infrastructure, and directly references how this can be achieved through designing more secure technical standards and protocols, promoting more secure software application development in the private commercial sector, and by patrolling Web content.

Third, we will examine the policymaking process that led to the *National Strategy*. This process can be characterized as open, but flawed. A presidential advisory board released 53 questions to the public for comment, then drafted an initial proposal which was discussed in several town hall meetings across the country, ultimately leading to the final version of the policy. It was heavily influenced at every stage by large private corporations, and from the outset of its implementation it came under heavy criticism for failing to allocate enough resources to the problem and for relying on a strictly

voluntary public-private approach. Implementation was further hindered by a high turnover rate at the top levels within the newly created Executive bureaucracy—the DHS's National Cyber Security Division (NCSD). As we will demonstrate, this policymaking process was inclusive of most of the major governing actors set forth in our political architecture (and that in itself is significant), however organizational conflicts between them, again contextualized in terms of who has authority at each specific layer, played a large role in derailing the policy's implementation.

Next, we will seek to clarify the current bureaucratic regime governing U.S. national cybersecurity policy. As will be explained, this regime had been headed primarily by the NCSD division within the DHS, however, following a weakened period of having conflicting roles with the newly created National Cyber Security Center (NCSC), the NCSD is now competing intensely to retain its governing authority with the Department of Defense (DoD) and the National Security Agency (NSA), particularly the military's CYBERCOM command center.

Finally, we will then attempt to tie all of this together by examining cybersecurity policy in action—namely, what actually happens in the face of a cyberattack. What becomes evident is the centrality of the private sector, particularly in preventing attacks; also, the reliance on software applications and technical protocols both in prevention and response, particularly network-monitoring tools and specific antivirus products; and finally, that the federal government's role is relegated primarily to being a coordinator among private actors. U.S. Computer Emergency Readiness Team (US-CERT) is vital to raising awareness about cyberattacks and for information-sharing, but ultimately, U.S. national cybersecurity policy thus far limits the federal government from taking more forceful measures beyond that point. The four-layer conceptual model again proves helpful in contextualizing both the problem stream and solution stream surrounding the issue by framing it in these terms.

Ultimately, by applying our four-layer model and its resulting map of political architecture to the issue of U.S. national cybersecurity policy, we will argue that its overriding policy design and policymaking process are reflective of how all four conceptual layers are important in their own right, and that this confirms the utility of the four-layer model in general. The acknowledged failures of U.S. cybersecurity policy have more to do with an implementation process characterized by institutional turmoil within the Executive Branch of the federal government than with a flawed policy design or policymaking process—and, in fact, this only serves to reinforce our argument that government alone does not have adequate governing authority to achieve their desired outcomes. Even the common criticism of the NSSC's policy design relying too heavily on public-private partnerships is not so much a flawed design element as it is a recognition of the Internet's decentralized reality where

numerous governing actors have authority at different layers. The lessons of U.S. cybersecurity policy reaffirm that the best way to create meaningful Internet policies that can be effectively implemented lies in creating policies that target the layer most appropriate to specific problems in order to produce intentional cascading effects at, what is often, another layer entirely.

In summary, the main purpose of this book will be three-fold: (1) to develop a new conceptual model that deconstructs the Internet into four policy layers; (2) to use this model in formulating a new political architecture that accurately maps out and depicts authority on the Internet—ultimately determining who governs at each layer; and (3) to use the case of U.S. national cybersecurity policy, post-9/11, in order to evaluate the usefulness of both. If we are to answer, "Who governs the Internet?," we need to know how to frame the question, how to answer it, and whether or not our method of framing and our answers are helpful. That is our goal in the following chapters.

NOTES

1. *Joint Project Agreement Between the U.S. Department of Commerce and the Internet Corporation for Assigned Names and Numbers* (September 29, 2006). Retrieved on February 28, 2013 from http://www.ntia.doc.gov/files/ntia/publications/signedmou290906.pdf.

2. "Program Statistics," *ICANN*. Retrieved on June 7, 2015 from http://newgtlds.icann.org/en/program-status/statistics.

3. Ingrid Lunden, "ICANN Applicants for New TLDs Revealed as Part of 'Reveal Day': The Full List," *TechCrunch* (June 13, 2012). Retrieved on February 28, 2013 from http://techcrunch.com/2012/06/13/icann-applicants-for-new-tlds-revealed-the-full-list/.

4. Milton L. Mueller, *Ruling the Root: Internet Governance and the Taming of Cyberspace* (Cambridge, MA: MIT Press, 2002).

5. Laura DeNardis, *Protocol Politics: The Globalization of Internet Governance* (Cambridge, MA: MIT Press, 2009).

6. B. Guy Peters and John Pierre, "Governance Without Government? Rethinking Public Administration," *Journal of Public Administration Research and Theory* 8.2. (April 1998): 223–243.

7. Robert A. Dahl, *Who Governs? Democracy and Power in an American City* (New Haven, CT: Yale University Press, 1961). See page 3: "if there are great inequalities in the conditions of different citizens, then there must also be great inequalities in the capacities of different citizens to influence the decisions of their various governments."

8. Carolyn J. Hill and Laurence E. Lynn, Jr., "Is Hierarchical Governance in Decline? Evidence from Empirical Research," *Journal of Public Administration Research and Theory* 15.2 (2005): 173–195.

9. H. George Frederickson and Kevin B. Smith, *The Public Administration Theory Primer* (Boulder, CO: Westview Press, 2003).

10. Hill and Lynn, 175.

11. Hill and Lynn, 174.

12. Candice Jones, William S. Hesterly, and Stephen P. Borgatti, "A General Theory of Network Governance: Exchange Conditions and Social Mechanisms," *The Academy of Management Review* 22.4 (October 1997): 911–945.

13. Mark C. Taylor, *The Moment of Complexity* (Chicago, IL: University of Chicago Press, 2001).

14. Richard Collins, *Three Myths of Internet Governance: Making Sense of Networks, Governance, and Regulation* (Chicago, IL: University of Chicago Press, 2009), 59.

15. Collins, 60.

16. Piotr Konieczny, "Adhocratic Governance in the Internet Age: The Case of Wikipedia," *Journal of Information Technology and Politics* 7.4 (October 2010): 263–283.

17. Henry Mintzberg, *Tracking Strategies: Toward a General Theory* (Oxford, England: Oxford University Press, 2007).

18. Lawrence Lessig, *Code and Other Laws of Cyberspace* (New York, NY: Basic Books, 1999). Lessig does not explicitly define "governance," but makes references to how code and "architecture regulates behavior" in cyberspace, and to "constraints on how you behave." See Chapter 7.

19. Michel Foucault, *Power/Knowledge: Selected Interviews and Other Writings, 1972–1977*, ed. Colin Gordon (New York, NY: Pantheon, 1980).

20. C. Wright Mills, *The Power Elite* (New York, NY: Oxford University Press, 1956). See pp. 3–4: "Whether they do or do not make such decisions is less important than the fact that they do occupy such pivotal positions: their failure to act, their failure to make decisions, is itself an act that is often of greater consequence than the decisions they do make. . . . Often they are uncertain about their roles, and even more often they allow their fears and their hopes to affect their assessment of their own power. No matter how great their actual power, they tend to be less acutely aware of it than the resistances of others to its use."

21. Bertrand Russell, *Power: A New Social Analysis* (London, England: Allen and Unwin, 1938).

22. Laura DeNardis offers a helpful definition of Internet governance referring to "policy and technical coordination issues related to the exchange of information over the Internet" in the context of architecting civil liberties into IPv6 protocol design. See Laura DeNardis, *Protocol Politics: The Globalization of Internet Governance* (Cambridge, MA: MIT Press, 2009).

23. Marcus Franda, *Governing the Internet: The Emergence of an International Regime* (Boulder, CO: Lynne Rienner Publishers, 2001).

24. Yochai Benkler, "From Consumers to Users: Shifting the Deeper Structures of Regulation," *Federal Communications Law Journal* 52 (2000): 561–563.

25. Lawrence Lessig, *The Future of Ideas: The Fate of the Commons in a Connected World* (New York, NY: Vintage Books, 2002).

26. Lessig, *Code*.

27. David Bell, *An Introduction to Cybercultures* (New York, NY: Routledge, 2001), 2.

28. Jack Goldsmith and Tim Wu, *Who Controls the Internet? Illusions of a Borderless World* (New York, NY: Oxford University Press, 2006).

29. Franda.

30. John Perry Barlow, *Declaration of the Independence of Cyberspace* (1996). Retrieved on January 5, 2006 from http://www.eff.org/~barlow/library.html.

31. Richard Barbrook and A. Cameron, *The Californian Ideology* (1996). Retrieved on August 23, 2012 from http://www.hrc.wmin.ac.uk/theory-californianideology.html.

32. Langdon Winner, *Autonomous Technology: Technics Out of Control as a Theme in Political Thought* (Cambridge, MA: MIT Press, 1977).

33. Andrew Shapiro, *The Control Revolution: How the Internet is Putting Individuals in Charge and Changing the World We Know* (New York, NY: Public Affairs, 1999).

34. Lessig, *Code*.

35. Tim Wu, "When Code Isn't Law," *Virginia Law Review* 89.4 (June 2003): 679–751.

36. Collins.

37. *National Strategy to Secure Cyberspace* (February 2003). Retrieved on March 5, 2005 from http://www.defense.gov/home/features/2010/0410_cybersec/docs/cyberspace_strategy%5B1%5D.pdf.