How Regulation and Competition Influence Discrimination in Broadband Traffic Management:

A Comparative Study of Net Neutrality in the United States and the United Kingdom

Alissa Cooper

St. Catherine's College

University of Oxford

Thesis submitted in partial fulfillment of the requirements for the degree of DPhil in Information, Communication and the Social Sciences in the Oxford Internet Institute at the University of Oxford

September 2013

93,137 words

How Regulation and Competition Influence Discrimination in Broadband Traffic Management: A Comparative Study of Net Neutrality in the United States and the United Kingdom

Alissa Cooper
DPhil in Information, Communication and the Social Sciences
St. Catherine's College, Trinity 2013

Abstract

Telecommunications policy debates concerning the contentious issue of net neutrality have revolved around a number of broadband network operator behaviors, including discriminatory traffic management – differential treatment of network traffic associated with different Internet applications for the purpose of managing performance. Some stakeholders have advocated for regulatory intervention to prevent network operators from discriminating to the detriment of independent application innovation. Others would prefer to rely on competition between network operators to discipline operator behavior.

Fixed-line broadband markets in the United Kingdom and the United States have differed substantially with respect to discrimination, competition, and regulation. The UK has experienced intense competition and pervasive discriminatory traffic management without triggering regulatory activity. The US has seen much less discrimination, limited competition, and regulatory threat followed by regulatory intervention. This thesis uses elite interviews, participant observation, and documentary analysis in a comparative study of these two cases between the mid-2000s and 2011 to determine why network operators take up discriminatory traffic management (or not) and how competition and the regulatory environment affect traffic management outcomes.

This thesis demonstrates that network operators take up discriminatory traffic management primarily to control cost, performance, or both. Competition promotes rather than deters discrimination because it drives broadband prices down, encouraging operators to manage high-volume applications whose traffic incurs high costs. Regulatory threat can be sufficient to counteract these desires, but in its absence and without concerns vocalized by interest groups, discriminatory approaches endure. Telecommunications regulators intervene to safeguard nondiscrimination when they conceive of their remits as encompassing social and industrial policymaking, are ambivalent about litigation risk, and are driven by their leaders' reputational agendas, as in the case of the Federal Communications Commission. With a narrower perception of its remit and more concern for its organizational reputation, Ofcom exemplifies the characteristics that inhibit traffic management regulation.

Acknowledgments

I am extremely privileged to have had the support of so many wonderful people in undertaking this research. One could not ask for a more caring and helpful family in pursuing doctoral work. I am forever indebted to my husband Marshall, who has endured so much with me and for me while I completed this project, and whose boundless interest in discussing arcane issues of telecommunications policy never ceases to amaze me. I never would have started or finished this thesis without the encouragement and assistance I received from my father, Dr. Mark Cooper, a legend in his own right who continues to inspire me every day. Both Marshall and Mark have read every word of this thesis and it is all the better for it. And of course I could not be more thankful for my mother Carol's willingness to do absolutely anything in her power to help me along the way.

I owe my sincere thanks to my advisor, Dr. Ian Brown, whose incisive assessments of my work and encouragement to persevere were invaluable. I am also grateful for his having introduced me to Dr. Alison Powell during my first week at Oxford. My collaboration with Alison inspired the structure of this thesis and her friendship has encouraged me during challenging times and delighted me at all others. I am also extremely thankful to have received scholarships from the Oxford Internet Institute and St. Catherine's College, without which this thesis would not have been possible.

This thesis would also not exist without the 70 busy people who were willing to spend their time offering frank, vivid insights about traffic management in interviews. I am especially thankful to the team at the ISP where I conducted participant observation for allowing me to both join their team and gain a window into their daily work.

The analysis of the UK broadband market in this thesis owes much to my endless conversations with Andrea Soppera, who taught me more about the broadband industry than I ever could have asked for. I was gratified for the help and feedback that I received from Prof. Barbara van Schewick, whose pioneering work was a dependable guide throughout every phase of my doctorate. I am also thankful to my colleague Gautam Hans for his proofreading assistance.

Last, but certainly not least, I am truly thankful to Leslie Harris and all of my colleagues at the Center for Democracy and Technology (CDT), where I worked full-time prior to arriving in Oxford and part-time throughout the duration of this project. The day that I started at CDT in 2006 we embarked on the organization's first net neutrality project and my interest in the topic has been nurtured every since through close collaboration and thoughtful conversations with my colleagues. Leslie was the ideal boss, accommodating and supportive at every turn. Working at CDT helped me keep a grip on the reality of Internet policy while experiencing the joy of a workplace where the density of brilliant and talented people is second to none.

Contents

Figures

Tables

Abbreviations

APA	Administrative Procedure Act of 1946 (US)
BEREC	Body of European Regulators for Electronic Communications
BIS	Department for Business, Innovation & Skills (UK)
CAT	Competition Appeal Tribunal (UK)
DCMS	Department for Culture, Media & Sport (UK)
DNS	Domain Name System
DOCSIS	Data Over Cable Service Interface Specifications
DPI	Deep packet inspection
DSL	Digital subscriber line
FCC	Federal Communications Commission (US)
FTC	Federal Trade Commission (US)
GB	Gigabyte
Gbit/s	Gigabit per second
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext Transfer Protocol Secure
IPv6	Internet Protocol version 6
ISP	Internet service provider
Kbit/s	Kilobit per second
Mbit/s	Megabit per second
MP	Member of Parliament
NPRM	Notice of Proposed Rulemaking
NRA	National regulatory authority (EU)
OECD	Organisation for Economic Co-operation and Development
Ofcom	Office of Communications (UK)
Oftel	Office of Telecommunications (UK)

OJC	Official Journal of the European Union C series (Information and Notices)
OJL	Official Journal of the European Union L series (Legislation)
Quango	Quasi-autonomous non-governmental organization
SMP	Significant market power
STM	Subscriber traffic management
TCP	Transmission Control Protocol
TSR	Telecommunications Strategic Review (UK)
U.S.C.	United States Code
VoIP	Voice over Internet Protocol

Chapter 1. Introduction

As global broadband Internet adoption has expanded, few telecommunications policy issues have been as contentious as that of net neutrality. The issue has engaged regulators, legislators, and courts throughout the world; precipitated massive advocacy and media campaigns; and triggered scholarly critique from across a wide range of academic disciplines. The debate has revolved around a number of existing or potential broadband Internet service provider (ISP) behaviors, including discriminatory traffic management – differential treatment of network traffic associated with different Internet applications for the purpose of managing the performance of the network. Some stakeholders have advocated for regulatory intervention on the basis that this sort of traffic management and other forms of discrimination give network operators too much power over which applications will succeed or fail. Others would prefer to rely on competition between network operators to discipline their behavior, fearing that regulation would unnecessarily constrain the future development of broadband.

The responses of regulators and network operators to these arguments have varied across countries. The fixed-line broadband markets in the United States and the United Kingdom provide a particularly insightful comparison, as the two countries have exhibited rather opposite extremes in three crucial respects: competition, discrimination, and regulation. The UK fixed broadband market has been among the most competitive in Europe and has seen extensive use of discriminatory traffic management without triggering significant regulatory intervention. The US has seen substantially less discriminatory traffic management, limited competition, and a policy environment characterized by regulatory threat, culminating with the Federal Communications Commission's Internet openness rules enacted in late 2010. Thus the two countries' contrasting experiences provide a means to understand the relationships among competition, the regulatory environment, and discrimination (or the lack thereof) in the provision of traffic management.

The regulation-versus-competition debate – and normative legal and economic suggestions about the appropriate role for regulatory intervention more generally – have dominated the net neutrality discussion within academic and policy circles. But the practical realities of how network operators have gone about managing network traffic and how regulators have responded are not consistent with the most prominent arguments about how markets and regulation should work. Examining why the broadband Internet is operated as it is and how regulatory influences affect operational decisions is critical to understanding the broader societal and economic implications of broadband policy choices.

To develop that understanding, this thesis uses a qualitative comparative study of broadband traffic management on fixed networks in the US and the UK from the mid-2000s to 2011. Combining insights from elite interviews, participant observation, and documentary analysis, it provides a variety of substantive, positive, empirical contributions to both academic and policy discourse concerning net neutrality and telecommunications regulation more broadly. It identifies both the operational and institutional influences that fuel ISPs' traffic management decisions, including the significant role of regulatory threat (or the lack thereof). It explains how intense competition can promote rather than deter discrimination, as evidenced in particular by the UK case. And it identifies the features of telecommunications regulatory agencies that make the most salient contributions to particular traffic management policy outcomes, focusing on both formal statutory constraints and informal pressures and tensions. Together these findings advance the current understanding of how discrimination, competition, and regulation influence and are influenced by each other.

The bulk of this thesis is dedicated to elucidating these findings. The remainder of this chapter provides terminology and concepts necessary to understand traffic management and reviews relevant telecommunications history in both countries. Chapters 2 and 3 review the literature and theory related to net neutrality and regulation more broadly. Chapter 4 discusses the research methods used in this thesis. Chapters 5 and 6 explain network operators' traffic management decisions in the US and the UK, respectively. Those findings are further synthesized in Chapter 7 in light of prominent arguments from the net neutrality literature concerning discrimination, competition, and innovation. Chapters 8 and 9 focus on regulatory environments, examining the roles of reputation and relationships in the context of traffic management. Chapter 10 situates the prior chapters' findings in the context of existing academic scholarship and future policy debates.

1.1 Traffic Management in the Context of Net Neutrality

The net neutrality debate has focused on a number of existing or potential network operator behaviors that can affect the performance of an Internet user's network connection. At the network level, performance can be described by the levels of congestion, delay, or jitter (variation in delay) that network traffic experiences and whether these metrics remain stable over time. At the human level, performance might be judged on quality of experience: whether streaming video is choppy, whether voice calls have clear audio, or how long it takes for web sites to load or files to download. While measures adopted to mitigate security threats or to prevent the transfer of illegal or unwanted content have been acknowledged within the policy debate (FCC 2009; FCC 2010b), the key focus has been on performance, and that is the focus in this thesis.

The ISP behaviors that have received significant attention include blocking access to particular content or services, charging application providers fees for prioritized traffic delivery, and differential treatment of different applications for the purpose of managing network performance. This thesis addresses the last of these practices, described hereafter as "application-specific traffic management" or "discriminatory traffic management." (The term "traffic management," common in the British policy discourse, is used here rather than the more general "network management" term from the American discourse, although as the terms are used by many commentators they often describe the same scope of practices.) This section defines application-specific traffic management and situates it within the larger scope of ISP behaviors that have been the subject of academic and policy debates. The terms defined in this section provide a novel synthesis of definitions offered within academic and policy discourses with the practical set of choices that ISPs have available to them for managing their networks.

1.1.1 Engineering Choices That Affect Performance

Internet service providers face many engineering choices that can affect network performance. These choices fall broadly into two categories: capacity management and traffic management.

Capacity Management

Capacity management involves decisions about deploying network infrastructure. Network capacity is comprised of two basic components: physical links (wires, cables, or fibers) and the network equipment that connects them. To increase network capacity, existing links may be replaced with faster links (for example, replacing a 100 Mbit/s link with a 1 Gbit/s link), or additional links may be deployed. If sufficient capacity exists on the links but the equipment is not able to make use of it, equipment upgrades can also provide additional capacity (for example, upgrading cable equipment to make use of existing cable channel capacity for Internet service, as many cable operators have done in recent years (Kumar 2008)). One question that has been explored in the net neutrality context (and in computer networking more generally) is whether relying on capacity upgrades alone can suffice to meet a network's performance goals without having to introduce traffic management measures (Bauer, Clark, and Lehr 2009; Bell 2003; Lessig 2006; Odlyzko 2009).

Increasingly, broadband networks also support storage capacity in the form of caches and content delivery networks. Storage capacity can be used within networks to reduce the distance that bits need to travel to reach end users, thereby improving performance. Likewise, ISPs make choices about how much capacity to dedicate to interconnecting networks and how close to end users they allow other networks to interconnect (Faratin et al. 2008).

Traffic Management

Traffic management concerns the treatment of network traffic at some set of management points (routers or servers, for example) within an ISP's network, where traffic flows in and out of the management points (Felten 2006). In the abstract, traffic management involves three components: (1) one or more subsets of traffic to be managed, (2) the trigger for applying differential treatment to those traffic subsets, and (3) the differential treatment to be applied. (This framework is a more general version of the recognition-manipulation-notification framework used by Mueller (2011) to describe the capabilities of deep packet inspection technologies.)

The first component relies on some criteria that can be used to identify a subset of traffic. An ISP may seek to manage traffic associated with a particular source, destination, application, user, or set of users, for example.

The second component is the trigger for applying the traffic treatment, which may take a number of forms. Traffic management might be triggered based on the time of day, so that it is only applied during peak usage periods, for example. It might be triggered by a particular network condition such as a utilization threshold or a specific level of packet loss in one or more parts of the network. Or the trigger might be a usage threshold, such that traffic management is applied when a user reaches a pre-determined volume limit. A usage threshold might also trigger other consequences instead of or in addition to traffic management – charging fees, suspending user accounts, automatically upgrading users to higher tiers, or sending warning letters concerning the possibility of any of the above.

The final component is the traffic treatment – what is done to the identified traffic when the trigger is met. There are three types of traffic treatment that have been at the center of net neutrality debates:

- **Blocking**, where identified traffic is prevented from being delivered (Dischinger et al. (2008) and Beverley et al. (2007) provide examples). Obviously blocking is sometimes used for reasons other than achieving performance goals, for example to prevent access to politically disfavored content or to competing services.
- **Prioritization (or de-prioritization)**, where identified traffic is sent sooner or later than other traffic at the management point regardless of which traffic arrived first (Kanuparthy and Dovrolis (2010) provide examples). Prioritization is only relevant in situations where more traffic arrives at the management point than it can process and send at once.
- **Rate limiting**, where the identified traffic is limited to a specified sending rate (Weinsberg, Soule and Massoulie (2011) provide examples). This can be accomplished at a management point in the network by buffering traffic if it arrives more quickly than the specified rate or dropping packets of identified traffic to achieve the specified rate. Rate limits can be assigned per user or in the aggregate. They may be set in absolute terms (for example, 256 Kbit/s), or relative to overall demand (for example, 2% of peak bandwidth).

Other forms of traffic management, including compressing or transcoding traffic, have not garnered as much attention as the types of practices listed above.

Several treatments may be combined, and different treatments may be used to achieve similar goals. For example, one way that ISPs have sought to protect the performance of delay-sensitive applications such as voice over Internet Protocol (VoIP) has been to rate limit peer-to-peer traffic on the assumption that peer-to-peer file transfers are less delay-sensitive (Zachem 2008). An alternative (or complementary) approach is to prioritize VoIP traffic so that if congestion does appear in the network, the delay-sensitive VoIP packets will be less likely to be affected (Mooyaart 2012).

1.1.2 Discriminatory Engineering Choices

As discussed in the next section, the concept of discrimination in the carriage of communications is as old as communications networks themselves. But what does it mean for engineering choices that broadband providers make to be discriminatory?

Whether traffic management is considered discriminatory (or "application-specific," where "application" is used as shorthand for "application, content, or service") depends on the criteria used to identify the traffic to be managed. Application-specific traffic management is based on criteria associated with particular *uses* of the network (FCC 2010b; van Schewick 2012). That is, traffic management is application-specific if traffic is selected to be managed because it:

- has a particular source or destination (bbc.co.uk, for example),
- is generated by a particular application (a BitTorrent client, for example),
- is generated by an application that belongs to a particular class of applications (a class of video chat applications that includes Skype, Google Talk, WebEx, and FaceTime, for example), or
- uses a particular application- or transport-layer protocol (Session Initiation Protocol or User Datagram Protocol, for example).

Traffic might be identified based on packet payloads (using deep packet inspection or other content-aware network devices), network or transport layer headers (port numbers, for example), heuristics (the size, sequencing, and/or timing of packets), or a combination of these characteristics (Allot Communications 2007).

Application-agnostic (or "nondiscriminatory") traffic management, by contrast, is based on criteria associated with particular *users* of the network. For example, an application-agnostic traffic management policy may target all users signed up to a particular service tier, or all users who have consumed a particular amount of data over an interval of minutes, hours, or days.

Notably, equating discrimination with application-specificity implies a narrower understanding of "discrimination" than one might obtain from the dictionary definition of the term. In some sense, many traffic management decisions involving scarce resources may be said to require some form of "discrimination" because arbitrary decisions must be made at network management points about which packet to send next when there is a choice of more than one, or which packet to drop when space is not available to store them all (Felten 2006). Routers that follow a first-in first-out sending policy or drop the most recently received packets could be said to be "discriminating" against packets based on their arrival times. But as it is used in net neutrality debates and in this thesis, "discrimination" means the much narrower set of differential decisions that are based on traffic being associated with particular uses of the network.

The distinction between application-specific management and application-agnostic management can also be applied to capacity management. For example, when network or storage capacity is allocated specifically for use by one application or class of applications, or when an interconnection agreement is struck for a particular kind of application traffic (VoIP, for example), those decisions can be considered application-specific. Increasing the capacity available for all Internet traffic is generally considered to be an application-agnostic approach, since capacity is made available for any uses of the network that users desire. Although questions about how regulatory proposals might address discriminatory caching, content delivery, or interconnection have arisen within discussions of net neutrality (see, for example, Ryan (2011)), the focus here is exclusively on the traffic management aspect.

A few examples out of the many application-agnostic engineering choices that ISPs in the US and the UK have made since the early 2000s include:

- Increasing network capacity so as to offer customers 2 Mbit/s peak rates instead of 1 Mbit/s;
- Instituting tiers of service that offer customers different speeds: 2, 8, or 20 Mbit/s; and

- Imposing volume limits (1 GB/day or 250 GB/month, for example) and charging fees or reducing the speeds of users who reach those limits.

Examples of application-specific engineering choices that ISPs in the US and the UK have made include:

- Rate limiting peer-to-peer traffic;
- Prioritizing VoIP and gaming traffic; and
- Preventing peer-to-peer file sharing applications from exchanging traffic.

Unlike several other practices that have garnered attention within net neutrality circles, traffic management techniques have been in use for years in many countries, including the US and the UK. As such, traffic management provides an observable phenomenon for study.

1.1.3 Distinguishing Traffic Management from Other Practices at Issue

By contrast, the less widely used practices that have nonetheless generated significant debate involve network operators charging fees to application providers for prioritized traffic delivery. Such schemes involve using a traffic treatment – prioritization – as the direct basis for a product offered to suppliers of applications. This kind of product offering is generally considered as distinct or outside of the umbrella of traffic management (FCC 2010b). Its distinguishing feature is that it involves up-front negotiations between ISPs and the providers of independent applications and services. Traffic management as conceived within the US policy debate (and to a lesser extent in the UK) tends to be confined to practices taken up by ISPs of their own accord without prior discussion with the developers of the applications that the management may affect.

As noted above, ISPs may decide to block traffic for reasons other than managing performance, including to foreclose competing applications or to limit access to disfavored content. In certain circumstances the lack of a performance rationale may be obvious, such as when Madison River, a small US-based DSL provider, began blocking its customers from accessing over-the-top VoIP services in 2005 (before agreeing to an FCC consent decree that prohibited the blocking). Indeed, as a general matter, it would be difficult for any ISP to argue that blocking VoIP traffic, which is known to be low-volume, noticeably improves the performance of non-VoIP uses of the network, and it certainly does not improve the performance of VoIP itself. Practices construed as traffic management, on the other hand, tend to be accompanied by some form of performance rationale.

In many cases, it may be impossible for outside observers to separate an ISP's motives to manage performance from its incentives to increase the profitability of its network or deter competition (Marsden and Cave 2007). For example, when an ISP that also sells an IP-based video-on-demand service decides to rate limit or selectively block peer-to-peer file-sharing traffic, it raises the question of whether the ISP is limiting peer-to-peer traffic so as to drive more customers to its video service, improve the performance of non-peer-to-peer applications, or both (FCC 2008). Similarly, when an ISP offers broadband products with different combinations of prioritization schemes to different users, it shows how a basic performance-boosting technique – prioritizing delay-sensitive traffic over delay-tolerant traffic – can be incorporated into a larger differential pricing strategy. As Wu (2003, 154) has noted, "[e]ven if the goal itself is legitimate, the method of achieving that goal may be suspect."

Although separating performance incentives from economic incentives can be difficult, the term "traffic management," as it is used in this thesis, describes practices roughly circumscribed by the two rules articulated above: they involve no up-front negotiation with application providers, and they are accompanied by some performance rationale. These two constraints do not provide a strict boundary, but they attempt to separate traffic management from the larger space of practices at issue in net neutrality debates, including blocking or discrimination for purely anti-competitive purposes, blocking to restrict speech or access to information, and schemes that involve ISPs charging application providers for delivery or prioritized delivery of their traffic.

1.2 Background and Context

The debate about network operators discriminating between applications in the provision of Internet service was first catalyzed by the transition from dial-up to broadband Internet access speeds at the turn of the 21st century. Offered initially by cable operators using their existing cable television plant and telephone companies using digital subscriber line (DSL) technology, the availability of speeds tens to thousands of times faster than dial-up spurred the development and popularization of new applications on the network (Bauer, Clark, and Lehr 2009). Media-rich web browsing, VoIP, peer-to-peer file sharing, interactive gaming, and video streaming proliferated, each with unique performance requirements and characteristics. VoIP applications, for example, generate only small amounts of data but require low-latency transmission to ensure call quality, whereas peer-to-peer file sharing applications, which can transfer large amounts of data, can perform acceptably despite variations in connection speed. Some new applications, including VoIP and video streaming and downloading, also presented potential competitive threats to similar services already offered by the telephone and cable companies.

Observing the effects of new applications on traffic loads, customer experiences, and competitive prospects, network operators began to seek new ways to manage traffic on the network (FCC 2009), including mechanisms that applied different treatment to different applications. At the same time, US regulators, legislators, and judges were beginning to grapple with how to apply and adapt existing telecommunications policy to broadband Internet service, a process that formed the precursor to the US net neutrality debate in the mid-2000s and the EU debate that followed. This section provides a brief chronology of relevant policy events on both sides of the Atlantic to provide the context that forms the basis of this thesis.

1.2.1 United States

Regulatory Obligations from Dial-Up to DSL

From the advent of telecommunications regulation in the US, telephone companies were subject to a variety of obligations as "common carriers," a concept established under English common law and incorporated into the regulation of US transportation and communications providers in the late 19th century. Common carriage obligations for US telephone companies were established both to temper monopoly power and to ensure that the public interest was served even where operators lacked monopoly power (Speta 2002b). These obligations were formally specified in regulation as Title II of the Communications Act of 1934, which created the FCC. Title II included a nondiscrimination clause that prohibited "unjust and unreasonable discrimination in charges, practices, classifications, regulations, facilities, or services" (47 U.S.C. 202(a)). In contrast to Title II, services classified under Title I are subject only to the FCC's much more general authority to carry out its regulatory duties, and are not subject to common carriage obligations (Nuechterlein and Weiser 2005).

The net neutrality debate grew out of 40 years of regulatory skirmishes over the extent to which common carriage obligations should apply to data services offered by telephone companies. In the mid-1960s, computers were increasingly being integrated with communications networks to offer new data processing and transmission capability. With concerns that AT&T as the nation's monopolist telephone carrier might seek to leverage its market power at the physical network layer into the nascent data processing industry ("vertical leveraging" in antitrust parlance), the FCC launched its first *Computer Inquiry* (known as *Computer I*). This resulted in a separation between "pure communications," regulated as common carriage services, and "pure data processing," classified under Title I (FCC 1971). Around the same time and after decades of litigation concerning the attachment of independently produced devices to the telephone network (Hush-a-Phone v. United States, 238 F.2d 266, 1956), the FCC instituted the *Carterfone* principle, establishing the rights of users to attach the devices of their choice to their home networks (FCC 1968).

As the decades passed, the distinction between common carrier services and data services persisted, albeit under different names. The *Computer II* proceeding in the 1980s retained the regulatory classifications but introduced the terms “basic” and “enhanced” services to distinguish pure transmission capability from services sold to the public that involve some sort of data or content processing (FCC 1980). The categories were again renamed with the passage of the 1996 Telecommunications Act, the largest overhaul of communications law in generations. “Telecommunications services” replaced basic services and remained under Title II, while “information services” replaced enhanced services and remained under Title I (47 U.S.C. 151 *et seq.*). Although the classification was not made immediately explicit, Internet services, offered at the time primarily by dial-up ISPs such as AOL and EarthLink, were generally considered to be information services (Nuechterlein and Weiser 2005). Telephone companies were prevented from discriminating against independent dial-up ISPs since these ISPs’ services used the telephone network’s transmission facilities, which were regulated under Title II.

Computer I had effectively excluded phone companies from the data services market through structural prohibitions. In *Computer II*, the Commission relaxed those prohibitions slightly, allowing large common carriers to enter the data services market through wholly separate subsidiaries (a policy known as “structural separation”). The FCC also added an unbundling obligation, requiring common carriers to separate out the raw transmission capability of their networks for sale on a nondiscriminatory basis to any buyer who wanted to use the network to provide enhanced services. A further option was created in *Computer III*, which instituted “non-structural” obligations regarding network information disclosures and nondiscrimination rules. Phone companies at this point had the choice of offering enhanced services under either the structural separation regime of *Computer II* or the non-structural separation regime of *Computer III* (Cannon 2001). Regardless of their choice, however, all wireline carriers that owned common carrier facilities and provided enhanced services were still bound by the

Computer II requirement to purchase their own transmission capacity at the same price available to all competitors (FCC 2005b).

Taken together, the *Carterfone* principle, the relaxation of the separation requirements, and the codification of service classifications first introduced in the *Computer Inquiries* formed the regulatory environment into which broadband Internet service offered via DSL was born and popularized in the late 1990s and early 2000s. Because DSL broadband was offered by wireline carriers that also operated as common carriers, DSL providers were still bound by the unbundling, disclosure, and (in some cases) separation requirements of *Computer II* and *Computer III*, all of which remained in place until 2005. The classification of DSL broadband under the updated Communications Act was less clear, but in cases where the owner of the DSL facilities (the phone company) was selling DSL transmission capacity to independent ISPs who would then offer broadband Internet as a retail service (as was common at the time), the telcos were clearly subject to Title II nondiscrimination requirements (Nuechterlein and Weiser 2005).

Cable "Open Access"

Relative to the telephone companies, the cable broadband industry had enjoyed a short and favorable regulatory history by the time the net neutrality debate began. Although cable television was regulated under Title VI of the Communications Act, cable broadband was introduced into the marketplace without explicit regulatory classification, and therefore without regulatory obligations. As cable broadband gained popularity in the late 1990s, debates about its appropriate regulatory regime played out primarily in two forums: agency reviews of proposed cable mergers and municipal decisions about whether to award cable franchises.

On both fronts, the cable industry scored early victories. In both the AT&T-Tele-Communications, Inc. and AT&T-MediaOne merger proceedings, despite requests from competing ISPs and consumer advocates to impose "open access" obligations that would require the merged entities to provide competing ISPs with access to their networks (similar to the unbundling rules for phone networks), the FCC declined to condition the mergers with open access provisions (FCC 1999; 2000a). AT&T became the largest cable operator in the country. Although the FTC subsequently imposed a limited open access condition on the far smaller cable company resulting from the AOL-Time Warner merger (FTC 2000), the dominant cable broadband business at the time was free of such obligations.

Around the same time, most municipalities debating the open access issue declined to condition their cable franchise agreements with an open access requirement. The most notable exception was the city of Portland, Oregon, which did impose such a condition on AT&T's franchise bid. But even Portland's decision ultimately resulted in a victory for the cable industry. In order to rule on AT&T's appeal of the Portland decision, the Ninth Circuit Court of Appeals needed to determine the regulatory status of AT&T's cable broadband service, as the Communications Act relieves telecommunications service providers of the need to obtain local franchises (AT&T Corp. v. City of Portland 43 F. Supp. 2d 1146 (D. Or. 1999)). Having an individual court decide the regulatory classification of a single cable broadband service urgently demonstrated the need for a nationwide policy on the matter, and the FCC was spurred to action. It launched the *Cable Modem* proceeding (FCC 2000b), which culminated in 2002 with the *Cable Modem Declaratory Ruling* in which the agency classified cable broadband as an "information service" not subject to common carriage obligations (FCC 2002). The Commission also declined to extend to cable broadband operators the sort of unbundling rules that applied to telephone companies.

Net Neutrality Beginnings

The *Declaratory Ruling* was soon challenged and eventually landed at the Supreme Court, where the FCC's classification decision was ultimately upheld (National Cable & Telecommunications Assn. v. Brand X Internet Services 545 U.S. 967 (2005)). Broadband Internet service offered over cable networks was officially classified as an information service not subject to common carrier obligations. Shortly thereafter, in its *Wireline*

Broadband Order (FCC 2005b), the FCC ascribed the same classification to wireline broadband Internet service (offered via DSL), creating parity with cable and ending the telephone industry's decades-long struggle to free its provision of data and broadband Internet services from regulation under Title II. As a result of the FCC's move away from unbundling in favor of inter-platform competition, most US broadband customers have had at most two choices of fixed-line broadband Internet service provider, one telephone company and one cable provider, since the mid-2000s (Government Accountability Office 2006; FCC 2012). The competitiveness of the dial-up era would not be carried over to broadband (Noam 2009).

To counterbalance the deregulatory steps taken in the *Wireline Broadband Order*, the FCC simultaneously issued its *Broadband Policy Statement* (FCC 2005c), laying out a framework of principles based on the highly influential "Four Freedoms" that FCC Chairman Michael Powell had articulated in a speech delivered in February 2004 (Powell 2004). These principles declared broadband consumers' rights to access the legal content and services of their choice, to connect the legal devices of their choice to the network, and to enjoy competition among broadband service providers. In making all four principles subject to "reasonable network management" in the *Policy Statement*, the FCC formally recognized broadband traffic management for the first time, although it declined to provide any guidance as to how it might judge the reasonableness of such management.

The deregulation of Internet service that resulted from the *Brand X* decision and the *Wireline Broadband Order* and the decline in operator competition since the dial-up era catalyzed concerns that had been brewing among academics, advocates, and Internet companies about the potential for broadband ISPs to discriminate against independent sources of content and applications. In 2003, Tim Wu had coined the term "network neutrality" to describe a policy prescription meant to ensure that the Internet "does not favor one application (say, the world wide web), over others (say, email)" (Wu 2003, 145). Spurred on by remarks made in the press by telco executives about their desires to introduce such discriminatory treatment on their networks (Mohammed 2006; O'Connell 2005; Reardon 2006), half a dozen bills were introduced in Congress between 2005 and 2007 to enshrine various versions of nondiscrimination rules into law. A highly public and vitriolic policy debate ensued, but net neutrality legislation was not enacted. The Federal Trade Commission conducted its own inquiry into the issue, concluding in 2007 that new regulation was unnecessary and potentially harmful (FTC 2007).

The FCC's response to the debate was piecemeal at first, with nondiscrimination provisions included in a number of telecommunications mergers (FCC 2005d; FCC 2005e; FCC 2006), a major spectrum auction (FCC 2007b), and a generic inquiry into broadband providers' practices (FCC 2007a). The agency stepped up its involvement by launching an investigation and issuing an order against Comcast in 2008 as a result of accusations that the ISP was engaging in discriminatory traffic management (FCC 2008). The FCC's authority to issue that order was challenged in court, resulting in the order being vacated (Comcast Corp. v. FCC 600 F.3d 642. (2010)). In late 2010, the agency nonetheless adopted formal net neutrality rules in its *Open Internet Order* (FCC 2010b), which has also been challenged in court. Thus the US broadband business in recent years has been characterized by both regulatory threat and regulatory uncertainty, with the attention of policymakers, regulators, and the public focused on both instances of discrimination and questions about the FCC's authority to prevent or prohibit them.

Within this climate, forays into discriminatory traffic management on US fixed-line networks since the mid-2000s have been limited (see Chapter 5). With the exception of some cable operators' use of peer-to-peer-focused management tools – which came to an end in relatively short order as a result of attention from the FCC and the public – the networks of the largest US broadband providers have been largely free of discriminatory traffic management. The combination of these facts with relatively limited levels of broadband competition and a culture of regulatory threat make the US case a compelling contrast to the UK.

1.2.2 United Kingdom

Drive Towards Liberalization

The roots of the current telecommunications regulatory regime in Britain – and the net neutrality debate – provide a stark contrast to those of the US. For most of the 20th century, telephone service in the UK was provided under state control by the Post Office. When British Telecommunications (BT) was created and then privatized under the Telecommunications Act of 1984, it accrued nondiscrimination obligations on the basis of its market power (c. 12, art. 8(1)) and oversight in the form of the Director General of the Office of Telecommunications (Oftel), who was responsible to the Secretary of State for Trade and Industry. While nondiscrimination obligations under the American regulatory regime were never solely predicated on demonstrations of market power, UK telecommunications policy cabins the regulation of discriminatory conduct within the confines of its competition framework, in line with broader European regulation.

As liberalization was being taken up across the continent at the turn of the 21st century, the major project of European telecommunications policy was to build a harmonized regulatory framework that would support competition across the region's telecommunications industries. The result was a new EU regulatory framework for electronic communications that came into force in 2003 (OJL 108/33, 2002). One of the key new features of the 2003 framework was the identification of specific markets in which national regulatory authorities (NRAs) could consider imposing ex ante regulations on communications service providers found to have "significant market power" (SMP), defined as a position of economic strength in which a provider could behave independently of competitors, customers, and consumers. The remedies available to NRAs in cases of SMP included transparency, nondiscriminatory provision of service (in the traditional common carriage sense of providing like treatment for like service while allowing for differential service tiers), and mandatory access to network facilities, among other obligations.

The bulk of the framework was transposed into UK law with the adoption of the Communications Act 2003, legislation that also created the Office of Communications (Ofcom) as the new converged communications regulator. One of Ofcom's first and largest tasks upon its founding was a complete review of existing telecommunications regulation, known as the Telecommunications Strategic Review (TSR). In keeping with broader European policy efforts, one of the key goals of the TSR (and of Ofcom more broadly) was to promote competition in telecommunications services. As Ofcom noted in launching its first TSR consultation, "despite nearly 20 years of regulatory activity intended to promote competition" (Ofcom 2004b, 2), BT remained in a position of significant market power in a number of telecommunications markets, including wholesale broadband markets (Ofcom 2004c).

The undertakings that BT agreed to at the conclusion of the TSR helped to change that. Local loop unbundling (LLU) had been possible in the UK (and required by EU regulation (OJL 336/4, 2000)) since 2000, but only in 2005 did it become a widely viable possibility for competitive operators. As part of the undertakings, BT agreed to functionally separate its access network division (which became Openreach) from the rest of the company and to provide equivalent wholesale prices, terms, and service guarantees to all ISPs, including its own ISP, BT Retail (Ofcom 2005a). Where the US had forsaken intra-platform competition, the UK embraced it, at least as far as the telephone network was concerned. Cable companies, viewed as key challengers to the incumbent BT, were never required to unbundle their networks.

Ofcom's move to functionally separate BT was among the most aggressive steps that any regulator in the world had taken to stimulate competition for broadband access. It worked. For years, Britain has enjoyed one of the most competitive broadband markets in Europe. In 2010, more than 70% of households were served by at least four wholesale broadband providers and many were served by dozens of retail providers (Ofcom 2010e).

Within the context of this intense competition, discriminatory traffic management has been pervasive among most of the UK's largest fixed broadband providers and has become more prevalent over time (see Chapter 6). By 2010, five of the six largest providers were putting limits specifically on peer-to-peer traffic at peak network usages times, and a number of providers had targeted other kinds of applications for traffic management.

Importing Net Neutrality

With its resources dedicated to policing dominant providers and stimulating competition between wholesale and retail network operators, net neutrality concerns had little visibility within Ofcom (and most other European regulators) during the mid-2000s. Only after the public debate about net neutrality in the US became highly charged did European policymakers begin to consider whether similar concerns existed in Europe and how they might be dealt with under the existing regulatory framework. After much discussion during the review of the EU telecommunications framework that began in 2006, a compromise was adopted in the revised framework in 2009, giving NRAs new powers to (1) require ISPs to be transparent about their traffic management practices and (2) set minimum quality of service requirements on network operators to prevent degradation in service quality (OJL 337/11, 2009). Ex ante prohibitions on discriminatory conduct in the absence of significant market power were not included.

To assuage the Parliament's fears that the framework did not go far enough, the European Commission also adopted a declaration affirming the importance of "preserving the open and neutral character of the Internet" and declaring its intent to monitor and report back to the European Parliament and Council about the implementation of these new provisions and their interaction with the "net freedoms" of European citizens (OJC 308/2, 2009). That process of monitoring and investigation, led by the newly restructured Body of European Regulators for Electronic Communications (BEREC), began in 2010 and was ongoing at the conclusion of this study.

In the meantime, the UK government had published its digital strategy document, *Digital Britain* (BIS and DCMS 2009), in which it endorsed the ability for ISPs to charge content providers for guaranteed levels of service quality. Ofcom had likewise remained skeptical during the framework review about the need for explicit net neutrality regulation. In public statements made throughout the mid-to-late 2000s, Ofcom officials emphasized their existing authority to deal with complaints associated with any provider's abuse of market power, but declined to acknowledge the potential harms of discriminatory traffic management in a competitive environment. They acknowledged that discriminatory traffic management was taking place and accepted it as a performance necessity (see Chapter 6).

Once the review was completed and the transposition of the revised framework into UK law had begun in 2010, the regulator launched a consultation to develop its official approach to net neutrality and traffic management. Ofcom's initial and enduring view was that the combination of retail competition and transparency would suffice to deter any harmful discriminatory conduct, and that minimum quality of service requirements were therefore unnecessary. While the regulator pledged to monitor and report on traffic management practices going forward, the UK government (with new parties in power) took a more aggressive stance, actively coaxing the broadband industry to establish self-regulatory principles around transparency and nondiscrimination (DCMS 2011a). Both processes were ongoing at the conclusion of this study.

1.2.3 Summary

The contrasts between the regulatory circumstances and market outcomes in the US and the UK from the mid-2000s to 2011 are clear. The competitive UK market has seen discriminatory traffic management flourish, calling into question the argument that competition can safeguard nondiscrimination in the absence of regulation. The regulator has been reluctant to intervene. The US, by contrast, has been characterized by limited competition, significant regulatory threat, and limited discriminatory traffic management.

These outcomes are summarized in Table 1 and they form the basis of the inquiry in this thesis.

	UK	**US**
Competition	Significant	Limited
Traffic Management Regulation	No	Threat
Discriminatory Traffic Management	Significant	Limited

Table 1. Comparison of traffic management outcomes in the UK and the US.

The net neutrality debate has been fueled by the claim that nondiscrimination provides the foundation for Internet innovation, and that allowing discriminatory traffic management gives network operators too much power over which applications and services will succeed or fail. The debate raises a heretofore unanswered empirical question: what causes broadband networks to be operated in a discriminatory or nondiscriminatory fashion? This thesis provides answers to that question, drawing from the divergent experiences of the US and the UK. The next two chapters present the detailed research questions and hypotheses used to structure the inquiry and explore how insights from the existing literature concerning net neutrality and regulatory theory apply.

Chapter 2.
Net Neutrality Literature Review

2.1 Introduction

Academic interest in the topic of discrimination on broadband networks was triggered by US regulatory events, beginning with calls from advocates, independent ISPs, and academics for cable open access in the late 1990s and early 2000s. The earliest work was legal scholarship (Lessig 2001; Speta 2000a; Wu 2003; Yoo 2004), often grounded in the law and economics tradition. Although the topic has since been addressed from a wider variety of disciplinary perspectives, the majority of scholarly attention to net neutrality has been generated by legal scholars evaluating economic aspects of net neutrality and economists evaluating the economic implications of discrimination and its potential regulation.

The bulk of this literature operates, often implicitly, within the confines of the traditional "public interest" theory of economic regulation. This theory emerged more than a century ago when regulatory agencies and commissions, created by legislatures and governments but independent of them, were first being established to oversee nationwide transportation and communications networks. The public interest paradigm assumes that regulation arises in order to protect consumers from the impact of market failures, including monopoly, information asymmetry, and externalities, among others (Baldwin, Cave, and Lodge 2012; Carpenter 2010; Posner 1974; Trebing 1981; Viscusi, Harrington, and Vernon 2005). Regulators are assumed to be benevolent, apolitical servants of the collective good or national welfare (Lowi 1979), operating with objective, self-evident standards for determining how best to meet the public's needs (Baldwin, Cave, and Lodge 2012; Stewart 1975). Strict adherents to this view have presumed that "the public interest in regulation would be identified automatically as the residue of the struggle among the conflicting demands of rival private parties" (Bernstein 1955, 126).

The notion of a regulator intervening to protect the public from market ills holds tremendous power as an ideal for how regulation should function in the economy. As a result, the concept of service to the public interest is ritually invoked in the rhetoric surrounding regulatory agencies (Wilson 1980), including in the statutory language that directs telecommunications regulators. The Telecommunications Act of 1996 created a requirement for the FCC to use its regulatory powers "in a manner consistent with the public interest, convenience, and necessity" (47 U.S.C. 1302(a)). Similarly, Ofcom's "principal duty" as described in the Communications Act 2003 is "(a) to further the interests of citizens in relation to communications matters; and (b) to further the interests of consumers in relevant markets, where appropriate by promoting competition" (c. 21, s. 3(1)).

As an explanatory theory of regulatory intervention, the public interest theory suffers from a number of well understood shortcomings (discussed in the next chapter). Above all, it lacks predictive power; it is more of a normative paradigm for how regulation should function than an explanation for why regulation arises (Majone 1994; Trebing 1981).

Nonetheless, the concepts and tensions inherent in public-interest-oriented views of regulation have persisted in policy and academic discourse, including discourse concerning net neutrality. The key debates in the net neutrality literature surround whether ISPs have economic incentives to discriminate and under which competitive circumstances, whether consumers and the economy might benefit or suffer as a result, whether regulatory intervention is warranted, and what form it should take. The existence and effects of market power (and, to a lesser extent, externality problems and information asymmetries) are central points of contention.

Given its intellectual grounding, the net neutrality literature offers predictions about ISP behavior derived from economic theory and observations of technological constraints. It treats legal, political, and social aspects of the issue normatively, or not at all. In short, many scholars have offered recommendations about what telecommunications regulators should do, but few have attempted to explain why they do what they do.

In economics and political science, scholarly reaction to the public interest theory over the last half-century has resulted in a shift toward positive theories that may be used to explain and predict what causes regulatory intervention in the economy and its associated market effects. Making a positive contribution to the net neutrality literature, as is the goal of this thesis, requires drawing from this theoretical work and identifying how it applies in light of the observed behavior of network operators and regulatory agencies.

One of the key paradigms that has shaped positive regulatory theories in recent decades is institutionalism. Although institutionalism itself comprises a broad spectrum of thought from across economics, political science, sociology, and organization theory, a common tenet across institutionalist traditions is that "institutions matter" in shaping human behavior and decision-making (Baldwin, Cave, and Lodge 2012; Black 1997). As defined in the pioneering work of Douglass North (1990), institutions consist of formal and informal constraints on behavior. Formal constraints include laws, regulations, and administrative mechanisms, while informal constraints comprise a potentially much broader category of norms and conventions that may be socially or culturally constructed and serve to enforce formal constraints (North 1990). Institutional perspectives are often cast as responses to models of human and organizational behavior based on traditional rational actor assumptions from economics.

Implicitly or explicitly, students of regulatory activity and theorists of regulatory behavior from across the social sciences have explored how a variety of formal and informal institutions influence regulatory and market outcomes. These approaches can be broadly classified according to the aspects of the regulatory environment in which they vest explanations of regulatory decisions: institutional design, external forces, internal characteristics, and nation-specific factors. Institutional design choices are in some ways the most observable aspects of regulatory space: agency governance structure, jurisdiction, independence from other branches of government, and procedural rules concerning transparency and public participation. The most significant body of work emphasizes the relationship between regulatory agencies and their external environments: how interest groups (including regulated industries), legislatures, courts, executive branch officials and departments, and the media affect agency decision-making. In contrast to theories that treat agencies as black boxes, internally focused scholarship takes the internal characteristics of agencies and those they employ as its central focus. Finally, some scholars have drawn attention to how national traditions, culture, and the diffusion of ideas create informal contours that shape regulatory and market outcomes.

The design of the research inquiry in this thesis reflects the narrowly conceived notion of regulation in the public interest (and against abuse of market power in particular) found in the net neutrality literature; the positive regulatory theory concept of regulation as a product of a multifaceted institutional environment; and observations of existing broadband conditions in the two countries of study. It asks not only why broadband firms have made particular choices in particular competitive environments, but also how institutional constraints have shaped those choices and the choices of regulators. Specifically, the thesis is guided by two research questions:

> Question 1: Why do network operators take up discriminatory traffic management (or not)? In particular, how does competition in the market for broadband service influence network operators' traffic management decisions?
>
> Question 2: How does the institutional setting – the formal and informal constraints that comprise the regulatory environment – influence traffic management outcomes?

This chapter explores what existing scholarship concerning net neutrality has to offer in seeking answers to these questions and proposes a research hypothesis based on a synthesis of the literature and observed characteristics of the US and UK broadband environments. The next chapter does the same for the regulatory theory literature.

2.2 Map of the Literature

In addition to the substantial body of normative legal perspectives on net neutrality referenced above, a small but growing number of studies have responded to calls to apply formal economic modeling to net neutrality questions (for example, Chen and Nalebuff (2006), Economides and Tåg (2012), and Hermalin and Katz (2007)). While these types of studies can provide novel contributions, it can be difficult to rely on theoretical models to draw conclusions about a space as dynamic and complex as Internet service: results are often ambiguous, and model parameters are often unrealistic (Chirico, Haar, and Larouche 2007). Indeed, as Faulhaber (2011, 10) notes in surveying the scholarship in this area, "the theoretical literature contains papers to support whichever position you may favor, and it is unlikely that we may expect a definitive answer from this research." These limitations have led to calls for empirical analysis to help assess existing normative and theoretical claims (Brito et al. 2010; Candeub 2012; Faulhaber 2011; Schuett 2010; van Schewick 2010).

Empirical work has been minimal, however. From a qualitative perspective, a few scholars have undertaken surveys of broadband providers' published terms of service and policies to analyze the restrictions that they levy on their Internet service offerings (Li and Losey 2009; Sidak 2006; Wu 2003; Wu 2007). Researchers in computer science and other technical fields have conducted limited investigations as to whether discriminatory treatment is occurring (Beverly, Bauer, and Berger 2007; Dischinger et al. 2008; Dischinger et al. 2010; Kreibich et al. 2010; Tariq et al. 2009; Weinsberg, Soule, and Massoulie 2011; Zhang, Mao, and Zhang 2009). Mueller and Asghari (2011) and Asghari, Van Eeten, and Mueller (2012) used data collected by Dischinger et al. (2010) to show correlations between regulatory events or other external factors and levels of detected differential traffic treatment, representing the lone body of work that links political or regulatory activity to empirical data. Researchers in technical fields have also offered a handful of descriptive articles that reflect on policy debates and normative claims from a technical perspective (Crowcroft 2007; Felten 2006; Jordan 2009; Jordan and Ghosh 2010; Odlyzko 2009; Peha 2007).

The literature review below reflects the distribution of methodological orientations for studying net neutrality – predominantly normative, with reference to theoretical or empirical work as appropriate. Because the public statements of regulators and network operators are one key basis of the analysis later in the thesis, this chapter focuses on academic sources. Much of the literature is US-centric, although the European literature grew in the wake of the EU telecommunications framework review. Contributions from elsewhere are noted as appropriate.

Because discriminatory traffic management is just one of several practices that have been the subject of policy debates, there is very little academic work that focuses exclusively on traffic management. However, many net neutrality arguments generalize to discriminatory traffic treatment of all kinds, including traffic management.

2.3 Technical and Economic Arguments

When it comes to explaining or predicting the circumstances under which discriminatory traffic management arises, the net neutrality literature has primarily focused on operators' economic motivations and incentives to discriminate under differing competitive conditions. The bulk of this section is therefore dedicated to the economic literature, prefaced by a discussion of technical arguments.

2.3.1 Technical Arguments

The earliest proponents of regulatory intervention to safeguard nondiscrimination were legal scholars who based their arguments on technical principles articulated about the Internet's early architecture (see, for example, Lemley and Lessig (2001), Wu (2003), and van Schewick (2007)). In this view, the technical design of the Internet, based on the "end-to-end" arguments that inspired its early architecture, provided a nondiscriminatory foundation that allowed application innovation to flourish. As conceptualized by Saltzer, Clark and Reed (1984), end-to-end design argues for implementing specialized functionality at the Internet's endpoints unless it can be completely and correctly implemented within the network. Based on this and later conceptions of the end-to-end arguments (van Schewick 2010), legal scholars have asserted that the Internet's nondiscriminatory character was baked into its technical architecture, since ISPs' efforts to treat traffic differentially in the middle of the network would violate end-to-end design (Frischmann and van Schewick 2007; Lessig 2001; Lemley and Lessig 2001; van Schewick 2010; Wu 2003).

Scholars in this camp argue that the Internet's end-to-end design allows parties operating at the edge of the network to introduce their innovations to large audiences with great speed and low barriers to entry, inducing survival-of-the-fittest competition that determines which innovations succeed or fail based purely on their own merits (Lemley and Lessig 2001; Lessig 2006; Werbach 2005; Wu 2003; Wu and Lessig 2003). Nondiscriminatory connectivity allows for "widespread experimentation by a large and diverse group of innovators who independently select whether to realize their innovative ideas" (van Schewick 2010, 351), thereby increasing innovation overall (Wu 2004). Discrimination, it is argued, up-ends this innovative model by turning network operators into "gatekeepers" (Lennett 2009, 100) with ultimate control over whether new Internet applications and services succeed or fail. As Lemley and Lessig (2001, 938) explain, "[a]n architecture that creates powerful strategic actors with control over the network . . . threatens innovation." A number of formal models have demonstrated some of these effects, showing how an ISP prioritization regime may boost revenue and growth for established content providers at the expense of upstarts on the fringe (Bourreau, Kourandi, and Valletti 2013; Guo, Cheng, and Bandyopadhyay 2012; Reggiani and Valletti 2012).

A contrary line of argument emphasizes the benefits of discrimination and discriminatory traffic management in particular. A number of scholars argue that discrimination improves performance by giving ISPs the ability to prioritize latency-sensitive traffic (Brito et al. 2010; Hahn, Litan, and Singer 2007; Owen and Rosston 2006; Peha 2007; Sandvig 2007), thereby making broadband products more valuable to consumers. Others make the further claim that discriminatory traffic management not only enhances performance, but is essential to ensure the continued smooth operation of broadband networks (Crocioni 2011; Hazlett and Wright 2011; Prüfer and Jahn 2007; Renda 2008; Singer 2007).

From a performance perspective, prioritization schemes are only useful when networks experience congestion. While those who are skeptical of discrimination claim that congestion is best dealt with by expanding capacity (Lennett 2009; Lessig 2006; Odlyzko 2009), Faulhaber and Farber (2010, 324) argue that discrimination is a far more reasonable approach to dealing with instantaneous congestion, since "[a]dding capacity to a network takes time, while congestion must be dealt with immediately." Some argue that discrimination may be the most cost-effective way – or the only technically feasible way – of managing congestion on the network (Speta 2002b; Yoo 2005). However, as Peha (2007, 8) notes, whether discrimination or capacity expansion is a better strategy for controlling congestion depends in part on whether "processing or communications gets cheaper at a faster rate," since discrimination requires processing while capacity expansion requires bandwidth.

Changing network usage patterns that create performance problems may be one key driver for discriminatory traffic management (Clarke 2009; Hahn, Litan, and Singer 2007; Marsden 2010; Shelanski 2007; van Schewick 2010). As Internet users employ new and different applications in differing quantities and at different times of day, network operators may feel pressure to respond by managing the particular applications responsible for growing traffic volumes. Changes in usage patterns may violate the assumptions that formed the basis of original network designs – that traffic would be asymmetric and connectivity intermittent (Bauer, Clark, and Lehr 2009; van Schewick 2010) – or they may vastly outpace even the most aggressive capacity planning schedules (Clarke 2009). In their study using measurements of traffic differentiation in 75 countries, Asghari, Van Eeten, and Mueller (2012) found a positive correlation between country-level bandwidth scarcity and differential treatment of BitTorrent traffic.

Contrary to arguments about performance benefits, some scholars claim that discriminatory treatment of traffic is an inefficient way of managing network resources (Frischmann and van Schewick 2007; Lennett 2009; Marsden 2010). Discriminating against bandwidth-intensive applications does not give the developers of those applications any incentive to improve their bandwidth efficiency since the application traffic gets restricted nonetheless (Cooper, Jacquet, and Soppera 2011). In fact, it may create further inefficiency by inducing an “arms race” between users and application developers on one side and ISPs on the other (Lehr et al. 2006; Sandvig 2007). As users and developers continually seek ways to evade the traffic detection and resulting discrimination, network operators will need to continually evolve their techniques in response, resulting in significant use of engineering resources on both sides that could be better spent (Lehr et al. 2006; Marsden 2010). Even if their own user experiences would improve, individual Internet users may oppose discriminatory schemes if they are perceived as unfair (Krämer and Wiewiorra 2013; Wiewiorra 2012).

The theoretical economics literature provides models demonstrating how discrimination can lead to bandwidth inefficiency. Economides and Hermalin (2012) create a model in which discrimination increases demand for the services that receive higher priority treatment on the network, causing the portion of the network dedicated to those services to “re-congest,” thereby reducing overall welfare. Krämer and Wiewiorra (2012) and Wiewiorra (2012) find a similar result in the case where the volume of congestion-sensitive services is significant, although they prefer to allow discrimination because of its potential to stimulate network investment (as discussed in the next section).

Several scholars differentiate between what they view as beneficial and harmful discriminatory practices. Authors in this camp see little value in deliberate degradation of certain traffic or content but argue that offering quality-of-service guarantees or prioritization to specific applications could improve application offerings for consumers (Frieden 2006; Jordan 2009; Jordan and Ghosh 2010; Marsden 2010). Peha (2007) additionally considers that discrimination against “unfriendly” applications – those that do not reduce their sending rate even when they encounter congestion – is beneficial to users of other applications that would otherwise be starved for network resources in the presence of unfriendly traffic. But he agrees that introducing artificial degradation or scarcity in the network can be anti-consumer to the extent that it alters the application landscape or increases customer fees. All of these arguments are aimed at highlighting uses of discrimination that can improve application performance without negative consequences for application providers or consumers. In some instances the same discriminatory techniques may be viewed positively or negatively depending on the purpose for which they are used (Mueller 2007; Peha 2007; Sandvig 2007).

Lastly, while they may not relate purely to technical motivations, commercial and governmental pressures related to the policing of illegal or unwanted content may justify the installation of the same equipment that network operators can then use for discriminatory traffic management. As Marsden (2010, 81) argues, "Once you have the equipment in place, abandoning net neutrality becomes a no-brainer. The motive is simple: do it because you can."

2.3.2 Economic Rationales For Discrimination

Technical arguments must ultimately be secondary in profit-seeking ISPs' decisions to discriminate. They must justify their choices economically, and there has been no shortage of debate about potential justifications in the literature. Within the realm of traffic management, arguments about the economic benefits of discrimination for ISPs tend to focus on ways in which discrimination can make broadband service more valuable for subscribers, which in turn allows ISPs to expand their revenues derived from subscription fees. Scholars posit that discrimination can help to create diverse consumer broadband product offerings that may better suit the needs of broadband users overall than the more limited set of Internet service offerings that would be available under a nondiscrimination regime (Litan and Singer 2007; Marcus 2008; Renda 2008; Yoo 2004; Yoo 2005). These arguments rest on the fundamental economic observation that differential pricing has the potential to be welfare-enhancing:

because different users value different applications, network operators should be able to tailor their service offerings to match these varied preferences (Weisman 2010; Yoo 2004). Some observers point to existing examples of products that are differentiated in this way (Litan and Singer 2007; Renda 2008; Valcke et al. 2009), including the UK ISP Plusnet's service offerings that are differentiated based on application-specific prioritization, as evidence that demand exists for diversified product offerings. As Faulhaber and Farber (2010, 331) conclude, "[i]t is for customers to decide how much network neutrality . . . they want."

Discrimination may also reduce costs. Since network operators tend to pay interconnecting backhaul and transit providers based on the amount of traffic exchanged, rising usage levels can yield increased operational costs, creating an incentive for operators to clamp down on high-volume applications (and to cache content locally) (Marsden 2010; van Schewick 2010).

From a welfare perspective, scholars argue that discriminatory management of network traffic creates the potential for allocating bandwidth optimally to the applications that need it most, maximizing welfare for all users of the network (Cave and Crocioni 2007; Hahn, Litan, and Singer 2007; Hermalin and Katz 2007; Krämer and Wiewiorra 2012; Sidak 2006; Sidak and Teece 2010; Weisman and Kulick 2010). Discriminatory traffic management may be more efficient than conducting fine-grained per-user traffic management or usage-based pricing because the transaction costs involved in the fine-grained schemes may outweigh any benefits derived from them (Crocioni 2011; Hahn, Litan, and Singer 2007; Yoo 2006). With the popularity of flat-rate pricing plans, and the fact that operators may be constrained to setting a single price for Internet service offered across a large region (Economides 2008), operators may be more willing to manage applications in the network than to charge for increased usage (Levinson and Odlyzko 2008). In essence, bandwidth-intensive applications can serve as proxies that are easier to manage than bandwidth-intensive users (Shelanski 2007; Yoo 2006). Frischmann and van Schewick (2007) reject this claim, contending that the social costs of imprecise usage-based pricing are less than those of application-based discrimination.

Although it does not speak to traffic management as defined in this thesis, much of the net neutrality debate has centered on models where ISPs charge application providers for access to their customers or for prioritized delivery of their traffic, viewing the Internet as a two-sided market (Cañon 2009; Choi and Kim 2010; Economides and Tåg 2012; Musacchio, Schwartz, and Walrand 2009; Weisman 2010). These schemes obviously imply additional revenue for ISPs. In the theoretical economics literature, a variety of such models of two-sided pricing have been formally developed, with mixed results that show positive, negative, or ambiguous welfare effects on ISPs depending on the model design and parameters (Bourreau, Kourandi, and Valletti 2013; Cañon 2009; Cheng, Bandyopadhyay, and Guo 2011; Choi and Kim 2010; Economides and Hermalin 2012; Guo, Cheng, and Bandyopadhyay 2012; Guo, Cheng, and Bandyopadhyay 2013; Hermalin and Katz 2007; Jamison and Hauge 2008; Krämer and Wiewiorra 2012; Lee and Wu 2009; Musacchio, Schwartz, and Walrand 2009; Reggiani and Valletti 2012).

In assessing whether ISPs will actually take up discrimination, the central questions debated in the literature are (1) whether operators with market power have incentives to discriminate, and (2) whether competition in the Internet access market diminishes those incentives.

2.3.3 Incentives in Market Power Situations

The idea that a provider with market power may seek to leverage its dominance is familiar in telecommunications. A dominant provider of telecommunications facilities may have the ability and incentive to exploit its control of the network to foreclose competition in adjacent markets, reducing service options to consumers (Nuechterlein and Weiser 2005). In the context of the net neutrality debate, those adjacent markets consist of services that may directly compete with network operators' offerings (most commonly voice or video), or other complementary services (search, social networking, and so forth). By raising prices or reducing the quality of independent applications, a dominant network operator may be able to foreclose competition and choice in those markets while preserving or increasing monopoly prices in the market for Internet access (Economides 2008; Greenstein 2007; Knieps and Zenhausern 2008; Marcus 2008; Nuechterlein and Weiser 2005; Peha 2007).

This is the core concern of many of those who advocate for some form of regulatory intervention to safeguard net neutrality: that absent sufficient competition, Internet service providers will act on these incentives to discriminate (Atkinson and Weiser 2006; Crawford 2007; Economides 2008; Herman 2006). They point to historical examples of monopoly providers of telephone service that sought to exclude competitors from complementary device and data services markets, spurring regulatory intervention (Herman 2006; Lemley and Lessig 2001; Weiser 2008). Even those who oppose regulatory intervention in specific markets (on the basis that sufficient competition exists) acknowledge the legitimate possibility of exclusionary behavior by dominant operators (Cave and Crocioni 2007; Knieps and Zenhausern 2008).

The notion that market power creates incentives for network operators to discriminate is not universally held, however. Applying the "one monopoly rent" economic theory, under certain conditions a monopolist network operator could earn the same profits by raising the price of its Internet service as it could by monopolizing an adjacent market and charging monopoly prices in both (Nuechterlein and Weiser 2005; van Schewick 2010). Because there is only a single monopoly profit to be had, the operator would have no reason to try to exclude competitors from the adjacent market (Nuechterlein 2009; Yoo 2006). To the contrary, the operator would want to ensure that as many applications as possible are "cheaply, innovatively, and efficiently supplied" (Farrell and Weiser 2003, 101), so as to drive as many customers as possible to purchase Internet service at the monopoly price (Baumol et al. 2007; Marsden and Cave 2007; Speta 2000a). This process is known as "internalizing complementary efficiencies" of independent applications, or "ICE" (Farrell and Weiser 2003, 101). It provided the foundation for arguments against cable "open access" in the US (Speta 2000b) and has received some support in the theoretical economics literature concerning net neutrality (Chen and Nalebuff 2006). When ICE holds, the private interests of ISPs are aligned with the public's interest in maximum availability of independent applications on the network, and harmful discrimination is not expected to arise.

However, even those who acknowledge the effects of ICE have focused on identifying the conditions under which the logic of ICE breaks down and whether those conditions may be met in the context of Internet service. As a general matter, monopolist operators may be able to increase their profits without excluding rivals altogether from complementary markets, or by discriminating against independent applications without directly blocking or otherwise excluding them from the network (van Schewick 2010). Furthermore, there may be specific market circumstances under which ICE fails to hold. Of the many exceptions to ICE identified by Farrell and Weiser (2003) in the open access context and van Schewick (2010) in the net neutrality context, two are most relevant for understanding incentives to discriminate for traffic management purposes: "primary good not essential" and "incompetent incumbents" (a third, "option value," is discussed in a later section).

In cases where a network operator with market power in the Internet service market also offers complementary services that do not require the purchase of Internet service (that is, the "primary good," Internet service, is not essential), they may have incentives to discriminate against or exclude competitors in the complementary market if that market is subject to network effects (van Schewick 2007; van Schewick 2010). Consider the case of an ISP with market power that also offers a video streaming service nationwide to all Internet users, in competition with independent video streaming providers. Since consumers outside the ISP's service area can choose between the ISP's video service and independent video services, the operator cannot recoup all monopoly profits in the video streaming market merely by raising the price of its Internet service. Instead, it can discriminate against independent video streaming services on its own network, with the goal of reducing the nationwide user bases of the independent services. If the strength of the network effects and the size of the economies of scale associated with the video streaming business are large enough, the discriminating ISP may sufficiently reduce demand for the independent services to drive them out of business, “tipping” the video market in the ISP’s favor (Frischmann and van Schewick 2007; Herman 2006; van Schewick 2007).

In the “incompetent incumbents” case, network operators may simply fail to realize the benefits of ICE. Allowing innovative independent applications to flourish is not always obvious or intuitive to platform providers (Wu 2003; Wu 2004; Wu and Lessig 2003), particularly those that have historically provided monolithic products that bundle transmission and service layer capabilities, such as telephone services. As van Schewick (2010, 364) notes, “there is a danger that if network providers are allowed to optimize the network in favor of specific applications they will optimize the network in favor of uses that create observable value that they can appropriate over uses that create less observable and appropriable benefits.”

This second exception hints at a broader criticism of ICE – that it only addresses the benefits of nondiscrimination that dominant providers can appropriate for themselves, ignoring the effects of a large swath of benefits that the Internet produces that ISPs are not capable of capturing for themselves by reflecting them in the prices of the services they offer. There is broad agreement that the Internet is characterized by “spillovers” (or externalities, as that term is often used) – effects of economic activity that are not mediated through the price system and impact others than those engaged in the activity (Frischmann and Lemley 2006; Greenstein 2007; Hogendorn 2012). The question is how the size and nature of those spillovers influences ISPs’ decisions to discriminate and vice versa.

Hogendorn (2012) argues that the bulk of the Internet’s benefits can be attributed to spillovers, which are derived from three main sources. First, the Internet is a General Purpose Technology (GPT), a single generic technology in wide use across the economy that provides the foundation for many further activities and the creation of new markets. Importantly, many of these activities are noncommercial, socially valuable productive uses – maintaining personal connections, educating, debating, and so on (Crawford 2008; Frischmann 2005;

Frischmann and van Schewick 2007; Ganley and Allgrove 2006; van Schewick 2010). Because GPTs are so widespread, their benefits are not very appropriable by anyone, and that is particularly true in the case of non-market productive activities.

Second, the Internet supports a diverse array of services that exhibit network effects, although the magnitude of these effects is unknown. Many of these are direct effects – the more users that join the network, the more valuable it becomes for all. As a consequence, an ISP that throttles BitTorrent at peak times, for example, may end up discouraging users from using BitTorrent altogether, reducing the number of peers available on the BitTorrent network to exchange content with all other users. These effects can also be indirect, as in the video streaming service example above: with fewer users, a video service provider may decide to offer less content or may be less able to acquire content licenses, reducing the benefits of the service for everyone. Where an ISP sees a compelling reason to discriminate against an application (throttling BitTorrent for congestion management purposes, for example) it may be further encouraged to do so if it cannot appropriate (or even estimate) the gains from supporting application-based networks at their optimal size (Brennan 2012).

Finally, the Internet is an innovation-spawning technology that not only serves as a platform for new functions and capabilities to be developed in the future, but also increases current productivity and serves as an input to non-Internet based innovations. Many commenters have expressed the intuitive notion that Internet innovation is extensive and beneficial to society (Lemley and Lessig 2001; Lennett 2009; Lessig 2006; van Schewick 2007; Werbach 2005; Whitt 2004; Wu 2003; Wu 2004; Wu and Lessig 2003). Hogendorn (2012) and van Schewick (2010) draw on empirical and theoretical economics research to show that innovation spillovers can be both large (compared to the direct gains accrued to innovators) and not easily appropriable, or even perceivable, by the innovators themselves or others.

Thus, in the view of Hogendorn, Frischmann, and van Schewick, spillovers associated with the Internet are extensive and not easily appropriable by network operators, opening the door for ISPs to make the suboptimal economic choice to discriminate. In response to these arguments, some observers have claimed that empirical evidence of spillovers and their effects specifically in the Internet context is thin (Becker, Carlton, and Sider 2010; Brito et al. 2010; Candeub 2012; Faulhaber 2011) and that proponents of the spillovers arguments fail to balance their analysis by evaluating other kinds of externalities and costs, including negative externalities that arise on the Internet in the form of congestion and positive externalities that may accrue from discriminatory business models (Rosston and Topper 2009; Shelanski 2007; Sidak and Teece 2010; Weisman and Kulick 2010). In essence, these authors argue that the magnitudes of the costs and benefits involved are "industry and company specific" (Rosston and Topper 2009, 10) and since no studies exist that attempt to quantify them for existing Internet service markets, the extent to which ISPs are incentivized to discriminate is unknown. Unsurprisingly, many of these authors also disagree about the extent to which competition affects incentives to discriminate, as discussed in the next section.

2.3.4 Effects of Competition on Incentives to Discriminate

Given historic concerns about market power in telecommunications and the arguable proposition that market power provides incentives for ISPs to discriminate, much of the net neutrality debate in the literature and in the policy sphere has turned on whether competition can deter ISPs from discriminating.

Scholars attribute a range of effects to competition, from limited anti-foreclosure effects to effectively complete prevention of discrimination. Competition among ISPs is said to prevent independent application providers from being foreclosed from the applications market, since such providers "will have access to the majority of broadband customers even in the unlikely event that any one network operator decides to block access" to an application (Sidak 2006, 472). In this view, also put forth by Litan and Singer (2007) and Yoo (2004), competition among ISPs promotes competition among applications but does not entirely deter discriminatory conduct. For many proponents of net neutrality regulation (for example, those who espouse the “tipping” argument explained above), lack of foreclosure is not a sufficiently nondiscriminatory outcome because it still allows ISPs to dampen application innovation.

A more broadly held view is that competition generally reduces ISPs’ incentives to discriminate because discrimination degrades the quality of consumers’ Internet connections, causing them to change broadband providers (Becker, Carlton, and Sider 2010; Cave and Crocioni 2007; Chirico, Haar, and Larouche 2007; Faulhaber and Farber 2010; Hahn, Litan, and Singer 2007; Nuechterlein 2009; Shelanski 2007). As Becker et al. (2010, 502) explain, “[c]ompetition among broadband access providers . . . enables consumers to switch providers if they are not satisfied with the service from their current provider. . . . As a result of this competition, attempts by a broadband access provider to limit access to Internet content would likely result in the loss of subscribers that prefer unrestricted access, which, in turn, provides a competitive constraint that limits incentives for such actions.” Notably, the telecommunications regulatory framework in the EU is built on this fundamental premise: competition exerts pressure on network operators because consumers switch if they are unhappy with their service.

Belief in the disciplining power of competition is by no means limited to those who oppose regulatory intervention to safeguard net neutrality. In fact, some proponents of net neutrality regulation argue that competition not only reduces incentives to discriminate, but practically eliminates the threat of discrimination (Crawford 2007; Herman 2006; Jordan 2007; Lemley and Lessig 2001; Lessig 2001). Indeed, the very premise of open access regulation in the US – the intellectual precursor to net neutrality – was that competition among rival ISPs would ensure nondiscriminatory access to content and applications (Bar et al. 2000; Cooper 2003; Lemley and Lessig 2001).

It is perhaps easier for supporters of net neutrality regulation to envision that competition would unequivocally deter discrimination since they also tend to argue that the US broadband market is not sufficiently competitive to achieve this result. The competitiveness of particular broadband markets is one topic where many scholars cite empirical data about market share and calculate standard measures of market concentration (usually the Herfindahl-Hirschman Index), although interpretations of these measures are contested. Many observers have argued that the US lacks sufficient competition to deter discrimination, based on FCC and other data that shows that most Americans have had two or fewer choices for broadband (Crawford 2011; Economides 2008; Herman 2006; Jordan 2007; Lennett 2009; Lessig 2006; Odlyzko 2009). Others have claimed that fierce rivalry between cable and telephone companies, combined with budding new entrants, is more than sufficient to provide a check on market-power-based leveraging (Becker, Carlton, and Sider 2010; Brito et al. 2010; Faulhaber and Farber 2010; Hazlett and Weisman 2009; Sidak 2006; Sidak and Teece 2010; Yoo 2006), or that the appropriate market to analyze is the competitive national ISP market that application providers face, as opposed to the more concentrated local broadband markets that consumers face (Singer 2007; Yoo 2004; Yoo 2005). Some scholars emphasize differences between US and European market structures (Atkinson and Weiser 2006; Marcus 2008; Sluijs 2009; Wallsten and Hausladen 2009), noting that European regulators, as part of their required market review duties, have thus far failed to find retail ISPs that possess significant market power (Cave and Crocioni 2007; Cave et al. 2009; Chirico, Haar, and Larouche 2007; Crocioni 2011).

An opposing strand of research argues that even with intense competition between network operators, competitive pressure is not an adequate tool for deterring discrimination. A number of flaws in the logic and the mechanics of ISP competition, related to both ISPs and consumers, have been identified.

With respect to ISPs, the fact that an operator is subjected to competition does not necessarily cause its motives to discriminate disappear. Technical incentives (discussed in the previous section) are unchanged in the competitive case (van Schewick 2010). For example, competitive operators may be motivated to discriminate to manage network performance or to differentiate their consumer products (Wu 2003). Marsden (2007; 2008) argues that in markets where competition predominantly takes place between an incumbent and resellers of the incumbent's network, ISPs have little ability to differentiate themselves on features or price (since those are dictated by the incumbent's network offering to the resellers), so differentiation via discrimination may become widespread. Furthermore, the incumbent can still make discriminatory decisions that impact all ISPs that make use of the same network (as has happened in Canada (Mueller and Asghari 2011)), and its incentives to do so do not necessarily change by virtue of opening the network to competing providers (Marsden 2007; Wallsten and Hausladen 2009).

Hogendorn (2007) provides a theoretical model showing that under conditions similar to those in the broadband Internet context, opening access to the incumbent's network need not reduce discrimination, and may even increase it if competitive ISPs find discrimination to be the basis of profitable strategies in the marketplace. Wiewiorra (2012) likewise finds that a discriminatory regime is more profitable for competing ISPs, results in greater variety of content, and is more welfare-enhancing than a nondiscriminatory regime as long as there are enough content providers that would benefit from having their traffic prioritized. Bourreau, Kourandi, and Valletti's (2013) model similarly shows how two competing ISPs have independent and unilateral incentives to discriminate even under conditions where both ISPs would profit more from a nondiscrimination regime, creating a prisoner's dilemma.

There are a number of reasons why the traditional logic of competition may not work as predicted. If all broadband providers discriminate against the same application – a kind of "parallel exclusion" (Hemphill and Wu 2013) – consumers seeking a nondiscriminatory service offering will have no operator to switch to (Lennett 2009; Marsden 2007; van Schewick 2007; van Schewick 2010). Even if they do have choices, they may not know (or be able to find out) that operator discrimination is the cause of a performance problem they experience, or they may falsely attribute the problem to the application provider or the manufacturer of the device in use (Marsden 2007; van Schewick 2007; van Schewick 2010;

Wu 2003). The complexity of diagnosing performance problems on the Internet results in incomplete information.

Broadband service may also be characterized by an array of switching costs that deter consumers from switching even when they are bothered by discriminatory treatment of traffic (Bar et al. 2000; Economides 2008; Krafft and Salies 2008; Lennett 2009; Marsden 2010; van Schewick 2007; van Schewick 2010; Wu 2007). These may be direct costs, such as termination fees, installation fees, or the loss of discounts associated with particular broadband packages. The rise of bundling creates lock-in, where consumers may be less willing to give up a package of services due to dissatisfaction with a single service, especially if other providers' bundles are not directly substitutable for the consumer's current bundle (Prince and Greenstein 2013). Free services that ISPs offer to their customers – email, instant messaging, stock quotes – can create lock-in since switching providers would require establishing new accounts for these services. Even without bundled services, consumers may exhibit status quo bias, causing them to prefer the product they currently have over a product they could have even if they dislike the current product. Switching can also require a significant investment of time and effort to research and compare potential options, contact customer service, wait for an engineer to complete the new installation, and so forth. With complex or bundled products, consumers may decide that switching is not worth the effort.

Finally, for those who view discrimination not just as a potentially anti-competitive practice, but as a threat to application innovation generally, competition as effectuated via consumer preferences is insufficient (Lemley and Lessig 2001; Lessig 2001; van Schewick 2010; van Schewick 2012). Lessig notes that when operators close their networks to particular applications, that closure creates a cost to innovation that does not get fully reflected in consumers' broadband purchasing decisions:

> That cost is not borne directly by the consumer. In the long run, of course, if it is a cost, it is borne by the consumer. But in the short run, the consumer doesn't notice the innovation that the closed model chills. Thus the consumer does not completely internalize the costs imposed by a closed system. And hence the pressure the consumer puts on closed systems to open themselves up is not equal to the costs that such closed systems impose on innovation generally. (Lessig 2001, 162)

The argument that competition deters discrimination relies on the idea that consumers internalize the costs of discrimination and discipline ISPs by switching providers when those costs outweigh the switching costs. But as discussed above, consumers likely cannot appropriate – or even perceive – all of the social benefits that future innovations could provide. As a result, they will underestimate the benefits of switching (van Schewick 2012), failing to provide the discipline that competition is said to impose.

2.4 Normative Perspectives Concerning Regulation

A great deal of scholarship has focused on whether regulatory intervention into discriminatory behavior is warranted and what form it might take (Krämer, Wiewiorra, and Weinhardt 2012). This section presents a discussion of the arguments concerning prominent regulatory approaches suggested in the literature, many of which are not necessarily mutually exclusive: reliance on existing law, ex ante prohibitions on discrimination, transparency requirements, principles-based case-by-case enforcement, and regulatory threat.

2.4.1 Reliance on Existing Law

During the years before the FCC adopted the *Open Internet* order, some US commenters advocated for reliance on existing competition law to deter potentially harmful discrimination, rather than the introduction of any new regulation (Baumol et al. 2007; Becker, Carlton, and Sider 2010; Nuechterlein 2009; Rosston and Topper 2009). As the European net neutrality debate unfolded, scholars likewise argued that existing regulatory tools – general competition law, existing powers under the telecommunications regulatory framework for national regulators to impose ex ante obligations on providers found to have significant market power, and new powers to impose transparency and minimum quality of service obligations – were sufficient to deal with harmful discriminatory conduct (Cave et al. 2009; Cave and Crocioni 2007; Chirico, Haar, and Larouche 2007; Marsden and Cave 2007; Renda 2008; Valcke et al. 2009). These commenters argue that the kinds of anti-competitive conduct that have raised concerns in net neutrality debates are well within the purview of competition authorities (in the US) and regulatory agencies (in the EU).

As explained above, many scholars view discrimination as generally beneficial to consumers, the economy, and society. They argue that regulatory intervention to deter discrimination beyond what existing law provides would dampen broadband investment and competition because network operators would be constrained in their product offerings and potentially unable to charge application providers for priority treatment (Cheng, Bandyopadhyay, and Guo 2011; Jamison and Hauge 2008; Litan and Singer 2007; Yoo 2004; Yoo 2005). The network efficiency benefits associated with giving priority to applications that need it the most could be lost under overbroad regulation (Cave, Prosperetti, and Doyle 2006; Crocioni 2011; Hermalin and Katz 2007; Knieps and Zenhausern 2008; Litan and Singer 2007; Sidak and Teece 2010; Singer 2007; Yoo 2004). In this view, these costs of regulation cannot be justified when the marketplace has demonstrated little or no evidence of problematic discrimination, when the provision of broadband Internet service is still relatively new, and when its future direction is uncertain (Becker, Carlton, and Sider 2010; Cave et al. 2009; Cave and Crocioni 2007; Hazlett and Wright 2011; Owen 2011; Shelanski 2007; Sidak and Teece 2010; Weisman and Kulick 2010; Yoo 2006).

From an implementation perspective, relying on existing antitrust/SMP enforcement is said to have several advantages. Given how difficult it would be for ex ante rules to clearly separate beneficial discrimination from harmful discrimination, and how quickly technological and market developments could render them out of date, relying on case-by-case antitrust evaluation is viewed as a more sensible approach (Becker, Carlton, and Sider 2010; Owen and Rosston 2006; Speta 2011; Yoo 2004; Yoo 2005). Although regulatory proposals tend to target particular network conduct – blocking or applying differential treatment to certain content or applications – the history of telephone regulation provides evidence that enforcing such rules could lead to detailed price regulation, since providers can evade non-price-based conduct restrictions through discriminatory pricing (Owen 2007; Yoo 2005; Yoo 2006; Yoo 2007). Finally, antitrust authorities would provide "a referee inclined towards calm objectivity and a rigorous adherence to economic principle" (Nuechterlein 2009, 58), as opposed to a highly politicized FCC where both the enactment and the enforcement of rules is often subject to political gamesmanship (Faulhaber and Farber 2010; Faulhaber 2011; Marcus 2008; Owen 2007; Weiser 2008; Yoo 2004).

2.4.2 Ex Ante Prohibitions

At the opposite end of the spectrum are those that argue in favor of new ex ante rules that would prohibit some or all forms of discrimination. They argue that discrimination threatens innovation at the network edge, free expression, and human interaction (Crawford 2007; Crawford 2008; Frischmann 2005; Frischmann and van Schewick 2007; Herman 2006; Jordan and Ghosh 2010; Lennett 2009; Lemley and Lessig 2001; Meinrath and Pickard 2008; van Schewick 2010; Wu 2003). Because discrimination need not be anti-competitive for it to interfere with either innovation or free expression, remedies based on competition law principles are viewed as inadequate for preventing the full range of harmful conduct (Crawford 2008; Economides 2008; Frischmann and van Schewick 2007; Hemphill 2008; Lessig 2006; van Schewick 2012). Furthermore, ex post enforcement would result in an arbitrary, fragmented landscape with different kinds of discriminatory conduct allowed or prevented depending on which cases get adjudicated (Economides 2008; Herman 2006), leaving network operators with flexibility to discriminate.

In Europe, commenters similarly posit that specific gaps in the European framework (both pre- and post-review) could also fail to prevent harmful discriminatory practices (Read 2012; Valcke et al. 2009). As noted previously, discrimination may be widespread among network operators without SMP (Marsden 2007). Until the framework review, independent application providers had limited standing to bring complaints to national regulators for resolution (Marsden 2010) (new provisions may have improved this situation (Marsden 2013), but expanded dispute resolution authority has yet to be thoroughly tested). The viability of regulatory intervention in situations where multiple broadband providers (rather than a single provider with SMP) are able to jointly leverage their market dominance for anti-competitive purposes has also been the subject of much debate (Cave and Crocioni 2007). All of these shortcomings raise the question of whether ISPs would be able to discriminate without sanction in the absence of further regulatory intervention in Europe.

As a general matter, proponents of regulation claim that the costs to innovation of relying on existing legal frameworks outweigh any potential costs of regulatory intervention. Taking a wait-and-see approach leaves independent application developers with far too much uncertainty about whether they will be discriminated against, dampening future prospects for innovation (Jordan and Ghosh 2010; Lennett 2009; Peha 2007; van Schewick 2007; van Schewick 2010; van Schewick 2012; Wu and Lessig 2003). Case-by-case adjudication proceeds too slowly to keep pace with application innovation (Economides 2008). While ex ante prohibitions on discrimination may prevent certain traffic management techniques that could increase bandwidth efficiency in the short term, the long term benefits of such rules for innovation and application diversity outweigh these costs (van Schewick 2010). Whether such rules would hamper network investment is debatable, but in any event regulators and legislators have other means (tax incentives or subsidies, for example) of stimulating investment (Atkinson and Weiser 2006; van Schewick 2010).

2.4.3 Transparency Requirements

Between the two extremes of leaving regulation as-is and adopting ex ante rules lie a number of middle ground approaches. A policy of requiring ISPs to publicly disclose the details of their traffic management practices, whether combined with additional regulation or not, has enjoyed widespread support. Transparency is viewed as a means to improve competition, because it aims to give consumers information they can use to make choices in the marketplace (Candeub and McCartney 2010; Chirico, Haar, and Larouche 2007; Crocioni 2011; Lennett 2009; Marsden 2007; Marsden 2010; Valcke et al. 2009; Wu 2007). In surveying Internet users, Wiewiorra (2012) found that more information about traffic management practices would discourage users from otherwise assuming that their ISPs are purposefully degrading performance when the network slows down in the normal course of Internet use. Sluijs et al. (2011) found experimentally that increased transparency (even if only offered to or understood by a subset of consumers) increases the quality of the broadband products that operators offer.

Even in the absence of competition, disclosure requirements can still fuel public pressure on network operators to change their practices (Atkinson and Weiser 2006). Weiser (2008) and Marsden (2010) have suggested that disclosures be further used to form the basis for regulatory agency monitoring or enforcement.

2.4.4 Principles-Based Case-by-Case Enforcement

Another middle ground approach would involve a regulatory agency adopting principles or guidelines concerning nondiscrimination and then enforcing them on a case-by-case basis. The generalized arguments for and against case-by-case antitrust enforcement apply here as well, but with important differences.

The virtue of a focused principles-based framework is that it offers a credible threat of enforcement – potentially even in cases of discrimination that would not be viewed as anti-competitive under existing law – while providing some flexibility for operators to experiment with discriminatory practices (Atkinson and Weiser 2006; Bauer, Clark, and Lehr 2009; Greenstein 2007; Sluijs 2009; Weiser 2003; Weiser 2008). The implications of this approach for ISPs depend on the specific principles adopted, but in the abstract the expectation would be that ISPs would avoid discriminatory practices in clear violation of the principles, while possibly pursuing other discriminatory approaches where enforcement appears less likely. As Lessig explained in describing the *Madison River* adjudication:

> [T]he most important action that this government has taken to preserve the Internet's end-to-end design was the decision by Chairman Michael Powell to commit the FCC to enforce what he referred to as the Internet's four "Internet Freedoms." . . . Those principles were relied upon by the FCC when it stopped DSL provider Madison River Communications from blocking Voice-over-IP services. That enforcement action sent a clear message to network providers that the Internet that they could offer must continue to respect the innovation-promoting design of end-to-end. (Lessig 2006, 1)

The FCC's reliance on the *Policy Statement* to bring its enforcement action against Comcast could also arguably be construed as a version of principles-based case-by-case enforcement (although the agency's authority to enforce the principles and the level of detail of the principles themselves have both been called into significant question). Using crowd-sourced network testing to detect differential treatment of BitTorrent traffic, Mueller and Asghari (2011) have shown that the issuance of the FCC's order in that case corresponded with nearly a 50% drop in tests showing differential traffic treatment on the Comcast network, and supplemental data shows a similar correlation in timing for other fixed ISPs in the US ("BitTorrent Manipulation in Selected Countries" 2012). The correspondence would appear to imply a deterrent effect.

Several scholars support a similar case-by-case approach to that described above, but would use co-regulation to establish and enforce principles or guidelines (or, as Weiser calls them, "cooperative norms") (Marsden 2010; Watal 2011; Weiser 2008). A co-regulatory strategy would potentially have different implications for ISPs since both the development of principles and their enforcement could be more dynamic and benefit from greater industry input than a strictly regulatory approach. Evidence from Japan appears to indicate that co-regulation can have a deterrent effect on discrimination. At the behest of the telecommunications regulator, Japanese ISPs collaboratively developed packet-shaping guidelines that expressed a preference against discriminatory traffic management without sufficient justification. After it was adopted, survey data collected from ISPs indicated a drop in application-specific throttling (Jitsuzumi 2011).

2.4.5 Regulatory Threat

Scholars generally stop short of recommending that an ongoing threat of regulatory intervention be maintained as an explicit strategy for deterring discriminatory conduct. As a regulator, continually emanating a sense that intervention may be imminent could be viewed as an abrogation of democratic process that undermines the agency's legitimacy and stifles beneficial market activity (Weiser 2009). Nonetheless, several commenters have recognized the impact that such circumstances can have on ISPs. Crocioni (2011, 4) noted that the "extent and form" of traffic management in Europe "may be influenced by . . . the potential threat of regulation." While the regulatory status of cable broadband was in flux in the early 2000s in the US, cable operators began removing application restrictions on their networks and making public commitments to operate in a nondiscriminatory manner (Wu 2003; Wu and Lessig 2003). In assessing the contentious US legislative debate about net neutrality in 2006, Felten (2006, 11) observed that "ISPs, knowing that discriminating now would make regulation seem more necessary, are on their best behavior," although regulators and legislators had not yet dug into "the difficult issues of line-drawing and enforcement." The implication of all of these observations is that once it is clear that regulators are paying attention but before they have taken any significant action, the possibility for regulation to be imposed suppresses ISPs' willingness to take up potentially controversial discriminatory practices.

Wu (2003; 2004) conceptualizes regulatory threat as a counterbalance to the "incompetent incumbents" exception to ICE. In this view, the threat of regulation serves an educational function: it inspires operators to consider whether discrimination really is in their best interests. It may also serve to balance out the positive reinforcements that ISPs receive about discriminatory practices from equipment vendors or financial analysts. The more prominent the regulator makes its interest in the subject known, the more likely that operators will give their traffic management decisions thoughtful consideration (Marsden 2007).

There is a further exception to ICE that may be characterized as regulatory threat, although Farrell and Weiser (2003) referred to it as “option value.” In the open access context in which Farrell and Weiser were writing, they noted that operators of closed platforms may be unwilling to open them to competitors for fear that later regulation would prevent them from returning to a closed model. In the traffic management context, this argument may be framed in terms of a network operator that is already engaged in discriminatory activity without having been restrained by prior regulation. Such an operator may be more likely to continue to perpetuate its behavior than an operator whose network has always been offered in a nondiscriminatory manner. In essence, it is more difficult for an operator to justify to regulators a switch from offering a nondiscriminatory network to a discriminatory one than vice versa. As Wu (2003, 155) notes, application-specific management “may become obsolete: adopted at a certain time for a certain reason that no long matters.” Operators engaged in discriminatory traffic management, even upon concluding that such management is no longer needed to efficiently run the network, may perpetuate the discrimination for fear of losing the option to re-introduce discrimination in the future.

2.5 Conclusion

The literature reveals a robust debate about the technical and economic motivations and incentives of ISPs to discriminate for traffic management purposes. Application-specific traffic management may be justified as efficient engineering, but its longer term implications for application design, innovation, and network engineering may argue against its use. Discrimination can be used to the advantage of ISPs, but whether they will engage in it or not may depend on their competitive circumstances, the existence of spillovers, network effects, and their ability to assess their own interests. The notion that competition disciplines operators from engaging in harmful discrimination is widely supported in academic scholarship and is central to the European telecommunications regulatory framework, but its detractors have identified a number of deficiencies in the logic and mechanics of competition that call this premise into question.

The normative debate surrounding regulatory intervention evokes familiar themes associated with the public interest theory of regulation. Those who identify market failures – or conceptualize economic and social concerns beyond the bounds of traditional market failure analysis – recommend regulatory intervention of various kinds. The broad academic consensus in favor of regulatory requirements to increase transparency suggests at least implicit recognition of existing information asymmetry and support for the notion that regulators can help rectify it. Beyond that, the literature reveals a wide divergence of opinions about the merits of regulatory intervention of different forms. Limited empirical, experimental, and anecdotal evidence suggests that case adjudications can send signals to ISPs about avoiding controversial discriminatory practices and that the threat of regulation tempers ISPs from dramatically changing course in how they manage traffic, but most regulatory proposals rely on normative arguments rather than empirical evidence.

Drawing from both the net neutrality literature and observations about market and regulatory circumstances in the US and the UK, the following hypothesis was developed at the beginning of this study in response to Research Question 1:

> Question 1: Why do network operators take up discriminatory traffic management (or not)? In particular, how does competition in the market for broadband service influence network operators' traffic management decisions?
>
> Hypothesis 1: The threat of regulation limiting how operators can manage traffic acts as an informal constraint on operator behavior. This threat is at least as important, if not more so, than the competitiveness of the market.

Simple observation of the facts about the competitiveness of the US and UK markets indicates that competition does not appear to be a deterrent to application-specific traffic management. The much clearer difference between the two regulatory environments lies in the threat of regulation. Where that threat has been minimal (in the UK), application-specific traffic management has arisen. Where the threat of regulation and the uncertainty of what it may require of operators has been high (in the US), operators have largely refrained from application-specific traffic management.

To understand potential reasons for these differences and the disparate traffic management outcomes that have resulted in the two countries, the next chapter explores what the regulatory theory literature from across political science, economics, law, and organizational studies has to say about why regulation comes about and in what forms.

Chapter 3.
Regulatory Theory

3.1 Introduction

For more than half a century, scholars have grappled with questions of "regulatory origin" (Fiorina 1982, 37): why regulatory agencies get established and why regulation gets imposed in particular sectors and in particular ways. The resulting body of work has diversified over time as government regulation of the economy has expanded, contracted, and changed. Establishment by government of specialist bodies to regulate aspects of the economy began during the late 19th century with the development of nationwide railways and communications markets. In the US, such "economic regulation," aimed at markets deemed susceptible to natural monopoly, continued to expand throughout the early-to-mid 20th century with the development of commercial aviation, trucking, and energy utilities. In Britain, the regulatory agency model was increasingly substituted by public ownership during the same period (Ogus 1994). The 1960s and 1970s saw a proliferation of regulatory agencies on both sides of the Atlantic that concerned themselves with cross-industry "social regulation" of health, safety, and the environment. Since then, scholars have been grappling with the tension between calls for deregulation – ridding the economy of formal regulatory oversight, and even of regulatory agencies themselves in some cases – and the expansion in modes of economic governance that go beyond traditional command-and-control regulation.

Experience with regulatory agencies across all of this time and many industry sectors has yielded a wide variety of theories of regulatory decision-making that have been tested using a diversity of methodological approaches, from historical analyses to econometric modeling to qualitative case studies. As noted in Chapter 2, support for the public interest theory waned substantially during the mid-20th century as scholars came to appreciate its multiple shortcomings. First, it assumes that the public interest is somehow uniform and easily identifiable, rather than vague, indeterminate, and comprised of conflicting views, as reality would suggest (Katzmann 1980; Noll 1985). Although legislators often charge regulators with vague public interest mandates (Lowi 1979; Stewart 1975), there is no single operationalized definition of the public interest that can be used to judge whether particular actions are serving the public interest or not (Mitnick 1980). The public interest theory also assumes that if an authoritative notion of the public interest existed in any given regulatory realm, regulators would have the tools at their disposal to identify it and act to protect it. In reality that is not necessarily the case (Mitnick 1980; Noll 1985).

The public interest theory does not account for how other interests affect regulatory outcomes. It assumes that the interests of affected groups and politicians – and even regulators' own self-interest – are exogenous to choices made about regulatory intervention. In reality, regulatory agencies are embedded within political systems where competing interests can drastically alter outcomes (Mitnick 1980; Stewart 1975; Wilson 1980). There must be some mechanism or advocate that allows the public interest to be identified among the field of interests presented to the regulator (Viscusi, Harrington, and Vernon 2005). The fact that regulators may be motivated by financial gain, jurisdictional desires to expand their turf, or other personal interests is ignored entirely (Baldwin, Cave, and Lodge 2012; Mitnick 1980).

Finally, there is ample evidence that refutes the theory. Regulation often arises in the absence of market failure and often does not arise despite market imperfections (Posner 1974). Regulation has been shown to have supported prices above cost in certain competitive industries and has had little effect on prices in certain monopolistic ones (Posner 1971; Stigler and Friedland 1962; Viscusi, Harrington, and Vernon 2005). A related problem is that the public interest theory does not account for any costs associated with regulation; it assumes that regulatory agencies are internally efficient and can costlessly deliver the best solution to the public (Noll 1985).

Realizing these drawbacks, scholars have looked to a wide variety of other sources of explanation of regulatory behavior. The field of regulatory theory is disciplinarily wide and diverse, with few thoroughly developed theories of behavior but many promising directions. It was for years dominated by US perspectives given the pervasiveness of regulation there, but has since garnered substantial attention in Europe. It reflects the familiar social scientific tension between theories grounded in traditional rational actor assumptions from economics and those that relax these assumptions or cast institutions and their effects as more central than utility-maximizing behavior. This section surveys the explanations of regulatory behavior found in an interdisciplinary cross-section of literature from political science, economics, law, and organization studies in the US, the UK, and the EU. The explanations are divided into four categories: institutional design choices, external forces, internal characteristics, and nation-specific factors. Where appropriate, perspectives are provided about how these explanations might be applied in the context of telecommunications or net neutrality.

3.2 Institutional Design

Perhaps the most fundamental and observable factors contributing to an agency's behavior lie in its institutional design. The kinds of institutional choices that define agency structures and processes are many and varied, including the scope of the agency's jurisdiction, the extent of agency independence from other branches of government, agency governance structure, and procedural rules concerning transparency and public participation in regulatory processes (Horn 1995). These choices are highly interrelated.

Regulators generally draw their authority from legislative delegations that determine the scope of the issues within their purview. The limits (or lack thereof) on this scope are important for understanding why an agency chooses to take up a particular regulatory agenda or not. One key consideration relates to whether the agency is delegated a policy-setting function, or whether it is designed purely as an implementer or enforcer of policies set by government (Ogus 1994); the latter characterization has a strong tradition in the UK but is less common elsewhere. Particular agencies may be delegated specific authority because they are presumed to behave in predictable ways – thus an environmental protection regulator and a governmental civil engineering department will presumably approach the same problems from quite different perspectives, and the delegating authority may prefer one over the other (Spence 1999). Mismatches between agency expertise and delegated tasks can create suboptimal regulatory outcomes (Baldwin and McCrudden 1987). Thus the choice of jurisdiction may have important implications for the agency's agenda, the alignment of its outcomes with legislative or executive preferences, and its overall effectiveness.

Although the independence of regulators is a subject of substantial academic attention, precisely defined ways of determining or evaluating independence are elusive. For example, scholars have suggested that formal independence from government is derived from employment conditions of key decision-makers (the standards and processes for their appointment and removal), agency jurisdiction (because legislative oversight is more easily accomplished for narrow, industry-specific regulators), location outside the executive branch structure, freedom from supervision within the executive hierarchy, the clarity of the division of powers between the regulator and the rest of government, and potentially other factors (Bernstein 1955; Gilardi 2002; Hanretty and Koop 2012; Horn 1995; Majone 1994; Prosser 2010). More independent agencies are assumed to have greater autonomy in setting their agendas and making their decisions. Even agencies that are not statutorily independent from government may have enough de facto independence to choose their own paths (Maggetti 2007). The model that is common for both the typical American "independent commission" (of which the FCC is one) and the British utilities regulators (of which Ofcom is one) tends to score high on most measures of formal independence (Ogus 1994).

Regulatory agencies come in many forms, from the multi-headed commission to the single-headed executive agency. The governance structure can have a profound impact on agency decision-making and is highly related to the question of independence. Commissions are said to offer more flexibility and independence than other structures, with commissioners serving on fixed and staggered terms that outlast the term of the President or Prime Minister;

appointments conducted on a bi-partisan basis; and removal of commissioners possible only on grounds of misbehavior (Horn 1995; Majone 1994). Commission action may be impaired in other ways, however: regulation by commission can result in incoherent compromises, horse-trading, mismanaged bureaucracies, and severe delays while trying to reach agreements (Katzmann 1980).

By contrast, the heads of executive agencies or ministerial departments may be more politically vulnerable since they can be removed based on the policy preferences of the President or minister in charge (Majone 1994). But by virtue of being able to act unilaterally within the agency they may be able to avoid problems of incoherence, mismanagement, and delay. Although many EU regulators were modeled on US regulatory commissions, Ofcom's governing structure – a board that includes executive staff (including the Chief Executive) and ministerial appointees – is a novel variant that potentially avoids the pitfalls associated with both the single-headed and commission models (Prosser 2010).

Finally, decisions about policies and procedures to allow for transparency and public participation in agency processes can have a profound impact on regulatory outcomes. Public participation is said to compel regulators to take societal interests into account (Stewart 1975) and to insulate them from political pressures since they need to be seen to be responding to the interests that participate (Croley 2008). However, the mere imposition of requirements that agencies publish their activities and consult with stakeholders does not automatically achieve these results, and may impact regulatory decisions in other ways. It may be that only well-organized or well-financed interests participate, that open processes reduce the impact of expert input, or that the resulting compromises fail to meet the needs of the public (Baldwin and McCrudden 1987; Ogus 1994; Rothstein 2004). This may argue for flexible rules concerning transparency and participation, along the lines of those established by Ofcom; in the absence of strict statutory requirements of the sort imposed on US regulators by the Administrative Procedures Act, the agency has established its own rules for consultation and openness (Lunt and Livingstone 2012; Prosser 2010), in line with a broader movement towards "better regulation" in Europe (COM(2001) 428 final; COM(2002) 275 final).

To an outside observer, the FCC and Ofcom appear very different when it comes to institutional design. A new FCC chair is nearly guaranteed to be appointed when the presidency changes parties; the agency's remit is comparatively broad (Shaffer and Jordan 2013); and it operates under strict requirements for transparency and public participation. Ofcom appears far less politicized, with its own discretion to rely on in determining the extensiveness of its consultations. The questions from the perspective of this study are whether those characterizations ring true and how they influence the decisions that the two regulators have made with respect to traffic management.

3.3 External Forces

Institutional design choices are deeply intertwined with the "mosaic of forces" (Horwitz 1989, 8) that external entities impose on regulatory agencies. Initial attempts to theorize about the influence of external forces on regulatory decisions tended to focus on a single external actor, with significant attention in the literature devoted to interest groups, legislatures, and courts, respectively. As those theories have matured, scholars have widened their focus to evaluate how a broader set of external forces combine with interest group pressures, legislative control, and judicial oversight to produce regulatory outcomes. This section examines the specific bodies of work on interest groups, legislatures, and the judiciary before reviewing broader theories of external influence on regulators.

3.3.1 Interest Groups

The earliest attempts to model the impact of interest groups on regulatory outcomes started by applying traditional economic assumptions to the regulatory process (Baldwin, Cave, and Lodge 2012; Black 1997). Under the "economic theory of regulation," all actors are assumed to be maximizing their utility (material wealth, or votes as a means to material wealth in the case of elected officials). All parties are expected to be well informed and to learn from past interactions. Regulation itself is assumed to be costless. With these assumptions in place, regulation is treated as a product whose allocation is governed by the same laws of supply and demand as market goods (Posner 1974).

The most prominent contributor to the economic theory was George Stigler. Stigler (1971) proposed that, given the market for regulation as modeled above, the greatest beneficiaries of regulation would be large, concentrated, well organized interests – that is, the industries subject to regulation. The consequence of this is regulatory "capture," in which "regulation is acquired by the industry and is designed and operated primarily for its benefit" (Stigler 1971, 3). Consumer interests might be present in the market for regulation, but they would likely be more diffuse and less well organized than the regulated industry itself, and therefore less able to obtain favorable regulation. A politically motivated regulator would maximize its gains from fulfilling industry's wishes.

The application of the capture theory in the net neutrality context would be complex. Articulated with reference to price regulation of natural monopolies, Stigler's theory focuses only on obtaining regulation, not remaining free from it (Wilson 1975). It does not easily accommodate the case where the potentially (or previously) regulated industry – Internet service providers – prefers not to be regulated at all, or where multiple industries have a web of relationships with the regulator (Prosser 1999). One could alternatively view Internet application providers as the industry capable of capturing the regulatory agency, but the regulation they seek to obtain would not apply to their own industry.

In an important predecessor to Stigler's work, Bernstein (1955) relates agency capture to the maturity of the regulatory agency. He depicts a life cycle of regulatory commissions in which their early existence is characterized by antagonism with the regulated industry and a crusading, pioneering spirit that is bolstered by strong political support from those who sought to create the agency in the first place. But as they mature, they find it increasingly difficult to extend regulation beyond the limits acceptable to the industry, settling into a role of protecting the industry and the status quo – the capture phase. This model could yield important insights for understanding net neutrality regulation globally, given differences in regulator maturity between countries.

The capture theory has been criticized because it fails to explain the existence of regulation to which industry is opposed (Posner 1974; Wilson 1980). The wave of social regulation that occurred on both sides of the Atlantic in the 1960s and 1970s, together with the deregulation of transportation and communications industries over the objections of monopoly incumbents, created substantial doubt about the universal applicability of capture (Derthick and Quirk 1985; Noll and Owen 1983). The capture theory is also based on neoclassical assumptions concerning rationality and utility, some of which behavioral economists have called into question (Simon 1955; Thaler and Sunstein 2009).

Furthermore, the capture theory fails to account for the interests of others outside the (potentially) regulated industry and the regulators, as discussed in the net neutrality case. Recognizing this, several scholars attempted to extend Stigler's theory to provide a fuller account of the interest group landscape, drawing on Olson's (1965) seminal work on collective action. Peltzman (1976) argued that to make regulation politically feasible, regulators would seek compromises that would require the regulated industry to share the benefits of regulation with other politically salient interest groups, including competing companies and potentially consumers. Becker (1983) suggested that regulation would particularly favor groups that were efficiently organized and put the least burden on other groups. Empirical studies in cases with a small number of affected interests have demonstrated that organized interests do often succeed (Noll 1989). Observers might reasonably claim to see these theories at work in the FCC's *Open Internet* rulemaking, whose compromises among organized producer and consumer groups have been openly acknowledged (Clyburn 2010).

When political scientists look at economic theories of regulation, they see specialized instances of broader theories of pluralism and interest group politics (Moe 1987a). Most influentially, James Q. Wilson emphasized the need to understand regulators as operators within the wider political system. Wilson (1975; 1980) identifies specific relationships between the diffusion of a regulatory policy's costs and benefits and the character and strength of interest group influence over regulatory outcomes. Policy issues yielding benefits only to a concentrated interest group but whose costs are more diffuse may well be subject to capture dynamics, whereas when a regulatory battle yields concentrated costs and concentrated benefits (on differing groups), regulation would emerge as the product of compromise or coalition-building between the interests. As the discussion of the net neutrality literature revealed, there is little consensus about the level of concentration or diffusion of the costs and benefits associated with net neutrality regulation across the population of application developers, broadband providers, and consumers, making it difficult to draw conclusions according to Wilson's taxonomy without further empirical evidence.

3.3.2 Legislatures and the Executive

The economic theory of regulation is distinguished by its purposeful omission of political institutions. As Moe (1987a, 475) has observed, "[t]he implicit claim is that institutions do not matter much . . . . There is little reason for accepting this claim. . . . if decades of political research testify to anything at all, it is that public policy cannot be understood without systematic attention to the nature and dynamics of political institutions." Moe and his American contemporaries therefore turned their attention to the impact of political institutions, particularly the US Congress and President, on regulatory outcomes. Accounting for differences between presidential and parliamentary systems, more recent work in Europe has followed similar lines of development, examining the relationship between elected politicians and independent regulatory agencies at both the EU and national level. This literature focuses on the extent to which regulatory agencies are controlled by politicians, how that control is exerted, and to what effect. The last of these questions is clearly most salient for this thesis, although it is the least developed.

The literature is characterized by two divergent views about the extent to which regulatory agencies are controlled by politicians. The first of these holds a traditionally bureaucratic view (following Niskanen (1971)), arguing that agencies are mostly autonomous and operate outside the bounds of political oversight. Some scholars, particularly early US commenters and more recent European ones, view this state of affairs as an explicit goal of legislators and government officials, along the lines of the public interest theory. They argue that politicians delegate authority to agencies to allow regulatory decisions to remain above the political fray, or because the technical complexity of regulatory issues is better dealt with by a dedicated expert agency (Baldwin and McCrudden 1987; Lowi 1979; Majone 1999; Thatcher 2002a). Particularly in Europe, with the shift from public to private ownership of utilities that began in the 1970s, delegation of authority to regulatory agencies has been viewed as an explicit means for politicians to make credible commitments that regulation will not fluctuate with the whims of the electoral cycle (Gilardi 2005; Majone 1999; Thatcher 2002a).

Others are less emphatic about bureaucratic autonomy being an explicit regulatory design goal, but they nonetheless conclude that agencies are able to maintain their own discretion for a variety of reasons. There may be little cause to believe that arcane regulatory issues have political salience – if an issue is too obscure to sway voters, it is not likely to be worth a politician's time (Spence 1997; Wilson 1980). Regulators may be capable of cultivating their own political legitimacy directly with interested parties, giving them the ability to challenge politicians with whom they disagree rather than being controlled by them (Carpenter 2001). Even if politicians have a desire to control agencies, they may lack the technical expertise or inside information to do so (Baldwin and McCrudden 1987; Niskanen 1971; Spence 1997). Leveraging these advantages, agencies will seek to maintain themselves or even expand their turf regardless of the preferences of politicians (Katzmann 1980; Niskanen 1971).

Reacting to some of these arguments, and focusing specifically on the US Congress, Weingast and Moran (1983) put forth an influential alternative view: the theory of congressional dominance. According to this theory, congressional committee members with oversight over regulatory agencies are assumed to have both the incentives to control agency behavior and the instruments to exert it: budgetary authority, oversight hearings, and threats of restrictive legislation. Weingast and Moran provide evidence by analyzing Federal Trade Commission behavior, showing how the FTC pursued controversial policies when encouraged to do so by Congress before becoming more conservative following a shift in the legislature's political orientation.

Scholars in the congressional dominance camp argue that congressional oversight is present even when it is not overt. Linking congressional dominance with the power of interest groups, McCubbins and Schwartz (1984) argue that Congress prefers to rely on "fire alarms" – concerns raised by constituents – to conduct regulatory oversight. By crafting procedural rules to ensure that interested parties can be informed about regulatory processes and have avenues to complain, Congress need not invest its own resources to determine when a regulatory agency is in need of a course correction (McCubbins, Noll, and Weingast 1987). These procedures may constrain agency action to such an extent that agencies are able to operate on "autopilot," with no need to seek further legislation – and no need for legislators to expend further resources and political capital – even as new policy issues arise. McCubbins, Noll, and Weingast (1987) use the FCC's regulatory approach to cable television as an example of the workings of an agency on autopilot, showing how the agency started by erecting barriers to the development of cable (at the behest of one of the agency's key existing constituencies, the broadcasters), moved on to craft concrete policies reflective of the political climate at the time, and eventually brought cable television fully within its remit, all without ever needing new legislation.

Rejecting both the bureaucratic and congressional dominance views, Moe (1985; 1987b) suggests a more president-centered theory of agency control. He argues that the White House uses its power of appointments to populate agencies with leaders who are conducive to Presidential control (Moe 1985; Moe 1987b) and loyal enough to the President's interests to continue to provide the President with inside information about agency business once appointed (Moe and Wilson 1994). However, the President's inability to foresee how an appointee will react to both emerging policy issues and influences from professional agency staff call Presidential control into question (Spence 1997). Presidents may choose agency heads based on politics rather than substantive policy (Heclo 1977; Wilson 1989).

Beyond appointments, executive agency officials including the President can centralize regulatory control in the executive by imposing procedural rules. Scholars often point to President Reagan's executive order requiring the Office of Management and Budget to review cost-benefit analyses of all major regulations and the somewhat less stringent requirements for cost-benefit analysis imposed by the Thatcher Administration as examples of this behavior (Baldwin and McCrudden 1987; Kagan 2001; Moe and Wilson 1994; Ogus 1994). It is far from clear, however, how effective these kinds of measures are in aligning agencies with executive preferences (Spence 1997).

The question from a traffic management perspective is whether the bureaucratic view, the congressional dominance theory, or theories of executive control provide explanations for the activities of the FCC and Ofcom. Does Ofcom operate as autonomously from government as other British regulators appear to have in the past (Thatcher 2002b)? Are the FCC's decisions controlled by members of the House and Senate Commerce committees, their constituents, the President's appointment choices, or other factors? Whether the primarily US-centric theories developed here can be generalized to other kinds of political systems is a significant question for further research.

3.3.3 Judicial Review

As much as political scientists have been attentive to the relationship between legislatures, the executive branch, and agencies, administrative legal scholars have focused on the relationship between regulators and the judicial institutions meant to hold them accountable. The body of work examining these relationships has been particularly responsive to increases, whether real or perceived, in the frequency and intensity of judicial review of agency decisions across many sectors.

In the US, the doctrine of judicial review is rooted in the Administrative Procedures Act (APA) of 1946, which provides a general authorization for courts to evaluate whether agency decisions are supported by substantial evidence and are not arbitrary or capricious (5 U.S.C. 500 *et seq*). In the decades since the passage of the APA, these standards have been significantly strengthened through case law and acts of Congress. Analytical requirements on agencies have expanded, such that agencies have been required to show that they have a "reasoned explanation" for specifying rules, not merely substantial evidence to back up their decisions (McGarity 1992). Substantive requirements have likewise become stricter, in particular by application of the "hard look" doctrine adopted by the lower courts in the 1970s (Breyer 1986). The hard look doctrine obliged the reviewing courts to "examine carefully the administrative record and the agency's explanation, to determine whether the agency applied the correct analytical methodology, applied the right criteria, considered the relevant factors, chose from among the available range of regulatory options, relied upon appropriate policies, and pointed to adequate support in the record for material empirical conclusions" (McGarity 1992, 1410). These standards are considered to be significantly more intensive than the original APA requirements.

The years since the passage of the APA have also seen a substantial enlargement in the population of interests entitled to participate in the regulatory process and seek judicial review. The expansion of these rights was articulated in a number of court decisions beginning in the 1960s as "a judicial reaction to the agencies' perceived failure to represent such interests fairly" (Stewart 1975, 1728). While the courts may have been aiming to make the regulatory process more available to public interest groups, the net result was an opening of the process to anyone interested in challenging regulatory decision-making (Horwitz 1989) and heightened attention within the judiciary as to whether agencies had given all interests "adequate consideration" (Stewart 1975).

British legal doctrine has no similar statutory basis for judicial review as the APA, but has instead a well developed common law basis for court review of governmental actions on the basis of illegality, irrationality, or procedural impropriety, supplemented by agency- or department-specific statutes that allow for appeals of regulatory decisions on their merits (Baldwin and McCrudden 1987; Bishop 1990; Wade and Forsyth 2004). As in the US, an uptick in the incidence and substance of judicial review in the UK has been observed in recent decades (Black and Walker 1998; Ogus 1994), despite a perceived general aversion among UK officials to the litigiousness of regulatory decision-making in the US (Moran 2003; Ogus 1994; Prosser 1997; Prosser 2004). Scholars have attributed these developments to shifting perceptions about the efficacy of the judiciary and its relationship to other branches of government. With heightened political divisions throughout government, the judiciary has emerged as an arena for adjudicating political disagreements (Baldwin and McCrudden 1987). Whereas the mere appeal of a regulatory decision might have once been considered a symbol of incompetence on the part of the responsible regulator or minister, frequent adjudication has merely come to reflect common interplay between an emboldened judiciary and other branches of government. Within telecommunications specifically, the increased incidence of appeal – and of Ofcom being overturned in court – is unmistakable, and possibly attributable to the widely held perception that Ofcom's predecessor Oftel was never properly held accountable, as Oftel was never successfully challenged in court during its 19 years of existence (Walden 2012).

Scholars have theorized (and in some cases sought to empirically verify) a number of implications for regulatory agency behavior based on observed and perceived increases in the strength and frequency of judicial review. The overarching claim is that agencies become more conservative as a result of heightened judicial review. The fact that courts may overturn them on an ever-expanding and ill-defined set of grounds creates uncertainty and excessive deliberation about how the agency should construct its rules and explanations for them (Baldwin and McCrudden 1987; Melnick 1992; Pierce 1995). The adoption of "adequate consideration" standards exacerbates this constant second-guessing, since agencies need to be able to anticipate every potential grounds for appeal that any interested party might launch and build appropriate defenses into its rules and rationales for them (Horn 1995; McGarity 1992).

This conservatism is said to manifest in a variety of ways. Rules can take much longer to develop (C. Scott 1998), and in some cases the combination of onerous rulemaking processes and extensive reconsideration of rules that have been remanded to the agency can be essentially debilitating, leading to the so-called "ossification" of the rulemaking process (McGarity 1992; Melnick 1992; Seidenfeld 1997). Agencies may also become less likely to reopen settled rules for fear of spurring a lengthy legal process, even when changes in the marketplace necessitate regulatory updates (Breyer 1986). As a general matter agencies become disinclined to experiment with radical or innovative ideas in promulgating new rules, preferring to rely on safer ground that has already been tested under judicial review (Baldwin and McCrudden 1987; Horn 1995; McGarity 1992). The overall effect may be one of fewer rulemakings altogether (Melnick 1992; Pierce 1988).

Increased judicial oversight may also cause agencies to seek alternative procedures that are not statutorily subject to appeal, including informal rulemaking, guidance, policy statements, and the like (Baldwin, Cave, and Lodge 2012; Pierce 1995; Seidenfeld 1997; Spence 1997). Scholars point to a number of examples where agencies have shifted their processes in response to a pattern of lengthy (and at times unsuccessful) court challenges: the US National

Highway Transportation and Safety Administration moving to post-hoc recalls and adjudications rather than affirmative promulgation of vehicle safety rules (Mashaw and Harfst 1986); the US Environmental Protection Agency preferring informal dealings with industry to implement the Clean Air Act rather than formal requirements (Melnick 1983); and the UK Commission for Racial Equality shifting to bargaining and rulemaking after judges proved hostile towards the agency's formal investigations of businesses (Baldwin and McCrudden 1987).

Whether claims of conservatism, ossification, and circumvention are manifest in practice, remain unchanged by developments in case law, and are relevant to telecommunications regulation are all significant open questions. Although certain agencies at certain points in time may appear to have been hampered by judicial oversight, empirical evidence of US regulators writ large suggests that they continue to produce more and lengthier rules (Coglianese 2002) and that most rulemakings are completed fairly promptly (Yackee and Yackee 2010). The US Supreme Court's jurisprudence regarding deference to agency expertise, first in *Chevron* (Chevron U.S.A. Inc. v. National Resources Defense Council, Inc. 467 U.S. 837 (1984)) and more recently in *Mead* (United States v. Mead Corp. 533 U.S. 218 (2001)), have further complicated evaluations of the impact of judicial review. While *Chevron* is said to restore judicial deference to the expertise of regulators, scholars dispute the evenness of its application by judges (Baldwin, Cave, and Lodge 2012; Melnick 1992). Whether *Mead*, which held that *Chevron* deference only applies in cases of formal agency decision-making, increases or decreases agency discretion has also been the subject of debate (Baldwin, Cave, and Lodge 2012; Yackee and Yackee 2010). In the UK, scholars and government officials have noted the increased incidence and intensity of judicial review (Black and Walker 1998; Prosser 1997) and the fact that the telecommunications regulator is largely unique among economic regulators in that it can be appealed on the merits of its decisions (National Audit Office 2010), but the implications of these facts are unclear.

Such implications may be weighty in the case of traffic management and net neutrality. The UK government was so alarmed by the effects of increased appeals on Ofcom that it launched three separate consultations in the span of four years about revising standards for judicial review (BIS 2010; BIS 2013; DCMS 2011c). No legislative changes resulted from the first two and the third was ongoing at this writing. The FCC notoriously gets challenged on every major decision it makes, including its most significant pronouncements related to net neutrality (one of which was based on the *Policy Statement*). Whether judicial review has inspired conservatism, ossification, circumvention, or other characteristics in the two regulators is a vital line of inquiry in this thesis.

3.3.4 External Signals

Trying to understand agency behavior on the basis of a single cause – the regulated industry's preferences, the party in power in Congress, or the standards for judicial review – may be appealing for its analytic simplicity, but it is unlikely to adequately reflect reality (Moe 1987a). Recognizing this, Roger Noll coined the term "external signals" to describe a theory that explains agency behavior as the product of multiple external forces:

> [A]gencies try to serve the public interest but have difficulty identifying it, because the public interest is such an elusive concept. Consequently, they judge the extent to which their decisions satisfy the public interest by observing the responses of other institutions to their policies and rules. . . . Among these sources of performance indicators are the courts; congressional committees that decide upon the budget and legislative program of the agency; the relevant budget examiners in the Office of Management and Budget or corresponding state or local government agencies; the press, whose primary locus of concern is the regulated industry and who may criticize the agency if performance there deteriorates; and the constituent interest groups participating in agency procedures, who, if dissatisfied with an agency decision, can appeal to the courts or can take their case to the politicians or the press. (Noll 1985, 41)

The external signals theory suggests that agencies are judged in multiple "theaters," including the marketplace, the hearing room, and the political arena (Noll 1971). They seek to minimize criticism and conflicts within and between each one (Joskow 1974; Mitnick 1980) while maximizing positive feedback. Succeeding in this endeavor involves creating a base of support among multiple external groups, allowing the agency to "shelter itself from changing political environments" and maintain its autonomy (M. K. Olson 1995, 388).

The external signals theory is attractive because it provides the opportunity to draw together the diverse strands of regulatory theory based around individual external forces and understand how they integrate with one another in a given regulatory context. The theory has been applied empirically to show that energy regulators respond to a combination of signals from environmental advocates and markets (Joskow 1974); that Congress and consumers jointly influence how the Food and Drug Administration (FDA) chooses regulatory instruments (M. K. Olson 1996); that different kinds of FDA drug approvals respond to different external signals (M. K. Olson 1995; 1997); and that the Environmental Protection Agency responds to a combination of marketplace and bureaucratic factors (Magat, Krupnick, and Harrington 1986). The external signals approach is a promising fit for understanding traffic management regulation given the complexity of the regulatory space for broadband.

3.4 Internal Characteristics

Theorists who focus on the influence of political institutions and external signals have noted that internal factors relating to agencies and the individuals they employ may be as important, if not more so, in understanding regulatory behavior (Carrigan and Coglianese 2011; Magat, Krupnick, and Harrington 1986; Moe 1987a). Scholarly attention to regulatory agencies from the perspectives of public administration, organizational sociology, economics, and management has shed light on how agencies function from the inside, how bureaucrats (or "agency officials," as they will be called here since "bureaucrats" does not appropriately describe Ofcom staff) behave, and the motivations behind both. This literature is vast and operates both at the level of the organization and of the individual employee. Given that the design of this thesis (involving interviews and participant observation) is focused on the experiences of individual regulatory employees as opposed to whole internal organizational structures, this section highlights the theoretical developments that focus on the link between agency officials and regulatory decisions.

The literature on bureaucracy and public administration reveals a tension – typical for any field that has engaged the attention of economists – between conceptions of agency officials as purely self-interested utility maximizers and those who suggest motivations based on a combination of self-interest and concern for other values (DiIulio 1994). This distinction harkens back to debates about the public interest theory of regulation – do officials concern themselves with the agency, its mission, and what it is intended to accomplish on behalf of the public? Scholars have pointed to a variety of motivations to explain the behavior of individual officials, some self-interested, some "public spirited" (Golden 2000, 12) and some that blur the line between the two.

Perhaps the most fundamental interest that agency officials have is "survival": remaining in office (Russell and Shelton 1974, 49). According to the economic theory of regulation, the desire to stay in office drives officials to make decisions that please the regulated industry (Niskanen 1975; Peltzman 1976; Stigler 1971), but under alternative conceptions it could imply appealing to the public or forming a cross-cutting political coalition that will continue to support the official's decisions (Russell and Shelton 1974; Wilson 1980). Some scholars argue that agency heads are imperialistic, constantly seeking to expand their jurisdiction as a means of survival (Niskanen 1971; Tullock 1965). The implication is that they will seek to please legislators and executive branch officials (in the Office of Management and Budget, for example) that control their budgets and authority.

Whether imperialism drives most regulators is debatable, however. Reflecting on case studies of nine US regulatory agencies, Wilson (1980, 376) was struck by the "defensive, threat-avoiding, scandal-minimizing instincts of these agencies." Officials at such agencies prefer a posture that gives them a low profile and reduces the chances of having agency decisions reversed or stirring controversy among interest groups – what Leaver (2009) terms "minimal squawk" behavior. Furthermore, evidence suggests that in reality many agencies have declined budget increases or volunteered to cut departmental units (Wilson 1989).

Since agency executives usually serve limited terms, survival may be less important to the self-interested official than post-tenure job prospects. A "venal administrator" might operate his or her agency so as to find the most lucrative employment afterward (Horn 1995; Noll 1985). As a result, executives may be reluctant to antagonize or alienate industry segments that might provide future employment, leading to regulatory decisions that favor the industry or temporary resolutions or lengthy procedures that postpone weighty decisions until after the executive's term (Hilton 1972; Noll 1985; Russell and Shelton 1974). However, Mitnick (1980) questions the extent to which venality motivates agency executives given that all high-ranking agency officials are all likely to be highly valued in the private sector at the conclusion of their agency service. Agency executives obtain valuable expertise, experience, and inside knowledge of the agency (Carpenter 2010), which together comprise "a non-transferable, personal capital asset which can only be realized in a future non-official position" (Russell and Shelton 1974, 48). For the professional staff of the agency, private sector jobs may not even be desirable give that "the material and nonmaterial satisfactions of public service may equal or exceed those of private employment" (Wilson 1989, 86).

Agency officials who make their own professional development a key priority will likely be responsive to the norms and teachings of their professional cultures, even if those norms drive behavior that is at odds with the goals of the agency (Golden 2000; Noll 1985). For lawyers and economists – two common breeds in the telecom world – the ways in which they are trained and evaluated are profession-specific: lawyers are considered successful when they win cases, and economists when they maximize welfare (Katzmann 1980; Wilson 1989). The impetus to fit these moulds can skew officials towards particular approaches to their work; for example, lawyers may focus on trivial cases they know they can win, or they may devote more resources to their cases than are strictly necessary (Katzmann 1980; Wilson 1989). If officials' post-agency job prospects depend on the judgments of their professional counterparts in the private sector, they may be more responsive to how professional culture shapes those judgments than to the mission of the agency (Horn 1995).

Underlying both the concern for future employment prospects and adherence to professional norms may be the more fundamental pursuit of personal status, esteem, and reputation (Carpenter 2010). Success in public service is often equated with perceptions of influence or popularity among the agency's key audiences (Wilson 1989). These less tangible benefits may motivate officials to favor regulatory work that reflects well on them, avoids controversy, or upholds their integrity (Russell and Shelton 1974).

As an agency head, building one's own reputation and maintaining the reputation of the agency may or may not be mutually reinforcing; executives of a more public spirited nature may favor the latter over the former. Agency reputations tend to be formulated around a multitude of specific traits: efficiency, uniqueness of service, or expertise, for example (Carpenter 2001; 2010). Regulatory decisions can be perceived by numerous audiences – organized interests, politicians, or the media – as either reinforcing or detracting from these reputations (Wilson 1980). Focusing on maintaining agency reputation may lead to the same kinds of conservative behaviors as focusing on "survival," since reversals of agency decisions or perceptions that agencies are shifting course can call agency reputation into question (Carpenter 2004). While "external signals" accounts of regulation likewise emphasize the feedback that agency officials receive from their external audiences, reputation provides a further refinement by helping to explain why agencies are more responsive to some signals than others at different times (Gilad 2012).

Having a reputation requires having an audience. But much of what defines the day-to-day workings of an agency, while not for public display or judgment, can influence regulatory decisions. As in organizations of many types, employees of regulatory agencies develop distinctive beliefs about their jobs, patterns of conducting their work, and shared expectations about how they relate to the agency as a whole. Scholars have variously labeled the collection of these items as an organization's "culture" (DiIulio 1994, 283), "sense of mission" (Wilson 1989, 95), "moral factor" (Barnard 1938, 72), or "essence" (Halperin 1974, 28). Particularly when an organization's goals are vague – as is often the case with regulatory agencies – its members develop distinctive ways of pursuing their work and understanding internal relationships (Wilson 1989), often with reference to their interpretation of how the agency is meant to serve the public. Far from being on "autopilot," agency officials act purposefully to conform to the mores of their colleagues and the agency as a whole (DiIulio 1994). A strong sense of culture in an agency can yield regulatory behavior that is consistent and cohesive (Kaufman 1960), although it may also make organizations resistant to taking on new and different tasks (Wilson 1989). Conversely, when subdivisions within an agency develop their own distinct cultures, conflicts between them can lead to dysfunction.

Finally, regulatory outcomes will depend on the personality and temperament of executives and staff. Agency leaders who favor negotiation and compromise will make different decisions than those who see themselves as advocates for a particular executive branch agenda or those who value their own personal autonomy in making choices and acting decisively (Wilson 1989).

As this section has demonstrated, the literature on organizations as it relates to regulatory officials offers few particularly strong theoretical directions to guide inquiries into specific regulatory decisions. Assigning causality at the individual level provides little predictive value without a deep understanding of the specific officials involved in any particular regulatory decision. This work does suggest, however, that between the need for "survival," pecuniary motivations, professional norms, the quest for status, and organizational culture, the internal workings of agencies and those they employ should not be overlooked in understanding regulatory outcomes.

3.5 Nation-Specific Factors

One of the criticisms of the economic theories of regulation is that they minimize or ignore the influence of culture, tradition, history, and ideas on regulators. While these aspects may be more difficult to measure or understand than individual or political preferences, an important collection of scholarly work seeks to make them more central to explanations of regulatory behavior.

Under a nation-centric view, industrial history and culture are said to circumscribe both the potential problems and the potential solutions perceived by national regulators (Dobbin 1994; Dyson 1983). Thus when national regulators are faced with new decisions, they are implicitly drawing on policy paradigms that have accrued national endorsement over time: notions of community sovereignty and mistrust of government in the US, or notions of champion firms and club-like regulatory governance in the UK (Dobbin 1994; Moran 2003; Vogel 1978). These same national traditions also help to constitute the repertoire of regulatory responses that firms perceive (Dobbin and Sutton 1998; Hancher and Moran 1989; North 2005), causing firm behavior to be driven and shaped by subjective national models of policymaking (North 1990). By emphasizing the importance of distinctive national patterns in regulators' approaches to problems, relations with industry, and conceptions of acceptable solutions, the nation-centered view challenges the notion that universal laws of utility maximization, efficiency, and the relative power of interest groups fully constitute policy outcomes.

What derives from "the experience of history, the filter of culture, and the availability of existing resources" (Hancher and Moran 1989, 280) are distinctive national "policy styles" that govern the interaction between the government's problem-solving approach and the relationship between government and other actors in the policy process (J. J. Richardson, Gustafsson, and Jordan 1982, 13). Scholars have discerned a number of measures to distinguish national policy styles from each other, including whether the government policymakers operate consensually or adversarially, with an incrementalist or rationalist approach, in anticipation of problems or in reaction to them, and along many other dimensions (Hancock 1983; Heidenheimer, Heclo, and Adams 1990; J. J. Richardson 1982; J. J. Richardson, Gustafsson, and Jordan 1982).

A number of scholars have sought to use the national policy style concept as a way of understanding cross-national variations in government regulation of particular industries (Kagan (2003) provides one list of studies). Vogel (1986) found national "regulatory styles" to be a key differentiator between British and American environmental regulations. The case studies presented by Kagan and Axelrad (2000) contrast the adversarial American style against the styles of a number of other nations by focusing on the experiences of multinational corporations operating in regulated industries such as chemicals and manufacturing. In each case, the national policy style is found to be a central differentiator in how regulation is created and in how firms respond to it. Generalizing, American regulatory style is said to be adversarial, with policymaking processes that provide ample opportunity for interest group participation but also breed suspicion and result in litigious, contentious relationships between regulators and regulated industries. Regulators in Europe and those operating under parliamentary forms of government are said to have a more consensual style, cooperating more closely with businesses which themselves tend to be less mistrustful of government bureaucrats (Kagan 2003; Vogel 1986; Wilson 1989).

Related to the nation-centric and policy styles approaches to regulation is the notion of the "power of ideas" in shaping national regulatory strategies (Baldwin, Cave, and Lodge 2012, 49). Ideas-based approaches focus on the role of the intellectual climate and the influence of experts in the policy process. When the intellectual or expert classes seize on new ideas – deregulation or environmental conservation, for example – the spread of these ideas fuses with the state's latent administrative and political biases to produce state-specific policy outcomes (Derthick and Quirk 1985; Hall 1989; Wilson 1980). Thus ideas and institutions serve in mutually reinforcing roles, with the diffusion of ideas providing some explanation for change in both the structure and output of institutions (Hall 1993).

The point of ideas-based arguments is not that intellectual paradigms are nation-specific, but that their interpretations can be. Policymaking in the European Union is particularly ripe for isomorphism and idea diffusion given the harmonized legal standards across Member States. Indeed, regulatory agencies have sought to emulate each other where they have found their neighboring counterparts to be successful (Gilardi 2002; Gilardi 2005). These dynamics have been further reinforced by the creation of regulators' groups, which foster the "cross-national fertilisation of ideas" between regulators in the same sector (Thatcher 2002a, 137). Thus idea dissemination is not the limited province of think tanks and academics that manage to infiltrate the policy process (Derthick and Quirk 1985), but may further rely on the intellectual exchange between regulators.

In a cross-country qualitative comparison such as this thesis, it would be difficult to ignore the influence of culture, tradition, and history on the specific regulatory decisions of interest. The questions that nation-specific regulatory theories raise for traffic management concern how the American and British regulators and firms filter their choices through these lenses, whether conventional notions about each country's regulatory style hold true in the case of net neutrality, and whether any particular ideas found institutional settings receptive enough to powerfully shape regulatory outcomes.

3.6 Conclusion

While the ideas behind the public interest theory clearly continue to influence scholarly thinking about regulation, and about net neutrality in particular, more recent work on regulatory theory has broadened and diversified in many different directions, some of which are in conflict with each other. Taking an interdisciplinary view requires considering and combining ideas from fields where the units of study, assumptions, methods, and solution spaces are at odds. This makes the development of meaningful hypotheses and generalizable findings in a specific case (such as traffic management regulation) uncomfortable at times. If there is one overarching lesson from surveying such a diverse corpus of work, it is that no single discipline or perspective has thus far successfully captured the nuance and complexity of regulation to the point of offering predictive power across a sizable range of agencies or industry sectors, let alone countries.

What is notable about the array of theories surveyed here is the extent to which they are interrelated. Institutional design choices clearly affect (and may be affected by) how external forces are brought to bear on regulatory agencies. The same could be said for national traditions and history, which may not only create path dependency in choices about how regulators relate to the rest of government, but may also influence the types of individuals that pursue work as regulatory officials, which in turn influences internal regulatory culture. Teasing these aspects apart to find causality between any particular factor and a specific regulatory outcome is a complex process.

With those difficulties in mind – and recognizing that observable facts about the relationships and inner workings that comprise telecommunications regulation in the US and UK were limited at the launch of this study – a simple hypothesis was crafted about the relationship between each country's institutional setting and its traffic management outcomes (the substance of Research Question 2). It focuses on a single type of institutional factor:

> Question 2: How does the institutional setting – the formal and informal constraints that comprise the regulatory environment – influence traffic management outcomes?
>
> Hypothesis 2: National regulatory styles are a key determinant of traffic management outcomes. Consensual regulatory regimes are more likely to produce regulatory outcomes that do little to constrain network operator behavior; adversarial regulatory regimes are more likely to restrain network operators from discriminating for traffic management purposes.

This simple hypothesis does not capture the complexity of regulatory arenas, but instead focuses on one of the most plausible differences between the US and the UK: regulatory style. Acting according to the UK's traditional regulatory logic of negotiation, Ofcom's reticence to regulate traffic management practices may be influenced by close relations to the broadband industry. By contrast, the adversarial nature of US regulation may help to explain both the FCC's willingness to act and the seemingly endless debate about the agency's authority. These differences may relate to particular institutional design choices while also relating to the broader logic at play in each nation's approach to regulation of the economy. The hypothesis does not foreclose explanations based on institutional design choices, but focuses on how those choices are manifest in the relationship between industry and the regulators.

Compared to national regulatory styles, the application of the external forces and internal characteristics theories to the observable facts in the US and the UK was much less clear at the launch of this study. Similar interests with similar agendas were affected in both countries, but the regulatory outcomes were obviously different. Agreement about the effects of judicial review was far from clear. Generalizing across both countries on the basis of primarily US-based theories concerning the influence of legislatures and the executive seemed imprudent given the differences in the two countries' political systems. Internal characteristics were viewed as fruitful potential explanations for the observed results, but theories based on the motivations or personalities of agency personnel offered little predictive value or means to shape the research inquiry.

Following the discussion of research methodology in Chapter 4, the second half of this thesis will explore whether the two hypotheses put forth provide explanations for the realities of traffic management decisions in the US and the UK, or whether other perspectives that have been discussed in the last two chapters better elucidate the experiences of the two countries.

Chapter 4.
Research Methodology

4.1 Introduction

In light of the predominantly normative and theoretical nature of existing academic scholarship about net neutrality, this thesis seeks to make a substantive, positive contribution by leveraging empirical evidence about the behavior of network operators and telecommunications regulators since the mid-2000s. Network operators around the world have been making decisions about how to manage their networks while regulators have grappled with the question of whether their intervention is warranted to safeguard nondiscrimination. The choices that both sets of parties have made – what they did, and, more importantly, why they did it – offer a rich basis for understanding the merits of existing normative arguments and theoretical modeling assumptions.

A close investigation of these decisions indicates that the outcomes do not necessarily align with what some of the most prevalent theories would have predicted, particularly as they pertain to the respective roles of competition and regulation as deterrents against discriminatory behavior. Given the limited number of countries in which national regulatory policy concerning broadband traffic management has received significant attention, I chose a cross-national comparative approach so as to draw out explanations for divergent regulatory and marketplace outcomes. This study makes use of the two prominent cases of the UK and the US, neither of which appears to fit the predicted paradigm.

This thesis seeks explanations for these discrepancies using qualitative methods – elite interviews, documentary analysis, and participant observation – to gain an in-depth understanding of the traffic management decisions that large broadband providers have made and the ways in which regulators have responded in both countries. By exploring qualitative empirical evidence pertaining to specific decisions that broadband providers and regulators have made throughout the last decade, the first-order goal is to elucidate which factors have been most determinative of traffic management outcomes and how those factors square with the existing body of regulatory and economic theory concerning net neutrality. More broadly, this research seeks to identify the mix of institutional factors that shape the regulatory landscape in cases where traditional principles of competition and regulation fall short.

4.2 Comparative Approach

For a number of reasons, a qualitative comparative approach is a natural fit for achieving these goals. The comparative method is a research strategy especially well suited to problems that exhibit a small number of cases and many variables (Lijphart 1971). Thus cross-national research often takes a comparative form given the limited number of nations in the world that exhibit a particular set of traits or circumstances. Events or policies affecting entire nations also tend to emerge from complex processes, fitting the "many variables" mold.

Both of these characteristics are exhibited in the case of regulatory policy pertaining to Internet traffic management. The issues of discriminatory treatment of Internet traffic and regulatory responses to it have garnered substantial consideration in a limited number of countries. Where these issues have arisen, a diversity of actors, institutions, and events have come together in a mixture of different ways to determine the ultimate market and policy outcomes. Comparative research thrives when the subject of study requires deep knowledge and familiarity with how the parts of each case contribute to the whole (Ragin 1994). The complexity of factors contributing to traffic management outcomes warrants this kind of in-depth analysis.

The comparative method is also highly suitable in cases where traditional controlled experimentation is infeasible, as is often the case when studying policy issues on a national scale. Instead, as Heidenheimer et al. (1990, 1) note, "the fact that different countries often adopt alternative strategies for dealing with similar problems represents a kind of natural experiment." Studying multiple countries side-by-side helps to dissect causes and effects in a way that is not possible from studying a single country in isolation, especially when one country chose a different course than the other (Rose 1973), as is the case with UK and US net neutrality policy.

4.2.1 Case Selection

Both scholarly debate and policymakers' attention to net neutrality originated in the United States in the early 2000s. A number of events conspired in the mid-2000s to elevate the profile of the issue, sparking legislative and regulatory initiatives and resulting in extensive grassroots mobilization and mass media coverage. Since then, regulators, policymakers, and other stakeholders around the world have engaged with the issue to various degrees. Major regulatory and legislative activities have taken place in a number of countries, including Canada, Chile, Finland, France, Japan, the Netherlands, Norway, Peru, Singapore, Slovenia, and South Korea. A number of other countries' regulators have hosted consultations or requested public input on the issue, including in Croatia, Germany, Sweden, and the United Kingdom.

I selected the United States and the United Kingdom for this study so as to create a rich basis for comparison. The choice of these two cases follows comparative approaches broadly based on John Stuart Mill's "Method of Difference" (Skocpol and Somers 1980). The Method of Difference involves contrasting cases that are similar in some respects but where the emphasis is on key differentiating characteristics, and where the phenomenon to be studied exists in one case but not the other (Skocpol 1979). This structure is especially well suited to research that takes one or more explanatory problems as its focus and seeks to draw out causal relationships, as this thesis seeks to do.

As a general matter, the US and the UK are similar when it comes to broadband: both are advanced economies that have had consistently above-average broadband penetration since the mid-2000s compared to their OECD counterparts (OECD 2011). But as explained in Chapter 1, the two countries are rather opposite extremes in three crucial respects relevant to broadband traffic management: competition, discrimination, and regulation. Conducting the comparison allows for analysis of the effects of these important differences, while their similarities with each other and with other advanced economies provides a basis to draw generalized conclusions outside of the two countries' national boundaries.

4.3 Data Collection

One of the key contributions of this thesis is the synthesis of empirical evidence collected via participant observation, elite interviews, and documentary analysis. Because of timing constraints at the site where participant observation was conducted, participant observation formed the first phase of data collection. This was followed by dedicated periods of interviewing and documentary analysis in the US and the UK, respectively. Each of the latter two phases began with some preliminary documentary analysis and the development of a list of potential interview candidates. As candidates began to accept invitations to interview, I developed individualized interview protocols, and as interviews were completed, I began to transcribe them. Transcription and documentary analysis continued for a number of months after the final interview was completed in each country.

4.3.1 Participant Observation

For a period of five months in 2010, I spent several days per month at a UK ISP observing a selection of meetings, conference calls, and interactions between members of a corporate traffic management research team. Participant observation was arranged as part of my joint work with the research team on a project at the company that sought to characterize the effects of traffic management on network traffic and on company expenditures. I was working with team members as they explored different traffic management options for the company and observing meetings and interactions between the team members while on site at the ISP. Thus my role was as participant-as-observer (Gold 1957) – I simultaneously contributed to the project team while developing relationships with informants and observing their activities and interactions.

Participant observation is a powerful tool for investigating the way that individuals think, feel, and act within a natural setting. It is particularly well suited in cases where the processes and interactions of interest may be obscured from public view or are poorly understood (Waddington 1994). For this study, participant observation provided an intimate perspective of traffic management decisions as they were being made, complementing the more reflective insights obtainable via the other two methods. It was also a valuable means of learning the language of telecom company culture, a necessary step in comprehending the mental models of study participants (Burgess 1995). Because participant observation comprised the first phase of the research, the insights and vocabulary that I gained at the research site were used to shape the framing and language used later on in developing interview protocols, particularly the UK protocols since industry jargon can be highly country-specific.

My process of data collection involved taking detailed field notes at the conclusion of each day at the research site. My observations focused on the people involved in researching, discussing, and setting internal policy about traffic management; the processes involved in deciding traffic management strategy from both technical and policy perspectives; the substantive discussions of different traffic management strategies; and the feelings and emotions of the informants. Immersion in the field for the purpose of observation is a naturally inductive process (Waddington 1994); as this was the first data collection phase, I aimed to gather a diversity of initial observations that could inform and refine my hypothesis formulation and research direction as I moved into the later phases.

Because I was in the role of participant-as-observer, it was never my aim to act as a completely passive observer of the team's work. I contributed to the team by analyzing existing data about network traffic in the presence of traffic management and creating models to predict future traffic growth and its associated costs under a variety of hypothetical scenarios in which the company's traffic management solution was altered or removed. Thus, my work for the team was primarily technical and descriptive, although it did involve conceptualizing arguments for and against the reduction or removal of application-specific traffic management. While the primary consequences of this work for me as a researcher were purely educational (particularly since I began the research process with a lot to learn about traffic management and the UK broadband industry), it also allowed me to develop an appreciation of the team's opinions about traffic management, including that there may be valid reasons to either continue or discontinue application-specific traffic management. I recognized the potential for these opinions to bias my research going forward, and sought at every turn thereafter in the data collection process to question whether newly collected data contradicted or supported what I had observed and discussed with the research team. I also interviewed five employees from the company as part of the interviewing phase, only one of whom was also a member of the research team, so as to obtain an in-depth understanding of the diversity of views about traffic management within the company.

My points of contact on the research team were fully aware that I was conducting participant observation and provided their consent to be observed as part of my thesis work. They understood the value of being able to contribute to a larger study of traffic management industry-wide and thus were receptive to questions about how traffic management was handled within the company. During the process of writing the thesis, I consulted with the team leader numerous times to ensure that I was accurately reporting information that I had learned from observation and to verify that I was not revealing any information that was sensitive for the company. Members of the research team are not directly quoted in this thesis; the terms "the observed research team" or "the observed network" are used to indicate information reported based on participant observation.

I was cognizant of the potential impact that my own participation had on the people and processes that I was studying. My aim was to remain as unobtrusive as possible while seeking to connect with informants on a human level so as to be able to observe them and their work environment in their natural states (Schwartz and Schwartz 1955). Furthermore, because participant observation was conducted at only one ISP, in one country, and with access to a handful of employees, the extent to which inferences and conclusions were drawn solely from the observation data was limited. Throughout the study I viewed participant observation as a means to triangulate my findings and as a vivid way to understand the day-to-day tensions and constraints that can contribute to traffic management decisions.

4.3.2 Elite Interviews

I opted to use interviews as a primary method of data collection to elucidate the decision-making processes of broadband providers and regulators. Interviews create the possibility of "seeing through others' eyes," providing researchers with a window into the depth and breadth of perspective of experts in the field (Bryman 2008, 465). Because traffic management and regulatory decisions have been occurring over a number of years, I chose interviews as an effective strategy for both reconstructing the circumstances leading to firms' and regulators' choices and illuminating each constituency's perception of the other as they have developed over time.

Interviews were vital to this research because all of the motivations and tensions behind traffic management decision-making are not necessarily matters of public record. As Useem (1995, 20) notes, "to answer some of the most vexing questions about business, we often have little choice but to enter that world directly." While public documentation provided a window into the formal, packaged rationales and justifications used by corporations and governments for their traffic management choices, interviews were used to elicit more raw, private, and unvarnished accounts of how those choices were constructed.

I conducted 70 semi-structured interviews for this thesis. These were "elite" interviews in the sense that participants were selected based on their expertise and standing in their professional fields (Marshall and Rossman 2010). The participant pool was comprised primarily of senior managers, corporate executives, and mid- or high-ranking government officials. Participants were drawn from two populations: current and former employees of broadband network operators, and current and former employees of regulatory or policymaking bodies, with a small number (three) of current and former employees of network equipment vendors also included. Six of the interviewees had relevant career experience both in industry and as a regulator or civil servant.

Any inquiry into organizational decision-making faces inherent limitations related to the biases and potential hesitancy of participants to provide frank views. I employed a number of mitigation strategies to help ensure that the data collected reflected an accurate depiction of the pertinent events and the reasoning behind them. I spoke with multiple employees of nearly every organization included in the study to ensure that different accounts of the same events could be corroborated. By further triangulating my findings with documentary analysis and participant observation, I sought to insulate my conclusions from the biases of any particular individual or company.

All interviews were conducted for anonymous attribution – participants are attributed by their field of occupation but neither their names nor the names of their employers are disclosed. Every quote in this thesis that is given with anonymous attribution is a quote from an interview. The guarantee of anonymous attribution was made explicit at the beginning of every interview and, in some cases, reaffirmed throughout if a participant wanted to revisit our agreement. While in a handful of cases participants were noticeably hesitant to engage candidly, it was far more common for participants to openly and forthrightly discuss their experiences and perspectives, often sharing details not available in the public record. In short, I found that rather than needing to constantly seek ways to break through the company line, it was more often the case that participants approached the interview as an opportunity to provide their own unique perspectives of past events, which was perfectly suited to the goals of the study.

United States

In the US, I invited 52 participants (all via email) and 37 agreed and were interviewed. My initial list of potential interview candidates was developed based on my own prior experience in the field and by gathering the names of key individuals that were publicly involved in net neutrality discussions either in the press or at the FCC. Because gaining access to elites can be notoriously difficult (Thomas 1995), where appropriate I leveraged my own professional ties or enlisted the help of former colleagues to make introductions to potential interviewees, as other scholars have suggested (Hirsch 1995; Ostrander 1995; Thomas 1995). I also formulated the invitations so as to draw in the potential participants by suggesting how their unique perspectives would contribute to the study.

As the interviews progressed, I used snowball sampling to generate further potential candidates. I asked every interviewee for recommendations for other individuals to interview, often basing my query on the interviewee's specific connections or knowledge (by asking about co-workers at the same company or counterparts at other companies, for example).

Interview participants generally had expertise and experience with traffic management either from an engineering perspective or from a policy or legal perspective. Only one interviewee had a strictly business background, but a number of others, particularly at the senior executive level, had substantial strategy or business management responsibilities. With limited exceptions, interviewees had at least a decade of experience in broadband or telecommunications and many had been in their line of work for several decades. Several current and former board-level executives and high-ranking FCC officials participated in the study.

Interviewees' employment experience was split fairly evenly between cable operators, telephone companies, and the FCC. Private sector employees came from six of the largest broadband Internet service providers in the country. The employment experience and expertise of the participants are shown in Figure 1, with overlaps represented (for example, a participant with employment experience both at the FCC and at a cable operator is reflected in both categories).

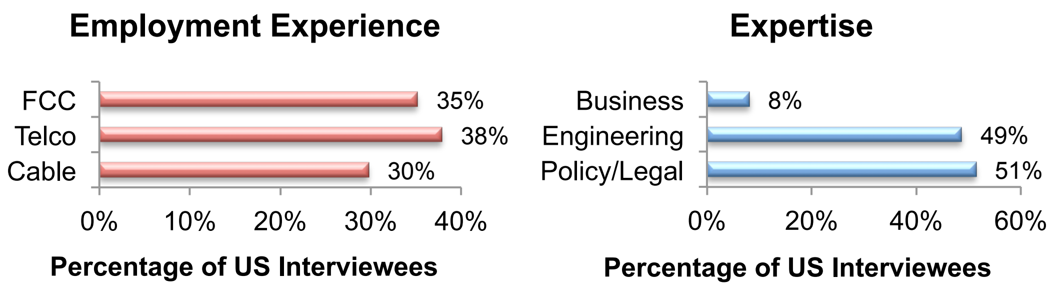

Figure 1. Employment experience and expertise of the 37 US interview participants.

Of those who were invited but did not participate, seven never responded, four declined to participate, and four agreed but were not able to accommodate an interview given scheduling constraints. If a response was not received within one to two weeks after the initial email was sent, a follow-up email was sent. In two cases, recipients of the follow-up email agreed to participate.

Interviews were conducted in person whenever possible, usually at the participant's workplace, and by phone or Skype (audio only) otherwise. In two cases, multiple employees of the same organization were interviewed together at the participants' request and despite appeals to interview individuals separately. In one case, a preliminary phone conversation was conducted with a participant and later followed up with an in-person interview. Thus 35 interviews were performed in total with the 37 participants.

All interviewees agreed to have the interview audio recorded except for one participant for whom detailed handwritten notes were used rather than audio recording. The recorded interviews totaled more than 25 hours of audio, with an average interview duration of approximately 44 minutes.

United Kingdom

In the UK, 60 participants were invited to participate and 33 accepted and were interviewed. My recruitment strategy was the same as in the US, but because I had fewer past professional contacts to begin with and because sources of public discourse about net neutrality that identify individuals are fewer in the UK, I relied more heavily on snowball sampling to develop my list of candidates.

Interviewees came from policy, engineering, and business backgrounds, with many of the business experts also having significant technical experience or knowledge. As in the US, most participants had worked a decade or more in the telecommunications industry and the participant pool included current and former senior level executives.

Participants were drawn from industry, government, and Ofcom. While in the US a strict focus on the regulator (the FCC) was used in choosing participants from the regulatory sector, there were several reasons for including government interviewees in the UK. The regulatory regime under which Ofcom operates is in large part determined by the broader EU regulatory framework, which underwent significant revision between 2007 and 2009. Net neutrality became a significant topic of discussion over the course of that regulatory review and has continued to draw attention from EU policymakers. Both Ofcom and the UK government participate actively in the EU policymaking process, and therefore it was important to incorporate the views of government officials as well as those from Ofcom. Separate from the EU activity, the UK government has been involved in coordinating self-regulation related to Internet traffic management, providing further impetus to include government participants in the study. Because responsibility for telecommunications policy has at various times resided with the Department for Trade and Industry, the Department for Business, Innovation & Skills, and the Department for Culture, Media & Sport, participants from all three departments were included.

As with US participants, more UK interviewees had private sector experience than public sector experience, with a fairly even split between those with engineering or business backgrounds and those with policy or legal backgrounds. Private sector participants came from six of the nation's largest broadband Internet service providers plus one smaller provider. The employment experience and expertise of the UK participants are shown in

Figure 2, again with overlaps represented in cases where participants belong to more than one category. Since there is at present only one large cable company in the UK, cable and telco participants are not broken out separately.

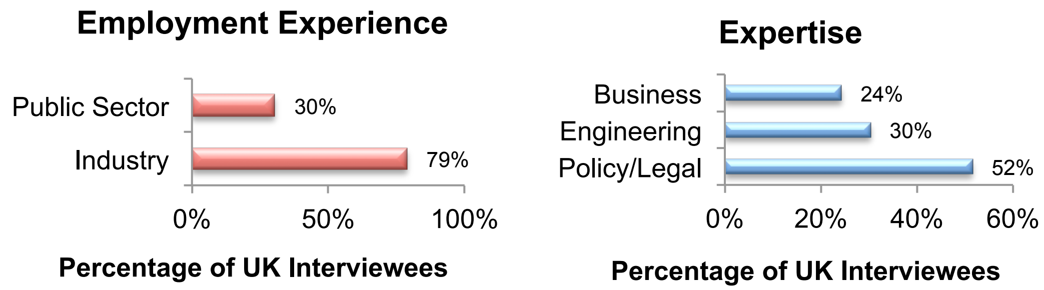

Figure 2. Employment experience and expertise of the 33 UK interview participants.

Of those who were invited but did not participate, nine never responded, eight agreed to participate but we were unable to resolve scheduling conflicts, five referred me to colleagues of theirs who did participate, and five declined. Five participants did not respond to the initial email but agreed to participate after receiving a follow-up email.

One interview was conducted via Skype (audio only), and one was conducted at a conference hotel, but otherwise all of the interviews were conducted in person at or near the participant's workplace. All participants agreed to have their interview audio recorded except one for whom detailed notes were taken instead. The average interview length was 60 minutes (33% longer than the US average) and the total interview time was 34 hours.

Interviewing Process and Structure

I conducted five pilot interviews (separate from the 70 core interviews) to become accustomed to interviewing and to trial interview questions. Four of these took place during the research design phase and one occurred in between the US and UK interview phases. None of the data collected from the pilot interviews was incorporated into the study's broader analysis.

The process of interview protocol development was evolutionary. I began with a generic interview questionnaire structure. For each participant, I altered, added, or removed individual questions based on research conducted prior to the interview about the participant's background, professional endeavors, and expertise.

One participant requested the interview questionnaire in advance and asked that we only discuss those specific questions. In all other cases, the interview protocols served merely as guides. In many instances the protocol questions were put aside for a time in order to explore in depth a particular experience or topic raised by the interviewee before returning to the protocol to ensure that all key topics had been covered, as recommended by Hirsch (1995). This flexibility contributed to further evolution in the interview protocols, where I would incorporate questions originally asked spontaneously in one interview into future interview protocols. In this way, I combined inductive and deductive approaches to data collection, at times steering participants towards topics drawn from the literature and my hypotheses, and at other times focusing on themes that emerged from the interviews themselves.

I drew heavily on the techniques of responsive and depth interviewing (Rubin and Rubin 2004). I strived as much as possible for the interviews to be conversational and grounded in participants' specific past experiences. I encouraged participants to tell their own stories about processes and events in which they had participated (Rubin and Rubin 2004). I found that participants often raised sensitive topics on their own when I inquired about related or tangential topics. For example, network operators' use of deep packet inspection (DPI) was one potentially sensitive topic given some of the negative press attention that it had garnered. I did not ask directly about DPI, but many participants brought it up of their own accord when I asked about their interactions with network equipment vendors (some of whom sell DPI equipment).

All interviews began with basic questions about the participant's job title, general work experience, and specific work experience relevant to Internet traffic management. For industry participants, I would then ask at a generic level about how traffic management decisions are made at the company or how changes in the way that traffic is managed come about. I started with these more basic and generic questions so as to ease into the conversation and build rapport, as recommended by Rubin and Rubin (2004).

Depending on the response to the generic traffic management question, I would probe on the influence of specific factors on the traffic management decision-making process. The entire selection of question topics was based on the classification of technical, economic, social, and legal influences provided by Lessig (1999; 1998) to describe the framework of formal and informal constraints on individual and organizational behavior. Further refinements in the technical, economic, and social realms suggested by van Schewick (2010) provided a framework for structuring questions. These included an individually tailored selection of questions about network usage, architecture, and engineering; the costs of different solution approaches and their potential to help providers reduce costs; the effects of competing providers' approaches and the role of competitive differentiation; interactions with network equipment vendors; and the intersection between traffic management and initiatives to police copyright infringement.

I also asked participants about specific regulatory events or incidents that had received attention in the press and the impact that those had on traffic management and internal decision-making processes. For those that had experience engaging with regulators, I would ask a series of questions to learn about the nature of the participant's relationship with regulatory bodies and his or her perceptions of the function and impact of the agencies' work on traffic management. I placed a particular focus on formal and informal constraints that were most relevant given each interviewee's professional experiences and how these constraints shaped participants' perceptions of the regulator as adversarial or consensual.

Interviews with public sector officials covered a similar selection of topics but were focused more on interviewees' perspectives on traffic management in the industry at large and on their experiences during particular regulatory events. I asked participants to reflect about why regulatory agencies had pursued particular courses and the nature and frequency of interaction with the broadband industry. We also discussed how a number of institutional factors – the authority and style of the agency chair, the strength of the agency's jurisdiction to act, the ability for companies to challenge agency decisions in court, and formal and informal agency accountability mechanisms – shaped regulatory outcomes, broadly pursuing themes inspired by a variety of theories discussed in Chapter 3 related to institutional design, external forces, and internal characteristics.

I concluded all of the interviews by asking about any topics not previously discussed (which often elicited substantial further commentary from the participant), requesting suggestions for others to interview, and thanking the participants for their time. I took time as soon as possible after each interview to document my own reactions to the substance and flow of the conversation in my interview journal.

My goal was to reach data saturation, and in each country I felt that was accomplished with respect to the both the dominant themes that I planned ahead of time to investigate and many of those that emerged during the process. By the end of each round of interviews, I was often hearing similar answers and observations in response to similar questions.

Data Handling

I transcribed all interview recordings using F5, a software package. Transcription was primarily word-for-word, although speech fillers, verbal stumbles, and interview portions that were far outside the scope of the study were not included in the transcripts. I approached transcription as a vital part of the analytic process, a means to gain intimate knowledge of my data, and a mechanism for early-stage recognition of emerging themes. I included my own mental notes in the transcripts (delimited with square brackets) at points where participants were particularly insightful or where I felt that I could draw a conclusion relevant to my research questions or hypotheses.

Since all quoted portions of transcripts were anonymously attributed, I did not seek to have interviewees expend time and effort to review their own transcripts or approve the use of individual quotes. The full interview recordings remained available to me throughout the process of writing the thesis and I frequently revisited the recordings to obtain complete assurance that quotes were being reported accurately. Some interviewees used profanity, and I chose to report their words exactly as they were spoken. The interview recordings will be deleted upon publication of this thesis.

4.3.3 Documentation

In this study documents served as an important source of data to help reconstruct the realities of industry and public sector postures related to traffic management and to help understand the nature of the relationship between regulators and industry. The documentary corpus was composed primarily of official company and government documents. These documents are of value not because they necessarily give objective accounts of the traffic management landscape, but rather "precisely because of the biases they reveal" (Bryman 2008, 521). The goal of combining documentary analysis with the other methods was to provide a basis for interrogating the messages that companies and governments use in public forums against the information they reveal in other ways. It is the aggregated insights drawn from all the data sources that together allow for drawing explanatory conclusions.

United States

The series of FCC proceedings that were relevant to traffic management and discrimination provided a structured document set from which to sample documents for analysis in the US portion of this study. I chose the three most relevant proceedings as the basis for document sampling: the *Broadband Industry Practices* inquiry, the *Comcast* proceeding, and the *Open Internet* proceeding. The *Open Internet* proceeding provided six separate opportunities for public comment: comment and reply comment periods for the initial Notice of Proposed Rulemaking, the regulatory reclassification of broadband, and a further inquiry into specialized services and mobile broadband. I sampled comments and reply comments from the first and last of these, but not from the reclassification inquiry as it was primarily concerned with interpretations of statutes and legal precedents.

For each proceeding, I included all of the comments and reply comments filed by the six broadband companies whose employees I had also interviewed. I also surveyed all of the "ex parte" filings made by each company in each proceeding to document meetings between company staff and FCC staff or to introduce further evidence into the proceeding record. Since the ex parte filings are much shorter than the comments, I reviewed each ex parte filing and included it in the corpus as long as it addressed traffic management in some way (ex parte filings often document wide-ranging meetings that span many topics).

To cover the perspective of FCC officials, I included in the corpus all of the Commission's notices and orders in the three proceedings. I also reviewed the relevant portions of Commission statements from a number of other agency activities, including the *Wireline Broadband* proceeding, the *Cable Modem* proceeding, and all of the large cable and telco merger proceedings from the late 1990s through the mid-2000s. Since individual commissioners often file their own statements when the Commission begins or ends a proceeding, I included in the corpus all of the commissioners' individual statements from this larger set of proceedings and the three key proceedings.

Several public hearings were held as part of the *Comcast* proceeding, therefore I also included commissioners' statements and testimony from those hearings. During the *Open Internet* proceeding, the FCC convened a series of meetings of its Technical Advisory Group, which included engineers representing a diversity of stakeholder interests. Summaries of those meetings were made public and included in the corpus. On a number of occasions throughout the 2000s, commissioners articulated key policy proposals in speeches (for example, Michael Powell's "Four Freedoms" speech in 2004), and I included those in the corpus as well.

In addition to the broadband provider and FCC documents, I included a selection of FCC filings from other parties, technical documentation from broadband providers and equipment vendors, press articles about salient events related to specific traffic management or regulatory incidents, online forum and wiki postings from broadband customers discussing traffic management, and broadband provider terms and conditions. These documents were primarily selected on the basis of having been referenced in an FCC filing that I had previously reviewed or having been mentioned by an interviewee. Therefore the document sample reflected a small amount of snowball sampling.

United Kingdom

The UK corpus was comprised of documentation produced by industry, Ofcom, and government, plus a selection of supporting documents from a variety of sources.

For each of the six large broadband providers whose employees I had also interviewed, I collected their filings from the EU telecoms package consultations, Ofcom's public consultation on net neutrality, and the UK government's consultation on the implementation of the telecoms package. Because the volume of commentary dedicated specifically to traffic management in these consultations was far less than in the FCC proceedings, I supplemented this set of documentation with relevant material from the companies' web sites, their terms of service documents, and corporate presentations from conferences and to investors.

I also reviewed the companies' Key Facts Indicators (KFIs) – disclosures about traffic management made as part of an initiative launched in 2010 within the Broadband Stakeholder Group (BSG), an industry self-regulatory body. I included the BSG's responses to the consultations listed above, as well as the responses of the UK Internet Service Providers' Association (ISPA) and the European Competitive Telecommunications Association (ECTA), two trade associations that represent subsets of UK broadband providers.

From Ofcom I included speeches and topic briefings relevant to net neutrality that were published prior to the publication of its net neutrality discussion document, the discussion document itself, and the agency's report published at the conclusion of the discussion. I also reviewed minutes from Ofcom board meetings where net neutrality or traffic management was discussed.

UK government documents in the corpus included the government's submissions to EU consultations on the telecoms package reform and on net neutrality and the government's approach, consultation, and response regarding its implementation of the telecoms package. Where other documents such as briefing papers or reports related to the Digital Economy Act touched on net neutrality, I included those as well.

As in the US, I collected a variety of further supporting documents. These included relevant press articles, consultant and analyst reports, blog posts, and postings by broadband users in online forums. Because of the volume of traffic generated by the BBC's iPlayer service and the BBC's somewhat outspoken role in the UK net neutrality discourse, I also included a selection of writings from both technology and policy staff at the BBC.

4.4 Data Analysis

Based on my three data collection methods, my data set was comprised of interview transcripts, my interview journal, documents, and field notes. I performed thematic analysis (G. W. Ryan and Bernard 2003) on each of these corpuses in turn, followed by further rounds of analysis to triangulate across the data sources and to compare between the two countries.

I used my research questions, hypotheses, and their theoretical underpinnings to derive an initial set of themes (represented as codes) at the beginning of the analysis. The initial codes were focused on the technical, policy, economic, and social rationales for various traffic management approaches; the role of the market, competition, and being able to switch broadband providers; the role of existing regulatory frameworks; and the relationship between industry and government.

The majority of the codes, however, developed throughout the analysis, again reflecting a combined deductive-inductive approach. I used a hierarchical code structure and went through a process of code expansion as more nuance emerged within particular themes. For example, many of the documents discussed the considerations that go (or should go) into traffic management decisions. Over the course of the analysis, this code class was expanded to include specific sub-codes concerning technical parameters (congestion, latency, efficiency, bottleneck locations on the network), customers' responses to traffic management, regulatory constraints, the capital and operational expenditures necessary to support traffic management, and a number of other considerations. Through each round of the analysis I revisited the code structure and made adjustments to reflect the patterns emerging based on new data sources. I also coded for document structure and language in some instances, noting which arguments were presented prominently, keeping track of specific examples, and recording the use of recurring metaphors.

I began the coding process using the NVivo software package, but due to instability issues with the software, I ultimately built my own relational database in MySQL using NavicatLite to store all information related to documents, codes, and coded text. I allocated separate database tables for the US and UK data, allowing me to run separate or joint queries on the data set as necessary for my analysis.

4.5 Ethical Considerations

The interviews and participant observation conducted for this thesis at times involved discussions about sensitive corporate and regulatory information. So as to protect the interviewees and informants from any potential harm that could result from their participation in the study and the publication of the thesis, I took a series of precautions. I made clear to every participant at the outset of our interaction that the thesis would be made public and that any quoted material would be anonymously attributed. I obtained the consent of interviewees before recording the interviews and did not record in cases where they declined. A number of interviewees requested that the recording be stopped for certain portions of the interview, or that particular words that they had spoken not be quoted, and I obliged all of those requests. With every quote that I did include in the thesis, I thought carefully about the selection of the quoted words and the characterization of the speaker ("cable executive" or "former FCC official," for example) so as to ensure that their words could not be used to inadvertently identify them by readers with background knowledge of the events discussed. All speaker characterizations were chosen to encompass broad groups that include a large number of people so as to eliminate the chance that any speaker could be uniquely identified.

I also took caution when reporting company-specific information that had previously not been publicly disclosed in the press or by the companies themselves, regardless of whether the information was reported in the thesis as a direct quote or not. I did not attribute this information to individual companies, but rather to the industries to which they belonged (US cable companies or UK DSL operators, for example). The goal of the thesis was not to reveal specific information about particular companies, but to characterize traffic management trends among different parts of the broadband industry in each country. As such, describing traffic management practices by industry sector strikes the appropriate balance between protecting potentially sensitive corporate information and contributing novel findings about industry-wide traffic management practices.

Chapter 5.
Traffic Management Decisions in the United States

5.1 Introduction

Since it became popular among the American public, fixed-line broadband Internet service has been primarily offered without application-specific network management. While some cable operators adopted specific techniques to manage peer-to-peer traffic, those practices largely ceased in the wake of the FCC's *Comcast Order* (FCC 2008). The large telephone companies did not make similar forays early on, and as the net neutrality debate intensified over time, they had even less incentive to pursue network management solutions that involved managing particular Internet applications.

This result was not solely the product of engineering designs or regulatory imperatives, but of the intermingled forces of both. The cable operators that deployed peer-to-peer management techniques were responding to a very specific engineering challenge that years of emergency bandwidth capacity upgrades showed no signs of fixing. But the context in which they made their solution choices was highly shaped by their unique regulatory history – or lack thereof. Compared to the telcos, the cable broadband operators had existed in a highly favorable and recently established regulatory environment prior to their adoption of peer-to-peer management solutions. For some operators, this environment imbued organizational paradigms for traffic management decision-making with a flexible, freewheeling character. It was under these circumstances that the solution that was ultimately denounced by the FCC was chosen and deployed. Only in response to the *Comcast Order* did these operators begin to institutionalize broader internal review and oversight of their traffic management decisions.

The telcos' networks were based on different access technologies with different characteristics. Their need to intervene to manage specific Internet applications was less urgent and they therefore largely refrained from doing so. But the telcos also had a radically different orientation toward their regulatory environment. They had fought for decades to have common carriage obligations lifted from their data and broadband services offerings, finally achieving that goal in 2005. They had good reason to tread cautiously so as not to upend these hard-fought gains. They had also accrued thorough regulatory review structures – policies, personnel placements, and means for internal cross-departmental dialogue – that they readily applied to internal traffic management debates. Throughout the entire history of their broadband offerings, the telcos' traffic management decisions were grounded in the combination of the unique engineering constraints of DSL and fiber mediums and deeply embedded organizational processes to address regulatory risk.

This chapter analyzes in detail how the commingled forces and technology and regulation shaped traffic management decisions and decision-making in the United States. The combined decisions of the cable and telephone operators have resulted in fixed broadband networks on which over-the-top Internet applications have not generally been managed, in stark contrast to the United Kingdom.

5.2 Regulatory History and the *Policy Statement* in Context

As discussed in Chapter 1, by the mid-2000s the telephone companies had been engaged in nearly 40 years' worth of uphill battles to relieve themselves of regulatory obligations on data services, with each successive round of the *Computer Inquiries* resulting in the removal of additional regulations. Cable operators, by contrast, had never been regulated as common carriers during the short existence of cable broadband service. The two industries' contrasting experiences provided the backdrop for a number of events in the summer of 2005 that profoundly altered the landscape of broadband Internet regulation in the United States.

First, the Supreme Court handed down its decision in the *Brand X* case (National Cable & Telecommunications Assn. v. Brand X Internet Services 545 U.S. 967 (2005)), classifying broadband Internet service offered over cable networks as an information service not subject to common carrier obligations. Shortly thereafter, in its *Wireline Broadband Order* (FCC 2005b), the FCC ascribed the same classification to wireline broadband Internet service (offered via DSL), creating parity with cable and ending the telephone industry's decades-long struggle to free its provision of data and broadband Internet services from regulation under Title II. To counterbalance this deregulatory step, the FCC simultaneously issued its *Broadband Policy Statement* (FCC 2005c), laying out a framework of principles declaring broadband consumers' rights to access the legal content and services of their choice, to connect the legal devices of their choice to the network, and to enjoy competition among broadband service providers. In making all four principles subject to "reasonable network management," the FCC formally recognized broadband network management for the first time.

The immediate operational effects of the *Policy Statement* were few. For the major fixed-line network operators, the *Policy Statement* largely reflected the status quo of industry practices at the time – it was "how the market had evolved in the US anyway," as one telco employee described it. Broadband offerings at the time were not characterized by significant limits on the content or services that customers could access or the devices they could attach to the network. Providers' previous restrictions on certain applications, in particular virtual private networking, had largely been lifted in the years before the *Policy Statement* was adopted (Coalition of Broadband Users and Innovators 2003; National Cable and Telecommunications Association 2003). As a result, there were few immediate consequences of the *Policy Statement* – none of the large operators needed to make drastic changes to the way they ran their networks in order to feel confident that they were in compliance with the *Policy Statement* in the months after it was adopted.

While the *Policy Statement* may have been simple to operationalize at the start, there was significant uncertainty about how it would be interpreted and enforced going forward. The ways in which the cable and telephone industries approached this uncertainty were shaped in significant part by their respective legacies of regulatory oversight at the Commission. The *Policy Statement* did not arise out of a vacuum, but rather was a by-product of a series of struggles in the courts and at the FCC. The same can be said for the respective mindsets that the cable and telephone companies took to interpreting the *Policy Statement* and making traffic management decisions as time wore on.

The greatest uncertainty was about whether the *Policy Statement* was enforceable. When the FCC published the *Statement*, Chairman Martin's accompanying press release explained that "policy statements do not establish rules nor are they enforceable documents" (Martin 2005a). Commissioner Copps, the FCC's most forceful advocate for nondiscrimination rules, was more equivocal, but still left some uncertainty as to his opinion of the *Statement*'s enforceability:

> While I would have preferred a rule that we could use to bring an enforcement action, this is a critical step. And with violations of our own policy, I will take the next step and push for Commission action. A line has been drawn in the sand. (Copps 2005)

The question of enforceability eventually would prove to be the most pivotal of the US net neutrality debate.

Perhaps only second to enforceability in its lack of clarity was the meaning ascribed to "reasonable network management." The *Policy Statement* provided no further explanation as to what "reasonable" might mean in the context of network management, nor which practices or kinds of practices constituted network management functions in the eyes of the Commission. The network management provision was, after all, in a footnote. The focus of the entire effort in developing the principles lay elsewhere, as explained by a former FCC staffer who had been at the Commission during the drafting of the Four Freedoms:

> We thought at the time, in 2004, that the real action in that four freedom document was in the first freedom, right? This is the consumer entitlement to use applications and content of their choice. You know if you'd asked us then where we were going to have to adjudicate or where we were going to have to enforce, it was all going to be around that. It's only when I guess Kevin Martin looked at the *Comcast* case that all of those anxieties around the way consumers would access lawful content were poured into the network management bucket. . . . But we didn't think at that time that that's where the enforcement focus would be.

Network management was not the focal point of the Commission's action, and therefore was not well specified.

Thus, although most network operators could view the *Policy Statement* as largely mirroring existing industry practices, they were still faced with significant uncertainty as to how the Commission might react and whether enforcement action was a possibility should they decide to change their practices, particularly their traffic management practices. As the next section shows, the approaches that the cable and telephone industries took to this uncertainty were highly influenced by their respective regulatory histories. The cable broadband industry's short and successful record of dealings at the Commission allowed for the creation of corporate environments where risk-taking and deployment of new traffic management approaches seemed justified under broad interpretations of "reasonable network management." The telcos' hard-fought deregulatory status created the opposite effect: caution so as to prove that fears about lack of FCC oversight were unfounded, and a sense of the need to examine thoroughly the potential regulatory consequences of any new practices on the network.

5.3 Industry Approaches to Traffic Management

By the mid-2000s, US broadband subscribership was growing steadily. Nationwide, year-over-year growth in subscribership was 30-50% (see, for example, data collected by the FCC (2004)). Traffic volumes, both on average and during the peak usage period (between when work and school let out and when people went to sleep), were likewise growing steadily, with reports of 25-35% year-over-year growth (Erman et al. 2009; J. K. Smith 2008).

On many networks, a disproportionate amount of traffic growth was attributable to peer-to-peer file-sharing traffic. Popular peer-to-peer applications are characterized by their ability to efficiently transfer large files unattended: users can configure their peer-to-peer clients to download specific files in advance and then leave those clients running in the background without the need for user interaction. This means that even if the fraction of a network's broadband subscribers using peer-to-peer applications is small, the proportion of peer-to-peer traffic can be large. Furthermore, peer-to-peer networks only function if peers both upload and download, because peers are the sources of all files to be shared. As a result, peer-to-peer traffic tends to be roughly symmetric, with comparable amounts of traffic being sent in the downstream and upstream directions.

Few public studies of the impact of peer-to-peer traffic on broadband networks in the US in the mid-2000s exist, but evidence from other countries suggests that peer-to-peer traffic was accounting for approximately 60% of residential downstream broadband traffic in that time frame (Cho et al. 2006; Plissoneau, Costeaux, and Brown 2005). Because residential broadband is generally offered as an asymmetric service with downstream bandwidth many times that of upstream bandwidth, peer-to-peer traffic was having an even more disproportionate impact in the upstream direction (Martin and Westall 2007). Several cable operators interviewed for this thesis reported peer-to-peer upstream traffic levels in the range of 80% of total traffic on cable networks in the mid-2000s.

The cable and telephone industries reacted differently to the challenges that the growth of peer-to-peer traffic presented. Numerous stakeholders have pointed to differences in access technology to explain the differential responses – that the cable companies, with more shared infrastructure closer to end customers and less upstream bandwidth capacity than either DSL or fiber networks, saw more intense effects from peer-to-peer traffic (Glover et al. 2008; Sandvine 2010b; Teplitz et al. 2010). Interviewees often explained cable's more drastic efforts to control peer-to-peer traffic in light of these differences.

However, technology differences were far from the only factors that distinguished the cable and telco approaches to managing peer-to-peer and network traffic in general. The cable companies' more cavalier approach to regulatory ambiguity created greater flexibility as to how traffic management decisions were made and opened the door for adoption of peer-to-peer-specific management solutions. While the technical characteristics of the telcos' infrastructure put less pressure on them to deal with growing bandwidth demands in novel ways, they also took far more care in considering the regulatory and public perception consequences of application-specific management, ultimately shying away from it for the most part. These distinctions in approach are explored in the next two sections.

5.3.1 Cable Industry

Managing Peer-to-Peer Traffic

There is no doubt that peer-to-peer traffic was putting serious strains on cable broadband networks in the 2004-2005 time frame, particularly in the upstream direction. Cable interviewees described how freshly upgraded network links would reach 80-90% utilization and require new upgrades in a matter of months. In some instances bandwidth requirements doubled every year for several years.

Cable broadband relies on significant shared infrastructure in the access network. Dozens of subscribers in the same neighborhood may share the same local fiber optic node, and hundreds of subscribers may share the same Internet Protocol port further up in the network. Popular peer-to-peer file-sharing applications place a particular strain on this kind of architecture because they were designed to maximize the amount of data being exchanged at any one time, in part by opening many simultaneous connections to other peers. The result in the mid-2000s was that with even a small number of avid peer-to-peer users sharing a particular network link, bandwidth was quickly becoming saturated as it was rolled out. One cable provider described the urgent tenor of the situation as follows:

> The example that comes to my mind is the case in which there was a university town . . . and the kids had come back to school somewhere around late August/early September time frame. And [the network engineers] saw this huge influx of traffic, even much more so than they would have expected after the kids have left at the end of the spring semester. I think there were emergency crews out there putting out new fiber nodes, splitting networks, and doing it all on an emergency basis because the network had quickly gotten bogged down with a traffic load that it couldn't handle. And it was a traffic load that it could have handled with the same population of people, say, three months before. And I think that that was the first time that, at least in my experience, the reality of sudden congestion or sudden rise of traffic started to become very real.

The upstream contention problem was acute for cable. Performance of other applications was suffering as a result. Some operators noted that upstream contention was taking a disproportionate toll on real-time applications such as VoIP that require minimal latency in both directions in order to provide reasonable call quality. Martin and Westall (2007) found that just 15 active BitTorrent users on a cable link shared among 400 total users could cause VoIP call quality to fall below a usable performance threshold. Operators were viewing the growth of upstream traffic as unsustainable from a capacity planning perspective. They were in search of a fix.

One solution that gained significant traction involved deploying deep packet inspection (DPI) technology to limit the number of upstream peer-to-peer connections at particular points on the network. This was the approach that Comcast pursued that eventually became subject to regulatory scrutiny. Sandvine, a network equipment vendor, supplied the DPI equipment that Comcast deployed in an "out-of-line" fashion: Internet traffic was copied and sent to the Sandvine devices for inspection, rather than having the devices sit directly between end users and the rest of the Internet (Zachem 2008). The equipment was configured to identify protocols associated with heaviest traffic usage, all of which were peer-to-peer protocols, including BitTorrent. When the number of upstream-only connections associated with any of the protocols reached a pre-defined threshold in a particular location on the network (indicating potential congestion on the upstream), the device would send a TCP "reset" packet to both sides of the peer-to-peer file exchange, causing the exchange to cease. This solution approach was viewed by many observers as discriminatory because it involved selecting traffic associated with particular Internet applications for different treatment from all other traffic (Ammori et al. 2007a; Copps 2008; Goldberg, Kumar, and Monahan 2007).

Comcast first trialed the Sandvine equipment in May 2005 and moved to network-wide commercial deployment that lasted from 2006 to 2008 (Zachem 2008). According to numerous interviewees, a number of other cable operators pursued this or very similar application-specific approaches within a similar time frame. The limited existing network-based research into detectable interference with BitTorrent traffic supports this conclusion as well (Dischinger et al. 2008).

In some cases, the configurations of the devices on the network and the identified protocols differed. For example, the cable provider RCN used Sandvine equipment to manage peer-to-peer traffic in two different ways that did not involve TCP reset packets (Kiddoo 2010). First, RCN redirected requests for peer-to-peer downloads from other networks onto its own network to avoid having to pay the transit costs associated with transferring the data from another network to one of its own customers. Second, the company limited the number of simultaneous upstream peer-to-peer connections that could occur within each geographic market in which the ISP operated. Because the Sandvine devices were installed "in-line" – directly between subscribers and the rest of the Internet – these connection limits could be enforced without the need to send TCP reset packets. Additional upstream requests above the pre-determined threshold could simply be blocked by the Sandvine device. Interviewees likewise confirmed that deployments varied between operators and that some providers chose to use the in-line approach to peer-to-peer management rather than having individual connections terminated via a TCP reset. The out-of-line approach was much easier for subscribers and other external parties to detect because the TCP reset packets had an easily identifiable signature.

According to interviewees and available public information, application-agnostic approaches to managing network traffic, such as usage-based billing or tactics based on per-user traffic volume, do not appear to have been widely deployed. One solution, as described by a cable interviewee, took a less application-focused approach, but could still be considered "application-specific" under the definition of that term provided in Chapter 1. The solution involved identifying traffic based on characteristics other than the application with which it was associated. In this approach, routers on the cable network would identify the network to which each packet was destined. If the destination network was another residential network (as opposed to a network used primarily to host web sites, for example), that was the first clue that the traffic was likely peer-to-peer. If the packet's size and transport protocol were also determined to be characteristic of peer-to-peer file transfers, the packet was then assumed to be part of a peer-to-peer exchange. Traffic flows containing such packets were then rate-limited in the upstream, helping to control upstream utilization.

This approach did not involve identification and differential treatment of specific applications, but it did discriminate on the basis of the destination of the traffic. Furthermore, some of the effects were likely similar to approaches that did single out specific peer-to-peer applications. Thus this example fits with the broader trend of large cable operators pursuing application-specific traffic management solutions. And as described by the cable interviewee, it was deployed in tandem on the network with another solution that did target specific peer-to-peer applications using Sandvine equipment.

Motivations Behind Application-Specific Management Decisions

By the time applications-specific solutions were being deployed, there was a growing debate in the press and in Congress about net neutrality and the potential for discriminatory conduct in the wake of the FCC's deregulatory activities. Why would cable operators take up application-specific traffic management solutions under these circumstances?

One feature that all of the approaches described above had in common and that made the TCP reset approach particularly compelling was that they were viewed as highly targeted, pre-packaged solutions that could alleviate, on relatively short order, what was becoming an emergency bandwidth problem for cable operators. The theory behind the TCP reset approach was to narrowly focus on the network's pain point – high upstream utilization caused in large part by peer-to-peer traffic – while having a minimal impact on the peers whose file transfers were interrupted. Most peer-to-peer clients are designed to automatically re-start a download when it gets interrupted (assuming another peer elsewhere has made the same file or file portion available). Assuming that the downloading peer would re-start by connecting to a peer on a different ISP or in a less congested part of the cable operator's network, the TCP reset approach would have the effect of moving upstream peer-to-peer traffic out of the trouble spot while introducing only a small delay in the time it took for the downloader to finish the file transfer. Uploaders were assumed to be mostly indifferent as to whether their uploads were interrupted or how quickly they finished.

Thus the TCP reset solution appeared likely to solve cable operators' specific problem immediately and persistently without significantly degrading the performance of peer-to-peer transfers. The other approaches described above – in-line peer-to-peer management and management based on traffic characteristics other than application – were similarly designed to narrowly focus on peer-to-peer traffic, although their performance effects may have been more variable since they involved dropping packets or preventing connections for some period of time rather than terminating entire peer-to-peer connections. To operators, all of these approaches seemed more focused on solving the existing problem than other conceivable nondiscriminatory approaches. Usage-based billing, for example, would have required accounting infrastructure and customer education without guaranteeing that it would actually cause heavy users to change their behavior or how long such a change might take. Capacity upgrades were ongoing, but they were perceived to be insufficient to keep up with traffic demand given their cost. By comparison, the approaches that the cable operators pursued appeared to offer more certainty, more immediate results, and more lasting efficacy.

Furthermore, solutions based on DPI were extremely attractive to the operators' marketing teams, which saw great appeal in using DPI not only for traffic management, but to understand how their customers were making use of their Internet service and how that usage might be parlayed into new product opportunities. At several ISPs, marketing teams purchased the equipment with little oversight from engineering staff. As one interviewee explained:

> [T]he people on the engineering side, if you ask them about what the experience of . . . installing [the DPI] was, there wasn't any of the typical consultation with engineering. It was, "We've decided to buy this system, you'll be installing it. Only you guys will know about it. Here's the budget, please get it done as soon as possible." And [engineering] looking at it and being like, "What the fuck is this? How do we configure this? What is this?" They had no vote in it at all, no influence.

Thus in some cases the fact that the DPI platforms were capable of providing marketing units with intelligence about network usage was an even more compelling reason to make a DPI purchase than any reason related to traffic management. As Marsden (2010) has argued, DPI's multiple uses make it an attractive investment. Nondiscriminatory traffic management solutions by their very nature could not have provided the same kinds of application-level insights as DPI-based platforms.

Finally, cable operators took a somewhat cavalier approach to their regulatory circumstances. Cable interviewees emphasized the streamlined structures of their organizations, describing operations as "bootstrapped" and "lean" with a "fly-by-the-seat-of-your-pants" culture. The cable operators had (compared to the telcos) relatively few personnel resources dedicated to developing interpretations of ambiguous regulatory policy and applying those interpretations to internal traffic management decision-making. In some cases regulatory staff exercised an extremely light touch, as exemplified by the following exchange with a cable public policy executive:

> Interviewer: [W]hat was the timing of your role – at the end of something you would come in and say that we should think about the public policy implications of this? Or was it kind of earlier? . . . How did it integrate into a decision to roll something out on the network? Or did it integrate at all?
>
> Respondent: Generally no. Other than ensuring that the business units and the engineering teams were acquainted first with the Four Freedoms as articulated by Michael Powell. And we always felt comfortable in asserting that we were operating our network consistent with those freedoms. And then in the *Policy Statement* articulated by Kevin Martin and his Commission, you know that consumers could go anywhere on the Internet that they wanted, use any content, connect any device that didn't interfere with the network . . . . But other than ensuring that those principles were known and understood and part of our daily practice, we didn't have much occasion to discuss specific network management decisions.

Even in cases where regulatory and public policy staff had the opportunity to develop internally their conceptions of "reasonable network management," those conceptions were sometimes quite broad. For example, some cable interviewees expressed that after the *Policy Statement* took effect, they assumed that any practice that served the purpose of providing a good user experience, or that was not anti-competitive, should be considered "reasonable."

The mere existence of the carve out for network management, and in a statement of questionable enforceability no less, gave some operators the confidence to pursue application-specific approaches. One cable interviewee explained this in discussing his company's dealings with Chairman Martin's staff regarding the reasonable network management footnote:

> [T]hat footnote was very important to us and we worked with them on that. And also I don't think it was in the document itself, but maybe it was in his press release or his statement where he specifically said, "this isn't enforceable." So those two things gave us a fair amount of comfort. Everything that was in there we felt that we were already doing at the time. So we were thinking that this set a framework, put things to rest.

During the *Open Internet* proceeding, a number of cable operators provided clear articulations of just how broadly they conceived of the term "reasonable." Cox exhorted the Commission to "establish a presumption that properly disclosed network management practices are reasonable, rebuttable only by evidence that the management tools are an artifice for anti-competitive conduct" (Wilson et al. 2010, 32). Time Warner Cable argued that the definition of "reasonable network management" should consist of any practice employed to reduce congestion, prevent the transfer of unlawful or unwanted traffic, or "any other network management practice that is intended to improve service quality or performance rather than to achieve any anti-competitive objective" (Teplitz et al. 2010, 71). Where cable operators were internally assuming such standards under the *Policy Statement*, the application-specific approaches seemed justifiable.

In short, the cable industry was facing a significant upstream engineering problem that even an accelerated schedule of capacity upgrades was not able to fully handle. A solution presented itself that was not only narrowly targeted, but had additional benefits in the eyes of marketing teams, some of whom made the actual equipment purchases. All of this came to pass in the absence of rigorous oversight from policy and regulatory staff.

Where policy and regulatory oversight did exist, interpretations of existing regulatory ambiguity and the likelihood of regulatory intervention gave operators tremendous leeway in choosing traffic management solutions. Such an approach to regulatory risk made sense in the context of the cable industry's brief and successful regulatory history. Cable broadband operators had never been subject to nondiscrimination obligations under Title II. Just weeks after Comcast began its Sandvine trial in 2005, the Supreme Court affirmed this regulatory classification. The cable industry could logically conclude that discriminatory traffic management approaches, including the TCP reset solution, were well within the bounds of what the FCC might consider "reasonable." That conclusion ultimately proved to be untrue and has shaped broadband Internet regulation ever since.

5.3.2 Telephone Industry

Absence of Application-Specific Management

By and large, the telcos refrained from pursuing the kinds of application-specific traffic management approaches that cable companies had deployed. This difference in approach was partly motivated by differences in access technologies – the urgent need to find a palliative for peer-to-peer upstream congestion simply did not materialize on DSL and fiber networks. But the telcos' decisions were likewise shaped by their specific regulatory history and obligations, just as the lack of such factors shaped the decisions of the cable companies.

The telcos were seeing the same kinds of broadband traffic growth in the early and mid-2000s as the cable companies saw, including surges in peer-to-peer traffic, but their networks bore the brunt of the growth differently. Residential DSL networks offer each subscriber a dedicated link in the access network, pushing the point at which multiple subscribers' traffic gets aggregated further up in the network where greater bandwidth is available to manage the aggregate load. As a result, DSL subscribers were much less likely to see the performance of their access network connections suffer when a small number of subscribers in the neighborhood chose to make aggressive use of peer-to-peer applications. Furthermore, DSL offerings tended to be less asymmetric than cable offerings, with the ratio of advertised upstream to downstream bandwidth around 1:4, whereas for cable it tended to be closer to

1:10 (Greenstein and McDevitt 2011). Thus the specific impact of the growth of peer-to-peer on upstream bandwidth was more muted for the telcos.

In addition to offering DSL broadband, beginning in 2003 the nation's two largest telephone companies began deploying new broadband products that supported video, voice, and Internet services over fiber networks: AT&T's U-Verse and Verizon's FiOS. These services provided tens to hundreds of times the access network capacity offered by DSL or cable at the time, giving the operators a much larger bandwidth buffer with which to handle traffic upsurges and relieving them of the urge to pursue fine-grained management of individual Internet applications.

In the absence of the kind of emergency bandwidth crunch that the cable companies were experiencing, the telcos were left without a meaningful impetus to begin managing peer-to-peer or any other specific Internet applications on their DSL networks. Overall, there is little evidence that the large telcos were discriminating against Internet applications in any significant way from the mid-2000s through the end of the decade. All telco interviewees emphasized this point.

It may be tempting to explain away these claims by way of interviewees' wariness to admit to engaging in practices that could reflect poorly on their companies, but their claims provide a sharp contrast to those from cable employees, many of whom openly discussed discriminatory practices even in cases where those practices had not attracted public attention. Telco interviewees made convincing arguments, often citing specific corporate executives, policies, or decision instances that governed traffic management choices. One telco executive put it this way: "What I can say, and as a company [we] have stated this, is that we do not impede any public Internet traffic. Our opinion is a bit is a bit."

In the case of fiber deployments, network operators did take steps to manage the interaction between their own IP-based video offerings and Internet traffic, but did not apply those management techniques to Internet applications themselves. The video services offered as part of the fiber deployments were designed to compete with cable television service, and as such the telcos sought to match the user experience of a customer clicking through cable channels. But unlike cable, the fiber networks were engineered to have capacity shared between Internet traffic and linear video and/or video-on-demand traffic. To ensure that video traffic would not squeeze out Internet bandwidth in a household with multiple people watching different HD or SD programs in different rooms, the network operators adopted management strategies that involved prioritizing their own video streams over Internet traffic and always reserving a portion of capacity dedicated to Internet traffic (AT&T 2012; Verizon 2012). These management strategies had the potential to impact a subscriber's Internet service when multiple video streams were in use, but they involved no specific management of Internet traffic. This was a deliberate choice that the network operators made.

Motivations Behind Traffic Management Choices

These choices were shaped as much by regulatory circumstances as by technological design. Compared to cable in the mid-2000s, the telcos arguably had a higher risk of having an enforcement action brought against them on the basis of a claim of discriminatory conduct. Until August 2005, the telcos were operating in many cases under the nondiscrimination obligations of Title II and in all cases under the remaining obligations from the *Computer Inquiries*. It was far from clear what these obligations meant for traffic management – the *Madison River* enforcement action (FCC 2005a) was the only remotely relevant case concerning Internet applications, and it dealt with explicit blocking of all traffic associated with a particular application, providing little guidance about whether more nuanced approaches to traffic management might be considered discriminatory.

With the *Wireline Broadband Order* (FCC 2005b), the remaining obligations were lifted and replaced with the *Policy Statement* of questionable enforceability. But just months later, Verizon and AT&T – comprising 70% of the nation's DSL market (Noam 2009) – each agreed to abide by the *Policy Statement* for two years as conditions of their respective mergers with MCI and SBC (FCC 2005d; FCC 2005e). While they may have had no more clarity than the cable operators about what “reasonable” network management was, they had far more certainty that if they were seen to be violating the principles of the *Policy Statement*, the FCC could take action against them.

In AT&T’s case, the FCC extended that commitment even longer and imposed further net neutrality conditions when the company acquired BellSouth in late 2006 (FCC 2006). Although the new conditions, which lasted until the end of 2008, did not speak to traffic management specifically, they required that the merged entity “maintain a neutral network and neutral routing in its wireline broadband Internet service,” (FCC 2006, 154) further increasing the risk that the FCC might take action should it learn of “non-neutral” practices.

As a result, the risk of regulatory enforcement could never be confidently disregarded. This risk did not completely quash the impetus to pursue traffic management approaches that involved differential treatment of different services – otherwise the video prioritization and bandwidth reservation techniques employed on fiber networks never would have materialized. But it did result in the institutionalization of corporate practices designed to mitigate regulatory risk. These organizational processes were integral to the telcos’ decisions not to pursue application-specific management, and they reflect a stark contrast in approach to regulatory risk as compared to the cable operators.

Within the telephone companies, regulatory review was thoroughly embedded in the process of traffic management decision-making. As one telco engineer explained, “our policy folks always have their antennas up” on how developments in the policy space might affect activities on the engineering side of the company. One policy executive likewise explained that as a policy team, “ideally you come in as early as possible or practicable” to review new product developments. Another policy-focused interviewee described how “every time they rolled out a new feature . . . it was pretty much par for the course that [the engineers] would bring me in and we would talk to them and advise them on whether we thought there were issues." Internal regulatory oversight is one feature that all the large telcos appear to have had in common.

This oversight was not only manifest in the usual exchanges between regulatory or policy teams and other corporate divisions, but also in bureaucratic forms that were developed even before net neutrality became an issue of public debate. One company convened an internal policy council that brought together regulatory, policy, marketing, legal, and technical staff to discuss traffic management decisions and ensure company-wide understanding of new traffic management capabilities. Another operator had policy personnel embedded from the very beginning within the team developing new broadband products so that a policy perspective could be provided as engineering, architectural, and business decisions were made along the way. Regulatory and policy oversight was enshrined in corporate organizational structures.

Might decades of regulation provide some explanation for why the telcos invested this effort in regulatory review? Year after year of operating under a considerable regulatory regime engendered a sensitivity to regulatory risk within these companies that shaped them profoundly, right down to the ways in which policy personnel were organized internally. In some cases, the risk of regulation appears to have penetrated even further, down to the very core of corporate philosophy. A number of telco employees spoke about how corporate philosophy, as articulated by senior executives, circumscribed the set of traffic management choices available to engineering teams. The precise interplay of the forces of regulatory risk and corporate philosophy is difficult to discern – which came first, the chicken or the egg? But as one former senior telco executive explained, there is no doubt about their combined contribution to the decision not to manage Internet applications:

> It's been [this company's] philosophy for as long as I've been promoting and pushing this philosophy that our job is to build capacity to satisfy the customer. And when you take that and you couple it with the regulatory – I'm going to call it overhang but it's probably not the right word – of net neutrality concerns, threats, we have stayed as far away from attempting to modulate network traffic as possible. If we're going to for some reason have congestion in the network, everyone's going to suffer from that congestion until we can physically get it fixed, either through true hardware or optimization techniques to get everybody back together. So . . . we have done no traffic shaping, traffic blocking, if you will.

Such philosophies were naturally not uniform across the industry or its executives. Former AT&T CEO Ed Whitacre famously ignited controversy when suggesting that rather than getting a "free lunch," large applications providers should be required to pay for preferential traffic treatment (O'Connell 2005), and his rhetoric was echoed by others in the industry, including Verizon senior vice president John Thorne (Mohammed 2006) and Qwest CEO Richard Notebaert (Reardon 2006). But there was clearly a gap between this rhetoric and the reality of how the telcos actually chose to manage fixed network traffic. When it came to traffic management techniques that had no revenue potential attached, the notion of staying away from modulating particular applications clearly pervaded the corporate mindset in some cases.

In sum, the telcos' initial and enduring approach to traffic management – "a bit is a bit" – can be explained by the combination of technology, regulatory risk, and, in some cases, corporate philosophy. In some instances they had application-specific tools available that they used to manage the performance of their video services, but they chose not to extend those tools to over-the-top services. While the technical imperative to implement application-specific traffic management responses was lacking, that did not prevent the telcos from injecting regulatory oversight into traffic management discussions early and often, in stark contrast to large cable operators.

5.4 The *Comcast Order* and its Aftermath

Cable operators – or some cable operators, at least – were about to drastically revise their approaches, however. Comcast did not disclose its use of the Sandvine TCP reset solution when it was deployed on the network. In May 2007, Robb Topolski, a Comcast customer, wrote about his peer-to-peer connections getting reset in the DSL Reports online forum (Topolski 2007). Over the following months, a number of parties conducted investigations into interference with peer-to-peer traffic and published their findings online, culminating with an Associated Press (AP) article published in October 2007 that documented nationwide tests the AP ran to demonstrate Comcast's use of the TCP reset technique. The following month, public interest groups and others petitioned the FCC to take action in response. After months of dramatic public debate (discussed in Chapter 8), Comcast began developing FairShare, a "protocol-agnostic" approach to traffic management. In August 2008 the Commission ordered Comcast to cease its use of the TCP reset solution by the end of the calendar year. By January 2009, FairShare had replaced the TCP reset solution throughout the Comcast network.

Rather than managing peer-to-peer uploads, as the TCP reset solution had done, FairShare was designed to reduce the impact of heavy users during times of high network utilization. Comcast deployed software to monitor network utilization in the aggregate and on a per-subscriber basis (Zachem 2008). When a network segment was nearing a congested state, the software would determine if one or more subscribers had been using disproportionate bandwidth, and those subscribers' traffic would be assigned to a lower priority level than the rest of the subscribers sharing the network segment. Thus this approach mitigated the impact of heavy users on the rest of the user base without regard to which applications were in use or which applications might be the source of disproportionate traffic. And rather than terminating individual connections, it allowed all heavy users' traffic to be exchanged, but ensured that other users' traffic would get priority. When the heavy subscribers' usage returned to normal levels, their traffic would return to normal priority.

For other cable companies that had been using similar approaches to Comcast, the FCC proceeding and the public attention surrounding it by and large spelled the end of application-specific traffic management. RCN agreed to cease using its in-line peer-to-peer management techniques in 2009 as part of a class action settlement that arose as a result of the public attention that Comcast's traffic management had accrued (Kiddoo 2010). One cable executive explained the effects of the public debate as follows:

> [W]e were looking at pushing [TCP resets] further as the FCC and Comcast had their brouhaha. And it was determined that Comcast had to stop. And in that case, sort of as the technology policy intersects with Washington policy, the uncertainty that Washington caused put a hold on us deploying further, and eventually caused us to stop. So we stopped packet resets . . . at year end whenever that year was when the FCC had ordered Comcast to stop.

The Washington policy debate may have put an end to management focused on peer-to-peer applications, but that did not mean that the kinds of traffic demands that had spurred its adoption in the first place had gone away. Cable operators still needed to manage that demand, and by and large they turned back to capacity upgrades, opting to "let standard capacity planning go back in place," as one engineer put it. But operators did not necessarily find themselves back in emergency node-splitting mode.

Even as cable operators had been pursuing the TCP reset approach, a new technical standard for cable modem design, known as Data Over Cable Interface Specifications (DOCSIS) 3.0, was being finalized. By bonding together multiple cable channels for use in delivering Internet traffic, DOCSIS 3.0 could supply cable broadband networks with several times more capacity both upstream and downstream than what was available previously. It required that equipment both in subscribers' homes and in the cable network be upgraded, but once those upgrades occurred, increased bandwidth was immediately available.

At the same time, patterns of network usage were beginning to shift. Web-based streaming video sites like YouTube and Hulu were maturing and attracting increasing numbers of visitors. Cisco estimated that 20% of US Internet traffic in 2007 was YouTube traffic (Cisco 2008). The relative portion of network traffic attributable to peer-to-peer applications was shrinking as a result, at least in the downstream if not in the upstream (Erman et al. 2009). Heavy peer-to-peer use could still wreak havoc on an under-provisioned network link, but it was no longer the main driver of traffic growth.

Thanks to both of these developments – and with the TCP reset solution off the table – cable operators could return to more pure capacity-focused approaches to managing network load without heading straight back into the urgent upgrade cycles of years past. As one engineer enthusiastically explained, “DOCSIS 3.0 is the knight in shining armor that arrived onto the scene in the nick of time.”

This is not to say that other traffic management approaches were not discussed and tested. The most prominent example was that of Cox, which in 2009 ran a trial of a system that distinguished between what the company considered to be “time-sensitive” applications and “non-time-sensitive” applications and prioritized the former over the latter in the upstream direction (Wilson et al. 2010). During times of congestion, non-time-sensitive applications (peer-to-peer applications, newsgroups, and other similar bulk file transfer applications) would have their speeds reduced to ensure that the performance of time-sensitive applications (web, VoIP, streaming, and all other traffic not classified as non-time-sensitive) was unaffected. The trial ran for a limited number of months on a portion of the Cox network and was never deployed network-wide.

In sum, with the watchful eye – or raised eyebrow? – of the FCC looming over them, cable operators by and large abandoned application-specific management techniques, taking advantage of DOCSIS 3.0 to help them handle the evolving traffic profiles on their networks.

5.4.1 Broader Effects of the *Comcast* Proceeding

The effects of the *Comcast* proceeding reached far beyond the particular choice of traffic management strategy for some cable operators. Across the broadband industry, it swiftly raised awareness of the potential for regulatory backlash associated with application-specific traffic management. The telcos were by and large not engaging in practices that might earn them the FCC’s ire, and the *Comcast Order* reinforced their decisions to steer clear of any practice that might raise questions before the Commission. This was true even after the expiration in 2007 and 2008 of the merger conditions that had made FCC enforcement a more plausible threat for Verizon and AT&T. Comcast’s ordeal at the FCC was enough to give the telcos pause even in absence of those conditions.

As traffic demands evolved over the years after the *Comcast Order*, the telcos leveraged their existing organizational structures to ensure that proposals for new traffic management practices received sufficient regulatory review. For some cable operators, the *Comcast* proceeding crystallized a need to develop similar processes. Numerous cable interviewees noted that their legal and policy staff became more involved in traffic management decision-making during and after the proceeding. The ordeal "had the effect of bolstering government and legal affairs teams . . . at least their ability to impact network management decisions," as one cable engineer described it. Another noted the "huge amount of time" he began to spend visiting with regulatory and policy staff in Washington.

The organizational adjustments were not limited to increased interaction between policy teams and engineers. One interviewee cited a trickle-down effect, where increased caution on the part of policy staff caused the finance department to view new research or development work with skepticism, withholding budget money for engineers to experiment with new approaches to traffic management in the lab. Engineering thus became constrained with both regulatory oversight and tighter budgets. But the opposite phenomenon also materialized: where marketing teams had been driving the charge towards DPI-based platforms that could support both traffic management and network intelligence, there was a distinct shift in leadership back towards the engineering side of the business. One cable engineer recounted such a change vividly:

> [E]ngineering was sort of back in the driver seat on all of the major decisions related to what we do in the network. And if someone in engineering wanted to raise their hand to . . . our leadership and say, "Hey, this isn't the right thing to do," or "This is going to really blow up in our faces," or "This isn't solving, really, the underlying problem," we sort of have that veto authority then. Whereas before, [the marketing teams] were like, "That's great. Just install the shit. You know, you're the network plumber. Shut the fuck up." That was kind of the view. . . . We, from that point forward, had an equal seat at the table.

Clearly, there was a belief within some organizations that allowing the marketing units too much control over these kinds of decisions had contributed to a choice of traffic management technology that was much derided by the public and regulators.

The *Comcast* proceeding and its aftermath profoundly shaped broadband and its regulation in the US, and many of its other effects are dealt with in later chapters. As for the substance and process of traffic management decision-making, it put an end to what application-specific management did exist in the US. It put broadband providers of all stripes on high alert that the FCC was taking a hard line against discriminatory treatment of Internet traffic. And it prompted some organizational transformation within certain cable operators, introducing greater regulatory oversight into traffic management decision processes and rearranging the power relationships between different corporate units. As such it left an indelible mark on the way that US broadband providers, and cable operators in particular, select the methods they use to manage Internet traffic.

5.5 Conclusion

The period following the *Comcast Order* was tumultuous from a legal and regulatory perspective (see Chapter 8). Comcast sued to have the order overturned. Barack Obama was sworn in as President just months later, having campaigned for office on a platform that included explicit support for net neutrality. The FCC, under its new chairman Julius Genachowski, launched a proceeding to develop open Internet rules. While it was ongoing, the D.C. Circuit upheld Comcast and ruled that the *Policy Statement* was unenforceable, causing the FCC to explore reclassifying broadband Internet service under Title II. Leaders in Congress and at the FCC held meeting after meeting in search of a compromise set of rules that a wide range of stakeholders would support. The Congressional talks broke down, paving the way for the FCC to adopt the *Open Internet* rules without reclassifying broadband. Legal challenges immediately ensued.

Amidst all of this uncertainty and upheaval, traffic management practices remained largely unchanged. Network operators were leery of ruffling feathers in Washington, and regulatory and policy staff had increased insight into potentially controversial engineering changes. US broadband users rode out the decade largely free of application-specific management.

The experience of US broadband users in the first decade of the 21st century, characterized by a relative lack of application-specific traffic management, arose out of the interplay between technological and regulatory forces. The engineering specifications of different access network technologies contributed to the difference in the degree of urgency with which network operators considered managing specific Internet applications. But this urgency, or lack thereof, was couched within existing organizational conceptions of regulatory risk and organizational institutions for addressing that risk. The decisions that operators made were no more determined by the architecture of their networks than by the amount of time and resources they had spent to obtain a particular regulatory classification. The combination of these forces secured a relatively "neutral" traffic management experience for US broadband users. As the next chapter will demonstrate, the same cannot be said for the United Kingdom.

Chapter 6.
Traffic Management Decisions in the United Kingdom

6.1 Introduction

In contrast to the US, application-specific traffic management was pervasive in the United Kingdom during the 2000s. The UK's earliest broadband providers adopted traffic management techniques aimed at controlling peer-to-peer and other non-web applications when broadband was still a nascent product with few subscribers, and in the years that followed, those techniques became entrenched and diffused throughout the industry. By 2010, more than 75% of UK residential subscribers were subject to some form of application-specific traffic management. Many of the techniques employed were blunt instruments that severely constrained the performance of the affected applications for large portions of the day, as opposed to targeting specific instances of instantaneous congestion. Even as network conditions improved or evolved, application-specific traffic management techniques often did not.

The most common form of application-specific management observed in the UK – rate limiting peer-to-peer file-sharing traffic – was initially adopted for the combined benefits of improving performance and controlling variable costs (predominantly backhaul capacity costs). The significant competition that ensued after Ofcom and BT took steps to make local loop unbundling viable had two effects that reinforced the utility of this approach: a significant decline in retail prices, and mounting pressure to increase broadband speeds, which required further investment. To meet marketplace demand for increasing speeds at decreasing prices, operators were incentivized to do whatever they could to keep costs down, including rate limiting high-volume applications.

Throughout the decade, Ofcom's actions and rhetoric created a culture of flexibility around traffic management. Ofcom was infused from its creation with the culture of a competition regulator and its embrace of functional separation was one of the most interventionist steps it could have taken to stimulate broadband competition. Having taken that step, it demonstrated a deep commitment to the disciplining power of competition in the marketplace, including in its policy formulation around net neutrality and traffic management. At the same time, Ofcom officials were adamant that the management of peer-to-peer applications was essential for networks to function. This logic gave operators nearly as much freedom as they could have desired in managing applications as they saw fit.

Consumer concerns about application-specific traffic management (to the extent that they existed) were muted, in part because of the perceived association between peer-to-peer applications and illegal copyright infringement. Affected users would complain or inquire about their operators' traffic management practices in online forums, but these concerns rarely garnered the attention of media outlets, consumer groups, or Ofcom.

Because neither consumers, nor Ofcom, nor the industry drew attention to application-specific traffic management, during the mid-to-late 2000s there was no nuanced public debate in the UK about the relative merits of application-specific and application-agnostic approaches or about whether the approaches in place could or should be more targeted or consumer-friendly. As a result, the relatively blunt approaches to managing applications that some operators adopted early on became ossified, hardly evolving over the decade despite massive evolution in network performance and speeds.

This chapter analyzes how the UK's wholesale market structure, the intense effects of competition, and a regulatory culture of permissibility combined to foster a British broadband landscape in which application-specific traffic management became pervasive. Marketplace dynamics bolstered operators' incentives to adopt or perpetuate application-specific management while countervailing public pressure never materialized, yielding a vastly different traffic management experience than that of the United States.

6.2 Early Adoption of Application-Specific Management

Application-specific traffic management was a fixture in the UK dating back to the earliest years of broadband itself. Numerous early broadband operators turned to application-specific traffic management strategies that relied on port-based application classification or deep packet inspection equipment to identify and manage high-traffic applications on the network. Application-specific management was adopted as part of larger strategies that often involved application-agnostic approaches as well, most commonly tiers of service with explicit volume caps or fair usage policies (FUPs) that prescribed limits for how much data customers could consume.

Peer-to-peer applications were the most common and widespread applications targeted, with operators limiting the capacity available to them either on a per-user basis or in the aggregate (for example, allowing the applications to consume no more than 1% of the network's overall peak bandwidth). Other non-web applications, including newsgroups, were at times subject to rate limits as well. Providers commonly imposed these limits across their entire user base, and at times did so with minimal or buried disclosure. Different providers imposed the limits over different time periods, with some applying them only during a few evening hours or at times of congestion, and others applying them for longer stretches of the day or at all times. Some operators, most notably Plusnet, adopted more complex application-based prioritization schemes, offering broadband packages where many more application classes were identified (web, email, VoIP, gaming, etc.) and prioritized according to their time-sensitivity or other factors.

This section explores the key drivers for adoption of application-specific management in the early years of broadband, prior to the fundamental market changes that accompanied the roll out of local loop unbundling (LLU) beginning in 2005 and 2006. It analyzes the most common case – DSL operators' management of high-volume applications – and contrasts that approach to those of the cable operators and Plusnet, whose early strategies helped to shape the traffic management landscape for the remainder of the decade. This section shows how application-specific management commonly met DSL operators' needs to control both performance and costs while incurring little attention from the public, the press, or Ofcom.

6.2.1 Wholesale Network Model for DSL

Broadband Internet service arrived in the UK around the turn of the 21st century. Early broadband was predominantly provided by cable operators and the retail arm of BT. With the launch of BT's wholesale bitstream products in 2000, bitstream operators proliferated, led by a number of companies that would all eventually come under the TalkTalk brand (Carphone Warehouse, Tiscali, Pipex, Nildram, and numerous others). Because the cable networks' footprint is geographically limited, DSL has been the dominant form of broadband access since 2003 (Ofcom 2009b).

Backhaul Prices

There are several features of the bitstream cost model that played an important role in operators' decisions to adopt both application-specific and application-agnostic traffic management in the first place. Figure 3 shows the key components of a typical bitstream operator's network (and BT Retail's network). Bitstream operators would contract with BT's wholesale division to run their customers' access lines. Although these access network costs typically accounted for more than 50% of the total monthly cost of each customer's line, the costs were bandwidth-independent. Access links are provisioned at a specific line rate per customer, and fluctuations in bandwidth usage up to that limit based on each customer's activity do not incur any additional charges. Nor do such fluctuations impact other customers in the neighborhood since each customer is connected to the exchange on a dedicated access line.

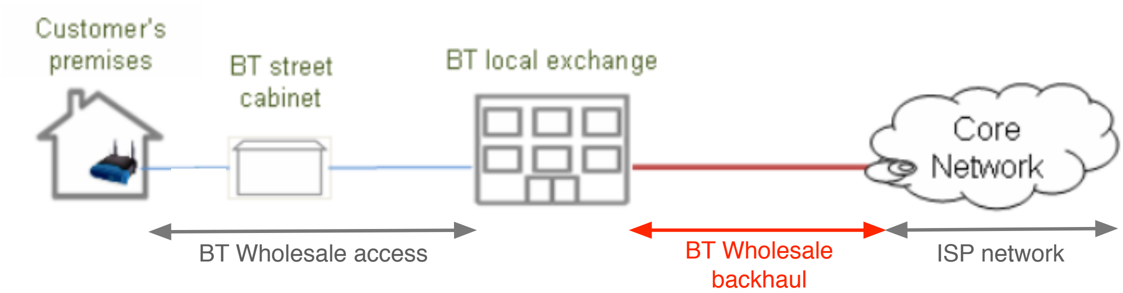

Figure 3. Typical UK bitstream ISP network in the early 2000s.

The backhaul component, in contrast, is both bandwidth-dependent and shared among many subscribers. Bitstream operators would purchase a certain amount of capacity from BT Wholesale to connect their customers from the local exchange to their own networks and on to the rest of the Internet. The amount of capacity purchased was based on a number of factors, including available budget, network utilization, and performance targets. If peak time bandwidth demand exceeded available capacity, operators could either purchase more bandwidth or allow performance to suffer as each customer's traffic competed for space on constrained links. In describing an early broadband network prior to the adoption of traffic management tools, one network engineer explained that "congestion was what actually controlled the amount of bandwidth and the experience the customer had."

The most popular wholesale products consumed by DSL operators in the early days were known as Datastream and IPStream with BT Central Plus. These products, whose prices were regulated by Ofcom based on a finding that BT had Significant Market Power in the wholesale market, were priced to average out the fixed and ongoing network costs that BT incurred in maintaining the nationwide wholesale network. As described by Plusnet's Dave Tomlinson (2008), UK operators saw a "huge price premium" for backhaul because the uniform national price was set to cover the cost of "links to some of the smallest exchanges in the country that perhaps only serve a couple of hundred customers, some of which are probably very uneconomical." Backhaul bandwidth was generally the most expensive bandwidth segment on UK operators' networks. Indeed, in the mid-2000s, the price per Mbit/s of IPStream (in the £120-£200 range for most ISPs' networks) was five to ten times the median price per Mbit/s that ISPs would pay for transit links that connected them with the rest of the Internet (Telegeography 2011).

The resulting backhaul prices also provided DSL operators with limited economies of scale in their purchase of capacity: a linear increase in capacity necessary to accommodate a growing traffic base was met with a roughly linear increase in the price of backhaul. For those using the IPStream product, for example, the price per Mbit/s when purchasing 622 Mbit/s of capacity on a BT Central Plus link was roughly the same as the price per Mbit/s when purchasing 155 Mbit/s on the same link once the annual rental for the link had been paid (BT 2006). Thus as traffic growth outpaced customer growth, operators were faced with backhaul costs that rose faster than revenues from customer subscriptions.

Retail Networks' Lack of Control

The fact that the wholesale network was run by a distinct entity from the retail networks also reduced the control that bitstream operators had over their networks in a number of ways. The lead time for deploying additional backhaul could be in the three-to-four month range, meaning that if operators' quarterly forecasts for traffic growth were too low, or if the process of getting new capacity installed incurred extra delays, there was no way to augment capacity on short order to mitigate performance problems. The kinds of "emergency" capacity increases described by US cable operators could not be deployed by UK bitstream operators.

UK ISPs were unabashed about the drawbacks of being beholden to BT's upgrade schedule. For example, the bargain bitstream operator FairADSL was explicit in explaining to its customers the source of performance problems on its network in late 2002:

> [W]e placed upgrade orders with both BT and our outgoing bandwidth providers in plenty of time to keep up with demand. Unfortunately due to ineptitude and bureaucracy on behalf of these parties, the bandwidth upgrade did not operate correctly and the BT Central Pipe upgrade has not been implemented in time. We are very annoyed about this, and extremely angry on behalf of our customers. (Jackson 2002)

ISPs felt the need to provide such explanations because the performance implications of delayed capacity upgrades or incorrect bandwidth demand predictions were obvious to subscribers. Their application performance suffered as a result. This issue was exemplified by the experience of the bitstream operator Nildram in January 2003, as reported on the ThinkBroadband blog (then known as ADSLGuide):

> Nildram, a well known ADSL ISP . . . is experiencing difficulties with a growing user base and lack of capacity on their central pipes. In the past few days, users of latency-sensitive applications such as online gaming and telnet have experienced problems at peak times when the central pipes have been overloaded. Orders for two more 155 Mbps pipes have already been placed with BT some time ago and were due to be on service by December 17th but Nildram now expect them not to be available until January 17th at the earliest. (Lahtinen 2003)

Degradations in service quality brought on by delayed capacity upgrades were a recurring theme in these early years.

Furthermore, when congestion arose operators had little visibility into its causes on the wholesale network prior to the interconnection point with the operators' networks. Because of the wholesale/retail split, operators lacked visibility into what was taking place on that first segment of the network where multiple subscribers' traffic came together, as one engineer explained: "the reality was that if [customers] had problems, you had very little that you could do with them because . . . all you knew is that they went through the wholesale network and somewhere along the line as far as [we were] concerned they magically connected to the Internet." In sum, operators lacked control over both the costs of their networks and the performance their users experienced, with little ability to diagnose or solve either problem.

6.2.2 Traffic Management to Control Cost and Performance

These problems came to a head soon after broadband became widely available. Subscribership grew rapidly in the early years, with well over 50% growth in both DSL subscribership and broadband uptake overall each year between 2002 and 2005 (Ofcom 2009b). With more subscribers, and more high-traffic subscribers, network operators were seeing increased complaints from customers about the speeds they experienced at peak times.

Keeping up with demand required increased investment in network capacity, with operators at times doubling their entire network capacity in a single upgrade. But given the pace of growth and the backhaul pricing structure in place, operators were searching for mechanisms to control the amount of traffic on their networks without needing to provision for an increasingly high and difficult-to-predict peak traffic rate. Backhaul costs were becoming a growing concern, and “writing a blank check for demand” to BT Wholesale (as one engineer put it) was not seen to be a sustainable solution.

In assessing their options, operators observed that a small minority of heavy users were creating the majority of the traffic, as has been common not just in the UK but across European and international broadband networks (Mooyaart 2012; Sandvine 2010a). Because high traffic meant high cost, these heavy users became an important target for traffic management. Volume caps and fair usage policies were introduced to help operators control heavy users’ traffic in coarse-grained ways. Some operators would send warning letters to users when they had exceeded daily or monthly volume caps, hoping to encourage them to moderate their usage. Others instituted overage fees, slowed users’ speeds, or disconnected them altogether when they exceeded their volume limits, attempting to create even stronger incentives for heavy users to moderate their usage while also increasing revenue or reducing traffic-related costs in the process.

While these approaches were aimed at reducing the bandwidth demand of a voracious minority, they did not give operators more precise ways of controlling or predicting that demand. Even operators that adopted a daily volume cap that resulted in slower speeds when breached could not know if, when, or how many subscribers would reach the limit on any particular day. In some cases users could choose to face the consequences of exceeding a cap – receiving a letter or paying a fee – without changing their behavior. Thus even with these approaches in place, operators were still looking for more fine-grained ways of predicting and controlling the demand for capacity.

Operators were aware that peer-to-peer traffic was quickly becoming the dominant type of traffic on the network, representing more than 50% of network traffic in some cases, as reported by interviewees and elsewhere (Ferguson 2002; Lunden 2005; Mooyaart 2012; Sandvine 2010). Figure 4 provides further evidence of the trend, showing the ratio of upstream-to-downstream traffic on one large UK network since 2004. Because peer-to-peer traffic tends to be roughly symmetric, a large presence of peer-to-peer traffic on a network implies a higher upstream-to-downstream ratio than would be observed in situations where web browsing or streaming are the dominant application types. The above-50% ratios in 2004 and 2005 demonstrate the dominance of peer-to-peer traffic.

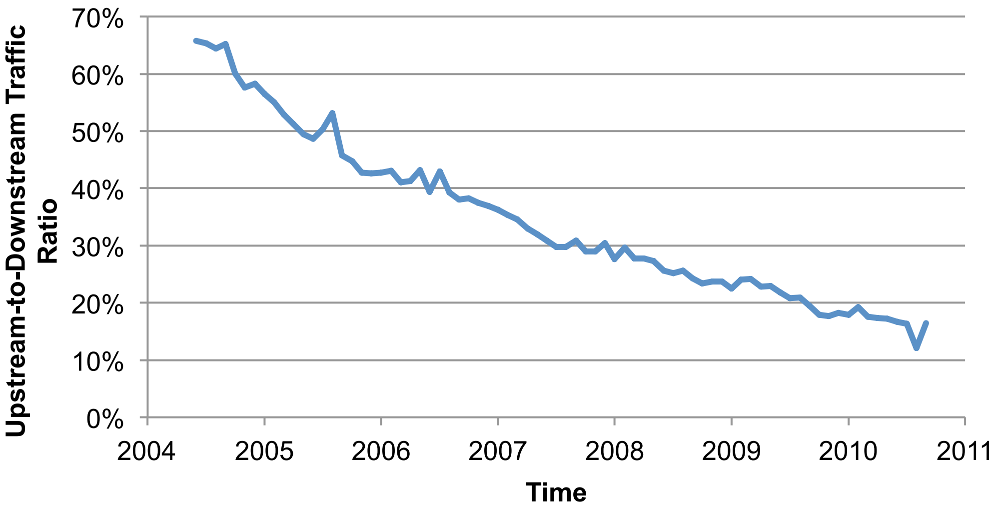

Figure 4. Monthly upstream-to-downstream traffic ratio on a large UK network. Reproduced from Cooper, Jacquet and Soppera (2011).

Using application-specific traffic management equipment to identify and control the capacity available to peer-to-peer file-sharing (and other high-volume applications such as newsgroups) seemed like a natural choice to operators under these circumstances. While application-agnostic approaches were broadly aimed at changing the behavior of a subset of customers, approaches focused on high-volume applications could very specifically and precisely stifle the prime sources of traffic on the network, giving network operators far more control over their bandwidth investments than they had previously. Management of peer-to-peer and newsgroup applications reduced the variability of bandwidth demand, as one engineer explained: "You could basically make your network grow, if you want, at a much lower rate than it would grow otherwise, you could control your costs that actually were real costs. That's why you have seen the retailers in the UK . . . the majority that are consuming BT Wholesale products have deployed that type of solution." In a market where operators were beholden to BT Wholesale's prices and upgrade schedule, these tactics gave them back some element of control.

But managing high-volume applications was not only about controlling the timing of bandwidth investments – it was also about reducing overall costs while safeguarding the user experience for the majority of subscribers. The use of application-specific traffic management as a way to reduce costs was a constantly recurring theme among the observed research team, and numerous interviewees were unequivocal about the effectiveness of peer-to-peer management as a cost reducer. The following is a sample of responses provided when interviewees were asked why this form of traffic management was deployed on ISPs' networks:

- "You save real money."
- "In order to save . . . from paying [BT Wholesale] loads of money in that interconnection point."
- "To save money! (laughing) Of course! Peer-to-peer traffic is a lot of traffic."

6.2.3 Operational Drivers of Application-Specific Management

While high-volume applications may have been prime contributors to congestion, it was the operators' view that they did not suffer its effects in the same way as other applications, providing a further rationale for management. Congested networks can indiscriminately delay applications of all types whether or not they are interactive. Suppressing peer-to-peer traffic and newsgroups helped to reduce the likelihood that more popular applications like web browsing would suffer from congestion. As one operator explained, "if you didn't apply some level of control, if we just allowed it to congest like we'd done previously, it was very clear you could see browsing issues, you could see performance of VoIP issues," because those were the kinds of applications where users would notice the delay caused by increased congestion. The Nildram example provided above illustrates this as well: users of latency-sensitive applications were the ones to notice the effects of constrained capacity.

Peer-to-peer and newsgroup applications were also viewed by operators as less time-sensitive and interactive than other traffic. By limiting the capacity available to these applications at peak times, file transfers would be delayed, but they would still eventually complete. Interactive applications like web browsing, VoIP, and gaming, by contrast, were viewed as having limited delay constraints, otherwise they could become unusable.

As with application-agnostic approaches, approaches focused on high-volume applications were appealing to operators because they were likely to have detrimental effects on only a minority of heavy users, to the benefit of the majority of users. On many networks, peer-to-peer file-sharing followed a similar power law distribution to broadband usage overall, with only 5% or 10% of a network's subscribers creating the peer-to-peer traffic that accounted for the majority of traffic on the network. Thus application-specific management was viewed as a safeguard for the majority of users' experiences. As O2 later noted in a filing to Ofcom, operators that implemented these kinds of strategies did so "in the interests of all users (and society) even though this may inconvenience a tiny minority of users who are consuming a disproportionate volume of scarce resources" (Telefónica 2010, 10).

As in the US, DPI-based solutions were also attractive to operators for the insights they gave about the causes of network problems and growth. Prior to the deployment of DPI, some operators used port- or IP-based information to guess at how much traffic could be attributed to various applications, but DPI provided the additional benefit of granular data about individual application usage, in the aggregate or on a per-customer basis. DPI gave operators a platform for understanding which applications were driving growth on their networks, which applications heavier users tended to use compared to light users, and a variety of other insights that could be used for capacity planning and design of new products (Cooper, Jacquet, and Soppera (2011) provide examples of such insights).

6.2.4 Alternative Approaches

Although DSL operators imposing limits on peer-to-peer applications was the most common type of application-specific management in the early years, alternative approaches pursued by Plusnet and the cable companies played important roles in shaping the overall UK traffic management landscape. Plusnet pioneered the explicit use of traffic management for product differentiation, while early experiences with complex application-specific tools in the cable industry created a lasting belief there that traffic management solutions should be simple to administer and explain to consumers.

Plusnet: Productizing Traffic Management

While other early adopters of DPI-based traffic management were quietly applying blanket restrictions to high-volume applications across their networks, Plusnet was spearheading an alternative approach that put traffic management at the center of its product offerings. In the early days of broadband, Plusnet bore the impact of BT Wholesale pricing perhaps even more so than most other ISPs because Plusnet had a policy of operating broadband as a profitable business. It would generally only buy capacity that customers were paying for and not more (Wyse 2008).

To accommodate growing demand while remaining profitable, the company initially experimented with an application-agnostic approach in which the traffic of the network's heaviest users was gathered together on the same link, forcing them to contend with each other rather than interfering with lighter users. However, based on a backlash from heavy users, conversations with other large ISPs where DPI was already employed for traffic management, and a desire to manage traffic in a more sophisticated and efficient way, Plusnet soon shifted to application-specific management that made use of DPI equipment (Imtech Telecom 2005).

As former Plusnet CTO Alistair Wyse (2008) has explained, Plusnet began by instituting an aggregate limit on peer-to-peer traffic capacity, but found it to be too much of a "blunt tool," treating heavy and light users of peer-to-peer applications the same. Instead, the company opted to develop specific usage profiles based on how their customers were using different applications across all hours of the day and to base its product offerings on those profiles. The result was a set of broadband packages that combined rate limits for high-volume applications with prioritization schemes that took effect when congestion occurred. For example, a lower tier package first offered in 2005 included (among other limits) rate limits of 100 Kbit/s for peer-to-peer and newsgroups, 2.5 Mbit/s for YouTube, and 2 Mbit/s for all other streaming from 6:00pm to 10:00pm. The rate limits differed during other hours of the day, with no limits from 12:00am to 4:00am. This was complemented by a prioritization scheme where VoIP and gaming received top priority, streaming and browsing were in the middle, and peer-to-peer and newsgroups received the lowest priority during times of congestion (Plusnet 2012a). Plusnet offered a series of higher-tier packages with higher rate limits and different prioritization schemes at higher price points, topped off by a package that prioritized VoIP and gaming but otherwise included no application-specific management (Plusnet 2012b).

Although Plusnet was a small ISP with a savvy user base, its "productization" and transparency around traffic management signaled to other ISPs that application-specific management could be a legitimate foundation for broadband product offerings. Numerous interviewees praised the Plusnet approach as a "really good solution" that was "more intelligent" than existing solutions and that other ISPs "definitely tried to learn from." By deliberately and openly segmenting its products based on specific application limits and priority schemes, it differentiated itself and set an example that other ISPs later looked to as they were considering traffic management options. Plusnet showed that to some degree and with a specific audience, customers would accept application-specific management. It could be used to sell broadband products, not just control costs and performance in the background.

Cable's Preference for Simplicity

Although the UK cable networks' architectures were obviously very different from those of the DSL providers, many similar factors played into the cable operators' traffic management decisions.

In the early days of broadband there were two primary companies offering cable broadband service, ntl and Telewest. They merged in 2005, began operating as ntl:Telewest, and were re-branded as Virgin Media in 2007 following the combined company's acquisition of Virgin Mobile. Prior to 2010, application-agnostic management aimed at heavy users was the primary approach used to control costs and performance. Like the DSL operators, cable operators adopted these strategies on the assumption that they would receive few complaints since only a minority of users would see an adverse impact.

Not long after launching broadband, the cable networks were seeing the same kinds of growth and traffic patterns as the DSL operators, with a small fraction of users and applications creating the majority of network traffic. As with the US cable companies (and unlike the DSL operators), one of the major expenses associated with surges in traffic was the need to manage demand in the access network, where subscribers shared capacity and a single heavy user could significantly degrade the experience for other users sharing the same Internet Protocol port. One cable product manager aptly captured how this dynamic influenced his company's traffic management decisions:

> [A] tiny, tiny percentage of customers, you know, single percents, were absolutely blasting the network. And not only costing us a lot of money but actually ruining the peak time experience for other customers on that port. So you actually solved for two things. You're able to better manage those heavy users and that's going to help our capital profile, which is of course an objective. (You know, we don't necessarily shout about that on [our] web site.) But it's not all about that. It absolutely protects customers who are normal users from congestion. And congestion on cable is obviously a very important metric.

Finding ways to control the need to invest in re-segmenting the network to reduce the demand on shared access infrastructure became imperative.

The cable operators' earliest approaches to managing traffic were application-agnostic measures aimed to limit the impact of the minority of heavy users on the rest of the user base. In 2003, ntl instituted a 1 GB/day download usage cap and began contacting customers who persistently exceeded this threshold (T. Richardson 2003). Telewest had a similar policy for its lowest tier service (Telewest 2004), but also began trialing a more fine-grained system. During certain daytime and evening periods, network routers would measure the volume of traffic that each subscriber was sending and receiving. If a subscriber exceeded a pre-defined threshold in either the downstream or upstream direction, his or her line speed would be significantly reduced (up to 75%) until the end of the period. This functionality, known as "subscriber traffic management" (STM), was a feature of the Cisco routers already in use on the network.

In 2006, before the two networks were fully integrated post-merger, application-specific measures were briefly added to the ntl network, reducing the capacity available to peer-to-peer and other high-volume applications at peak times. The rate limits for different applications were set differently in different parts of the country depending on the extent of the use of high-volume applications. In areas with large student populations, for example, capacity for these applications was more severely limited.

This deployment was short-lived, however. With Telewest staff primarily taking responsibility over the product management functions in the combined company, the decision was made to remove the application-specific measures and apply STM across the whole network. The technology used to do the application shaping was "incredibly complex," "expensive to run," and "had obviously been a little bit unloved on the ntl network" according to one interviewee. The fact that it operated differently in different parts of the country made it difficult to explain to consumers and to assess network capacity plans overall.

The STM functionality, on the other hand, was "pretty low maintenance and pretty easy," as the interviewee put it. It was viewed as a simpler approach that could be explained in the same way to customers everywhere. The technology was well understood within the combined company and did not require additional devices on the network beyond the routers that were already necessary to route traffic. STM was therefore viewed as a simpler and more cost-effective solution that helped to solve a largely similar problem that application-specific management was designed to address. The value of simplicity in traffic management was internalized at the company and continued to influence its decisions in later years.

6.2.5 Lack of Public Attention to Application-Specific Management

Role of Perceived Copyright Infringement

The widely held assumption that most, if not all, peer-to-peer file-sharing consisted of illegal copyright infringement created an important subtext in which both DSL and cable operators made their decisions about how to manage traffic from the very beginning. By the early 2000s, peer-to-peer applications had been widely associated with copyright infringing activity given the content industry's pursuit of a number high-profile lawsuits and messaging campaigns since the late 1990s. This association had two key effects on operators' traffic management decisions: it created an additional rationale for targeting peer-to-peer file-sharing with traffic management, and it gave operators some comfort that their customers would be unlikely to complain about traffic management practices aimed at peer-to-peer applications.

The extent to which illegal uses of peer-to-peer applications helped to legitimize their management varied from operator to operator, but across the broadband industry operators acknowledged its role in shaping traffic management decisions, even if that role was small. One former cable executive explained that when considering traffic management solutions, the thinking was that "if you're going to traffic shape, [de-]prioritize [peer-to-peer] first because it's illegal and it's all the heavy users." While others were not quite so emphatic, interviewees generally agreed that the perception of illegality made peer-to-peer applications a more justifiable target for rate limiting. One policy executive explained that peer-to-peer was simply "politically disfavored traffic."

Perhaps more importantly, the perception of illegality provided operators with some assurance that customers were unlikely to complain about peer-to-peer management, either by calling up customer service agents or contacting the press or regulators. They expected customers to keep quiet for fear of retribution for their infringing activity. This dynamic underscored the attractiveness of peer-to-peer management because it was unlikely to increase either the public attention paid to traffic management or the volume of customer service calls, which can create significant costs for broadband operators. One former DSL strategist explained this vividly, illustrating why such a customer service call would be unlikely to take place:

> You know everybody used to just choke the hell out of peer-to-peer at peak. Just almost squeak it to nothing. And that was okay because who's going to complain about that? Who's going to phone up [and say], "I'm trying to download these films for ages and I can't get them"? "Alright kid, would you like me to forward you to Warner Brothers?" You know, it's not going to happen.

From the operators' perspective, consumers had a clear disincentive to complain. Some operators took this one step further by intimating to heavy users that given the volume of traffic they were downloading, they were likely to be engaged in illegal activity. In this case, the operators themselves were raising the specter of legal consequences for infringing activity, further reinforcing the message to consumers that drawing attention to themselves by complaining about traffic management practices was a potentially risky activity.

When they did complain, users primarily limited themselves to online forums and bulletin boards where their identities could be protected and where they could learn from other similarly situated users about ways to evade traffic management limits or switch providers. Operators have long been aware of these forums and the complaints they contain; many operators even host their own or have staff who participate on the forums and address concerns as they are raised. But awareness of forum-based complaints never created any particular imperative against the use of peer-to-peer management. One engineer explained this as follows:

> I mean drop the peer-to-peer on the floor for two hours, I don't think that anyone is going to complain. And it represents quite a significant portion of traffic on the network. And tends not to drive too many calls when people phone up and say, "my BitTorrent isn't working." Generally they're not actually phoning up to do that. They might be bitching on a message board somewhere, but it's not such – the cost of that and the reputational impact it causes – it's a case still to be proven.

In short, the fact that peer-to-peer applications were widely associated with copyright infringement added an additional element of justification to some operators' traffic management decisions while also insulating them from the kind of backlash, both public (press or regulatory) and private (customer service calls), that could create financial and reputational costs.

Regulatory Focus Directed Elsewhere

With user complaints about application-specific traffic management maintaining a low profile, questions about consumer harm and discriminatory practices related to traffic management failed to materialize in the halls of Ofcom (and Oftel before it) during the early years of broadband. Instead, Ofcom was concentrated on creating a more competitive marketplace, where faster broadband speeds could be garnered at lower prices across more of the population.

UK broadband penetration was just 10% in 2004, about average across all of the OECD countries (OECD 2006). The provision of broadband by LLU operators was scant. Improving that situation was one of Ofcom's first priorities after it was created. The regulator laid out its plans for a revised broadband regulatory framework in 2004 (Ofcom 2004d) that, together with the broader set of regulatory changes being formulated as part of the strategic review of telecommunications, would allow LLU-based broadband to become a widespread service offering. That review on its own was a massive undertaking, and combined with the LLU market review, Ofcom had plenty of work to do to make broadband more competitive and widely available. Traffic management practices that were potentially discriminatory or anti-consumer simply were not on Ofcom's agenda.

Blunt Instruments of Traffic Management

With neither Ofcom nor consumers putting pressure on ISPs early on to change their traffic management practices, operators had the flexibility to deploy application-specific management as a means of controlling costs and performance in the background, rather than as a central product feature. One DSL product manager explained that peer-to-peer and newsgroup management "was put in out of necessity as opposed to anything in terms of differentiation of product," although there were occasional exceptions such as Plusnet.

This flexibility accorded to the ISPs meant that operators could be aggressive – "pretty draconian" or "pretty blunt" in their own words – in configuring when and by how much they reduced the capacity available to high-volume applications. Limits on high-volume applications would commonly apply to all customers regardless of the package to which they were subscribed. Some limits were in place 24 hours per day, while others applied for shorter but still significant portions of the day. Under the "just squeak it to nothing" approach described earlier, the targeted applications had their speeds reduced dramatically, often down to single-digit percentages of users' headline speeds. With few customers and no market or regulator providing incentives to make application-specific management more targeted to congestion, more consumer-friendly, or less blunt in general, operators generally sought as much cost and performance control as they could without blocking high-volume applications altogether.

6.2.6 Summary of Early Years

Traffic management, both application-agnostic and application-specific, was a feature of UK broadband since the inception of broadband itself. The structure of the DSL market – with retail providers as customers reliant on BT Wholesale's nationally averaged, linearly priced backhaul – spurred DSL operators to seek out solutions that would allow them to control both costs and performance. Because the majority of traffic could be attributed to heavy users and high-volume applications, those became logical targets. The cable operators saw similar usage trends, but early experience with a complex application-specific approach created a desire to maintain a simpler application-agnostic approach going forward.

The widespread management of applications on DSL networks was further rationalized on the basis that peer-to-peer and newsgroup applications were more amenable to having their performance reduced than other applications on the network while also being in use by a smaller segment of the subscriber base. The application awareness involved in conducting application-specific management benefitted operators by providing them with insights into how their networks were being used. Operators expected few customer complaints, publicly or privately, given consumers' perceived hesitancy to be associated with illicit activity taking place on peer-to-peer networks.

Taken together, all of these factors drove the adoption of traffic management strategies that would become the norm across the industry. By 2004, a selection of large and small DSL providers – serving more than a quarter of all broadband users – had taken up application-specific traffic management. Fairly blunt and unsophisticated management of peer-to-peer and other high-volume applications was most common. The cable networks and smaller operators like Plusnet differentiated themselves in ways that would have consequences in the market going forward.

While Ofcom was keenly focused on broadband issues, traffic management was not among them. The regulator's attention was diverted to larger efforts, the results of which contributed to a reshaping of the UK broadband landscape in the second half of the decade.

6.3 Competition Entrenched and Diffused Application-Specific Management

The UK broadband market entered a period of upheaval in 2005. Ofcom conducted a complete review of telecommunications regulation, culminating in a significant set of reforms adopted to spur competition in the broadband access market. Shortly thereafter, as net neutrality discussions in the US intensified, Ofcom began to formulate its position regarding discriminatory treatment of Internet traffic, beginning from the notion that the competitive framework in place should serve to discipline providers' behavior. At the same time, competition itself was intensifying some of the pressures that had led to adoption of application-specific traffic management in the first place, requiring network operators to invest in increasing capacity even as retail prices declined.

This section explains how these events created a climate in which application-specific traffic management became legitimized, entrenched, and diffused across the bulk of the UK broadband industry. Early broadband providers that had already been using application-specific approaches by and large left them in place untouched. Application-specific strategies had become so well accepted that BT briefly began rate limiting video traffic, including traffic generated by the BBC's popular iPlayer service, in 2009. Later adopters of application-specific management expanded the use of traffic management in new ways and rationalized their choices differently than early adopters had. The DSL operator Sky was a notable exception to these trends, showing that a major ISP could compete on price and performance without managing applications.

6.3.1 Ofcom's Shaping of the Traffic Management Landscape

One of Ofcom's first and largest tasks when it was founded was a complete review of existing telecommunications regulation, known as the Telecommunications Strategic Review (TSR). Announced in late 2003, one of the key goals of the TSR (and of Ofcom more broadly) was to promote competition in telecommunications services. As Ofcom noted in launching its first

TSR consultation, “despite nearly 20 years of regulatory activity intended to promote competition” (Ofcom 2004b, 2), BT remained in a position of significant market power (SMP) in a number of telecommunications markets, including wholesale broadband markets (Ofcom 2004c).

The undertakings to which BT agreed at the conclusion of the TSR helped to change that. Local loop unbundling had been possible in the UK since 2000, but only in 2005 did it become a widely viable possibility for competitive operators. As part of the undertakings, BT agreed to functionally separate its access network division (which became Openreach) from the rest of the company and to provide equivalent wholesale prices, terms, and service guarantees to all ISPs, including its own ISP, BT Retail (Ofcom 2005a). Separately, BT also made commitments to lower access network prices, maintain its bitstream prices at sufficient levels to allow LLU operators to compete, and make improvements in provisioning of service systems used by competitive operators (Ofcom 2005b).

Together, these commitments – and the oversight structure that Ofcom and BT agreed to use to enforce them – provided the foundation for an explosion in LLU, as depicted in Figure 5. From 2005 to 2006, the number of unbundled lines grew by more than 500%, instantly giving UK users increased broadband choices (Ofcom 2011a). By 2009, 35% of all broadband subscriptions were via unbundled lines and 85% of broadband customers had access to at least one LLU entrant (Nardotto, Valletti and Verboven 2012), in addition to BT, bitstream providers, and, in many areas, cable.

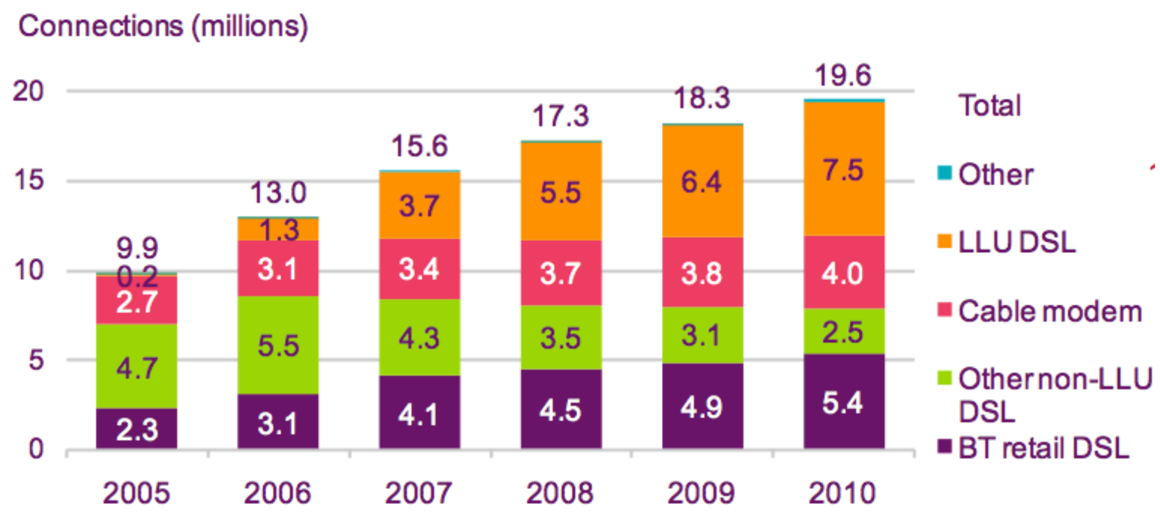

Figure 5. Number of UK residential and small business fixed broadband connections by connection type, 2005-2010. Reproduced from Ofcom (2011a).

Ofcom's Belief in Competitive Discipline

The TSR was a "huge exercise" (as one former Ofcom official put it) for the newly constituted regulator. The TSR also resulted in some of the boldest steps to promote competition that any European telecommunications regulator had taken since the beginning of liberalization – a "radical intervention," as Ofcom's former head of strategy described it (M. Brown 2010). Although other national regulators in Europe and elsewhere had intervened to spur the development of LLU and other means of competitive entry, Ofcom was the first to require functional separation of the incumbent (Tropina, Whalley and Curwen 2010). The former official characterized the undertakings as "a very heavy regulatory regime" with "nondiscrimination rules on steroids" to ensure that alternative providers would have the chance to compete.

It was obvious from the beginning of the review that one important measure of Ofcom's success would be whether competition emerged in the broadband market. In the years following the conclusion of the TSR, Ofcom repeatedly highlighted the emergence of competition and the associated consumer benefits of lower prices, higher speeds, and more choices, as evidenced by the agency's annual reports:

> Genuine competition in the fixed-line telecoms markets, at the deepest levels of infrastructure, creates a virtuous circle of new investment in emerging technologies and innovation in services and price competition. This benefits consumers and helps maintain the competitiveness of the UK economy as a whole. (Ofcom 2006a, 18)
>
> Towards the year-end, the UK passed a significant milestone with more than half of all households now on broadband, encouraged by cheaper prices, higher bandwidths and an increasing range of new services. . . . Ofcom's role in wholesale broadband regulation has been critical to the roll-out of Local Loop Unbundling, which passed the two million lines mark during the year . . . . (Ofcom 2007a, 4)
>
> Citizens and consumers are already reaping the benefits of competition. Over half of all households have broadband from one of more than 500 different providers. (Ofcom 2008a, 6)

Ofcom's belief in the promise of competition and its devotion to ensuring competitive entry had profound effects on the regulator's approach to traffic management in the context of net neutrality. As the net neutrality debate attracted increasing attention in the US in 2006, Ofcom began developing its own position as to whether regulatory intervention was necessary to prevent discriminatory practices, including application-specific traffic management. While in the early years of traffic management in the UK Ofcom had been focused elsewhere, the growing attention devoted to net neutrality abroad drew Ofcom's attention to the issue and created an imperative for the regulator to react.

One of the key tenets of its early positioning on the issue was that competition among broadband providers in the UK was sufficient to discipline providers' urges to discriminate. One former Ofcom official explained that "because it's a competition regulator it sees [net neutrality] as a competition problem." In its first public writing on the issue, Ofcom declared that "[f]or those operators without SMP in the relevant market, we consider that efficient working of a competitive market will address the risks posed to consumers from non-network neutral approaches. An effectively competitive market at the retail level, with relatively low barriers to entry, means that customers have a range of choice in their ISP" (Ofcom 2006b, 17). As a competition-focused regulator, Ofcom instinctively approached the issue with the belief in the ability of competition to restrain operators' behavior.

In public addresses over the years that followed, Ofcom officials reiterated this argument, emphasizing that the UK has “such strong retail competition” (Ingram 2006); that “competition itself can provide a constraint” (Kiedrowski 2007); that degrading application performance would make an ISP “a less attractive proposition to consumers” (D. Scott 2007); and that “[i]n a more competitive environment, there is less inherent problem with traffic management and prioritisation” because “[i]f network operators get these calculations wrong, consumers will switch to another provider” (Richards 2008). The belief in competition existed at the highest levels of the agency.

Underlying these arguments was not only a belief in the logic of competition, but also the need for Ofcom to prove that its interventions to promote competition were justified. To those who had crafted the extensive competition framework that emerged from the TSR, it was unconscionable that a problem could so soon emerge in public discussion that could not be remedied by the competitive market. One former Ofcom official described this internal agency logic:

> Providing open access is a pretty big competition intervention. You don’t do it for fun. Therefore, hang on a minute, surely, the threshold we would want to set would be so high that we can’t even begin to see there’s going to be a problem.

To Ofcom, it was essential to demonstrate that competition delivered on its promises. Having expended tremendous time and resources to achieve the post-TSR regulatory framework, it was not conceivable that discriminatory traffic management was a problem that required a solution beyond competition. Not only did Ofcom believe in competition, it was institutionally imperative that competition be shown to work for consumers.

Competition was therefore a constant refrain for Ofcom clear through to the end of the decade. Spurred on by the conclusion of the review of the EU telecommunications framework and burgeoning net neutrality discussions in Europe, Ofcom issued a public discussion document in 2010 in which it repeatedly affirmed the argument it had been making since 2006:

> Without market power there is a strong presumption that no anti-competitive effects and consumer harm will arise. This is because consumers will tend to punish attempts at exclusionary behaviour by simply shifting their business to an alternative provider who does not engage in the same exclusionary practices. At present, both the fixed and wireless retail broadband markets in the UK are considered effectively competitive. In principle, therefore, there ought to be sufficient choice of provider to discipline firms' behaviour. (Ofcom 2010c, 26)

Ofcom's Claims about the Necessity of Application-Specific Management

In the first several years after the conclusion of the TSR, Ofcom officials also took the position that application-specific management was essential to safeguarding network performance. In 2006, Ofcom explained to the European Commission that ISPs were "routinely degrading" peer-to-peer applications "given the implications this traffic can have on other users of the network" (Ofcom 2006b, 17). Ofcom officials argued publicly that peer-to-peer management was an "example of discrimination on today's internet that is essential for its smooth operation" (Ingram 2006), that without it "the net would be in a sorry state right now" (Ingram 2006), and that peer-to-peer traffic "would otherwise cause the network to gridlock" if its performance were not degraded (Kiedrowski 2007).

These arguments are remarkable in their insistence about the necessity of the specific technique that network operators had most widely adopted, especially given that when network operators took the occasion to discuss traffic management themselves, even they were not publicly framing the issue in such dire terms. Nildram, for example, explained to its customers in 2006 that high-volume use of peer-to-peer applications "sometimes has a knock on effect for other users on the network" (Ferguson 2006). When ntl briefly introduced peer-to-peer limits in 2006, it was "to maintain an excellent quality" service in light of the fact that peer-to-peer applications "may have a detrimental affect [sic] on other users' services" (Towny 2006). While the impact of peer-to-peer users on others was noted, operators were not claiming that the measures taken were absolutely necessary for the network to continue functioning. By comparison, Ofcom officials insisted that peer-to-peer traffic would render the network unusable unless specifically degraded.

Perhaps because large operators had been managing peer-to-peer traffic for several years, Ofcom took for granted that this was a necessary tool to safeguard performance. The regulator failed to acknowledged the cost-based rationales that accompanied operators' decisions, as if the tradeoff between the demands that peer-to-peer traffic put on the network and the ability to meet these demands by deploying additional capacity did not exist. Nor did Ofcom even hint at the possibility that application-agnostic traffic management approaches could achieve the same goals, or that application-specific approaches could be more or less tailored, invasive, or blunt depending on how they were designed. Instead the regulator appears to have observed what was going on in the marketplace and assumed that it was indispensible.

Ofcom's initial and enduring position was that traffic management, whether discriminatory or not, was acceptable and that problematic practices could be remedied by competition (or, as a last resort, by Ofcom's existing regulatory powers). By focusing on competition as a disciplining force and rationalizing peer-to-peer management as a performance necessity, Ofcom made it obvious to operators that their application-specific traffic management approaches would not meet with regulatory resistance. The regulator took no issue with the approaches that were already in place, and if operators did take up traffic management that consumers found problematic, they were expected to "be punished in due course through the market mechanism, not through anything that the regulator might do," as one former Ofcom official explained. The threat of regulatory backlash was essentially nonexistent.

6.3.2 Market Effects of Competition

After the TSR was completed, it was not long before increased competition started having profound effects in the marketplace, most notably on price and speed. The year 2006 became the year of the "free" broadband offer, with TalkTalk, Sky, and Orange all offering free broadband as part of a bundle with other services (Ofcom 2007b). From 2005 to 2008, the average cost of broadband fell by over 40%, with Ofcom attributing around half of that decrease to LLU take-up (Ofcom 2009b).

At the same time, operators were investing to rapidly increase connection speeds. Average headline broadband speeds more than doubled in 2006, from 1.6 Mbit/s to 3.6 Mbit/s (Ofcom 2007b). By 2009 they had nearly doubled again, to 7.1 Mbit/s (Ofcom 2010d). In short, the arrival of LLU made faster speeds available to a larger portion of the population at lower prices. Figure 6, reproduced from Analysys Mason (2010), demonstrates this aptly: as the fraction of broadband subscribers on LLU rose, prices declined sharply, making 8 Mbit/s service affordable where it had previously been prohibitive while also driving prices down for slower offerings.

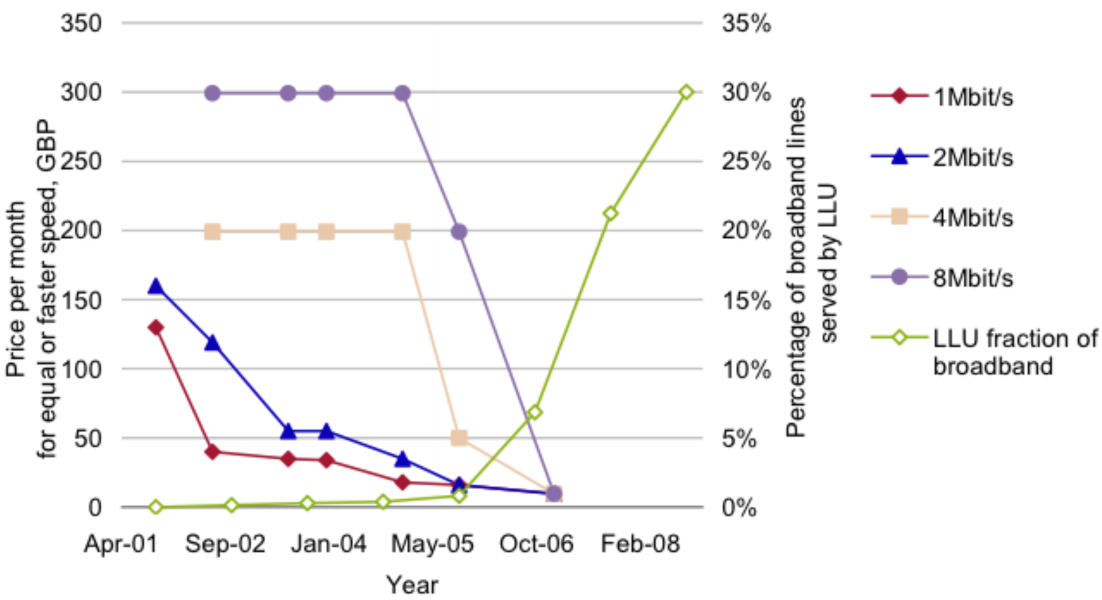

Figure 6. Lowest market price for a UK broadband product of specified speed, 2001-2008. Reproduced from Analysys Mason (2010).

As faster headline speeds became available, Ofcom increased the pressure on operators to continue to deploy new capacity and to deliver the speeds that they advertised. Faced with complaints from consumers and Ofcom's own Consumer Panel (2007), the regulator worked with industry in 2008 to create a voluntary code of conduct requiring signatories to disclose, among other things, estimated actual speeds in addition to headline speeds (Ofcom 2008c). That same year, Ofcom initiated a nationwide speeds testing initiative that used specialized hardware placed within thousands of residences to obtain statistically significant samples of broadband performance across the nine largest fixed retail providers (Ofcom 2009a). Since 2009, Ofcom has reported the speeds testing results annually and continued to add new tests, most of which focus on web browsing performance.

The simultaneous decrease in prices and pressure to increase speeds worked in tension. Operators felt pressure from Ofcom and the market to increase network capacity, which required up-front investment. But offering broadband for free, as operators felt pressure from the market to do, was not much of a strategy for raising the capital necessary for serious network expansion. One DSL product manager described this dynamic aptly in puzzling over a competing ISP's pricing structure: "It's just crazy! We look at our pricing and just can't work out how we do it. And then they're doing it at half the price again . . . this is crazy." These pressures were direct results of the newly introduced competition spurred by LLU.

6.3.3 Application-Specific Management As a Legitimized Practice

The climate of permissibility that Ofcom established around traffic management during its early forays into net neutrality policy combined with competitive forces in the marketplace to entrench application-specific traffic management where it already existed and diffuse it across the industry, even to some operators that initially resisted such measures. The extent to which application-specific management had been legitimized was demonstrated differently by different factions of the largest operators, but by the end of the decade they had all adopted application-specific management in some form, with the key exception of Sky.

Early Adopters Left Application Management Largely Untouched

For operators that had deployed application-specific management in the early-to-mid 2000s, the competitive dynamic and Ofcom's reluctance to intervene reinforced their original traffic management decisions. In the newly competitive environment, expending budget on bandwidth-independent costs made sense: every provider paid the same to hook up an access line, and more access lines meant more customers. But anything providers could do to reduce bandwidth-dependent expenses helped them to compete on price in the marketplace. As Everything Everywhere explained to Ofcom in 2010, "[w]ith both speed and costs driven by the level of peak time capacity provided, networks must optimise the balance between peak time speeds and the impact on retail charges to arrive at a competitive position" (Mooyaart 2012, 2).

This logic applied whether providers were unbundling exchanges (as was the case with TalkTalk, Tiscali, and Orange, for example) or not (as was the case with BT Retail). The cost model for LLU backhaul was more attractive than for bitstream. Operators could deploy their own backhaul and enjoy the associated economies of scale without being beholden to BT Wholesale's nationally averaged bandwidth prices (Analysys Mason 2008). But the competitive nature of the market post-2005 gave even LLU operators that were transitioning customers from bitstream a good rationale to keep their peer-to-peer management in place for all customers regardless of whether they were being served from an unbundled exchange or not. Peer-to-peer traffic was a large fraction of overall traffic; managing it down continued to help operators save on bandwidth costs. The market dynamics that resulted from the growth of LLU simply gave DSL operators more reason to keep their existing traffic management strategies in place.

Operators had so much flexibility to manage applications as they wished that a number of them made virtually no major changes in how they managed applications over the course of the decade even though the capacity of their networks changed radically during that time. One large operator, for example, had instituted a per-user rate limit for peer-to-peer traffic at peak times that in the early 2000s equated to about 3% of a typical user's overall headline speed. Since that rate limit remained largely unchanged – and it always applied to all users regardless of the package to which they were subscribed – by the end of the decade users' peak time peer-to-peer traffic limits were closer to 0.05-0.5% of headline speed depending on their chosen packages. Because there was little external pressure to increase the limit and plenty of cost savings to be had by keeping it as low as it was to begin with, this approach to traffic management did not evolve as the network evolved. At the ISP where participant observation was conducted, traffic management researchers' attempts to convince broadband product managers to relax or reduce peer-to-peer limits were met with significant, ongoing resistance. Some providers did reduce the number of hours per day during which applications were managed, but the mechanisms used for management and the limits in place did not necessarily evolve. Many of the "pretty blunt" and "pretty draconian" approaches that were instituted early on only became more so as time wore on.

A 2009 decision made by one early adopter, BT Retail, provided further compelling evidence of the extent to which application-specific management had been legitimized in the UK. That year BT Retail quietly began limiting video streaming to less than 1 Mbit/s at peak time on its low-end "Option 1" package, which at the time offered an 8 Mbit/s headline speed. BBC staff soon noticed this change and the impact that it was having on the popular iPlayer streaming service (Cellan-Jones 2009a). The dispute between the two corporate giants escalated in the press, with BT publicly seeking payment from video providers to carry their traffic and the BBC accusing BT of hiding its practices from consumers (Watson 2009). After intense scrutiny in the press, BT abandoned the practice in less than a year (the ramifications of this decision are further discussed in Chapter 7).

Despite reversing course, BT's mere attempt to manage traffic in this way was an extremely bold move. BT Retail was the nation's largest ISP. Option 1 was among its most popular broadband products. Video streaming was rapidly becoming one of the most popular applications on the Internet, with Sandvine (2010a) estimating 80% growth in peak-time streaming traffic in Europe from 2008 to 2009. The iPlayer was exploding in popularity (BBC 2010). The fact that BT believed, or even hoped, that imposing limits on such popular applications for so many subscribers would be accepted in the marketplace was a testament to the culture of legitimacy that had developed around application management.

Later Adopters Rationalized Application Management in New Ways

Not all operators entered the post-Undertakings era with application-specific management in place. In the cases of O2 and Virgin Media, the companies had strong reasons for refraining during the early years of serious LLU growth. By the end of the decade, however, they both developed rationales for adding application-specific management to their networks, further demonstrating the diffusion of application-specific management across the industry.

O2

O2 purchased the nation's only pure ADSL2+ network, Be, in 2005 and launched its own line of consumer broadband products in 2007. When it came to traffic management, Be had always been something of a "purist brand," as one interviewee described it, with no management or restrictions of any kind on its network; the company's motto was "life uninhibited" (Be 2005). In the early years after O2's acquisition, this philosophy persisted and was bolstered by the fact that the network was flush with capacity, obviating the need for the kind of cost and performance control mechanisms used by earlier traffic management adopters. With up to 24 Mbit/s service offered across 150 unbundled exchanges at purchase, O2 had a clear speed advantage over competing ISPs whose top-of-line products maxed out at 8 Mbit/s. It also had a smaller user base, yielding minimal contention on the network. As a result, there was no need to invest in traffic management solutions.

The situation changed as the network and the market matured. As with many other LLU providers, O2 offered both "on-net" LLU service and "off-net" broadband service that relied on bitstream access from BT Wholesale. For O2, the disparity in the performance of the two kinds of service offerings was more severe than its competitors since the LLU network provided faster-than-average speeds. Congestion on the off-net network was becoming a serious problem, resulting in the decision in 2009 to introduce application-specific management for its off-net product that combined prioritization of streaming, gaming, and other time-sensitive applications with rate limits for peer-to-peer and newsgroups (Ferguson 2009b), not unlike the approach taken with Plusnet's product offerings years earlier. That move improved customer satisfaction while slowing the pace of bandwidth investment necessary to meet demand.

Impressed with the results of application-specific management on the off-net network, O2 moved to shed all vestiges of Be's purist approach, adopting a strategy that put traffic management at the center of its product offerings. In 2010, O2 unveiled a wholly new product strategy that involved three different packages, each of which had different rate limits for peer-to-peer and streaming applications at different times of day (Ferguson 2010). The low-end package capped peer-to-peer traffic at 50 Kbit/s per user at peak and 100 Kbit/s off-peak and included an 800 Kbit/s limit on streaming video. The only limit on the high-end package was a 250 Kbit/s limit on peer-to-peer at peak times. The theory behind this approach combined the usual rationales for application-based rate limits – cost and performance control – with the notion that traffic management could be used to segment the market, attracting customers who were willing to knowingly buy a product with degraded application functionality at a lower price point.

This was a bold step from O2, which was the sixth largest fixed ISP in the UK at the time. Even after BT Retail had met with resistance and ultimately abandoned its effort to impose streaming video limits on its lower tier packages the previous year, O2 accepted the risks associated with taking the same approach. In doing so, it became the largest ISP to use traffic management as a product differentiator – not just as a means of background cost and performance control – and to have the public accept that strategy as legitimate. Smaller ISPs like Plusnet had operated on a similar model for years. But the fact that O2 – a highly recognizable brand with a fixed network supporting more than 600,000 customers – was able to sustain a product strategy with traffic management at its center demonstrated how legitimized application-specific traffic management had become.

Virgin Media

Virgin Media, meanwhile, was benefitting from a tried and trusted application-agnostic traffic management system together with cost effective network upgrades that increased capacity without requiring significant new investment. The subscriber traffic management (STM) system that Telewest had deployed years prior continued as an effective tool for moderating the behavior of heavy users on a daily basis. Like O2, Virgin offered an "off-net" DSL product that relied on bitstream backhaul and had an application-specific priority scheme that gave higher priority to VoIP, gaming, and streaming than to other applications. But the core cable product that served the vast majority of the user base had no application-specific management.

The years that followed the TSR saw the launch first of 20 Mbit/s offerings and then 50 Mbit/s offerings made possible by the introduction of DOCSIS 3.0 technology on the network. Compared to the up-front investments made by the LLU providers and BT's investments in fiber, the capital expenditure necessary to deliver increasingly higher speeds on the cable network was small. As Virgin Media CEO Neil Berkett explained in 2009, "iPlayer was launched in January of 2008. All of the DSL ISPs sort of went how terrible is this, it is costing us capital. We felt it, but it didn't cost us capital" (Berkett 2009, 2). Moreover, the deployment of DOCSIS 3.0 that year required only "incremental tens of millions of pounds" from which the Virgin network derived "significant upgrades in terms of capability and capacity" (Berkett 2010, 10).

However, even the combination of the well-functioning STM system and relatively inexpensive capacity upgrades were not enough to defend against competitive pressures to minimize operating costs through application-specific management. In 2010, Virgin Media introduced aggregate rate limits for peer-to-peer and newsgroup applications across all of its packages, just as numerous DSL ISPs had done years prior. By applying limits in the aggregate, they affected all users similarly and thus did not suffer from the complexities of the ntl system that had been in use years prior.

Unlike in earlier years, Virgin Media's decision in 2010 was not a matter of finding the applications creating the most traffic on the network and ratcheting them down. Just two months after the application management was introduced, Virgin Media CTO Paul Buttery presented the graphs shown in Figure 7 to analysts, showing that the network had experienced 300% traffic growth in the previous three years and that streaming and web traffic combined to represent more than half of peak time network traffic, while peer-to-peer and newsgroups accounted for less than one quarter.

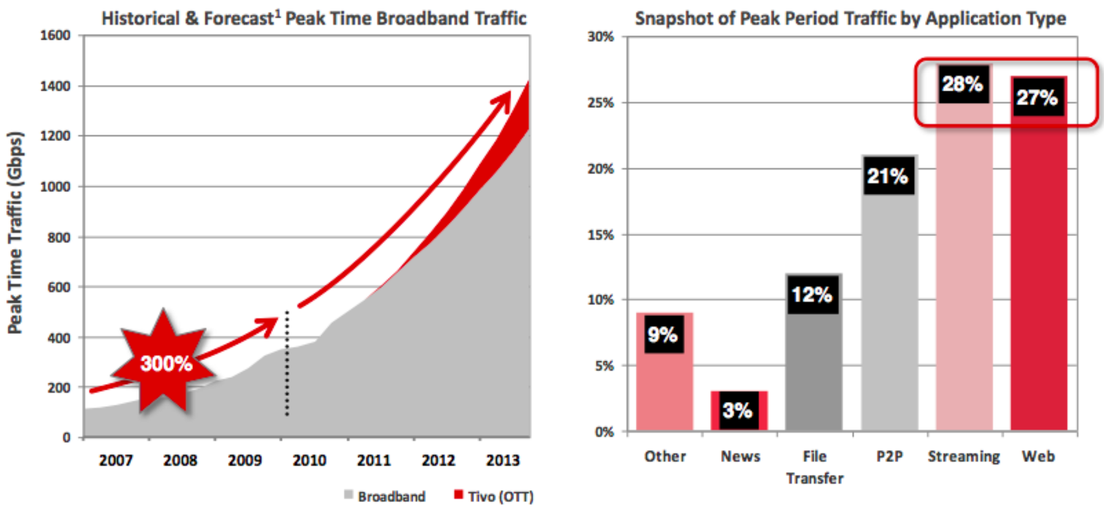

Figure 7. Peak period traffic growth and volume
by application type on the Virgin Media network.
Presented December 2010. Reproduced from (Buttery 2010).

Buttery explained this data as follows:

> [I]f you look at the things that are driving that growth – what we've got there is a snapshot of peak traffic, and we've broken it down into the major types there. If you'd been looking at that chart three years ago, that would have been a picture that was dominated by peer-to-peer. And what we're seeing now is because of the quality of our network and because of the drive towards over-the-top TV, people are taking more and more video streaming from us, and that streaming traffic is continuing to grow and grow. We've also seen web traffic rising. (Berkett 2010, 7)

Clearly, although traffic growth had been substantial, it was not being driven by the applications that Virgin Media started managing in 2010. When those applications were dominant years prior, Virgin had refrained from managing them. In the interim, application-specific management was legitimized by Ofcom even as competition forced providers to offer higher speeds at lower prices. The result was that in 2010, Virgin had the flexibility to take a step that, while not aimed at its biggest traffic source, helped to reduce peak time demand even further than the STM system could.

Although the idea of managing the highest volume applications may have no longer held, many of the other rationales that providers had long used to justify their traffic management decisions did. Peer-to-peer applications were still viewed as less time-sensitive than other popular applications, their management still seemed unlikely to draw customer complaints, and limiting the traffic of modest peer-to-peer users – those whose peer-to-peer usage was higher than average but not enough to trigger the STM thresholds – could still improve performance for other customers on shared links.

Unlike O2, the particular ideals upheld internally at the company created an additional incentive for Virgin to seek ways to manage traffic beyond STM. For Virgin, being able to claim that all of its packages were "unlimited" had always been viewed as extremely important from a marketing and product differentiation perspective. One of the key criteria that had been established through years of advertising campaign disputes before the Advertising Standards Authority (ASA) was that packages with monthly caps could not be advertised as "unlimited." As a result, Virgin was committed to not introducing monthly caps on any of its cable products. There was an ethos about unlimited products that ran deep in the company.

Having taken monthly caps off the table, and with the STM system already in place, Virgin had limited further options if it wanted to reduce peak time traffic demand. Adopting an application-agnostic solution similar to the one Comcast had developed after the FCC intervened was considered to have the same problems as the old ntl application-based management: because it was only triggered when links became congested, the traffic management would affect different customers in different parts of the network differently. Establishing an aggregate rate limit for peer-to-peer and newsgroup applications allowed the products to be described as "unlimited," with a traffic management explanation that applied the same to everyone. These public-facing considerations plus the opportunity to reduce capacity costs spurred the nation's second-largest ISP in 2010 to follow the path that so many of its competitors had adopted long before.

Sky As the Lone Major Differentiator

Similar to O2, Sky entered the broadband market by purchasing a network (Easynet) with large excess capacity, few existing customers, and many unbundled exchanges. Also like Be, Easynet had operated under a corporate philosophy that did not support application-specific management. This had been a key underpinning of the high-end unlimited broadband packages that Easynet offered under the UKOnline brand and was believed internally to be a powerful differentiator in the marketplace. At the point of acquisition, accommodating relatively few customers on an underutilized network made it simple to perpetuate the unmanaged approach from a demand management perspective.

But even as take-up soared and Sky surged from tens of thousands of users at launch to 3.5 million in 2011, that strategy remained unchanged. Sky held firm to its approach of dimensioning the network to accommodate demand, and as a result likely transitioned its backhaul network to higher capacity links sooner than its competitors. Sky offered an off-net product with application-specific management just as its competitors did, but that technology was not applied to the core on-net broadband offering that most of its customers purchased.

Unlike most of those competitors, Sky had a television product that provided it with a major source of revenue. Sky customers were not able to purchase broadband without a TV subscription until 2010 – the first year that broadband was profitable for the business (Darroch 2010). Although the exact figures are not public, it is widely assumed that Sky was able to continue to compete on both speed and price without traffic management or obvious performance problems because it subsidized the cost of broadband capacity with TV revenues.

This was in stark contrast to competitors such as Virgin, whose CEO acknowledged in 2010 that "the most gross margin for us is produced by broadband, followed by fixed line telephony, followed by our B2B and our mobile business, followed by video" (Berkett 2010, 9). Sky did not feel the need to manage cost in the same way as the rest of the market because it was able to cross-subsidize its broadband service. A number of much smaller providers such as Zen and Andrews & Arnold likewise refrained from managing applications, but they did so at a significantly higher retail price point than Sky given that they had no way to subsidize the cost of unmanaged traffic demand.

Sky's unmanaged approach belied Ofcom's claims about the necessity of peer-to-peer management while signaling the extent to which other operators' decisions were at least partially grounded in concerns about cost. By maintaining a network whose performance was good enough to continue to attract new customers, Sky showed that focusing on capacity upgrades rather than traffic management could allow its customers to achieve acceptable performance on the DSL network without the operator needing to take action against peer-to-peer applications. In other words, Sky's existence and growth demonstrated that there was nothing inherent about the impact of peer-to-peer applications on the network that required some response other than capacity upgrades.

Sky also held firm to the belief that lack of traffic management was a marketable quality to consumers, often emphasizing its lack of traffic management in marketing materials. Figure 8 shows an example from an advertising campaign that included the claim that Sky "never slows you down at peak times no matter how much you use" (Hermes Project 2011).

Figure 8. Sky Broadband advertising campaign from 2011.
Reproduced from Hermes Project (2011).

It is difficult to know how much the unlimited, unmanaged aspect of the product has contributed to Sky's growth, but having maintained that philosophy even as every major competitor decided otherwise, the company continued to demonstrate its belief in the virtues of running a nondiscriminatory network. As will be discussed in the next chapter, Sky provides an important practical example to test theoretical arguments about how to safeguard nondiscrimination. If the goal of a regulatory policy is to ensure that most users have access to a nondiscriminatory offering, then the growth of Sky indicates that supporting a competitive retail marketplace may be one strategy for achieving that goal. But if the goal is to ensure wide nondiscriminatory access to applications offered by independent application providers, the UK case should be more accurately viewed as a disconfirming example of the power of competitive forces.

6.4 Conclusion

In stark contrast to the US, the UK broadband landscape has been dominated by traffic management, with application-specific management becoming increasingly pervasive over time. The DSL wholesale market structure and the nature of high-volume applications provided incentives for DSL operators to limit those applications early on, while early experiences with complex application-specific management on cable networks inspired cable operators' preferences for simpler solutions. Over time, competitive pressure, the low risk of backlash from customers, and Ofcom's insistence on both the market's disciplining power and the necessity of peer-to-peer management firmly ingrained application-specific management as an acceptable practice industry-wide. As a result, both DSL and cable operators toward the end of the decade rationalized its adoption and extended its use in new ways. The key exception was Sky, which used lack of management as a product differentiator and found other ways to pay for the costs associated with its customers' unfettered traffic demand.

The implications of the UK experience are profound. The UK case reveals that in markets with a wholesale/retail split, the structure and bandwidth prices built into the wholesale market can have significant effects on traffic management decisions. It demonstrates how competitive forces – in absence of any countervailing pressure from customers or regulators – can intensify operators' incentives to manage applications. It provides a testament to the influence of the regulator not only in assuring the industry of its regulatory restraint, but also in bolstering competitive conditions. And it shows how much inertia can surround traffic management technology once it has been deployed in an environment where neither customers nor regulatory authorities demand – or even discuss – how the technology might evolve as networks evolve. The next chapter draws out these implications for the net neutrality policy debate by comparing and contrasting the US and UK experiences.

Chapter 7.
Discrimination, Competition, and Innovation in the Field

7.1 Introduction

Building on the examinations in Chapters 5 and 6 of how and why network operators made their traffic management decisions, this chapter analyzes the implications of UK and US experiences as they relate to key debates found in the net neutrality literature. It focuses on three central questions highlighted in the literature review in Chapter 2: (1) whether network operators use discriminatory traffic management to control performance and/or cost, to diversify their product offerings, or for anti-competitive purposes; (2) whether competition serves as a deterrent to discrimination; and (3) whether discrimination creates barriers to application development and innovation. Because discriminatory traffic management has been more prevalent in the UK than in the US, this chapter draws more evidence from the former than the latter, but reflects the experiences of US ISPs as appropriate.

Findings from the US and the UK are predictable in some cases and unexpected in others. Performance and cost were the driving factors behind many application-specific traffic management deployments, but interesting exceptional cases arose in the UK where operators sought to differentiate their products or use traffic management as a lever in business negotiations. While the competitiveness of the UK market yielded some nondiscriminatory options for consumers, relying on consumer switching behavior to provide more comprehensive competitive discipline was insufficient for a variety of reasons, including the presence of switching costs. Finally, the process of correcting errors in the technology used for application-specific management revealed the costs that application-specific management created for application developers and innovators. In unpacking each of the three debates from the literature, the evidence collected from the two countries revealed a wealth of nuance and detail that has heretofore been lacking from the academic discourse concerning net neutrality.

7.2 Rationales for Traffic Management

The net neutrality literature points to three core types of rationales to explain the use of discriminatory traffic management: to control performance and/or cost, to segment the broadband market, and to disadvantage competing applications. A number of conclusions can be drawn about the prevalence of these various rationales and the interactions between them in the cases of the US and the UK. Experiences in the two countries show that controlling performance and cost was a key factor driving adoption of discriminatory traffic management in most cases, but that the cost of traffic management equipment itself was as important for many operators' decision-making. In the UK, BT's brief management of video streaming demonstrated an additional cost-based rationale: using traffic management as a lever in a business negotiation about the costs of video traffic.

For a handful of ISPs, it was clear that no anti-competitive motivations were present given the competitive dynamics of those providers' service offerings. In many other cases, particularly where peer-to-peer and video streaming services were affected by traffic management, a number of factors complicate the assessment of whether anti-competitive motivations spurred the adoption of discriminatory strategies. This section analyzes how the rationales put forth in the literature manifest themselves in the US and the UK.

7.2.1 Performance and Cost

Both the academic literature and the net neutrality policy discourse are replete with arguments about the benefits of discriminatory traffic management for controlling performance, costs, or the combination of both. Discrimination is said to enhance performance by allowing ISPs to prioritize latency-sensitive applications over delay-tolerant applications (Bartlett et al. 2008; Brito et al. 2010; Hahn, Litan, and Singer 2007; Owen and Rosston 2006; Peha 2007; Sandvig 2007; Zinman et al. 2007). Some argue that discriminatory traffic management is a necessity that networks cannot function without (Everything Everywhere 2010; Ingram 2006; Kiedrowski 2007; Renda 2008; Singer 2007; Speta 2002a; Yoo 2005; Zinman et al. 2007).

Others acknowledge that achieving particular network performance comes at a cost, and that the necessity of traffic management is driven by the need to meet customer demand within the confines of ISPs' budgets (Crocioni 2011; Glover et al. 2010; Hazlett and Wright 2011; Mooyaart 2012). Whether or not they improve performance, limitations placed on particular applications can also help to control backhaul and interconnection costs by reducing overall traffic loads (Marsden 2010; van Schewick 2010).

Differences Between Network Technologies and Markets

As the previous chapters demonstrated, enhancing or maintaining performance was a key motivator for many ISPs that deployed application-specific traffic management (particularly peer-to-peer management) in the US and the UK, although these concerns manifest themselves in different ways in the two countries. In the UK, DSL operators became early adopters of peer-to-peer and newsgroup management solutions because these solutions provided a way to keep their networks from congesting that was less expensive than buying more bandwidth from BT Wholesale. Keeping the cost of backhaul bandwidth down without allowing the user experience to deteriorate was a prime motivator that endured even after the spread of LLU and the appearance of new entrants. Notably, the management of high-volume applications on both DSL and cable in the UK focused on limiting downstream transmission rates, since the majority of the capacity (and therefore bandwidth cost) was in the downstream direction.

The US cable operators that deployed peer-to-peer management solutions were also motivated by the need to improve performance, but the nature of both their problems and chosen solutions was different. Most cable operators were focused on reducing upstream congestion, and they therefore chose peer-to-peer management solutions that limited the number of upstream connections. They were less concerned about controlling the overall growth of their backhaul bandwidth than with managing localized congestion problems in their access networks, which were the most expensive portions of the network to be constantly upgrading (one partial exception was RCN, which cited transit costs as one of several rationales for its peer-to-peer management). The difference in approach between the two countries' operators was related not only to the physical differences between cable and DSL technologies, but to market dynamics: while both cable and DSL operators in the UK justified their decisions on the basis of downstream bandwidth costs, by and large neither cable nor DSL operators in the US did so.

Costs of Traffic Management Equipment

The academic and policy discourses tend to focus on the cost reduction benefits associated with traffic management solutions while ignoring the costs of implementing those solutions. In both countries, the cost of the equipment used to identify and manage specific application traffic (usually DPI equipment) had a significant impact on traffic management decisions, both for ISPs that adopted application-specific approaches and those that did not.

Where discriminatory traffic management survived for many years – on DSL networks in the UK – it had a strong business case to back it up. As one engineer explained, "the cost of managing and investing in [DPI] was a better return than continuing to buy bandwidth." The two choices were viewed as interchangeable: whichever one cost less was the one that could be justified. Another engineer estimated that the savings his network accrued from having deployed DPI for traffic management amounted to 15% of the company's broadband budget. This included not only bandwidth savings but equipment savings as well, as the server equipment used to connect the ISP to the BT Wholesale network was a significant expense. Reducing bandwidth demand by installing DPI boxes – which cost an order of magnitude less than the servers – meant fewer total servers that would need to be purchased in order to satisfy the bandwidth demands from the same number of customers.

Some US cable operators that adopted the TCP reset solution viewed it as "very cheap." Compared to the hundreds of millions of dollars being spent to upgrade to DOCSIS 3.0, the DPI equipment cost one company "[on] the order of single millions of dollars to deploy across [the] whole network." Some ISPs realized savings by purchasing lower-end DPI equipment with simpler processing technology than what might be found in higher-end switches or routers; one engineer described the equipment as "just a PC with 2 Gbit NICs [network interface cards] in it."

ISPs that chose not to invest in traffic management equipment, or that later abandoned it after a few years of use, perceived the cost/benefit trade-off very differently. Some of the difference in perception is attributable to differences in network size or architecture; the expense associated with DPI was naturally less for smaller networks or those with fewer aggregation points where DPI would need to be deployed to manage traffic effectively. For larger networks, 2 Gbit/s DPI equipment was not going to suffice. One engineer who had been involved in repeated unsuccessful attempts to build a business case for DPI described the difficulty of "trying to find a DPI solution that scales" to accommodate "networks handling a few hundred gigs a second." Instead of realizing equipment savings, ISPs operating at such scale required DPI equipment that cost more per gigabit than their routers or switches. Engineers from across the spectrum of large ISPs – both cable and telco, US and UK – described repeatedly reviewing proposals and business cases for application-specific management technology that never got deployed for having been too expensive. One former DPI salesperson explained the challenge from the vendor side: "It's hard to make a business case to a carrier to put $50 million worth of equipment in their network to [apply] QoS . . . when they can't charge for it." Relatively speaking, application-agnostic solutions, such as expanding bandwidth, Comcast's FairShare, and Virgin's subscriber traffic management (which made use of existing cable system features) came with price tags that were easier to justify.

Beyond network size and architecture, ISPs differed as to their fundamental outlook about the long-term value of application-specific management. In contrast to the idea that DPI and bandwidth investment were interchangeable solutions to the same problem, ISPs that abandoned or never deployed application-specific traffic management felt that the value to be gained from a DPI deployment – whatever its cost – could not be rationalized compared to the kinds of things that ISPs were accustomed to spending money on, including network expansion. One US engineer aptly summarized this idea:

> If we were going to install DPI everywhere right now, it would probably be – would certainly be more than $100 million, maybe $125 or $150 million depending on the kind of discounts you got. That's crazy, that's a crazy amount of cost. I think most people internally when we talked about things like that, [thought] if you're going to spend that much money, why not spend it on network capacity, customer care, promotions to get more customers – there's a million other things that would be higher priority to spend more on than that.

Whether because of equipment costs or other reasons, numerous large ISPs have successfully offered broadband without application-specific management. Thus, as a general matter, prioritizing latency-sensitive applications or limiting delay-tolerant ones is clearly not essential to offering a successful broadband service. The approaches taken by Sky, the large US telephone companies, and the large US cable companies since 2009 have disproven this argument. Achieving particular performance at a particular cost has certainly motivated those that chose application-specific management, but many ISPs have chosen otherwise.

Costs of Video Streaming

Cost-based rationales also played an important role in ISP decisions about video streaming. In 2009, within a highly charged environment that had seen multiple public disputes between the BBC and the UK ISPs over the costs associated with the BBC iPlayer video streaming service, BT Retail imposed an 896 Kbit/s limit on streaming video for the customers on its "Option 1" broadband package (sold at "up to 8 Mbit/s" speeds). This particular use of traffic management revealed how traffic management could potentially be used as a bargaining chip in negotiations concerning traffic costs.

Tension Between the ISPs and the BBC

Well in advance of the launch of the BBC's iPlayer video service in 2007, BBC executives and staff had been meeting regularly with executives from the large ISPs to discuss the potential impact of iPlayer traffic on their networks. While those meetings were usually productive and collegial, at times the tension between the ISPs and the BBC over iPlayer traffic growth spilled into the press. Several months prior to the iPlayer's official launch in 2007, a number of press reports cited ISPs (most notably Tiscali) threatening to place restrictions on iPlayer traffic unless the BBC agreed to pay for at least some of the bandwidth costs associated with the delivery of that traffic ("Net Firm Warns on Web Video Costs" 2007; Fluendy 2007; Murray-Watson 2007; Palmer 2007; Williams 2007). Although BT was mentioned in some of those reports, the company denied having any such complaints about iPlayer at the time (Williams 2007).

As the iPlayer continued to gain popularity, the private discussions continued, with the BBC's Director of Future Media and Technology, Ashley Highfield, occasionally highlighting them on the BBC Internet Blog (Highfield 2007a; Highfield 2007b; Highfield 2008a). In one post, Highfield detailed a 19-point plan for clarifying the relationship between video services and broadband costs, which included the following provision:

> Content providers, if they find their content being specifically squeezed, shaped, or capped, could start to indicate on their sites which ISPs their content works best on (and which to avoid). I hope it doesn't come to this, as I think we (the BBC and the ISPs) are currently working together better than ever. (Highfield 2008a)

The ISPs were not pleased about the perceived threat to "name and shame" them (Orlowski 2008). In the view of one ISP policy official, the BBC management was "very much trying to shape the agenda" on the issue. Another former ISP executive explained his internal response upon reading the blog post: "ultimately, I'll manage my network the way I want to, and if my customer doesn't like it, there's plenty of other suppliers. [My company] doesn't need [the BBC], as a publicly funded body, to start moving the market around." The debate continued in the press, with Tiscali railing against Highfield's "inflammatory comments" (Williams 2008) and Highfield responding by emphasizing both the "strong relationship" between the ISPs and the BBC and his company's unwillingness to pay for the delivery of iPlayer content (Highfield 2008b).

BT's Decision

BT executives had been involved in the private discussions – but not the public spats – concerning who should be responsible for bearing the costs of video traffic. In April 2009, BT instituted its video streaming limit on Option 1. The cost of providing the capacity necessary to support iPlayer at a decent quality clearly factored into this decision. At the time, iPlayer video was using an adaptive encoding that would sense the network conditions and choose from one of three bitrates: 1.5 Mbit/s, 800 Kbit/s, or 500 Kbit/s. By limiting video streaming to 896 Kbit/s, Option 1 customers could in theory still make use of the bottom two bitrates, but BT could potentially reduce the capacity of many streams by almost half by foreclosing the highest bitrate option. Furthermore, in the absence of sufficient capacity to meet all users' demands for 1.5 Mbit/s streams (or streams at higher rates offered by other video service providers), the rate limit was viewed as a fairness mechanism: "a means of ensuring that everybody got some experience" of iPlayer, as one engineer put it.

The public discourse surrounding BT's change to Option 1 points to a strong connection between capacity costs, business maneuvering between BT and the BBC, and the video streaming limit. Although an explanation of the limitation on video streaming was provided in BT's fair usage policy on its web site in February 2009 (Ferguson 2009a), it did not gain wide public attention until it was reported in June on the BBC News web site by Cellan-Jones (2009a). The limitation was particularly noticeable to the engineers working on iPlayer because it appeared to be knocking users down to the 500 Kbit/s bitrate, even though the stated policy was to allow streaming up to 896 Kbit/s, which should have allowed for streams at 800 Kbit/s. One BBC employee explained that an internal team would regularly "look quite closely at the traffic going out to various different networks to see what's going on," and this was what they noticed when looking at the BT network:

> [T]he one to look at . . . is the 1500 Kbit/s stream. And what you can see is as the traffic management kicked in, at about 6 o'clock, then the number of people who can get that stream goes right down. And you see the number of people who get the 500 Kbit/s stream. Because they said they were clamping it to 864 [sic], so the 800 should have been making it through. Actually what we found was that it wasn't making it through at all, so that was being impacted. And in fact, 50% of users were getting the 500 Kbit/s, which was kind of a bit rubbish.

After the initial news story, the BBC continued to report on the issue and began to engage with BT spokesmen about the ISP's response (Cellan-Jones 2009b; Cellan-Jones 2009c). This exchange culminated with a Financial Times report in which BT Retail's managing director, John Petter, explained that BT could not "give the content providers a completely free ride and continue to give customers the [service] they want at the price they expect" (Bradshaw 2009). At that point, the private discussions between the ISPs and the BBC had been going on for years, but this marked the first time that BT had publicly sought payment from the BBC (and other video providers as well) for the traffic costs associated with video streaming. In this way, the limitation placed on Option 1 video streaming could be viewed as a means of leveraging the ISP's control over the network in its business negotiations with the BBC. One interviewee explained that the ISP's traffic management decision was "about a relationship issue with the BBC." Perhaps the traffic management technology in use had a more severe impact than BT had been expecting, or perhaps Petter was particularly irked by the fact that the BBC itself first reported the story (Clark 2009). Regardless, there was clearly a link between the cost of iPlayer and other video traffic and BT's decision to impose the streaming limit. The traffic management decision was cast as a bargaining chip in discussions over who would bear that cost.

7.2.2 Market Segmentation and Anti-Competitive Motivations

Beyond performance and cost control, the net neutrality literature points to two other potential rationales for discriminatory traffic management: segmenting broadband markets and undermining competing services. Some commentators have suggested that by using different traffic management strategies for different broadband products, ISPs can offer a menu of products that may suit the needs of Internet users better than a suite of product choices that are all offered on a nondiscriminatory basis (Litan and Singer 2007; Marcus 2008; Renda 2008; Valcke et al. 2009; Weisman 2010; Yoo 2004; Yoo 2005). The market segmentation rationale implies that an ISP either offers some products with discriminatory traffic management and others without it, or that discriminatory traffic management is applied differently to different products, such that the applications that receive priority treatment vary between products, for example. Market segmentation also requires that these differences be advertised to potential customers – otherwise they would not have the information needed to select the broadband product most appropriate to their needs.

The use of discriminatory techniques to disadvantage competing services is a core concern that motivated calls for regulatory intervention (Atkinson and Weiser 2006; Crawford 2007; Crawford 2013; Herman 2006; Lemley and Lessig 2001; Weiser 2008). Since broadband providers often offer services that face potential competition from independent application providers, they may use traffic management to reduce the quality of the independent applications so as to drive more customers to their own service offerings (Economides 2008; Greenstein 2007; Knieps and Zenhausern 2008; Marcus 2008; Nuechterlein and Weiser 2005; Peha 2007). For a particular ISP's traffic management to be used in an anti-competitive fashion, that ISP must offer some service that potentially competes with an application whose performance is degraded as a result of the traffic management. Conversely, if traffic management solutions have a positive or neutral impact on independent applications that compete, the motivation for deploying those solutions cannot be said to be anti-competitive.

Obvious examples of potentially competing applications include VoIP and video streaming, which can offer substantially similar functionality and content as ISPs' own offerings of voice service, linear television, or video-on-demand. Because peer-to-peer applications are used to download movies and television shows, they may likewise serve as potential competitors to ISPs' own video offerings.

Assessing whether ISPs acted with anti-competitive intent would be a complicated task even for antitrust investigators with access to internal business documents and corporate correspondence. Such materials were not available for this thesis and are unlikely to be made available to researchers in the future. Nonetheless, some observations concerning anti-competitive motivations and their relationship to market segmentation can be made on the basis of the evidence collected from the US and the UK.

Market Segmentation in the UK

In some instances it is possible to conclude that traffic management has been deployed without anti-competitive intent. The UK operators Everything Everywhere, O2, and Plusnet have all offered broadband packages that limit the transmission rates of peer-to-peer applications or video streaming (or both) while not offering their own video services. These limits could not have been motivated by a desire to undercut independent video providers to the benefit of the ISP's own video offerings, since they had no such offerings. Instead, as the previous chapter demonstrated, O2 and Plusnet sought to segment their customer bases by offering different combinations of application-specific traffic management strategies on different broadband packages and explicitly marketing those differences to consumers. Everything Everywhere did not use traffic management as a product differentiator, but rather applied the same application-specific traffic management to all of its residential broadband customers: a combination of limits on peer-to-peer and newsgroup traffic and prioritization of VoIP and gaming traffic designed to help the company control performance and cost.

Everything Everywhere and Plusnet also provide obvious examples where undermining voice competition could not have reasonably been a rationale for employing the traffic management strategies they chose. Both companies used priority queuing systems where VoIP traffic was prioritized over other traffic. In Plusnet's multi-level priority system, VoIP was prioritized at the highest level (see, for example, Plusnet (2012c)). This is the opposite of what one would expect from an ISP that was attempting to disadvantage challengers to its own voice services.

Enhancing performance under constrained bandwidth provides a more sensible rationale in these cases.

Although BT's rate limiting of video only applied to Option 1 and not users of BT's other broadband products, it could not have been considered a market segmentation tactic. The limit was disclosed in BT's fair usage policy, but it was not advertised to consumers or disclosed in the terms and conditions to which new Option 1 customers had to agree. Consumers could not efficiently sort themselves into the appropriate product categories if they were unaware of the differences between the products.

Peer-to-Peer and Video Streaming Management

The dynamic wherein application-specific traffic management served to benefit VoIP was not necessarily restricted to cases where ISPs gave VoIP explicit priority at network bottlenecks, as Everything Everywhere and Plusnet did. As explained in Chapters 5 and 6, part of the reason that ISPs were drawn to the idea of limiting peer-to-peer and newsgroup applications was because the high volumes of traffic that they tended to generate could impair the performance of other applications on the network, and VoIP in particular. One US cable engineer described how the impact of peer-to-peer traffic on customers of Vonage, a leading US VoIP provider, became an impetus for adopting peer-to-peer management:

> [A]round 2005-6, Vonage started to get really big. We had constant calls from Vonage management, customers pissed off, that we were degrading Vonage's service. . . . [S]omething was going on on the network, we didn't know quite what, that was impacting VoIP and other real-time applications. And so we were sort of feeling our way through that. We didn't really know much about it. But we knew that if we did some stuff in the network we could make that experience better. Demonstrably better.

Similarly, after Cox's widely publicized trial of a congestion management system that imposed rate limits on high-volume applications during times of congestion, the company confirmed that the system "had a positive effect on upstream time-sensitive traffic such as gaming and over-the-top VoIP calls" (Wilson et al. 2010, 29). Thus imposing limits on peer-to-peer and other high-volume applications, even in the absence of explicit prioritization for

VoIP, was arguably having a positive effect on independent voice providers in competition with the ISP for voice customers.

However, unlike in the cases of Everything Everywhere, O2, and Plusnet, the analysis of other ISPs' competitive motivations for imposing limits on high-volume applications is complicated by the fact that they also offered their own video services. While it is possible that their traffic management approaches may have been deployed in part to benefit VoIP users, the same could not be said for the applications whose performance was being degraded. Whether peer-to-peer management was conducted for the purpose of disadvantaging potential video competitors has been widely debated in the US policy discourse, particularly with reference to Comcast's use of the TCP reset solution (Ammori et al. 2007b; Brill and Taubman 2008; Martin 2008d; McDowell 2007). While it is clear that operators in both the US and the UK took up peer-to-peer management to help control performance and cost, the data gathered as part of this study does not allow for conclusions to be drawn about whether they were additionally motivated to undercut applications that were perceived to be in competition with their own video offerings.

In the case of BT's video streaming limitation, BT did have its own digital video offering, BT Vision, at the time when the limitation came into force. Vision offered a selection of broadcast and premium television content, including a limited selection of catch-up iPlayer content. Given that the video streaming limitation on Option 1 affected some content and services that were highly comparable to those being made available through Vision, this instance of traffic management has perhaps the strongest potential out of all those studied to be construed as anti-competitive. As with the cases of peer-to-peer management, more data would be necessary to understand whether strategic behavior was involved in BT's decision to limit video streaming.

7.2.3 Summary

Many of the instances of application-specific traffic management observed in the US and the UK were motivated by a desire to control performance, cost, or both. The traffic management solutions that ISPs ultimately pursued in some cases reflected the differences between the two countries' market structures and the differences in physical technology between cable and DSL. The cost of traffic management equipment itself was an important consideration for all ISPs, whether they chose to adopt application-specific solutions or not. Given that many large ISPs have successfully offered broadband service without application-specific management, it is clearly not a prerequisite for offering broadband products that are acceptable to consumers.

BT's decision to limit video streaming and the highly charged context in which it was made reveal an added dimension to ISPs' cost-based motivations: the desire to use traffic management as a lever in business negotiations with content and application providers. While this situation may be viewed as highly specific to the special position that the BBC occupies in the UK video market, at a more general level it demonstrates the importance of the market context and the relationships between ISPs and application providers in understanding why networks are engineered in specific ways.

Finally, while the UK provides examples of the explicit use of traffic management for market segmentation and specific instances in which anti-competitive motivations were provably absent, in general more data would be necessary to draw general conclusions about the competitive interactions between operators' own video offerings and the peer-to-peer and video-focused traffic management solutions they deployed.

7.3 Relationship Between Competition and Discrimination

The second key debate concerning net neutrality relates to whether competition can deter broadband providers from discriminating against applications. One broadly held view is that because discrimination can impair application performance, ISPs in competitive markets will be reluctant to take up discriminatory strategies for fear of losing customers (Becker, Carlton, and Sider 2010; Cave and Crocioni 2007; Chirico, Haar, and Larouche 2007; Faulhaber and Farber 2010; Hahn, Litan, and Singer 2007; Nuechterlein 2009; Shelanski 2007). This premise is fundamental to both the EU telecommunications regulatory framework (OJL 337/11, 2009) and Ofcom's regulatory approach (Ofcom 2011b), and it is what drives regulatory interest in increased transparency about broadband operators' practices. Unsurprisingly, many ISPs have argued publicly that competition serves to deter discrimination (Atherton 2010; C. J. Brown and Boucher 2007; Casserly, Wallach, Alvarez, Walker, et al. 2008; ECTA 2010; Sky 2010; TalkTalk Group 2010a; Telefónica 2010).

Others question whether competition can reliably deter discriminatory conduct and whether it can do so sufficiently to safeguard application innovation. Even in competitive markets, consumers may not understand their choices or the relationship between discrimination and the service they experience (Marsden 2007; van Schewick 2007; van Schewick 2010; Wu 2003). They may encounter switching costs – financial, logistical, cognitive, or otherwise – that deter them from switching when they otherwise would have done so (Bar et al. 2000; Economides 2008; European Commission 2010; Krafft and Salies 2008; Lennett 2009; Marsden 2010; Skype 2010; van Schewick 2007; van Schewick 2010; Wu 2007). Furthermore, because consumers are unlikely to realize and account for the fact that future innovations may not materialize as a result of discriminatory conduct, their choices in the broadband market cannot be relied on to reflect a preference for nondiscriminatory networks (Lemley and Lessig 2001; Lessig 2001; van Schewick 2010; van Schewick 2012).

With a half-decade of experience under highly competitive conditions, the UK broadband market provides a rich case study for understanding whether either of these lines of argument are supported in practice and the nature of the relationship between consumer preferences and discriminatory traffic management. This section draws four conclusions based on evidence from the UK:

- It is unlikely that many consumers have switched their broadband provider because of application-specific management;
- Because application-specific management has not been a primary concern when choosing a broadband provider, competition only serves to safeguard the most popular applications;
- When ISPs are ambivalent about the harms that consumers experience as a result of traffic management, competition does not encourage nondiscrimination; and
- Barriers to switching broadband providers have been higher than for other telecommunications services.

7.3.1 Lack of Consumer Awareness and Understanding

For competition to serve as a check on discriminatory traffic management, consumers need to know that such traffic management is occurring. Consumers might be motivated to switch ISPs when they notice a decline in application performance (Casserly, Wallach, Alvarez, Waz, et al. 2008; TalkTalk Group 2010a). But identifying the cause of a particular performance problem can be difficult, as described by one ISP policy official:

> [T]here's a hell of a lot that goes on. How long is your [DSL] line? The line speed that you get depends on length of line and whether you've got your vacuum cleaner on at the same time. Then there's how much backhaul goes in. . . . Another thing is what the site on the Internet [that the person is visiting] can deliver and another would be the effect of the traffic management policy. So there's a whole bunch of stuff and it is quite difficult for customers to distinguish what makes what effect.

This assessment aligns with the findings of a 2012 study conducted by Consumer Focus, a consumer representation and advocacy organization created by the UK government. The study, the only one of its kind to focus specifically on traffic management, involved in-depth interviews with 32 broadband users who had varying levels of technical sophistication. Respondents in the study demonstrated no awareness that slow network speeds might be the result of traffic management or that they might experience different performance from different broadband providers. They instead attributed slow service to the network being busy or to their computer equipment being outdated (Kisielowska-Lipman 2012).

The potential for this confusion and lack of awareness is one of the reasons why policymakers have focused extensive attention on the transparency of traffic management disclosures. Transparency is a central feature of the EU telecommunications package (OJL 337/11, 2009), Ofcom's approach to net neutrality (Ofcom 2011b), and the FCC's *Open Internet Order* (FCC 2010b). In addition to helping customers shop for new services, accurate disclosures could potentially help users that experience degraded service to understand whether that degradation might be the result of traffic management. In theory, disclosures could also help consumers who value having a generally nondiscriminatory network even if their own application usage is unaffected (or benefits) from application-specific traffic management. They could use the information to choose providers that run nondiscriminatory networks.

In practice, most consumers have difficulty finding and understanding traffic management disclosures. In the Consumer Focus study, participants engaged in a product searching exercise in which they were asked to comparison shop for broadband using the real web sites and disclosures of the UK's largest ISPs. Participants had extreme difficulty finding the relevant information and understanding it when it was found. The study concluded that consumers do not understand the most basic of terminology – including the term "traffic management" itself:

> Consumers struggle to imagine how the term 'traffic management' could apply to an internet service, and most are unable to link the expression to the amount of usage on a network. Some of the vulnerable consumers and light users actually think 'traffic management' relates to information about road traffic. (Kisielowska-Lipman 2012, 18)

A London Economics study commissioned by Ofcom in 2011 found similar results (London Economics 2011). Participants were given usage profiles for hypothetical broadband users and asked to choose the most suitable broadband packages (with traffic management information included) to match the profiles. The study found that participants "chose the incorrect package for their usage profile in a large proportion of cases, irrespective of the type of and how the information is provided to them" (London Economics 2011, 2).

UK broadband providers are well aware of this lack of consumer understanding. One ISP official estimated that less than 1% of the UK population "understands that they're even being traffic managed," despite the fact that the majority of UK broadband users were subscribed to packages with application-specific management at the time: "They have no understanding whatsoever that they're being managed off the park between 6:00 and midnight every day."

Large ISPs engage in frequent market research that explores the impact of various kinds of information on consumers. One ISP product manager described this experience as follows:

> [W]e've done a lot of research on this . . . we got a research group to go out and speak to customers, what do they understand, should we call it traffic management, should we call it fair share, should we call it, you know, fair management? Is "threshold" right? Is "speed reduction" right? What sort of language – how transparent should we be, how interested are they? And I sat . . . behind one of those one-way mirror things just watching customers go, "Do they do that? Oh the bastards! I can't believe it! Why do they need to do that? God . . ." But then most customers don't actually know what speed they're on, you know. "Are you with [our ISP]?" "Yeah I am with [your ISP]." "What speed are you on?" "Oh I'm on the 7 meg product." It's like, we don't even have that product. And it's sort of head in your hands.

Notably, all of these observations came at a time when more information about traffic management was available, and in standardized formats, than at any previous time since the advent of broadband in the UK. Prior to the 2008 publication of Ofcom's first voluntary code of practice for broadband speeds (Ofcom 2008c), traffic management disclosures were generally not uniform and in some cases not even available. This began to change with the code of practice's recommendation that ISPs publish traffic management information on their web sites. In 2011, the broadband industry went a step further in adopting a uniform "Key Facts Indicator" that was used to display traffic management information in a standardized format across the web sites of many ISPs (Broadband Stakeholder Group 2011). If consumers were so devoid of traffic management understanding after all of those efforts, it is hard to imagine that consumer understanding was any better at any previous point in time.

In sum, consumers largely do not understand traffic management, whether because they cannot distinguish its effects from other factors that impact performance or because the information that exists about traffic management is not comprehensible to them. As a result, it is unlikely that many UK consumers have switched their broadband provider on the basis of traffic management.

7.3.2 Secondary Role of Traffic Management in Switching Decisions

Even if more consumers had understood the existence and effects of traffic management, it is not clear that traffic management would have been a primary factor in choosing a broadband provider. Employees from every large UK ISP interviewed for this thesis agreed that most consumers care far more about the price and overall speed of a broadband package in making a purchasing decision than any other factors. A survey of several thousand broadband users conducted by Ofcom in 2008 revealed similar findings: price was by far the feature that was most frequently compared when choosing among broadband providers (cited by 67% of respondents), followed by speed (45%), reliability (31%) and customer service (31%) (Ofcom 2009a). In Ofcom surveys conducted yearly between 2009 and 2012, consumers who had considered switching but chose not to consistently listed "no cost benefit" as a reason for not switching more often than "satisfied with provider," indicating that price was a top concern (Ofcom 2009c; Ofcom 2013b).

This is unsurprising given that broadband advertising and public policy put a strong emphasis on price and maximum or average speeds, with little emphasis on traffic management and its impacts. As noted in Chapter 6, the arrival of LLU inspired fierce price competition, with several operators introducing "free" broadband offers and many others emphasizing low prices in their broadband marketing. ISPs usually distinguished the products in their product lines by maximum download speed, price, and usage cap (Ofcom 2010b). Reflecting on marketing developments since 2005, one ISP official explained that "all that's happened is it's just gone faster and cheaper."

The focus of the UK government and Ofcom on maximum speeds reinforced the idea that it should be of prime importance to consumers. The government framed its goals for broadband in terms of download speed: originally 2 Mbit/s to every home (BIS and DCMS 2009) and more recently 25 Mbit/s to 90% of the UK population (Richmond 2011). Ofcom's findings from its speed-testing program likewise emphasized download speeds (Ofcom 2009a; Ofcom 2010d). In the 2008 survey mentioned above, Ofcom asked respondents about their perceptions of the impact that nine different factors might have on broadband speed – distance to the exchange, computer processing speed, and so on. Traffic management was not put forth as one of the choices (Ofcom 2009a).

This is not to say that the impact of traffic management (for those who understood it) was not a secondary concern, but it did not often rise to the level of importance necessary to trigger a switching decision. A policy executive working for an Internet application company whose product had been discriminated against in the UK and elsewhere explained this aptly:

> [T]he current competition between the ISPs or mobile operators can only be sufficient for the very few Internet services that are huge enough and indispensible enough, such as Google search, that consumers would switch tomorrow if they didn't have you. We're not even in that category. I think Facebook and Google are probably the only ones in that category, where if you didn't have them, the vast majority of users would just say, "you're crazy" and move with their feet. Even us at [my company], we're clearly not important enough in consumers' minds.

With consumers reasoning this way, competition among broadband providers might have safeguarded applications that consumers already valued very highly – applications whose performance was satisfactory enough for consumers to enjoy. But for applications that have yet to obtain deal-breaker status, or for new applications that consumers have yet to even discover, discrimination against them would not necessarily be eradicated through market discipline, and could therefore serve to prevent those applications from ever performing well enough – and becoming so important to consumers – that they could trigger a switching decision. This is a specific manifestation of the argument put forth by Lessig and van Schewick that consumer preferences would not adequately reflect the importance of application innovation: it is not that consumers do not care at all about less popular or established applications, but that they do not care enough to switch their ISP over them. These findings also comport with a number of formal economic models that show that discriminatory regimes can harm niche applications and their users while benefitting large, established providers and their users (Bourreau, Kourandi, and Valletti 2013; Guo, Cheng, and Bandyopadhyay 2012; Reggiani and Valletti 2012).

Although this weaker version of competitive discipline that protects only the most popular, established applications has not been the subject of significant attention from pro-competition advocates in academia or public policy, it has not been lost on all ISPs. At the ISP where participant observation was conducted, employees reported that when broadband product managers were confronted with the idea that eliminating discriminatory traffic management might improve the overall application development environment on the Internet, they were unsympathetic, preferring to hear about solutions that would make customers happier or cut costs. TalkTalk, in responding to Ofcom's consultation about net neutrality and traffic management, explained the virtues of competition as follows: "if access to . . . services / content *were important to customers* and they wanted unrestricted access, then any ISP which blocked or degraded access would see subscribers defect" (emphasis added) (TalkTalk Group 2010a, 6). The clear implication is that whether competition acts as a safeguard against discriminatory treatment of a particular application relates to whether that application is already valued by consumers.

7.3.3 Competitive Discipline When Unwanted Customers Are Affected

Despite most consumers not understanding traffic management or caring about it enough to switch ISP, a small minority of UK Internet users did highly value having a broadband connection that does not discriminate against specific applications. Because peer-to-peer traffic management has been so common, heavy peer-to-peer users provide a useful minority group to study. Estimates of how many UK broadband users are consistent users of peer-to-peer applications vary; on the network where participant observation was conducted, for example, significantly less than 1% of customers were considered heavy peer-to-peer users. Data collected for an Ofcom copyright infringement study in 2012 indicated that at least 4% of UK broadband users (about 800,000 people) had downloaded or shared files via peer-to-peer applications in the three months prior to the study, although the study did not differentiate between casual and consistent peer-to-peer users (Kantar Media 2012).

In interviews, network operators broadly observed that many consistent users of peer-to-peer applications chose to buy broadband from providers that had reputations for not throttling peer-to-peer traffic. As one network engineer described it, “when we didn’t have traffic management in place, when you Googled and looked on message boards and things, there was definitely a [theme] of ‘go to [my company], they’re doing nothing.’ And the percentage of users was ridiculous.” This small and sophisticated user population was capable of attributing the peer-to-peer degradation they experienced to traffic management and finding the providers that might supply better service.

Network operators were aware of this purchasing behavior, but they were somewhat conflicted about how to approach this particular customer segment. As discussed in Chapter 6, peer-to-peer usage tended to generate significant traffic-related expenses for ISPs, which created the motivation for instituting peer-to-peer management in the first place. A network that was known to allow peer-to-peer applications to run at maximum speeds would be a magnet for “the very expensive users,” “the problem types,” “the mad peer-to-peer guys,” as ISP employees described them. Did operators want to attract these customers?

For some, it was “a question mark.” One policy executive’s musings about his company’s heavy peer-to-peer customers revealed the inherent tension between wanting to maximize the size of the customer base and minimize expense: “To be fair, we’ve got quite a few actually. Because our policy is quite nice, actually. It’s not that stringent. So unfortunately, we’ve got quite a few. But, okay, fair enough.” For the ISP in question it was difficult to say whether these customers were desired or not, but clearly they came to the ISP for having known about the lax traffic management policy. Other operators were more willing to cast these customers off. As one engineer explained, "you want to manage [peer-to-peer traffic] and if anything the attitude has been, well, if they get a bit annoyed and they don't like that, I mean, okay . . . they can go somewhere else."

This uncertainty and even ambivalence towards heavy peer-to-peer users highlights the difference between predictions in the literature about the effects of competition and its practical effects in the marketplace. Rather than producing a generally nondiscriminatory market for broadband access with a handful of providers exceptionally deploying application-specific management (Litan and Singer 2007; Sidak 2006; Yoo 2004), the reverse occurred: only a few providers were willing to become magnets for "expensive" peer-to-peer users by not managing peer-to-peer traffic. The rest were willing, perhaps with some reluctance or equivocation, to let those customers go elsewhere. The small size of the affected user base and the large expense associated with its usage patterns meant that most ISPs either valued the benefits of peer-to-peer management over the harms it caused, or did not even view the associated loss of customers as a harm to their businesses. Even where many ISPs compete for customers, there can be no competitive discipline of discriminatory conduct if ISPs are ambivalent about or actively in favor of losing the customers harmed by that conduct.

This result contradicts both the stated views of policy stakeholders and previous research findings. One former Ofcom official explained how the switching behavior of consumer subgroups was expected to cause broadband providers to respond:

> [I]n most consumer markets, you have, anyway, always a lead group of people who switch and allow the other customers to benefit. And there's always a sizable chunk of the population where there's a lot of inertia. . . . Most people don't even understand their tariffs for most of these things. So, you know, you have to expect that some consumers are much more sophisticated and you are relying on those leaders to make sure that the rest of the population benefits from it, rather than expecting, you know, everybody to switch depending on their usage.

This logic is supported experimentally in the finding of Sluijs et al. (2011) that if only a subset of consumers are well informed about the quality of their own broadband connections, the average quality of broadband offered in the market increases. The idea is that as well-informed consumers (or "leaders" of some sort) switch to providers of nondiscriminatory broadband offerings, this will prompt other providers to offer nondiscriminatory choices as well. Some commenters have claimed that this group need not necessarily be large; as TalkTalk explained to the European Commission in 2011, "for transparency and switching to be an effective discipline on the behaviour of ISPs does not require all customers to understand the impact of traffic management and/or switch in respond [sic] to a harmful practice. In reality, if only a small number of customers respond to a particular traffic management practice by leaving the ISP then it will render that practice unprofitable" (TalkTalk Group 2010b, 9).

Perhaps the validity of this claim depends on just how "small" the population is that switches, but in the peer-to-peer case there are tens if not hundreds of thousands of users in the nationwide pool of customers significantly affected by peer-to-peer management. That some of them gravitated towards the handful of ISPs that were not managing peer-to-peer traffic (and took their "expensive" usage behavior with them) clearly did not render peer-to-peer management unprofitable for the rest of the large ISPs. Nor did it inspire a general trend toward nondiscriminatory offerings in the marketplace. In fact, the opposite occurred, with more large ISPs taking up discriminatory traffic management as time went on. Thus, at least in the case of peer-to-peer management, the decisions of a fraction of consumers did not create consequent benefits for the rest of the broadband customer population. More ISPs tolerated their departure than recruited them to stay.

7.3.4 Barriers to Switching

Finally, although this study did not include data collection concerning specific switching costs, it is worth analyzing the extensive consumer switching data collected by Ofcom. Indeed, judging based only on the amount of time and resources that Ofcom has spent studying consumer switching behavior and improving relevant regulations, it is reasonable to conclude that significant barriers to switching existed when the agency first took up the issue in 2006 and that barriers persisted in the years that followed. Between 2006 and 2012, Ofcom conducted three public consultations, ran two separate working groups with industry participation, hosted two workshops, and conducted or procured at least 15 separate studies with the involvement of seven different outside research firms on the topic of consumer switching.

Many of those studies involved surveying a representative sample of the UK broadband user population. Among consumers surveyed in Ofcom's survey research who have switched their communications provider, more consumers have consistently found it difficult to switch their broadband service than their telephone, TV, or mobile services (Ofcom 2006c; Ofcom 2009c; Ofcom 2013b). Notably, more consumers who have never switched perceive it to be difficult than those who have switched: 22% compared to 15% across surveys conducted annually between 2008 and 2012. Among those who have considered switching broadband providers but decided not to, between 15% and 33% listed "hassle" as a reason for not switching in various years, while between 17% and 20% listed contractual reasons (Ofcom 2009c; Ofcom 2013b). Thus there is some evidence of lock-in, cognitive barriers, and logistical costs preventing switching (Tambini 2012).

As opposed to the perceived difficulty of switching, the perceived likelihood of a disruption in service during a switch is far less than it is in reality. While 11% of broadband users in a 2010 survey who had not switched expected a disruption in service, more than twice that – 27% – actually experienced a disruption when they did switch, which again was the highest number out of telephone, TV, and mobile (Ofcom 2010g). The average period without service across all switchers was 12 days. This is an alarmingly high figure in light of the fact that by 2011 nearly three quarters of the UK population considered the Internet to be an essential information source (Dutton and Blank 2011), a figure that grew steadily over the course of the decade (Dutton, Di Gennaro, and Millwood Hargrave 2005). Giving up an essential service for multiple weeks at a time can surely be viewed as a significant barrier to switching.

7.3.5 Summary

Assessing whether competition has provided sufficient discipline against discrimination in the UK broadband market depends on what one considers to be sufficient. Sky and several smaller operators provided UK consumers with nondiscriminatory choices. Although some operators chose to manage video streaming and other popular applications, the predominant applications targeted – peer-to-peer and newsgroup applications – were in heavy use by only a small fraction of the population. Thus, if the goal of competition is to ensure that some nondiscriminatory choices exist and that consumers can choose among broadband providers that do not discriminate against applications they already value, that goal was met in the UK market.

If instead competition is intended to support a broadband market where nondiscrimination dominates and discriminatory conduct is relegated to the margins through consumer rejection, the UK market clearly came up short. Most consumers did not understand or concern themselves with application-specific management enough to make a switching decision on that basis. For those who did want to switch, barriers persisted throughout the decade, including a significant incidence of service disruption associated with switching even as broadband became more central to people's lives. The applications discriminated against were niche enough that many large providers did not feel the need to compete over how well or poorly those applications performed. Some ISPs even considered the loss of customers associated with application-specific management as a benefit, thus nullifying the power of competitive discipline.

Theories about the externalities associated with nondiscrimination assume that Internet users experience the benefits of innovation indirectly, not that a small number of users experience benefits directly on ghettoized networks. But whether an outcome like that in the UK is considered to be harmful to innovation relates at least in part to whether one believes that discrimination has a dampening effect on innovation and application development. That is the subject of the next section.

7.4 Impact of Discrimination on Application Development

The third and final topic from the literature to be explored – the interaction between discrimination and innovation in Internet applications – has been at the crux of the net neutrality discourse. Many advocates for regulatory intervention have argued that network operator discrimination erects barriers that make it more difficult for application developers to innovate and seamlessly deliver their products to Internet users. With discrimination taking place, the claim is that application providers would need to seek permission from network operators in order to get their applications to work properly or with enhanced quality (Economides 2008; Lennett 2009; Lemley and Lessig 2001; Wu and Lessig 2003). One possible consequence of application-specific traffic management is for developers to seek ways to evade it, resulting in an "arms race" where network operators and developers continually attempt to circumvent one another, using significant engineering resources on both sides in the process (Lehr et al. 2006; Marsden 2010; Sandvig 2007). Whether arms races occur or not, discrimination generally reduces application developers' incentives to innovate because it vests the control over which applications succeed or fail in the hands of network operators (Lemley and Lessig 2001; van Schewick 2010; Wu 2003).

The research conducted for this thesis did not focus on application developers, but observations and interviews with ISPs and others in both the UK and the US revealed an unexpected wealth of insights about the impact of application-specific traffic management technology on application developers and the landscape of application development. This section demonstrates how discriminatory traffic management creates difficulties for application developers, how attempts to overcome those difficulties have required a significant investment of resources from both developers and ISPs, and how ISPs sought to perpetuate this model despite acknowledging its detrimental effects on application innovation.

7.4.1 Difficulties with Traffic Management in Practice

Nearly every ISP included in this study that used deep packet inspection equipment for application-specific traffic management reported difficulties with the technology once it was deployed. The problems they encountered at times affected the applications that were the intended targets of the traffic management, but more often had adverse effects on other applications that were never intended to be rate limited, blocked, or otherwise managed. Thus application-specific management caused collateral damage for a wider space of applications than intended by the ISPs. Because the traffic management discriminated between applications, its errors were discriminatory as well, degrading the performance of some applications but not others.

The problems encountered with traffic management technology were diverse. Sometimes traffic was correctly identified, but the treatment applied to it did not function as expected, as discussed earlier in the case of the BT Option 1 video streaming limit. In some cases, applications were misclassified despite the fact that they used different application protocols from the ones targeted for management. One ISP executive described how VPN applications would sometimes be miscategorized, causing them to be de-prioritized and resulting in complaints from users trying to access corporate networks from home. After numerous user reports (and articles in the press (Bangeman 2007; Wellman 2007)) about how Comcast's TCP reset solution for peer-to-peer management appeared to be resetting Lotus Notes connections, Comcast explained that its systems had "inadvertently" affected the corporate messaging application (Casserly, Wallach, Alvarez, Waz, et al. 2008, 38).

ISP employees were generally less than pleased about the extent and frequency of this kind of misclassification error. One US engineer's assessment was that "there are a lot of sort of half-brained DPI implementations out there doing stuff that is probably absolutely weird and unexpected." Another gave a stark assessment of the DPI platform in use on his network:

> [I]t was completely shocking how poor the analytics and detection capabilities of the platform were. Like what a – it would be too strong to say what a Mickey Mouse sort of system it was – but the system had these classifiers so it would be able to characterize traffic as certain types of things. And they were often not just wrong, I mean *completely* wrong, where we'd look at data and be like, "this data can't possibly be correct." Then a day later we'd notice that the report looked different from after that conversation [with the vendor], like a day afterwards. [We'd ask them,] "What happened?" "Well, we found a bug in our classifier, so we changed it." . . . We joked that it's like the magic eight-ball on the network. It's like, "What kind of traffic is that?" "Ahh! Peer-to-peer!" "What kind of traffic is that?" "Web streaming!" It's a different answer every day.

As frustrating as these DPI deficiencies were to ISPs, the real victims of these errors were application developers whose products were being needlessly and inexplicably affected. Not only would application traffic be mistakenly rate limited or otherwise managed, but the effects of the mistakes would come and go as DPI vendors sought to improve their classifiers, making diagnosis of the situation from the application developer's point of view even more complicated. Similarly, at the ISP where participant observation was conducted, the team in charge of the traffic management policy was constantly experimenting with different application-specific rate limits and time windows in which those limits were applied. Application developers would have had no straightforward way to know whether these fluctuations had to do with their own products or interference from ISPs.

A more subtle but equally vexing type of problem arose from the fact that the same application protocols were often used by both applications that ISPs intended to manage and those that they did not. One interviewee described a situation where a popular video streaming application shifted its traffic from HTTP to the more secure HTTPS to support users of Nintendo Wii game consoles, which only accepted HTTPS traffic. This created problems for customers of an ISP whose traffic management policy involved throttling all HTTPS traffic, under the assumption that HTTPS was usually used for non-interactive activities (e-commerce, for example). The ISP did not intend to throttle streaming video, but that was the result.

A more common example involved the use of peer-to-peer protocols by gaming companies. A number of UK ISPs encountered problems because game updates and functionality would be distributed using BitTorrent and other common peer-to-peer protocols that the ISPs generally rate limited during peak hours. One ISP product manager explained that "every now and again we'd make a colossal error," the largest of which was accidentally managing World of Warcraft, a popular online game created by Blizzard Entertainment. The response to that mistake from customers was so severe that it "just felt like a catastrophe."

Similar incidents involving gaming applications were discussed at the ISP where participant observation was conducted. Researchers at the company would occasionally use these incidents to try to demonstrate to broadband product managers that peer-to-peer protocols had a diversity of uses, some of which had extremely loyal customers who would post angry messages in online forums when their gaming experiences were degraded. Such forum postings, and responses from ISP customer service agents, became common across the largest UK ISPs (Jackson 2011; "Re: BT Throttling P2P Traffic!!!!!!!!!!!!!!!!!!!!!" 2010; "World of Warcraft Latency and Lag Spikes" 2011). This particular type of misclassification was clearly widespread; customers of the Canadian ISP Rogers encountered enough problems with World of Warcraft and other games over an extended period of time to prompt regulatory authorities to intervene (Thompson 2011). The lack of sophistication of DPI classifiers had obviously detrimental effects specifically on gaming applications and their users.

Some of the attitudes expressed within the network operator community about these kinds of problems reveal a fundamental disconnect between the way that operators view the Internet as a platform for application development and the way that application developers and others might view it. If providers assume that "HTTPS is normally just used for banking," that "the problem is that some gaming companies use peer-to-peer signatures," or that "we have problems when . . . [gaming providers] start using the same protocols that the BitTorrents of the world use," as interviewees claimed, it puts the blame for the miscategorization on application developers rather than on ISPs for choosing to discriminate based on guesses about the applications associated with particular traffic. This is exactly the opposite view of application providers (and net neutrality advocates), who assume that they can use any open, non-proprietary application protocol and have their applications perform in a predictable way. The root of the problem is not their choice of application protocol, but the fact that ISPs have chosen to interfere with particular application protocols for traffic management purposes. Some ISPs, having made significant investments in DPI platforms, came to view the application behavior as abnormal rather than viewing their own interference with application traffic that way.

In some cases, the DPI equipment simply failed to classify traffic. One UK ISP interviewee described having “three months where it just didn’t work.” Another explained that his company had chosen a US-focused vendor, “so everything they do is focused on the US in terms of new protocols coming out, new products, etc. . . . so things like Spotify, iPlayer, all of that which are problems for us, they have no interest in. So it took probably, I would say, six to nine months for them to even acknowledge that yes, Spotify existed.” The Spotify streaming music service had launched in the UK prior to launching in the US.

While these problems did not necessarily degrade the performance of applications that were not targeted for traffic management, they further illustrate how network equipment could fail to live up to the promise of accurately distinguishing one application from another – and how long it could take to develop accurate classification capability for even a single application. Similarly, a class action settlement related to Comcast’s use of the TCP reset solution revealed that Comcast had misclassified Lotus Notes for up to six months (Brill 2009). Rogers customers in Canada encountered problems with gaming applications for well over a year (Rosen 2012). For a small application provider seeking to rapidly expand its user base, it is reasonable to believe that degraded performance over half a year or more could cause irreparable damage to its chances of success.

Many of the delays resulted because ISPs were beholden to vendor update schedules. One of the frustrated US engineers quoted above explained that he had asked to see the vendor source code that was causing the errors, but "that was like the crown jewels, we couldn't see that." Instead, he simply had to wait for the vendor to fix the problem on its own timetable. At a UK ISP where the operator's relationship with the vendor was more collaborative, engineers had begun developing their own code to recognize applications because "the vendors have release cycles, we get asked to do something next week, we can't say that it takes three months because we need to get the vendor to do it." Introducing a third party's technology as a factor that determined the performance of specific applications clearly complicated the task of running the network, sometimes so much so that operators chose to bypass the third party rather than wait for a fix.

Deploying and operating DPI for traffic management purposes was, in short, a "tortuous process," as one interviewee called it. The technology for classifying applications was immature and in some cases inadequate. Resolving problems often required lengthy waits as vendors cycled through multiple attempts at fixes and batched their product updates. All the while, application providers – many of whom were never intended to be the targets of traffic management in the first place – were faced with degraded and fluctuating performance for reasons entirely outside their control. They had built applications using their choice of available application protocols and under the performance assumptions that those protocols provided. By interfering with those protocols, ISPs created costs that those application providers were forced to bear.

No technology is perfect, and it would be unreasonable to expect that any traffic management system, whether application-specific or application-agnostic, could be deployed without operators later discovering bugs in need of fixing. But what differentiates application-specific management is the fundamental premise that network operators (and the DPI vendors they work with) can distinguish applications targeted for management from all others accurately enough to achieve some desired traffic management effect without prior coordination with application providers. Taken together, the classification errors of the sort described in this section represent not just transient software bugs, but an underlying weakness in the notion that any network-based classification engine can stay ahead of the multitude of ongoing, uncoordinated innovations and application changes taking place at the edges of the network. Furthermore, discriminatory traffic management turns network operators into arbiters of application-specific network performance. When the arbiters make mistakes, applications suffer the consequences differentially. As the next section demonstrates, operators have sought to correct these errors by involving application providers in their decision-making – creating a cost to application innovation that would not exist were it not for application-specific management.

7.4.2 Engaging Developers to Resolve Traffic Management Problems

ISPs would often be alerted to classification errors when customers posted complaints in online forums or sent emails to customer service, or occasionally when application developers notified the operators of performance problems. Several UK ISPs reacted to this information by attempting to engage developers of the misclassified applications in crafting solutions to better identify their traffic. In some cases, this engagement became extensive. One ISP product manager named five gaming providers – including four of the world's largest game makers – that his company was "proactively" working with out of a "constant sort of desire to make sure that we are not inadvertently traffic managing something that we shouldn't be." This engagement was an unwanted cost borne by application developers that was a direct result of the use of application-specific management.

As perceived by network operators, the willingness of application developers to help ISPs sort out these problems appears to have varied, both over time and from developer to developer. While some companies were described as being "happy to speak" to individual ISPs, others were "quite difficult to deal with." One interviewee noted that some developers seemed to react to user complaints with the "kind of attitude where 'it's not our problem, go and talk to your ISP'" because they were "fundamentally anti-traffic-management as a principle." Upon recognizing the latest in a series of problems related to the classification of World of Warcraft traffic, a Virgin Media support forum manager exposed the reticence of some gaming companies:

> We appreciate that some customers will have noticed a similar issue with the previous World of Warcraft update. The reason behind this is because gaming companies are not prepared to share the updates with Virgin Media or traffic management suppliers prior to its release and so the first time we see the new packets is when people start to use the new updates. We are trying to change this view point of the gaming companies however at present they are un-willing [sic] to work with us. (Wilkin 2011)

Application developers had multiple reasons to be skeptical. Sorting out the DPI problems, both from a technical and a logistical standpoint, could be resource-intensive. The product manager who had worked with the gaming companies described how the process to "get to the right team, the right person, set up the right call there, [and] get the [application] signatures" from a gaming company could take four or five months and how it could be "a real struggle to actually get to the right person who gets what we're talking about and who's happy to help." Having not planned to engage with network operators about these kinds of issues, application developers did not necessarily have the organizational structures in place to respond, and they may have been reluctant to invest in creating those structures to solve what they perceived to be the ISPs' problems. For some ISPs, the issue of misclassification took on enough organizational significance so as to require substantial institutional investment, at times involving the most senior level executives; making use of cross-functional teams comprised of product managers, engineers, and customer service agents; and creating ongoing engineering investigations, holding weekly conference calls, or calling emergency meetings with vendors to diagnose emerging problems and develop solutions. If engaging with an ISP was going to require even a fraction of that amount of time and resources, application providers may have preferred to brush off the ISPs' requests – particularly if that investment would need to be replicated for multiple different ISPs using different technologies and vendors. Small developers simply might not have had the resources to spend on that level of engagement, as one ISP policy executive readily admitted:

> I think to me the point is the friction point. There's a kind of administrative overhead involved in making these adjustments and allowing traffic through. If you're Google, it's fine. They've got enough people to go and sort it out. If you're one man in a garage . . . obviously that's a real problem. Actually, you look at, I think, Twitter, Google, Facebook, they all started [as] two people in a garage. And you don't want to kill off the ability to do that . . . basically you start up for almost zero cash cost.

Moreover, some application developers may have been reluctant to do anything to encourage further entrenchment or expansion of application-specific traffic management. By helping ISPs solve these problems, application developers would have been implicitly legitimizing a technology that network operators took up entirely of their own accord, that application developers never asked for and had no control over, and whose primary impact was collateral damage to the developers' products and users. They also had no guarantees that ISPs and vendors would not leverage their improved ability to identify the developers' products to intentionally degrade their performance if traffic management policies were to change in the future. From that perspective, depending on the size of the application user base affected by a particular misclassification, it would have been entirely legitimate for an application developer to refuse to engage with an ISP in the hope that the drawbacks of application-specific management would eventually cause the operator to shift to a nondiscriminatory approach. Being "fundamentally" against the notion that ISPs and DPI vendors should be determining the performance of specific applications makes perfect sense from a developer's perspective.

Nonetheless, some application developers did engage. For example, in responding to repeated messages from game players about latency problems associated with individual ISPs, a Blizzard technical support forum manager explained that, "[w]e have devoted tremendous time and resources to monitoring, isolating, and eliminating any possible contributing factor within our immediate control . . . . Of course, this hasn't stopped us from attempting to work with every local (and some foreign) ISP who has contacted us for further information" (Brianl

2011). Both diagnosing problems and helping ISPs to resolve them required a substantial investment.

ISPs viewed the path of seeking assistance from application developers as a reasonable way forward, and in the case of some UK ISPs as a strategy worthy of substantial expansion. The product manager who had been working with the gaming companies spoke about not having "anywhere near as many gaming companies on board as we want" and plans of trying to expand "to the whole gaming industry." He faulted his company for being too reactive to misclassification problems, explaining that "what it actually is going to need is a quite proactive, almost tiger team to go out and just establish [relationships]. Thinking, you know, Netflix is going to be massive, let's go and speak to them, let's go and speak to such-and-such game company. Let's really find a strategy for reaching out to these companies to make sure we stay ahead of it." Another wanted to see more direct engagement between application providers and DPI vendors:

> For me the key thing, really, is we've got to get application developers and providers with a much better relationship with the vendors, the Allots of the world and the Ciscos. And actually, well, accept that ISPs are going to do this. You're not going to stop it. So that way we can get a much more open dialogue of actually talking to each other about, well, "We're going to be launching this [application], this is what the traffic looks like, so you can take it into account." Because all that happens in the current circumstance is that we get aggravation from some customers, the gaming providers get aggravation, and the network vendors get aggravation from both of us, saying, "Can you help us sort it out?" So that doesn't seem like a – I wouldn't call that a virtuous circle.

Such are the visions of a discriminatory future: where a multitude of application developers are in constant contact with ISPs as their applications get updated and as new applications get introduced; where application developers expend resources to understand how DPI classifiers work so that they can help ISPs and vendors craft application signatures; and where developers engage directly with the technology companies whose equipment is used to degrade application performance. While these visions may never be fully realized, the seeds of this future were at least partially sown when application developers began to engage with ISPs to resolve misclassification issues, incurring the kinds of costs to application development and innovation that some academic and policy stakeholders had predicted. For application developers, particularly small upstarts with few resources to spare, it was a lose-lose situation with barriers in both directions: they could either sink time and resources into solving the problems that ISPs created for them for the sake of restoring their own products' performance, or decline to engage in the hope that ISPs would find their own solutions or shift to nondiscriminatory strategies before their product performance became a casualty of imperfect DPI technology.

7.4.3 Commitment to Discrimination Despite its Impact on Innovation

In some cases network operators remained committed to discriminatory traffic management despite being aware of its impact on application development and innovation. As acknowledged by the policy executive quoted above, application-specific management "creates a friction" that small application developers may not be capable of bearing. Another interviewee explained the situation of many UK operators as one where "control is taken away by the ISP." An internal report supplied by one UK ISP gave a more blunt assessment: "The deployment of DPI has made innovation with P2P technology almost impossible." Operators thus understood the consequences of their choices, but preferred not to acknowledge that understanding publicly. The policy executive explained this calculus forthrightly around the time when Ofcom was finalizing the publication of its official approach to net neutrality (Ofcom 2011b):

> T]o be honest with you, this point about "permission to innovate" stuff that [Ofcom is] going on about now: hey, if I was writing a biased document for ourselves, I wouldn't include it. Because it kind of goes against what we do. But, hey, I'm not going to disagree with it because you can't disagree with it.

The impact on innovation was known, but ISPs were not prepared to discuss it publicly, much less change their traffic management strategies to mitigate it.

Operators were also well aware of the "arms races" taking place with application developers. The report cited above detailed how such races had unfolded with developers of peer-to-peer applications:

> There are risks in imposing 'single-sided' bandwidth management practices. Any control choices taken by ISPs can be counteracted by control choices made by other stakeholders, which can result in arms races which are counterproductive for all parties involved. A classic example is traffic generated by Bit Torrent [sic] P2P file-sharing. Most ISPs are shaping this type of traffic during peak hours to mitigate issues of network overload. Unfortunately this has led to a proliferation of mechanisms to encode, encrypt and manipulate P2P which make it a lot harder for ISPs to detect P2P traffic. Every measure service providers introduce to exert their control over the network leads to developers trying to circumvent them.

The previous sections also demonstrate a more subtle incarnation of a race to circumvent, albeit one in which application providers were reluctant to engage. As exemplified by the cases of Virgin Media in the UK and Rogers in Canada, network operators would work to correct DPI classification errors that affected games and other applications they did not intend to manage, only to have those applications alter their application behavior and once again be misclassified. Rather than peer-to-peer or other application developers reacting to advances in traffic management technology, it was the operators and their vendors that were in the reactionary position, constantly needing to adapt their DPI as games and other applications changed without warning. In both cases, the presence of application awareness caused each side to waste resources adapting to the other. One ISP interviewee questioned the utility of this arrangement, just as an application developer might have: "it does frustrate me significantly that these guys know they are going to be launching this, so why don't we just prevent it, rather than deal with it afterwards?" His preferred solution was to have applications provide their signatures to vendors in advance, but to application providers, foregoing the use of application-specific management in the first place would have been the simpler prevention mechanism.

In short, operators knew that discrimination was harmful to application innovation and that it required ongoing investment from both sides to adapt to constant technological changes. They knew that the situation was not perfect, but they much preferred to retain the performance and cost control benefits of application-specific management than to forego them for the sake of application innovation. The policy executive reflected this balance aptly:

> If somebody came to us, or [if] Ofcom came and said, “we scan all the new developments in innovation going on and we find that there’s actually a bunch of new stuff coming here,” we’d probably respond to it. Because the last thing we want is somebody to see a brilliant new application, go home, and they can’t use it or the experience is terrible. That’s going to reflect badly on us. There’s no point in us bothering to do it. It’s about us having the knowledge, I guess, to not do that blocking. The alternative is we don’t do any traffic management at all.

It is precisely because it is impossible for any individual or company to predict in advance where new application innovations will come from or which ones will take hold that some stakeholders have argued in favor of nondiscriminatory networks. With nondiscriminatory traffic management, ISPs would not need to be able to predict which applications would succeed – or how their application signatures would be interpreted by DPI equipment.

7.4.4 Summary

The practical experiences with the use of DPI for application-specific management across a range of ISPs provide a wealth of detail to inform what has been a primarily theoretical debate about the impact of discrimination on application development and innovation. It is clear that discriminatory traffic management creates costs for application developers, whether because of intended performance degradation or inadvertent traffic classification. Not all applications suffer equally and those that make use of application protocols that are commonly targeted for management tend to bear the brunt of DPI’s faults. Some ISPs place at least part of the blame for these problems on application developers for using commonly managed application protocols and not sharing information about their application updates in advance.

Application developers have choices about how to resolve these problems, but all of them raise barriers to further application development and innovation. Some operators have sought developers’ ongoing assistance in improving traffic classification, and some developers have complied. The alternative is for developers to see the performance of their applications suffer on particular networks. Operators are aware that this puts application developers in a bind, but the benefits they derive from perpetuating application-specific management outweigh their concerns for how it may be affecting the environment for application development. They are willing to take certain steps to minimize the impact of application-specific management on developers, but not so willing that they would rethink their overall approach to traffic management.

7.5 Conclusion

The UK and the US provide important lessons for understanding how net neutrality policy arguments are manifest in practice.

Do operators adopt application-specific management to cut costs, improve performance, differentiate their products, or disadvantage competitors? Many operators were clearly motivated by the ability to control performance and bandwidth costs, but an important balancing factor was the cost of traffic management equipment itself, which some ISPs could not justify. ISPs in both countries have demonstrated that whether because of equipment costs or other reasons, it is entirely possible to run a successful broadband network without application-specific management, countering claims to the opposite in the academic literature. Whether those that did take up application-specific management did so for anti-competitive reasons is largely left as an open question.

Does competition deter discrimination? Certainly not in the UK in the way in which it has been envisioned in the literature, relegating discriminatory conduct to the margins. Instead, most consumers did not understand traffic management or use it as a basis for switching. Those who did do so comprised a group perceived to be small or insignificant enough that most ISPs did not seek to factor them into their product decisions, despite some consumers' complaints about traffic management. Competition may have effectively safeguarded the performance of the most popular applications, but not the nondiscriminatory application development environment as a whole.

Does discrimination create barriers to application innovation? It clearly created costs for certain developers whose products were differentially affected by application-specific management. Operators and the technology vendors who served them were not always able to meet the fundamental requirement that defines application-specific management: distinguishing one application from another. As a result, application providers either had their product performance degraded, expended resources to help the companies that were the source of that degradation, or both. Operators acknowledged these detrimental effects on application developers, but remained committed to application-specific management.

With a solid understanding of these implications, the next chapter returns to topics of regulation and policy, comparing and contrasting how the regulatory environments in the US and the UK produced distinct traffic management outcomes.

Chapter 8.
Regulatory Reputation

8.1 Introduction

The preceding chapters have demonstrated that when it comes to understanding why network operators have foregone application-specific traffic management, the subtleties of the regulatory environment were more influential than the presence of competition. In the US, the history between the regulator and the operators and the threat of regulation created through the FCC's initial forays into net neutrality policy created a deterrent against discrimination, while intense competition and regulatory permissiveness in the UK spurred adoption of application-specific management. Understanding why, as a general matter, a particular national broadband market is or is not characterized by widespread discriminatory conduct therefore requires understanding the roots of regulators' approaches. What is it about the FCC and the US regulatory environment that led the agency to intervene in formal and informal ways? What is it about Ofcom and the UK regulatory environment that caused the regulator to forbear from any significant intervention? The answers to these questions provide rich new data to allow the extension of theoretical frameworks concerning industry regulation. They are the subject of this chapter and the one that follows.

Because the FCC and Ofcom differ along so many dimensions, the key task of a comparative inquiry is to identify which of the differences help explain the regulators' approaches to traffic management. The first section below provides a brief sketch of the surface characteristics that differentiate the two agencies: governance structure, age, remit, core competency, politicization, and independence. These differences fall across the spectrum of categories of explanations for regulatory behavior outlined in Chapter 3: institutional design choices, external forces, internal agency characteristics, and nation-specific factors. The rest of this chapter and Chapter 9 focus on three particularly salient lenses through which the behavior of the FCC and Ofcom can be understood based on the evidence gathered in this study. The first is reputation, an internal agency characteristic that plays vastly different roles within each of the two regulators and is closely intertwined with institutional design choices about agency governance structure and the politicization (or lack thereof) of each agency. That is the subject of this chapter. The influence of judicial review (or more precisely the regulator's approach to litigiousness) and the presence or absence of interest groups are examined in Chapter 9.

8.2 Differences Between the FCC and Ofcom

The FCC was established in 1934 with the broad remit of "regulating interstate and foreign commerce in communication by wire and radio so as to make available, so far as possible, to all the people of the United States . . . a rapid, efficient, Nationwide, and world-wide wire and radio communication service with adequate facilities at reasonable charges" (47 U.S.C. 151). It is the agency that sets national telecommunications policy (Shaffer and Jordan 2013). It has both legislative-style rulemaking authority and judicial-style adjudication authority. It operates as a five-member commission (reduced from seven in 1983), with commissioners appointed by the President and confirmed by the Senate for five-year terms (reduced from seven years in 1986) with a three-term limit. No more than three commissioners share the same political affiliation and all Commission rules and adjudications require affirmative votes from the majority of commissioners. The chairman is designated by the President and is generally expected to serve for a single Presidential term, although on average FCC chairmen have served for just 2.7 years (FCC 2013). The FCC staff, by contrast, tend to be primarily civil servants whose average length of employment exceeds 15 years (U.S. Office of Personnel Management 2013). Since 2001, the FCC has been consistently staffed with approximately twice as many lawyers and other legal professionals as engineers and about ten times as many legal professionals as economists (Marcus and Schneir 2010; U.S. Office of Personnel Management 2013).

As with a number of other agencies created during the New Deal era, the FCC was established as an independent expert agency – "independent in the sense that, while it is subject to laws passed by Congress and court decisions, most of its actions cannot be directly overruled by the President through the administrative process" (Taylor 2006, 264). Despite this designation, for decades the agency has been broadly understood as a highly politicized, contentious organization (Napoli 1998). Whether observers attribute this politicization to the influence of Congress, the executive branch, industry sectors, or all of the above, all agree that a substantial portion of the Commission's decisions involve what Weiser calls "political deal-making and rewarding those with influence" (2009, 6) (see also Cowhey, Aronson, and Richards (2009), Geller (1974), Horwitz (1989), and Napoli (1998)). Among the major federal independent regulatory commissions, the FCC has the highest rate of commissioner dissents (Ho 2007).

Ofcom is different in many ways. It is far younger, having been created in 2003 as a converged regulator to replace five existing sector-specific regulators with oversight over broadcasting, television, radio, and telecommunications. It was deliberately organized as a commercial entity would be, with a chief executive and executive team who run the organization day-to-day and a board that oversees the executives. Board members, including the chairman, are selected by the Secretary of State to serve for individually specified terms that may be renewed. Although the majority of Ofcom's budget is allocated by the government, it operates much like a private sector entity, with a significant portion of both executives and staff drawn from industry and salaries offered on scales competitive with the private sector (House of Commons 2007; House of Commons 2009; House of Commons 2011b; House of Commons 2011a; Ofcom 2004a; Prosser 2010). Unlike the FCC, Ofcom employs more professional economists than lawyers (Marcus and Schneir 2010).

The Communications Act 2003 tasked Ofcom with two principal duties: "(a) to further the interests of citizens in relation to communications matters; and (b) to further the interests of consumers in relevant markets, where appropriate by promoting competition" (c. 21, s. 3(1)). The citizen-oriented duty was a very contentious addition to Ofcom's enacting legislation, however, and debate has continued since Ofcom's creation about whether the agency has taken too much of an economic regulator's approach while neglecting its mandate to protect citizens' interests (Gibbons 2005; Livingstone 2008; Lunt and Livingstone 2012; Prosser 2010; P. Smith 2006).

While Ofcom has been responsive to government, particularly as it relates to managing the agency's budget, its individual regulatory decisions related to telecommunications tend to be less politically charged than those of the FCC. The Ofcom board operates by consensus unless a member insists on a vote, in which case only the vote count but not individual members' votes are made public (Ofcom 2013a). Ed Richards, a former Tony Blair staffer, was named Ofcom CEO under Tony Blair's Labour government but has continued in his role for years after the Conservatives came to power. On matters of telecommunications policy in the EU, Ofcom and the UK government often support distinct (although not always contradictory) positions to each other.

As discussed in Chapter 3, some scholars have emphasized the linkages between the sorts of structural characteristics described here and regulatory outcomes. While these characteristics formed the foundations upon which the two telecommunications regulators constructed their varying approaches to net neutrality and traffic management, they do not provide complete explanations for the many puzzles of regulatory behavior presented by the FCC and Ofcom. For example, despite being organized as an evidence-based, politically insulated agency, Ofcom did not seek to intervene with traffic management regulation when presented with evidence of discrimination. The FCC, despite having a much less technocratic orientation than Ofcom, reached far deeper into the intricacies of traffic management. To understand these kinds of nuances, a better grasp of the informal aspects of the respective regulatory environments is necessary. The next section begins that analysis by focusing on reputation.

8.3 The Role of Reputation

Chapter 3 highlighted a body of scholarly work that has considered reputation as central to understanding regulatory outcomes. In particular, Daniel Carpenter conceptualized regulation as the product of "the imperfect human official motivated neither by neutral competence nor by monetary enrichment nor by raw empowerment, but by status, esteem, legitimacy, and reputation" (2010, 43). He and others sought to tease out how agencies could develop their own power and autonomy by cultivating organizational reputations with key audiences and how the influence of individual reputations within an agency could reinforce or detract from the agency's image (Carpenter 2001; Carpenter 2010; Russell and Shelton 1974; Wilson 1989).

Reputation provides a way for the differences in approach to traffic management between the US and the UK to be understood. Neither of these agencies are immune to political pressures, but the fact that those pressures are brought to bear within differing reputational frameworks resulted in substantially divergent responses, as explained in this section.

8.3.1 FCC

Despite its collegial structure, the FCC is an agency that is centered around its chairman. Chairmen have tremendous influence over the Commission's agenda, the regulatory tools used to address specific policy issues, and the speed with which proceedings get resolved. They cannot fully control which matters will come before them as a result of stakeholder petitions, technological advancements, or congressional directives, but they can choose how swiftly or how forcefully to react; as one interviewee put it, "the chair has the power of deciding what won't get done." From 1976 to 2003, chairmen voted to affirm FCC decisions 97% of the time, demonstrating the extent of their control over which work the Commission completes (Candeub and Brown 2008).

As a result of this institutional design that vests the chair with such significant control, the primary reputational concern that affects the work of the FCC is that of individual reputation, not of organizational reputation. Chairmen are judged on their individual merits by the Washington political class because their actions are viewed as being highly determinative of what the FCC accomplishes. The public's perception of the FCC and its long-term health are not main concerns; rather, "the FCC is composed of serial mini-administrations, each of which leave their stamp on media and communications policy" (Taylor 2006, 275). Although they lack the same means of controlling the agenda as the chair, individual commissioners also develop specific reputations in the course of their work at the agency, at times using not just dissenting votes but also concurring statements to speak to particular constituencies or provide distinctive views (Candeub and Brown 2008).

For chairmen, the drive to leave their stamp on the agency's work is reinforced by their short tenures. Unlike a corporation with a CEO who is invested in developing a successful organization in the long term, the FCC attracts a series of leaders who arrive knowing that they have a few years at most to accomplish their individual policy goals. As an organization, the FCC certainly has a reputation – mostly for incompetence, delay, short-sightedness, internal strife, and a raft of other deficiencies (Geller 1974; Hundt and Rosston 2006; Napoli 1998; Weiser 2009) – but improving how the agency is regarded as an organization is not generally high on any chairman's list of priorities. That individual reputation trumps FCC reputation is not meant to imply that chairmen pursue their work narcissistically, but that they focus their short tenures on the substantive policy outcomes that they most desire to be associated with, rather than on how those outcomes fit into longer range objectives for the FCC as a regulatory body.

This is not to say that the chair can ignore the opinions of the other commissioners, but the need for the chair to make significant, substantive compromises in order to achieve the broad outlines of his desired outcome in any given FCC proceeding is most relevant when fewer than five commissioners are voting (whether because one or more commissioner seats are vacant or as a result of a recusal). Most orders are adopted unanimously (Candeub and Brown 2008; Ho 2007), but on divisive issues, the chairman's key task is to find support among at least two of his fellow commissioners. With a full commission voting he can often find that support either from the other two commissioners of his party or by appealing to the personal interests and idiosyncrasies of the opposing-party commissioners (Candeub and Brown 2008). More significant compromise occurs in cases where fewer than five commissioners are voting, because any two commissioners who disagree with the chairman have enough leverage to spur a negotiation. His political skill matters.

The conception of the FCC as described here – as an agency largely controlled by its chairman, affected by how the chairman wishes to be perceived, but susceptible to policy compromises resulting from the vagaries of commissioner appointment schedules and recusals – explains much about the agency's approach to traffic management. The following sections illustrate this by examining the actions of the two chairmen who presided over the bulk of the FCC's involvement in traffic management regulation up to 2011: Kevin Martin and Julius Genachowski. Other FCC officials certainly played critical roles in shaping the traffic management landscape in the US, particularly Martin's predecessor Michael Powell and former Commissioner Michael Copps. The examinations of Martin and Genachowski presented here can be considered as case studies that illustrate the overall reputational dynamic that characterizes the FCC.

Kevin Martin

By the time Kevin Martin was named FCC chairman in March 2005, his credentials as a conservative political insider were well established. He had worked for George W. Bush's presidential campaign, including helping with the Florida vote recount after the 2000 election; for FCC commissioner Harold Furchtgott-Roth, by some estimates the most conservative FCC commissioner ever to serve (Ho 2007); and for Kenneth Starr, the independent counsel who led the investigation into the Monica Lewinsky affair. But rather than being marked by obvious partisanship or a characteristic conservative taste for deregulation, Chairman Martin's tenure at the FCC provides an extraordinary example of the extent to which an individual chairman's personal pursuit of a particular reputation as a regulator can leave an imprint on FCC policy, modulo the political exigencies of the time. By reinterpreting one broadly popular narrative about Chairman Martin – that he disliked the cable industry – this section demonstrates how both the power and the limitations of an agency driven by individual reputational concerns affected traffic management policy.

The disputes between Chairman Martin and the cable industry were many and varied during his time at the Commission. He denied cable companies' petitions for waivers to use low-end set-top boxes while granting similar waivers to telephone companies entering the TV business (Baumgartner 2008); he tightened cable's media ownership rules while relaxing them for broadcasters and newspapers (Davidson 2008; Labaton 2007); he repeatedly pressured the cable industry to offer "à la carte" or per-channel pricing (Davidson 2008), at times using questionable procedural maneuvers that sparked outrage in Congress (House Energy and Commerce Committee 2008; Weiser 2009); and, most relevantly, he authored the order that required Comcast to cease its discriminatory traffic management (FCC 2008). The assessments of Martin's tenure offered by interviewees from both government and industry echoed those that often appeared in the press during his time as chairman (Eggerton 2009; Lasar 2008; Ulaby 2008): that "from the beginning, [he] didn't like cable," that he had an "anti-cable approach," and that "cable companies were guilty until proven innocent" during his tenure. Speculation in the industry and in the press was that his apparent bias resulted from his personal objections as a conservative to what he viewed as overly indecent content on cable television, which was subject to far more lenient content regulation than broadcast (Davidson 2008; Lasar 2008).

Many interviewees offered his animosity towards cable as the key explanation for the *Comcast Order*. One cable policy executive explained that the relationship between Martin and the cable industry had "soured" and that "the Comcast situation was in large part a reflection on that relationship as opposed to the substance of the net neutrality debate." A former FCC staffer was more blunt: "he had a complete vendetta against the cable industry, he caught them with their pants down, and he wanted any reason to nail them."

At the time, the "attack-on-cable" explanation certainly seemed more plausible than any policy-based reasoning for what appeared to be a dramatic reversal of course and a significant regulatory intervention. When the Commission adopted the *Internet Policy Statement* in 2005, Martin explained that "policy statements do not establish rules nor are they enforceable documents" (Martin 2005). When compliance with the *Policy Statement* principles was included as a merger condition in the SBC/AT&T and Verizon/MCI mergers several months later, Martin stated: "I do not believe that all of the conditions imposed today are necessary" (Martin 2005c, 2). And when the AT&T/BellSouth merger was approved the following year with even stronger nondiscrimination requirements, Martin's distaste for regulatory intervention in this area was palpable:

> Some of the conditions will certainly provide additional consumer benefits. . . . Other conditions, however, are unnecessary and may actually deter broadband infrastructure investment. The conditions regarding net-neutrality have very little to do with the merger at hand and very well may cause greater problems than the speculative problems they seek to address. These conditions are simply not warranted by current market conditions and may deter facilities investment. Accordingly, it gives us pause to approve last-minute remedies to address the ill-defined problem net neutrality proponents seek to resolve. . . . Importantly, however, while the Democrat Commissioners may have extracted concessions from AT&T, they in no way bind future Commission action. . . . For example, today's order does not mean that the Commission has adopted an additional net neutrality principle . . . although AT&T may make a voluntary business decision, it cannot dictate or bind government policy. Nor does this order. (Martin and Taylor Tate 2006, 2)

Despite what seemed to be his clear opposition to intervention and his contention that the *Policy Statement* was unenforceable, when the allegations about Comcast's handling of peer-to-peer traffic surfaced in 2007 and petitions for action were filed with the FCC, Martin responded aggressively. He sought comment from the public, organized a series of hearings that elicited testimony from neutrality advocates and embarrassing revelations from Comcast, and made provocative public statements about Comcast's behavior (Martin 2008b). All of this activity culminated in the Commission's adoption of the *Comcast Order*, which bound Comcast to a significantly higher nondiscrimination standard than that in the *Policy Statement*. The *Order* required the company to transition to nondiscriminatory network management practices that "further a critically important interest" and are "narrowly or carefully tailored to serve that interest" (FCC 2008, 28) – a standard that the chairman openly admitted to have borrowed from the extremely stringent legal test known as "strict scrutiny" (Martin 2008b), which is applied when the government seeks to limit what would otherwise be protected individual speech under the First Amendment. Other than by fining Comcast, which the Commission declined to do, it is difficult to conceive how the result of the *Comcast* proceeding could have been more interventionist within the confines of an enforcement action. The former FCC staffer explained his perplexity:

> I've never seen him angrier than . . . AT&T/BellSouth and he wrote that blistering statement about how this is the most appalling thing and this net neutrality concept will ruin earth as we know it, and then eight months later or whatever it was, he had totally flip-flopped. . . . It was a measure of how eager he was to stick it to the cable industry that he would use this weapon, that you'd think would be the last one he'd want to use.

But Martin repeatedly denied that he was out to disadvantage cable or favor its competitors, arguing instead that his actions were motivated by a desire to deliver lower prices to consumers and to create regulatory parity between the cable and telephone industries so as to spur competition (Davidson 2008; Eggerton 2009; Lasar 2008; Martin 2007a). The latter goal was of particular relevance to net neutrality policy, as it was bound together with changes to the regulatory classifications of DSL and cable broadband. With the publication of the *Wireline Broadband Order*, which removed the remaining common carrier obligations on DSL broadband providers and thereby created regulatory parity with cable broadband, Chairman Martin celebrated the achievement of his long-standing goal of ending the "regulatory inequities" between cable and DSL: "As I have said on numerous occasions, leveling the playing field between these providers has been one of my highest priorities" (Martin 2005b, 1). In the press and before Congress he often reiterated his desire for both industries to be treated equally, including by enforcing broadband nondiscrimination in the Comcast case as had been done in the case of Madison River (Lasar 2008; Martin 2007a; Martin 2008c). This makes sense in light of how different the regulatory histories of the two industries had been, as explained in Chapter 5.

A deeper analysis of Chairman Martin's actions and statements combined with the perspectives of the regulatory and political actors who surrounded him reveals a more coherent explanation for his actions than those of a cable adversary or relentless equalizer: Chairman Martin wanted to do as little as possible to constrain ISPs' behavior ex ante, but wanted to retain the right to discipline them ex post if they contravened his vision of how Internet service should be offered. This approach combined the conservative ideals of regulatory restraint and strong enforcement. His pursuit of this approach was, by necessity, shaped by the political exigencies of the time, but at each step in the development of net neutrality policy in which he was involved, he sought as best he could to make the industry aware of the potential for FCC intervention – but not to intervene ex ante. A senior FCC official explained the message that Martin sought to send to the ISPs: "Hey, we are going to watch this. These are certain principles we're going to follow. And you've got to follow these guidelines, and if you get out of line, then people can complain, and we can take action against you." An analysis of each step in the process demonstrates both Martin's commitment to this approach and the effects of political pressure from other commissioners and Congress on the outcome.

At the time of adoption of the *Wireline Broadband Order*, only four commissioners were seated, two Republicans (Martin and Kathleen Abernathy) and two Democrats (Michael Copps and Jonathan Adelstein). To accomplish his objective of leveling the playing field between DSL and cable, the chairman needed the support of at least one of the Democratic commissioners. A former Republican FCC official familiar with the negotiation explained it as follows:

> It was 2-2, so we had more limited ability to adopt something, we really needed a third vote. . . . From the Republicans' perspective, it was we don't want net neutrality legislation and we don't want net neutrality rules, so if we could buy off the Democrats with a Policy Statement that's not binding, that's a win. We did intend for it to be a Policy Statement that people actually abided by, which is why you needed the footnote on reasonable network management. But it was mostly – they were statements that Kevin strongly believed should govern the Internet, he just didn't want there to be rules or legislation. So the hope was that this would stave off the need for either of those.

Thus the *Policy Statement*'s mere existence was a result of political compromise, which explains why a chairman generally inclined towards deregulation would support its adoption at all. But it was also a statement of policy that the chairman generally believed in – he simply did not want that policy to unduly constrain operators by obtaining the status of a formal rule. Thus, while Chairman Martin claimed that the *Statement* was not enforceable and that "regulation is not, nor will be, required," he also admitted that the *Statement* "does reflect the core beliefs that each member of this Commission holds" (Martin 2005a). It expressed how he thought broadband should be offered, even if he preferred that the Commission not intervene to make it so.

At both times when the telcos sought approval for their mergers in November 2005 and December 2006, there were again only four commissioners voting. Chairman Martin clearly disagreed with the idea of including nondiscrimination conditions, as that would have placed ex ante constraints on the largest DSL providers that did not exist for the cable companies. The lightest touch regulatory approach possible – of describing principles to guide how the industry should behave and hoping that they comply – only remains light of touch until those principles become enshrined in binding orders. Commissioner Copps, joined by Commissioner Adelstein, used the leverage of the 2-2 commission to enshrine them despite Martin's objections. It was clear that Martin's position had more to do with regulatory process than with policy: just months before he wrote his dissenting statement about the AT&T/BellSouth merger conditions, he stated in testimony to Congress his belief that the FCC had authority under Title I of the Communications Act to take action if violations of the *Policy Statement* were to occur (Martin 2006).

By the time the allegations about Comcast's peer-to-peer management surfaced in 2007, there had been several heated net neutrality debates in Congress, with a number of bills introduced that would have codified far more detailed nondiscrimination rules than those that existed in the *Policy Statement*. Defending against what he viewed as unnecessary Congressional intervention, Chairman Martin repeatedly emphasized to Congress and the public that the "Commission remains vigilant and stands ready to step in to protect consumers' access to content on the Internet" (Martin 2007, 7) (see also Martin (2007b; 2008a; 2008c)).

The fact that the *Policy Statement* provided no guidance as to what kinds of network management practices were to be considered "reasonable" created space for debate about whether Comcast's behavior violated the principles in the statement. Comcast's lack of transparency and shifting public statements about its practices (discussed further in Chapter 9) inflamed the public discussion. With Comcast potentially contravening in so public a manner the chairman's vision for how broadband should be offered and Congress threatening intervention that he did not want, Chairman Martin used an enforcement action (over the objections of the other two Republican commissioners) as a show of strength. A former FCC staffer explained the chairman's predicament:

> The FCC had principles in place and if Comcast could do what it did and the FCC, which nominally had these policies in place – whether they were rules or whatever, that's a legal game – if the FCC had these policies in place and yet was unable to do anything about it, that would have just shown that the emperor had no clothes, almost, right? So the FCC either – it sort of had a choice, it could either act like it had some power, or it could just reveal to the world that it had no power. And then what would have happened if it revealed that it had no power? [Congressman] Ed Markey would have pushed a bill through the House and [Senator] John Kerry and others would have picked it up in the Senate . . . .

While he may have declared the principles unenforceable and opposed the idea of making them binding in the absence of proof of specific infractions, issuing the order against Comcast provided an important deterrent to Congressional action while also proving that the regulatory threat embodied by the adoption of the *Policy Statement* had teeth. A former senior FCC official explained this preference for ex post enforcement over ex ante restrictions by analogy:

> There used to be highways – and I think there still are, out west – where there is no speed limit, but you can still be stopped for speeding. The speed limit is: you have to drive safely. And it's the difference between having a speed limit and having a principle of "you have to drive safely." And if you have a speed limit, you're on notice that you can't go faster than 55. If you have a principle that says you have to drive safely, you can get stopped and then you can be told, "what you're doing is not safe and if you continue to do it, then we're going to fine you." . . . I think to understand the commission's precedent, that's exactly what we did in *Carterfone* . . . we ultimately adopted rules, but the rules codified the principles that we had begun enforcing in the *Carterfone* decision. And that's exactly the same approach.

This was the embodiment of regulatory threat: putting the industry on notice, without too many specifics, and developing a record of adjudications over time so as to further refine what Commission policy should be. In a situation where, as the senior official explained, the FCC had not "fully vetted" what it meant for traffic management to be "reasonable," creating ex ante restrictions was seen as inappropriate: "it's much more difficult to write rules for the hard cases when you haven't had enough experience trying to figure it out." That there was highly relevant Commission precedent for principles-based adjudication was also important to the chairman – the *Comcast Order* refers to "our venerable *Carterfone* principles" (FCC 2008, 23) and uses *Carterfone* as an example in support of the claim that "the Commission has often relied on adjudications rather than rulemakings to enunciate and enforce new federal policy" (FCC 2008, 18).

Viewed through the lens of individual reputation and shaped by political pressure from his fellow commissioners and Congress, Martin's actions were fully consistent across time. They reflected his political skill and his personal desires to both refrain from ex ante intervention and bring the weight of the agency to bear ex post. They also had a marked effect on how traffic management in the US would evolve, as they set up a legal battle over the FCC's authority that Martin's successor, Julius Genachowski, had to navigate.

Julius Genachowski

The link between individual reputation and net neutrality policy outcomes is far more straightforward in the case of Julius Genachowski. Prior to being named as FCC chairman in 2009, Genachowski spent more than a decade as an executive and investor in the technology industry. This included a long stint at IAC/InterActiveCorp, whose web properties included Expedia, Match.com, Evite, and many more of the web's most popular sites, and advisory roles for Bandwidth.com, Truveo, and other start-ups engaged with VoIP and video (Carney 2006; Daver 2005). He was involved in the founding of multiple Internet-focused venture firms. In short, his professional experience was oriented towards businesses that harnessed the Internet and the web as platforms for novel services and applications.

He brought that orientation and his knowledge of the Internet industry with him when he joined Barack Obama's campaign team as a chief technology advisor ("Julius Genachowski" 2012). He urged the campaign to make aggressive use of social media and capitalized on his network of tech and venture contacts to raise hundreds of thousands of dollars for the campaign (Heilemann 2007). As he and his team developed Obama's tech policy platform, he made support for innovation a focal point. The platform was announced in late 2007, listing support for "the principle of network neutrality to preserve the benefits of open competition on the Internet" (Obama for America 2007, 2) as the first item on the future President's technology agenda. Long before he became FCC chairman, his desire for a policy of nondiscrimination on the Internet was clear. Once he became chairman in 2009, his commitment to reaching this goal took on significant reputational value; as one industry interviewee noted, "I think Genachowski wants to be viewed as more tech-friendly, so he's interested in things like . . . net neutrality."

It was no surprise, then, that the FCC issued a notice of proposed rulemaking (NPRM) concerning Internet openness shortly after Genachowski was appointed as chairman (FCC 2009). The proposed rules would have codified the four principles in the *Policy Statement*, added rules concerning transparency and nondiscrimination, and made each of the rules subject to "reasonable network management." In its treatment of a number of thorny policy issues, including those related to traffic management, the NPRM provided an early example of another characteristic that would come to define Genachowski's reputation: his centrist approach to policy (Crawford 2013; Gustin 2013; McDowell 2013). Of the "narrowly and carefully tailored" standard for traffic management that had been established in the *Comcast Order*, the FCC wrote: "We believe that this standard is unnecessarily restrictive in the context of a rule that generally prohibits discrimination . . . . We seek comment on our proposal not to adopt the standard articulated in the Comcast Network Management Practices Order in this rulemaking" (FCC 2009, 51). Genachowski was showing that he was taking a more flexible substantive approach to net neutrality policy than his predecessor, even as the Commission sought to formally enshrine in rules what had previously existed only in the *Policy Statement*.

The complicated path that led from the NPRM to the adoption of the *Open Internet* rules continued to bolster the chairman's reputation as a tech-friendly centrist and revealed his skills of political negotiation. In April 2010, Comcast's legal challenge to the *Comcast Order* was upheld in court, calling into question the FCC's authority under Title I to issue broadband nondiscrimination rules (Comcast Corp. v. FCC 600 F.3d 642 (2010)). The Chairman and his staff responded by proposing a compromise "third way" legal framework to resolve the authority question (Genachowski 2010a; Schlick 2010), conducting a public inquiry into the viability of reclassifying broadband service under Title II (FCC 2010a), and engaging with members of Congress who mounted an effort to craft legislation to resolve the issue (Genachowski 2010b).

Due to political opposition, none of these efforts ultimately resulted in changes to the regulatory classification of broadband Internet service, but the chairman pushed forward with rules nonetheless. One industry interviewee explained that the seeds of this pursuit had been planted years prior: "once [Internet openness] became a platform in the President's campaign and he won the election, now there was this political driver to do something on net neutrality

. . . obviously doing something on net neutrality was a mandate." Another made a direct connection between the chairman's reputation and his resolve: "for Genachowski, with the way FCC chairmen are evaluated, if he doesn't do a rule he's a failure." He had committed himself to a policy that he felt would protect "Internet pioneers with little more than a good idea and a no-frills Internet connection" (Genachowski 2010c, 1); the emergence of legal hurdles did not deter him. A former FCC staffer explained this inevitability in the political context of the Commission, where, crucially, all five commissioners were seated and eligible to vote on the *Open Internet* rules:

> I think the chairman believed when he started and believed when he finished that having high-level rules was the right thing to do and he was going to bring it to order. I guess the only way it would not have happened is if there weren't three votes for any particular conception of it. . . . But I don't think – there was never any wavering . . . as to, should we have high-level rules of the road. He always thought we should and he was going to drive it through if he possibly could.

The rules that were adopted were perhaps the strongest indication of the chairman's tech-friendly yet centrist approach. They bolstered protections for innovators by creating enforceable rules and adding nondiscrimination and transparency provisions that the *Policy Statement* had lacked. As a result of tireless negotiations, their adoption was ultimately greeted without objection from nearly every major industry stakeholder in the broadband industry (save Verizon) and the tech sector (Albanesius 2010; Cicconi 2010; National Cable and Telecommunications Association 2010). The rules created a presumption of reasonableness for application-agnostic traffic management but refrained from adopting the "unnecessarily restrictive" (FCC 2010b, 49) standard from the *Comcast Order* and committed to further elaborating the reasonableness standard on a case-by-case basis. In short, the *Open Internet Order* arguably delivered on the Obama campaign promise of protecting "the Internet's traditional openness to innovation" (Obama for America 2007, 2) by carving out a middle ground approach with enough flexibility to accommodate most ISPs and enough teeth to assuage application providers' concerns. The rules precisely reflected the reputation of the chairman.

Although Genachowski charted a more regulatory course than his predecessor, both chairmen took full advantage of the flexibility that the FCC afforded them to act in the presence of uncertainty. Martin preferred to establish policy through principles-based adjudication because "[t]o establish principles, you're able to say, 'this is okay' and 'this is bad,' but 'we're not sure about the middle and we'll be able to continue to judge it as we get more and more facts,'" as the former senior official put it. Genachowski's approach likewise recognized that specifying very precisely which behaviors should be permitted and prohibited was not likely feasible, and that recognition together with his centrist leanings led to the creation of "a set of high-level rules of the road" (Genachowski 2010c, 136) that could be further refined through case-by-case adjudication. To pursue the policy outcomes that would define their legacies, the chairmen had to act without complete certainty about the effects of discrimination or the prohibition of its various forms. A former FCC staffer summarized this approach:

> [Y]ou're never going to resolve all ambiguities in a regulatory proceeding. . . . If we could come up with a perfect rule and that institution, the FCC, operated perfectly, that would definitely be the better place to be because everyone has certainty about what is and is not permitted. But that's a fool's errand. So, let's state some high-level principles and let's adjudicate cases on the facts in a case-by-case way.

The drive to take bold action despite uncertainty in the marketplace was notably absent from Ofcom's treatment of the net neutrality issue, as was the dominant effect of individual reputation. Those contrasts are the subject of the next section.

8.3.2 Ofcom

The comparison between the FCC and Ofcom is particularly salient because the role of reputation was quite central for both agencies' traffic management decisions; the difference with Ofcom was that its strong organizational reputation dominated its decision-making. As explained in this section, Ofcom's early work was so highly regarded by private and public sector stakeholders that it took on significant policymaking influence, including in the context of the EU framework review. As political pressure compelled the agency to become more conservative, its roots as a competition regulator colored its approach to net neutrality, yielding a cautious, non-interventionist policy and focus on transparency that could reinforce the regulator's desire to be viewed as protective of consumers.

This section's analysis of Ofcom as an organization is not meant to suggest that the individual opinions and approaches of Ofcom officials were unimportant in the UK's net neutrality debate (in fact, the influence of key individuals is highlighted below). Rather, the claim here is that the organization's culture and image were highly determinative of the frame within which individual-level debates took place. They served to control the overall outcome, if not every nuance of Ofcom's traffic management decision-making.

Ofcom's Strong Organizational Reputation

During its formative years, Ofcom developed a strong organizational reputation and devoted significant attention to maintaining and promoting that reputation within the UK and elsewhere. It quickly came to be known for its effectiveness as a regulator characterized by technical expertise, rigorous analysis, evidence-based decision-making, and inside knowledge and understanding of the private sector (Lunt and Livingstone 2012). It would be difficult to find many stakeholders who would describe the FCC with the kinds of superlatives that interviewees – including representatives from regulated companies – used to describe Ofcom: "very rigorous in evidence-based policymaking," "a leader," "a heavyweight," "a very, very effective regulator," "a very successful institution," "a shining example," "the most competent authority," and "a leading light regulator." MP Peter Luff once remarked that Ofcom was "an organisation that has generally been held in very high regard by the industry concerned" (House of Commons 2007, Q57). Compared to regulators in other sectors and other countries, interviewees described Ofcom as "one of the leading regulators in the EU," "by far the best regulator, not just in telecoms, but in Europe," and "the best regulator in the world."

Several interviewees pointed to Ofcom's repeated top-ranking in the annual ECTA Regulatory Scorecard, which benchmarks the national regulators in Europe according to their effectiveness in spurring competition in telecommunications, as proof of Ofcom's effectiveness (European Competitive Telecommunications Association 2009). Ofcom's reputation clearly spanned international borders. One former Ofcom official recalled "flying every week sometimes to Brussels to speak at external conferences," while a former government official noted that "we had calls pretty much every week from regulators all over the world saying, 'We want to come to London, we want to hear about Ofcom,' or, 'Can you please send someone here to explain how it works . . . ?'" The agency commanded intense respect from its peers in other countries.

With this lofty status came a desire to continuously curate and preserve the organization's reputation. A phrase that interviewees used time and again when discussing Ofcom was "be seen" – whether Ofcom was "cautious about not being seen" to overstep its boundaries, that Ofcom "doesn't like to see itself" as interventionist, or that it needed to "be seen to be doing something" on consumer issues, the agency clearly had great concern not only for its work and the quality thereof, but for how it was being perceived by others. Indeed, one Ofcom official explained that a key function of Ofcom's international division was "protecting Ofcom's reputation overall and making sure that what Ofcom does and says is rightly understood and interpreted." In this way its corporate orientation shone through. Just as companies whose survival is ultimately in the hands of their customers invest significantly in their image and branding, so did Ofcom.

Ofcom built its reputation by being more than a technocratic implementer of powers delegated to it by the UK government and EU legislation. In its early years, it took bold, innovative steps, not the least of which was the Telecommunications Strategic Review that culminated in BT's Undertakings. The TSR was the first major overhaul of telecommunications regulation in decades – in Ofcom's own words, an endeavor to replace the existing "unsustainable" (Ofcom 2004e, 5) framework with a "new regulatory contract"

(Ofcom 2004e, 17) – and its relative success was the subject of envy within regulatory circles (Analysys Mason 2010). The TSR was just one reflection of Ofcom's commitment to holistic, strategic thinking. Because it was not bound by the need "to take each decision on its merits," as Ofcom's first head of strategy Kip Meek explained, the agency could take riskier, more decisive action than many other regulators (Crabtree 2004).

Impact of Reputation on Policy-Setting

Over time, Ofcom's depth of expertise, strategic focus, and outsized reputation created an unusual dynamic between the regulator and the government. Although Ofcom was statutorily designed to be the implementer of regulations to fulfill the government's public policy objectives, Ofcom's strategic and policy work reached much farther in reality. One former official who joined Ofcom upon its creation explained how this was implicitly blessed by government ministers:

> [In the] early days, you very much got the sense that ministers were, like, you know, "We created Ofcom, we're really proud of the fact that we've created Ofcom, now go away and do your stuff. Of course we will want to be kept in touch and occasionally kind of intervene just to tell you there's something that you're missing or there's a broader public policy objective at work here." But basically Ofcom was given a huge amount of latitude to make policy in a way that has never really been – that's not the way that European regulators have generally worked. It's not the way that Oftel worked. And it's not the way that the other European peer regulators in the sector worked, actually. And they always used to be amazed at the amount of stuff that we could get away with.

By the time the European telecommunications framework review began in 2007, Ofcom was taking the lead in representing the UK's policy voice during the negotiations in Brussels. One official explained that "Ofcom were in the driving seat," while a former government official from the Department of Business, Innovation & Skills (BIS) described all of the advantages that Ofcom had as a policy-setting body compared to BIS, the government department with authority over telecommunications:

> [W]e say they're the regulator and we're the policymaker, which is fair enough. But in some people's eyes that was sometimes blurred. And during the framework review, for example, Ofcom played a very prominent role. I mean, they are the regulator, they understand the issues probably better than policymakers do in various areas. They have more resources, they've always had more resources than I have. They have more money than we do, and I know that's not irrelevant. And so

> sometimes they would lobby in the [European] Parliament and they would say, "This is our view," and it would be interpreted as the UK view rather than the Ofcom view. So, you know, so there is a sort of policy dimension when your regulator, if you like, goes outside of the normal role of a regulator and sort of goes onto the public stage.

Ofcom had used its reputation and expertise to accrue policy responsibilities, which had important implications for the UK's position on net neutrality and traffic management during the framework review. As explained in Chapter 6, Ofcom's early positioning on the issue was staunchly pro-competitive and anti-interventionist. Its influence on the UK's position was palpable, as the former BIS official explained:

> [D]uring the framework review, we had a debate with Ofcom where they said that if blocking took place, that would be part of the normal competitive process. . . . We had discussions with Ofcom because it wasn't, at the time, it wasn't our position. We would have had to sort of come to an agreement on what our position [was]. They won. I mean I lost because ministers wouldn't back the position that I wanted to take as an official. Which is fine. But their argument at the time was that if blocking did occur, that was part of the normal process of delivering services. And so, you know, if one day you woke up and you found that you couldn't get Google, you would just go to another provider. So there wasn't any harm because it would be like taking another packet, going from one brand of cornflakes to another brand of cornflakes.

Because Ofcom spoke with such respect and authority, the agency's position that relying on competitive forces was the way to address issues of blocking (and, by extension, discrimination more generally) essentially became the UK government's position.

Reputational Transition

Ofcom's prominence within policy debates, both in telecommunications and in other areas, did not go unnoticed by those in industry and government, particularly in light of a long tradition of UK regulators acting with narrow remits as technocratic implementers of government policy (Ogus 1994). As one government interviewee remarked, "there's always been a bit of a question about is Ofcom sort of over-reaching." These concerns took on new significance leading up to the 2010 election. The Conservatives were solidifying their electoral prospects on a platform that included extensive cuts to government and regulatory agencies (also known as quasi-autonomous non-governmental organizations or "quangos").

In announcing his approach to quango reform, future Prime Minister David Cameron singled out Ofcom:

> OFCOM is the regulator for the communications industry, and it's clear that it has an important technical function. . . . But . . . OFCOM currently has many other responsibilities that are matters of public policy, in areas that should be part of a national debate . . . . These should not be determined by an unaccountable bureaucracy, but by minsters accountable to Parliament. So with a Conservative Government, OFCOM as we know it will cease to exist. Its remit will be restricted to its narrow technical and enforcement roles. It will no longer play a role in making policy. And the policy-making functions it has today will be transferred back fully to the Department for Culture, Media and Sport. (Cameron 2009)

With Cameron's intentions made clear, one former Ofcom official recalled a "huge change as the mood sort of darkened in Ofcom as the election approached." The Conservatives' electoral victory presented a significant challenge to Ofcom's status as an expert agency and its reputation as a strategic heavyweight. Combined with the government's austerity program, which for Ofcom meant instituting a 28% budget cut over four years beginning in 2010, the agency was faced with an urgent need to alter both its workload and its image to meet the expectations of the new government. Numerous interviewees noted how Ofcom CEO Ed Richards, being "politically astute" and "pragmatic," instituted a plan to adroitly re-shape both the agency's reputation and the allocation of its resources so as to appease the government without allowing Ofcom to be gutted entirely. Stepping back from any appearance of policy-setting was a key component of this plan, as explained by a former Ofcom official:

> [F]rom that very high, lofty position, it's a long way down. And certainly you feel threatened. . . . [I]t was definitely a shift to try to make us look more low key, more as implementers, rather than strategic. We were pretty much effectively told that the word "policy" was banned from being used in documents, okay? We won't talk about it, even if [it's] regulatory policy, because the Conservatives hated that. . . . Policy, for policymaking, they thought should reside in government, strategy should reside in government, we were there to be the technical implementers.

In some corners of Ofcom, the move to publicly rein in the agency's strategic work was welcome because it matched the staff's natural proclivities as evidence-based competition regulators. Some areas of policy that were arguably within their remit, including media plurality and the innovation and free expression issues bound up in the net neutrality debate, did not lend themselves to straightforward quantification or traditional competition analysis. One industry interviewee observed that this type of work "is not based on evidence. It's based on kind of belief, core beliefs. And . . . [Ofcom is] just sort of not constitutionally organized, I think, to do those kinds of things." A former Ofcom official described how this sentiment was manifest internally, explaining that Ofcom "wants to find numbers and these are not numbers issues. These are to do with what kind of society we want to live in. . . . The natural instinct is to say, '. . . that kind of social policy, we shouldn't really be engineering that.'" If Ofcom could be said to have an organizational "culture" or "essence" (Barnard 1938; DiIulio 1994; Halperin 1974), it naturally skewed towards quantitative, technocratic, and narrowly focused. Ofcom Board Chair Colette Bowe and other senior Ofcom officials acknowledged the agency's difficulty in approaching issues that were not amenable to concrete definition and quantitative analysis (Lunt and Livingstone 2012).

Aversion to qualitative assessment and ambiguity had already colored Ofcom's approach to net neutrality before being reinforced by the agency's departure from the policymaking limelight. Numerous former Ofcom officials explained how, in the 2005-2009 time frame, Ofcom was undecided about whether discriminatory traffic management was problematic, and as a result it took a cautious, reactionary approach, only engaging when circumstances demanded it (as in the framework review) and advocating against intervention when it did engage. One interviewee described "strong internal divisions about whether or not it was an issue," resulting in internal deliberations of this sort: "This could be kind of important. We're not entirely sure. Is it really an issue right now? We're not entirely sure. Is anyone going to be upset? We're not entirely sure. We'll come back to it in six months. And then people would tend to come back to it in 18 months." Because some aspects of net neutrality did not fit neatly within competition analysis, Ofcom delayed getting deeply involved, by which time the organization was aiming to reduce its perceived policy presence so as to assuage the government's concerns.

Impact of Reputation on Net Neutrality Policy

The inclination to stay out of the limelight had important implications for Ofcom's approach when it did decide to engage publicly on net neutrality ("net neutrality, again, is just a symptom, it's a symptom of a shift within Ofcom" was how one former Ofcom official put it). Some stakeholders perceived that "Ofcom is now much more narrowly focused; that it has become risk averse; that it is not contributing to some of the big debates, which previously it would have played a major part in, for instance, things like net neutrality," as MP John Whittingdale described it (House of Commons 2011b, Q2). In 2010, the agency issued a discussion document about net neutrality and traffic management, outlining its initial views on the subject and seeking public comment. On the document's very first page, Ofcom made clear that it intended to approach the issue narrowly: "The debate ranges widely including questions such as whether citizens have a 'fundamental right' to a neutral internet, or whether 'net neutrality' promotes economic competitiveness and growth. These are important questions, but also ones primarily for governments and legislators." (Ofcom 2010c, 1)

The most fundamental questions about discriminatory traffic management squarely concern economic competitiveness and growth (and "industrial" policy, also mentioned in the document as a matter for government (Ofcom 2010c, 8)), since discrimination potentially affects the competitiveness and growth of the Internet applications sector, but Ofcom declared such topics outside of its scope. As one official involved in drafting the document explained, "different people could come to a different view of the weighting that they want to attribute to the innovation aspect . . . that's beyond our pay grade." A government interviewee concurred: "it's not for the regulator to make that judgment about what is good for the overall economy or the good of consumers. They can make judgment as to whether something is to their detriment, but not what is necessarily to their good."

To the extent that this reluctance to engage on the broader innovation issue resulted from Ofcom's trepidation about the perception that it had overstepped the bounds of its authority, its approach marks a significant contrast to that of the FCC. Julius Genachowski pursued his goal of establishing net neutrality rules despite tremendous uncertainty about the FCC's authority to do so. Kevin Martin likewise pursued his own enforcement agenda without the greatest clarity about the agency's jurisdiction. At Ofcom, meanwhile, officials explained that they were "consciously very careful . . . not to tread on anybody's toes and make sure that it was very much linked to the transposition of the telecoms framework." The official involved in the drafting of the discussion document expressed that "what I thought was really important for Ofcom was not to get sidetracked into a debate where we really have no authority to just start making it up as we went along." The policy legacies that the individual FCC chairmen sought to leave behind trumped concerns about the scope of the regulator's remit and spurred them to take bold action to deter discriminatory conduct, whereas Ofcom's desire to re-brand itself as a technocratic implementer of government directives caused it to err on the side of caution, leaving the toughest questions about discriminatory traffic management unexplored.

Over the course of a year and a half of deliberations, Ofcom came to recognize the importance of some of the broader economic and innovation issues involved in net neutrality policy, discussing them in detail in its net neutrality position paper, *Ofcom's approach to net neutrality*, published in November 2011 (Ofcom 2011b). The regulator had clearly moderated its pro-competition rhetoric, claiming that "[w]here innovation is of particular importance, as here, a departure from the standard competition-based approach may . . . be justified" (Ofcom 2011b, 21) and signaling to the industry that "[o]ur stance as a regulator is that any blocking of alternative services by providers of internet access is highly undesirable, because of the potential effect on innovation" (Ofcom 2011b, 26). Some of this shift was attributable to the work of Steve Unger, the Ofcom executive in charge of the agency's net neutrality work. One industry interviewee recalled that after engaging with Unger "it was clear that he knew there was a problem" with allowing discrimination to get out of hand. After the document was published, Unger emphasized his appreciation of the innovation aspect at a Parliamentary hearing, noting that the net neutrality debate "is not primarily around competition" and that there is “a story there around innovation and it goes beyond conventional competition laws” (House of Lords 2012, 21).

Despite the shift in rhetoric, however, Ofcom declined to use any regulatory powers, including the power to establish a minimum quality of service created by the transposition of the EU telecommunications package, to prevent or prohibit discrimination of any kind (including outright blocking). Instead, it focused on promoting transparency to ensure that consumers would receive all the information they would need to select broadband services. For a regulator with an aversion to ambiguity in an environment where it did not want to be perceived as making strategic decisions, transparency was a safe and relatively easy issue to pursue. A former official involved in the work noted the agency’s reputational motivations in describing the press release that accompanied the release of the document:

> [The press release] led with consumer information. It was three quarters about consumer information, and that was mostly what was picked up in the press. Because you can’t be wrong on consumer information. And it was a tiny bit about blocking and innovation at the bottom, which I think only The Register picked up in any real degree. . . . But I mean if you looked at everyone else that picked up the press release, they only focused on the first three quarters. Because you can’t be wrong on that.

The focus on consumer information in the net neutrality context was part of a larger shift that Ofcom was making towards increasing its activity and presence in the area of consumer protection. The 2009 Ofcom Board Chair appointment of Colette Bowe, who had chaired Ofcom’s consumer panel since 2003, combined with Ed Richards’ desire to assuage the new government concerns about Ofcom’s over-reach led to significant changes in the organization’s structure and work. A new division was set up to focus specifically on consumer policy, Ofcom’s annual plans devoted increasing attention to consumer protection (compare Ofcom (2008b) with (2010a)), and the agency increased its visibility on a wide variety of consumer protection issues, including the accuracy of advertised broadband speeds, bill shock, and many others. As one Ofcom official explained, “[w]e understand how markets work and therefore how industry works, but the most important thing that we need to understand and that we struggle much more to understand is what do consumers think and how do consumers feel and what are their expectations and have they been harmed and can they switch . . . ? And that's where a lot of the energy . . . has gone." This was in some ways the latest evolution in the agency's quest to reconcile its statutory duties to further the interests of both consumers and citizens, its core competency as an economic regulator, and the political pressures it was facing (Lunt and Livingstone 2012).

This change was dramatic, as was its relationship to the agency's public image, as another formal Ofcom official explained:

> There's a very, very strong desire on the part of senior people at Ofcom now to be taking every opportunity to assert their pro-consumer credentials. And for people in the industry who have always seen the role of an economic regulator as being more of a, you know, kind of technocratic function, in many cases supporting the industry rather than, you know, rather than being focused purely on the consumer interests, it's been a bit of an unpleasant wake-up call to find that, you know, now, Ofcom is routinely sort of, you know, taking the opportunity to criticize individual firms for different courses of conduct and all this sort of thing. A lot of stuff that previously used to be done through more informal contact and in a more kind of light-touch way is now being done very publicly . . . and the obvious conclusion that people are drawing is that Ofcom has this, you know, wants to show that it's being tough and effective on behalf of the consumer.

Numerous interviewees stressed these reputational and publicity aspects – and the benefits to Ofcom – of the agency's shift towards consumer protection. One industry executive explained that Ed Richards was "more interested in doing things that will play well from a media point of view, which tends to be some of the consumer protection stuff," as opposed to technocratic analyses of telecommunications markets. Focusing on consumer protection allowed the agency to justify itself in public as a consumer champion and not be relegated to a pure technocratic function. As noted above, this dynamic was evident in Ofcom's work on net neutrality, where the tricky regulatory issues were discussed but not acted on, while the promotion of consumer interests received all the emphasis.

8.4 Conclusion

Reputational concerns clearly had profound effects on the decisions of both the FCC and Ofcom related to traffic management. In the case of the FCC, the agency's institutional structure endowed the chairmen with significant control over the agency's output, providing support for the theory that an agency with a commission structure endows those who lead it with the ability to act independently (Horn 1995; Majone 1994). This in turn put high reputational stakes on chairmen's decisions. Kevin Martin's endeavor to embody the combined conservative ideals of regulatory restraint and conscientious enforcement led the agency down a path of principles-based guidance followed by an adjudication that served as protection against further regulatory incursion by Congress. Julius Genachowski approached jurisdictional uncertainty and vitriolic debates with the tech-focused, centrist approach he had developed in the private sector. In reference to the *Comcast Order*, one industry interviewee noted that "had there been a different decision-maker with the same set of facts, I think it may have come out quite differently." That assessment would be equally true of many major decisions made by FCC chairmen (modulo political circumstances and less-than-full commissions) and it is a testament to the extent to which the FCC's work is shaped by the individual reputational concerns of the chairman.

The role of reputation within Ofcom stands in sharp contrast. Its organizational reputation grew so strong – on the basis of both the quality of its work and its attention to reputation management – that it was able to impress its competition-centric view on UK broadband policy writ large. But the broader economic issues bound up with questions about discriminatory traffic management did not fit neatly into its traditional competition framework, and as government pressure caused Ofcom to evolve its brand, the agency's inherent cultural aversion to qualitative assessment and ambiguity caused it to view regulatory intervention to safeguard nondiscrimination narrowly and skeptically. The agency turned the bulk of its attention to transparency as part of a larger internal push to develop its image as a consumer champion.

As discussed above, Ofcom is widely respected among many of its constituencies for the rigor of its analysis and depth of expertise, whereas the FCC as an organization is not highly respected. However, for policy issues like net neutrality that require taking potentially controversial positions and that cannot be fully accommodated by competition analysis, Ofcom's institutional structure presents significant barriers when compared to that of the FCC. While allowing important policy decisions to be subject to the whim of the chairman (at least in part) has the disadvantage of potentially creating incoherence and arbitrariness in the long term (as theorized by Katzmann (1980)), the fact that the FCC chairmen had such broader leeway to act was a key factor in establishing a nationwide nondiscrimination policy in the US. Kevin Martin declared a policy unenforceable and then enforced it. Julius Genachowski had rules enacted with near-uniform industry support despite serious doubts about the Commission's authority to do so. Both of these steps and the many others in between created the culture of regulatory threat that deterred many large ISPs from discriminating for many years. Ofcom, presented with evidence of widespread discrimination but unable or unwilling to take a broad view of that evidence in light of its reputational need to scale back its involvement in policy matters, conversely created a regulatory culture of permissiveness.

Each regulator's particular reputational focus and interpretation of its own remit was in part a reflection of its structural foundations. The reputational pressures and perceived remits themselves, however, were the key determinants of their decision-making, notwithstanding the presence or absence of discrimination or competition in either market.

Chapter 9.
Regulatory Relationships

9.1 Introduction

Much of this thesis has been devoted to analyzing the relationships between regulators and ISPs and the influence of those relationships on traffic management outcomes. This chapter elaborates a vital but heretofore unaddressed aspect of that relationship – litigation – while putting it in context of other stakeholders' interactions with the regulatory agencies. Together with the reputational motivations discussed in the last chapter, litigiousness and broader interest group pressures combine to explain the differences in approach to traffic management between the FCC and Ofcom.

9.2 Litigiousness

As discussed in Chapter 3, scholars have suggested that increases in the depth and breadth of judicial review of regulatory agency decision-making causes agencies to become more conservative, to get bogged down creating excessively thorough rules to stand up in court, and to seek alternative ways of regulating to circumvent the judicial process. The case of traffic management regulation reveals that these effects are moderated by the extent of the tradition of litigation in a particular context. The FCC is accustomed to being sued and exists within a long-standing litigious culture, whereas for the much younger Ofcom legal challenges are a more recent and influential phenomenon. This contrast and its relationship to reputation, uncertainty, and regulatory remit in both countries are explained in this section.

9.2.1 FCC

The FCC is subject to judicial review as outlined in the Administrative Procedures Act of 1946 and associated case law developed since then, including the seminal *Chevron* and *Mead* cases discussed in Chapter 3. Over the course of decades, nearly every major order that the FCC has issued has been challenged in court. As interviewees noted, the agency gets sued "every time" it makes a big decision, "for doing nearly anything controversial" and when it acts on "anything that's actually going to matter to consumers." Companies and trade associations have challenged the agency on media ownership, data roaming, broadcast indecency, pole attachment, universal service, number portability, and many other issues. Of the major decisions that relate to net neutrality and open access, nearly all of the FCC's major orders have been subject to litigation: the *Computer Inquiries*, the agency's decisions related to unbundled network elements and line sharing for local loop unbundling, the *Cable Modem Declaratory Ruling*, the *Comcast Order*, and the *Open Internet Order*. Whenever the agency is considering a weighty or contentious issue, the likelihood of litigation is high; a former senior FCC official explained that in telecommunications, the carriers' "standard is they're always going to sue." For those who work at the Commission, litigation is an inevitability.

Because legal challenges are always expected and have been for decades, they do little to deter the Commission from regulatory intervention when that is the desire of the chairman and the majority of the commissioners. That the potential for judicial review does not inspire conservatism at the Commission is especially evident in the context of the *Open Internet Order*. As one former FCC official who was involved in drafting the order explained, "I don't think the consciousness of being sued had a tremendous effect. It was more like consciousness of the uncertainty in the law about our authority and trying to make sure we were within the bounds of that authority." The Commission drafts its orders "knowing that they may come before a skeptical court," but it does not necessarily decline to act because of fear of a reversal. Another former official explained the balance between the inevitability of litigation and the compulsion to set a particular regulatory direction:

> I think we felt that whether we did a Title II approach or a Title I approach, they both would be subject to significant legal challenge. And you could poll whatever lawyers you want, you're going to find some on one side and some on the other side as to which is more risky. But either way we thought doing the rulemaking was a good idea because you never get certainty. Every order we do is challenged.

The inevitability of litigation is intertwined with the particular pressures that FCC chairmen face as a result of the agency's structure. Chairmen have only a handful of years to accomplish their goals in office and establish their legacies as regulators. As discussed in the previous chapter, their personal reputations are bound up with the Commission's achievements during their tenures. As a result, the fact that courts may overturn their decisions (often when they are already out of office) is not necessarily of significant import to them; their reputations are established based on the policies they enact, not whether those actions withstand judicial scrutiny. The former senior official noted that "if you don't proceed because someone tells you that they think that they can stop you, then you're never going to be able to proceed on anything."

Chairmen do not put litigation risk entirely aside, but nor do they allow it to divert them from their need to establish their reputations according to their own preferences during the short time that they have to do so. One former official explained that "there's an understanding that the people at the commission are judged in the way that they feel they will be judged . . . based on the rules that they create. Whether or not they get overturned, often a year plus later, is sort of not their problem." They may prefer not to be overturned, but as one industry interviewee observed, "they're going to say what they're going to say and make their decision. If it gets thrown out in court, at least they can go back and say, 'Look, we tried. And we disagree with the court and we think we knew what was in the public interest and we thought we had the authority, but the court did what the court did.'" FCC chairmen are known for the particular regulatory decisions they make, not whether the courts later find them to have overstretched their authority. The former official's conclusion was unequivocal: "of this I am totally confident having known well several chairmen and worked with them closely: they don't give a shit about being overturned by judges. That's not the thing that's going to drive their decision-making."

This dynamic was clearly in evidence during the *Open Internet* proceeding. The chairman was committed to enacting nondiscrimination rules. Both the Title I and Title II approaches seemed tenuous. Although he and his staff dedicated many months to creating a set of rules that they hoped would go unchallenged by industry, Verizon was "very clear all the way through that they didn't think this was right and they probably weren't going to support it. And they would certainly come in and talk but the unspoken premise was that they didn't see any huge advantage to them in agreeing to support something that they clearly did not think was good policy or in their interests," as one interviewee put it. In the end, as an official involved in drafting the order explained, "you just have to go ahead and do what you think is going to be right, and I guess it is a peripheral benefit that during the litigation process the rules are very likely to be in place. . . . I think when we looked at it, it was just, these are the rules we think are smart rules to adopt and we'll employ what we think is the best legal foundation for it, recognizing that neither [the Title I nor the Title II] way is perfect." The chairman and his staff knew that they were almost sure to get sued, but it was more important to establish the rules as a matter of policy than to avoid legal risk. As Crawford (2013, 62) noted, "Verizon sued. Someone always sues."

What so clearly distinguishes the FCC from Ofcom is that FCC chairmen have a mandate to pursue this kind of policy-setting and that they are comfortable doing so in light of significant legal uncertainty. Conducting public interest analysis, allowing gray regulatory areas to be refined through case-by-case adjudication, and setting national communications policy are all accepted as being well within the realm of the Commission's discussions, even if the courts at times disagree with the agency about its statutory authority to act. Setting policy is what FCC chairmen come to the agency to do.

Kevin Martin and Julius Genachowski took different approaches to net neutrality, but they were both more willing to intervene in an uncertain legal environment than Ofcom because they both led an agency where it was expected that chairmen would take bold steps to advance their policy agendas. One industry interviewee who had worked at the Commission in the mid-2000s explained how the drive to set policy trumped concerns about legal challenges both while he was there and during the *Open Internet* proceeding:

> I think what happens is that the people that are working on this are just trying to come up with the right policy and it's almost like they just outsource the litigation risk to the lawyers and it happens afterwards. . . . So it's always in the back of the policymaker's mind, but I would say it's not all that different between the open Internet order and, say, the data roaming order . . . or the LNP [local number portability] rules, which are a good example of how we knew we were going to get sued, but it's just, it's a good answer for consumers and you do it and you let the lawyers clean up the mess, essentially. You'll take their advice to make sure that you don't overreach, so you don't create risk where there isn't some, but you're stuck with the statute at the end of the day.

The combination of the inevitability of litigation and the expectation that FCC chairmen will pursue communications policy broadly, far beyond the confines of competition analysis, gives chairmen the freedom to make regulatory decisions whose value lies more in the direction they set than in their statutory longevity. As the former staffer who had worked on the LNP rules remarked about the *Open Internet* rules, "[S]ay they get overturned in the D.C. Circuit or wherever. There is value in having the policy said, because then you've narrowed the universe of contested issues to just the jurisdictional question. You've taken a step forward on the policy, maybe you failed on the jurisdiction, but at least you've gotten some rough consensus about what the policy should be." Those involved in drafting the *Open Internet Order* similarly recognized the value of rulemaking regardless of legal challenges; they explained how "in essence the rules are just about the default that you've created" and that "contending stakeholders benefit from having someone lay down a marker." FCC chairmen act in light of legal uncertainty (or near-guaranteed legal challenges) because they realize that setting a policy direction can have a long-lasting impact regardless of what happens in court.

The interaction between litigiousness, uncertainty, and the agency's remit at the FCC was in substantial contrast to Ofcom, where policymaking was eschewed and litigation was a more recent and intense phenomenon. That is the subject of the next section.

9.2.2 Ofcom

Unlike the FCC, where challenges to regulatory decisions have been so common for so long as to have few moderating effects on chairmen's agendas, a high incidence of appeal of Ofcom's decisions is a much more recent phenomenon. Ofcom is obviously a much newer agency, but beyond its relative youth are a number of factors that have contributed to the increased incidence of appeal: the legal framework, the competitiveness of the marketplace, and the increasingly international nature of telecommunications firms operating in the UK. Because the uptick in litigiousness is recent and has coincided with agency budget cuts, it has caused Ofcom to become more conservative in regulating. This trend had a noted effect on the agency's approach to net neutrality.

Judicialization of UK Telecommunications

Ofcom faces a substantially higher standard of judicial review than other economic regulators in the UK, a standard that is arguably higher than what the FCC faces under the doctrine of *Chevron* deference. In transposing the 2002 EU Framework Directive, which required that "Member States shall ensure that the merits of the case are duly taken into account" (OJL 108/40, 2002) in the context of judicial review, the UK established the grounds for merits-based appeals of all Ofcom decisions to the Competition Appeal Tribunal (CAT) (c. 21, s. 192-197). Unlike a traditional judicial review in Britain wherein the court is confined to scrutinizing the process by which a regulator made its decision, the CAT can challenge any aspect of an Ofcom ruling, including, for example, "whether each of sometimes more than a hundred variables in a model used by Ofcom to reach its decision was correct" (DCMS 2011b, 6). In effect, the CAT can act as a duplicate regulator, with the ability to take an Ofcom decision and "look at it again and decide anew," as one industry interviewee characterized it – "a second regulator making the decision from first principles." The CAT can receive new input that was not previously available to Ofcom and can supplant Ofcom's decisions with its own (National Audit Office 2010). While judges may vacate FCC rules, their power to tweak or supplant the regulator's decisions with their own is far less than in the UK.

The fact that appellants can essentially receive a full re-hearing of every decision that Ofcom makes has been widely linked to the increasing incidence of appeal. The UK government has emphasized that the frequency, cost, and resource-intensiveness of appeals have been on the rise since 2007 (DCMS 2011b). Figure 9 is reproduced from a 2011 consultation impact assessment and demonstrates the rising costs associated with appeals.

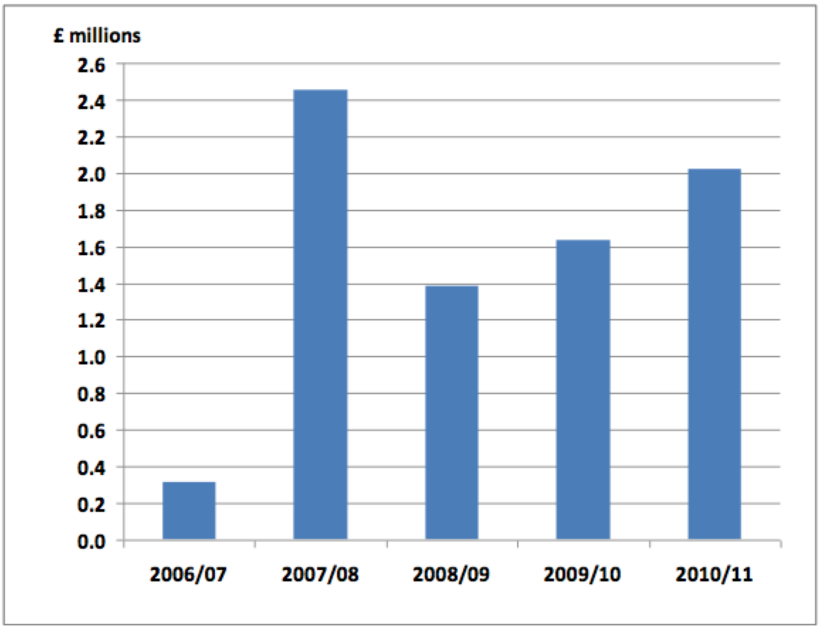

Figure 9. Ofcom's appeals costs, 2006-2011.
Reproduced from DCMS (2011b).

As one former Ofcom official noted, "the incentives on litigation are higher – much, much, much higher – under the current regime than they were under the previous regime," before the Communications Act 2003 was passed.

The change in the legal framework was accompanied by a number of other market developments that further increased litigiousness. As telecommunications markets became more competitive – in part thanks to Ofcom's early interventions discussed in Chapter 6 – firms began to realize that the costs of litigation could be low compared to the potential savings they could garner by reversing Ofcom decisions that advantaged their competitors. Interviewees from both Ofcom and industry were unabashed in describing this trend. One former Ofcom official explained that the companies had "nothing to lose and huge amounts to gain through appeals," while an industry executive noted that "the amounts of money at stake are so large relative to the costs of litigation, that people tend to revert to a strategy that says, 'well, it's worth a roll,' because unless the chances are completely hopeless, you won't be wasting your time." As the firms' margins were shaved thinner and thinner by new entrants into the market, they turned to litigation as a competitive tactic. Another former Ofcom official who was involved in the TSR described the effects of the competitive structure on the orientation of telecommunications firms towards litigation:

> One of the side effects of making the industry highly competitive and operating on very, very tight margins, is that people are looking for extra money everywhere. And it's kind of, "I've looked down the back of the sofa, I've shaken the piggy bank, you know, can we not get something out of BT or whoever it might be through bringing some really complicated dispute?" And so the art of crafting disputes and appeals against regulatory decisions so that you – the costs of mounting disputes or appeals is very low. The benefit can be enormous. Even getting the Competition Appeal Tribunal to agree to shave one pence per minute off a termination rate is worth millions and millions and millions. So why wouldn't you do it? So we kind of created a rather bizarre incentive structure, in my view, which means there will be more litigation.

Chapter 6 explained how Ofcom's major intervention to spur competition among broadband providers amplified ISPs' incentives to manage peer-to-peer traffic as a way to reduce costs in a highly price-competitive market. But peer-to-peer management was just one cost-saving behavior that competition incentivized; litigation was clearly another.

The uptick in litigation was also partly attributable to changes in the relationship between Ofcom and the industry. When Ofcom was still a new creation, one former official explained, it "had a tremendous amount of goodwill. The reason you got less court cases is because people didn't really want to mess with it." Over time, however, it became a "known quantity" where "you know you can bring a court case and keep it shut up for 18 months." At the same time, the key firms offering telecommunications services in Britain were becoming more international. This meant that they were both more accustomed to litigating in other countries and less likely to have concerns about the perception that by challenging the national regulator they were also challenging the interests of the nation. Both the fixed and mobile broadband markets saw entry from major foreign firms, including Telefónica, Orange, Deutsche Telekom, and Tiscali, while British firms expanded their presence overseas. Because they were all operating in an international context, the regulatory culture had shifted from one where "people used to . . . defer to government without anybody having to kind of say it too explicitly" to one where "no one cares what the British government wants," as one industry executive explained. If challenging the state in court had ever given them pause, the increasingly competitive and international nature of the market pushed firms in the opposite direction, emboldening them to use the threat of a lawsuit to their benefit. One former government official illustrated this brash use of legal threats:

> [A]ll the evidence seems to suggest that this has had a – some would say even a pernicious effect – on policymaking and regulation . . . the ability not only of operators to appeal against every decision that Ofcom ever makes, but also the way that they're up front about doing it. You know, like they'll say to the minister, "of course, if you do want to go down this road, then we'll be seeing you in court," you know, quite openly.

The "pernicious" effects of increased litigiousness were universally observed by interviewees from industry, government, and Ofcom: the agency became "conservative and slow." The time and resources dedicated to appeals had begun to hamstring the agency. Ofcom CEO Ed Richards questioned whether the UK had "ambled somnolently into a world where regulators are expected to make timely decisions to promote competition, but find it ever increasingly difficult to do so" (Richards 2010). The government was clearly alarmed: it launched three separate consultations in the span of four years on the subject of reforming the appeals system (the last of these was ongoing at this writing, with no legislative changes resulting from the previous two) (BIS 2010; BIS 2013; DCMS 2011c). In the 2011 consultation, the government warned against "a 'gridlock' of continuous overlapping appeals and market reviews" (DCMS 2011b, 8) and declared that "the resource allocated to defending market decisions before the CAT is unsustainable" (DCMS 2011c, 15). Notably, even if Ofcom's decisions were upheld by the CAT – which they were in the majority of cases (DCMS 2011b) – the agency still bore the costs of the appeals process. The perceived problem was not necessarily with the regulatory outcome, but with the costliness of achieving it.

With litigation consuming increasing staff time and budget even as the agency's overall resources were being cut for austerity purposes, Ofcom increasingly sought ways to limit its litigation risk so as to avoid further appeals. One former Ofcom official explained that as a consequence of the combined effects of budget cuts and increased litigiousness, Ofcom would "develop absolutely bullet-proof procedures for taking, you know, uncontroversial decisions . . . but will shy away from some of the more difficult, contested territory where it may not have to act." One government official further elaborated the quandary in which the regulator found itself:

> [I]t's quite difficult for Ofcom to make these big bang decisions. It means that the decisions have to iron-plated in order that they can't be appealed or that all bases are covered. At the same time, by iron-plating them, you end up with sort of a 400-page document where perhaps 30 pages would have done. And you have 400 pages of text that can be appealed . . . it's all rather chicken and egg, clearly . . . . I think it's the view of a number of people in government if not the government's formal position . . . that the system is gamed, the system is dysfunctional, and is leading to both regulatory log-jam and the sort of cartelization of certain aspects of the market because that's an inevitable consequence of an appeals system that means it's highly beneficial for the status quo to be maintained. So as a consequence of this, I think Ofcom has concentrated on slightly more winnable consumer issues. And that's why people say they're not brave anymore, they don't take these big bang positions.

In contrast to the innovative steps it had taken early on, the agency felt increasingly constrained by the need to spend time and resources fighting appeals and increasingly averse to taking any action that would risk further appeal. Thus, in addition to reputational pressures, litigiousness was yet another factor that caused Ofcom to both recede from high-profile intervention and focus more attention on consumer protection. Regulating the relationship between businesses and consumers, rather than amongst businesses (as would have been required in the net neutrality case), was perceived as less risky.

Effects of Litigiousness on Ofcom's Approach to Net Neutrality

In this climate, the issue of net neutrality was not one that was apt to spur the regulator to take bold action. A number of interviewees felt that, had Ofcom taken a more regulatory approach to the issue, an appeal would have been inevitable. "I don't see how they see they can make a decision on net neutrality that wouldn't be challenged by each and every player in the market" was how one described the situation. Another explained that to build a case for intervention Ofcom "would have to show examples of detriment because of blocking" and that "this would be something they would definitely get appealed on, so I just don't think they would be able to make that case."

As discussed in Chapter 8, neither the effects of regulation nor the consequences of failing to intervene to prevent discrimination were fully understood or quantifiable when Ofcom was formally considering its approach to net neutrality. The agency was already predisposed to avoid tackling issues involving ambiguity or any perception of policy-setting. The legal pressures added an impetus to either intervene on the basis of overwhelming supporting evidence, or not at all. The regulator clearly decided that the justification for regulating in support of nondiscrimination failed this test. The fear was not that Ofcom would be overturned by the CAT, but that it would be forced to face the CAT based on the industry's opportunistic exploitation of perceived evidentiary weaknesses.

While the FCC intervened with full knowledge that a court challenge would ensue and a degree of ambivalence about being overturned, Ofcom allowed the threat of litigation to constrain its decision-making. One former Ofcom official observed that at the agency, "the only way to actually know if you're actually going to win is if people don't want to fight you in the first place." Over the course of its short existence, Ofcom transitioned from a bold, risk-taking regulator to a more conservative, cautious one, grounding its definition of success in avoiding challenges from the industry rather than using regulation to revolutionize it. No regulatory agency wants to find itself in court, but whereas litigious culture did little to deter the FCC chairmen from pursuing their goals with respect to net neutrality, it did much to shape Ofcom's approach to the issue. The combination of litigation risk with internal and external pressures to recede from policymaking and focus on consumer protection served as a potent mix of influences to deter Ofcom from intervening ex ante, even in the face of pervasive discriminatory traffic management.

9.2.3 Judicial Review and the Regulatory Life Cycle

Notably, the experiences of the FCC and Ofcom illustrate an interesting twist on Bernstein's seminal "life cycle" model of regulatory behavior (1955). As explained in Chapter 3, in Bernstein's life cycle agencies begin with a pioneering willingness to antagonize the regulated industry, move into a capture phase where their processes become highly judicialized, and settle into an ossified old age where they do little more than maintain the status quo. What this model misses are the effects of the novelty of certain forms of adversarialism – namely, litigation – and how those effects change as novelty transitions into desensitization. Compared to the FCC, Ofcom is clearly an agency in its youth, dealing with judicial antagonism as a new phenomenon with the potential to substantially re-shape its character as a regulator. For the FCC, litigation has been a mainstay for so long that its threat no longer affects the policy course that the chairman aims to set. Thus, while the age of an agency may be important for understanding its behavior, the correlation between age and the will to regulate may not be as Bernstein suggests. Newfound litigiousness can cause an agency to adjust to avoid clashing with the industry, but once the agency has grown accustomed to the inevitability of litigation, legal threats from the industry may have little effect.

9.3 Interest Groups

Comparing the US and UK experiences with net neutrality regulation sheds light not only on the applicability of Bernstein's conception of industry capture, but of regulatory theories based on interest group influence more broadly. As discussed in Chapter 3, scholars have elaborated a diversity of theories to explain regulatory results as a function of interest group influence, focusing on the political salience of particular groups (Peltzman 1976), the relative concentration or diffusion of regulation's costs and benefits (Wilson 1975; Wilson 1980), and the extent to which some groups are more efficiently organized than others (Baldwin and McCrudden 1987; Becker 1983; Rothstein 2004). Despite their differences, one common assumption among these and other theories of interest group regulation is that at least one salient group will seek regulation to begin with. In the case of net neutrality regulation, however, some of the key differences between the US and the UK relate to the presence and intensity of interests advocating for regulatory intervention. As this section demonstrates, this more fundamental question about interest groups – whether they participate and actively make demands of the regulator – explains much about regulation in this comparative context.

9.3.1 Profile of Net Neutrality as a Regulatory Issue

Early in its life as a regulatory topic in the US, net neutrality became a contentious, high-profile issue in which stakeholders of many kinds became deeply involved and invested. The legislative debate in 2006 spawned a massive grassroots mobilization that involved dozens of rallies, delivered millions of petition signatures to Congress, and fueled the years of heated debate that followed (MoveOn.org 2013). Nothing of the kind took place in the UK, where net neutrality was initially viewed as an American problem resulting from lack of infrastructure competition (Atherton 2010; Sky 2010; TalkTalk Group 2010b) and where a "noisy debate" (TalkTalk Group 2010a, 9) never subsequently developed. As one former Ofcom official explained, net neutrality "really came to our attention very much driven by, wow, there's some really interesting developments going on in the US and it seems to have triggered a whole tsunami of lobbying and quite a lot of very interesting technical and academic analysis of the issues, but no comparable drivers in the UK."

Between 2006 and 2010, more than 10 times as many articles about net neutrality appeared in the mainstream press in the US as they did in the UK (Powell and Cooper 2011). More than 100,000 comments were filed with the FCC during the *Open Internet* proceeding

(Genachowski 2010c) while fewer than 100 were filed during Ofcom's consultation on net neutrality (Ofcom 2010f). There are many factors other than the intensity of each country's respective debate that contributed to the disparity between these figures, but they are illustrative of the differences in the profile of the issue among stakeholders and the public in the two countries.

Both public interest advocates and Internet applications companies were heavily involved in advocating for neutrality legislation and regulation in the US. Telecommunications, and net neutrality in particular, became a top lobbying priority for Google and other Internet firms as the debate developed (OpenSecrets.org 2013). Although their net neutrality lobbying budgets remained a fraction of those of the ISPs (Allison 2009), net neutrality was one of several issues that spurred growth in lobbying spending among Internet application providers in the mid-to-late 2000s.

Public interest advocates seized on the net neutrality issue in its early stages and played a critical role in urging a regulatory outcome that included strong nondiscrimination safeguards. The Electronic Frontier Foundation provided the most rigorous early proof of Comcast's use of the TCP reset solution (Eckersley, Von Lohmann, and Schoen 2007) and the *Comcast* proceeding that followed was initiated, sustained, and the outcome substantially negotiated by Free Press with help from other advocacy groups and Vuze, a small peer-to-peer service provider. Analysis of regulatory and mass media discourses shows that within the US net neutrality debate, arguments first made by public interest groups later appeared in the FCC's own pronouncements, demonstrating the spread of the groups' influence (Powell and Cooper 2011). These groups often joined forces with major Internet applications providers in their advocacy efforts (Erickson 2010).

There was no comparable coalition actively seeking regulatory or legislative intervention to safeguard nondiscrimination in the UK. Skype and Yahoo were the strongest private sector advocates for regulatory intervention (Sherwood 2010; Skype 2010), but they had little support from others. As alluded to in Chapter 7, the BBC leveraged its public presence to advocate for open Internet principles, but its influence over Ofcom's decision-making is generally complicated by its unique position as a public service broadcaster. A handful of advocacy organizations, consumer groups, and academics followed the net neutrality debate, but when compared to other digital rights issues, their advocacy around net neutrality was minimal (Powell and Cooper 2011). Several interviewees noted that there had been no incident or regulatory proposal in the UK akin to the Madison River VoIP blocking in the US or the HADOPI copyright enforcement law in France to galvanize a critical mass of activists around the net neutrality issue. One industry executive noted the absence of any sizable constituency to advocate for intervention:

> I mean, we have no kind of public interest groups. We don't fund public interest groups. So there's not lobbying dollars behind the case. Google doesn't seem to be that interested in really making an issue of it. So there's no one on the other side really saying . . . I mean there are Internet activists and so on in the UK, but they don't seem to be picking this issue as a huge issue.

This void gave Ofcom a powerful role in framing the debate, allowing it to set up the arguments to which a tiny cadre of neutrality advocates responded. Within such a structure, advocates often employed not only the terms but also the arguments put forth by Ofcom, and correspondingly refrained from demanding strong intervention. The most striking example of this pertained to Ofcom's often-repeated view that the combination of competition and transparency would keep discrimination in check and give consumers the ability to switch ISPs if discrimination got out of control. The Open Rights Group (ORG), the most active advocacy organization on the neutrality issue, essentially parroted Ofcom's argument:

> Thankfully in the UK there is currently a healthy amount of competition and choice this [sic] is currently preventing any company from trying to drop Net Neutrality. . . . There is a small amount of traffic shaping happening in relation to ports used by P2P software by some network providers, this should be made more explicate [sic] to the customers when they sign up. (Open Rights Group 2010)

ORG not only accepted Ofcom's framing of the issue in competition terms, but it also agreed with Ofcom that instances of discrimination against particular applications could be rectified through transparency. While this view was certainly not universal among the handful of interested advocates, it is revealing of the influence that Ofcom's framing had on the debate in the absence of a sizable constituency demanding stronger safeguards. The kind of discrimination that ORG implicitly endorsed is precisely what pro-neutrality groups and companies argued against in the US.

9.3.2 Emotional Investment

Beyond the differences in press exposure, lobbying expenditure, and advocacy constituency, perhaps the most significant difference in character between the US and UK debates was the intensity of the emotions involved. While the US debate could arguably be characterized as the most intensely emotional telecommunications policy debate of the 21st century, the issue was treated with ambivalence in the UK.

United States

Many interviewees observed that the level of emotion and personal investment that numerous stakeholders exhibited in the US net neutrality debate was atypically high. The public rhetoric was heated from the very beginning, with the CEO of AT&T claiming that applications providers would like to "use my pipes free, but I ain't going to let them do that" (O'Connell 2005) and stakeholders on either side launching slogan-oriented campaigns to "Save the Internet" or conversely to keep the government's "Hands off the Internet." From there, the issue became (in the words of interviewees) "noisy," "polarizing," "political," "overheated," "infected," and full of "extraordinary hostile partisan venom." It was a debate "where people were not willing to say the other side has a point" – a religious debate, in essence. And as one policy executive noted, in "[r]eligious debates you're never going to convince anyone. 'My religion is better than yours.' Then we go to war." By telecommunications policy standards, the rhetoric was vicious.

Comcast's behavior after it was accused of blocking peer-to-peer file transfers was a particularly potent catalyst for heightened rancor. Company spokespeople initially denied that the ISP was blocking or throttling any traffic (Svensson 2007; Wein 2007). Comcast then changed its Terms of Service and issued new, more detailed disclosures (Casserly, Wallach, Alvarez, Waz, et al. 2008), but continued to insist that it was only “delaying” peer-to-peer downloads and uploads, not blocking them (Albanesius 2007; Casserly, Wallach, Alvarez, Waz, et al. 2008; Cohen 2008). Interviewees recalled that the company “wasn’t very transparent,” was “really not being forthright,” and “dug a hole for themselves by not being more forthright from the beginning.” The company’s behavior clearly angered public interest advocates and Chairman Martin, all of whom took Comcast to task for both shifting its story and continually insisting that it was not blocking traffic (Martin 2008b; Sohn et al. 2008). One interviewee noted that, “[i]t was a really easy sound bite for the commissioners and for the public interest groups that were pushing. The hallmark of whether something is reasonable is whether it’s disclosed, whether it’s transparent. And they [Comcast] didn’t have that argument.” Another explained that Comcast’s behavior “bred a lack of trust, both to the public and the FCC,” further incensing stakeholders who were already embroiled in a hostile regulatory fight.

As the debate wore on, stakeholders became increasingly invested at a personal level in achieving their desired outcomes. By many interviewees’ accounts, net neutrality was a more emotional issue than any other telecommunications policy issue in recent memory. “[T]here was a surprising amount of emotion behind it,” one former FCC staffer explained, recalling that “I remember sitting in a lot of meetings thinking, ‘Wow, these people are really pissed off.’” Speaking to the religiosity of the debate, another former staffer observed that “there was more emotion in the debate than logic in some ways.”

It is difficult to identify the precise causes of this dynamic, but interviewees honed in on the fact that for stakeholders of many different kinds, the way that the Internet functions is utterly central to their lives, their work, and their businesses. Net neutrality therefore went to the very core of their personal and organizational motivations. One policy executive from an Internet application company explained how the substance of the issue got to the foundational interests of the businesses on both sides:

> [A] carrier that blocks your innovation process goes straight at what makes you go. So there's something very personal about that. Likewise, I think for the network operators, their world view is . . . "we're doing the heavy lifting to get the Internet to all of America," and I think there's something about the fact that government would micromanage that heavy lift that is sort of the analog to what I'm feeling about the innovation process. . . . And so in both cases you're going really at the core function, the core values of what the ISPs as well as what the Internet guys do.

This feeling that the regulatory outcome could have a foundational impact was at least as strong, if not more so, among individuals and public interest advocates. The following exchange from an interview with a former FCC staff member who was heavily involved in the *Open Internet* proceeding reveals the unique position that the Internet occupied in some stakeholders' lives and how the importance that the Internet held for them personally stoked the unprecedented emotionality of the debate:

> Respondent: There were some people for whom I think this was a very emotional issue. Yeah, I had people actually in tears in some of the meetings I had, which surprised me.
>
> Interviewer: Really?
>
> Respondent: Yeah, it's interesting, I thought about it at the time, because people don't come in here and talk about intercarrier compensation and break down crying – at least they haven't yet, we'll see when we get closer to the order on that, maybe that will happen. But I think there's the sense that the Internet is really this wonderful thing, I think people feel like it's part of them, it's part of us, there's a special connection to it that people don't typically have to special access or certain spectrum issues, and so when there's the possibility for changes that could affect it in ways people perceive to be negative – although I think there's a questionable basis sometimes for that perception – my sense is that it becomes very personal. They feel that there's some part of them or something that they're tightly enmeshed with is harmed or is in jeopardy of being affected. Very strong emotions, yeah.

The strength of people's personal connections to the Internet elevated the emotional importance of net neutrality and helped compel the regulator to act, in utter contrast to the UK.

United Kingdom

In comparison to the US, the net neutrality discussion in the UK was largely devoid of emotion – and even of demonstrated interest, in some respects. The British public never took interest in the issue on any significant scale. One former Ofcom official observed that "[p]eople didn't seem to be that agitated about it" when the agency issued its discussion document. One reason for this is that it was not viewed as being part of broader social policy or digital rights, and even the consumer protection aspects of the issue were not the ones of typical concern for consumers, as the former official explained:

> [C]onsumers felt here that they have a relatively good terms [sic] of broadband services and good quality. The one thing that has created a lot of consumer dissatisfaction here . . . is the inaccuracy of broadband speeds that are advertised or actually delivered. That is related to net neutrality to some extent. So there was a lot of consumer concern but it was a consumer issue, it wasn't a citizen issue. There was no sense that this was important for democracy or free speech or something like that, or not to that extent.

A narrow consumer or commercial issue such as broadband speeds advertising was capable of piquing the interest of the public, but there was no compulsion to get exercised over the broader policy questions inherent in net neutrality regulation. Again here, Ofcom's framing of the issue matched (and perhaps guided) the broader public's understanding of it. For example, in 2008, Ofcom CEO Ed Richards characterized the US net neutrality debate as follows:

> The question of paying more for better services, whether traffic prioritisation, higher speeds or higher usage limits, has also been caught up in the net neutrality debate in the US. In part this reflects a slightly atavistic sense that the internet ought to be free and egalitarian in all respects. (Richards 2008)

Characterizing concerns about discrimination as backwards-looking, principles-based worries reveals an inherently functionalist view of the Internet's value. The regulator as much as members of the general public did not accept the idea that deciding the questions of whether and how discrimination should occur would have a fundamental impact on society. Instead, the Internet was viewed as a technical implement wherein decisions about its functional intricacies would perhaps benefit from technocratic oversight, but need not be of broader societal concern.

Some interviewees placed the ambivalence over net neutrality in the context of more generalized apathy for activism in the UK compared to the US. One interviewee linked the lack of public outcry over net neutrality to “the psyche of the country,” claiming that Britons “queue and complain formally and properly, but they don’t get upset.” Another characterized it as cultural:

> We’re just not that “arsed” in this country. “Here, come on, join!” “Aw, I can’t be arsed, I’m watching TV.” Again it’s about vested interests – is it really going to impinge on my life? Not really. When it does, maybe I’ll moan about it, or I won’t do anything. It’s just cultural.

Identifying the links between British culture or psyche and public engagement on net neutrality is beyond the bounds of this thesis, but these observations help to explain the disparate treatment of the net neutrality issue between the two countries. The FCC was not only a regulator tasked with setting broad social policy, it was also confronted with a coalition of interests making passionate arguments about society’s broad interest in safeguarding nondiscrimination. In a context where no group made similar demands and where the Internet was viewed as a commercial good much like any other, Ofcom felt no comparable impulse to act.

9.4 Conclusion

The last two chapters demonstrate the most salient factors that differentiate the telecommunications regulation paradigms in the US and the UK. The US regulatory environment was characterized by individual reputational ambitions, accustomization to litigation, and an emotional public debate spurred on by a powerful coalition of Internet companies and public interest advocates. The combination of these forces created an environment in which regulatory intervention was viewed as appropriate and even necessary. In the UK, the reputational pressures on Ofcom as an organization, conservatism induced by increasing litigation, and the absence of interest groups demanding action contributed to regulatory forbearance. These factors took clear precedence over the presence or absence of discrimination and competition.

These explanations for regulatory activity fail to provide support for many of the earliest and most thoroughly developed theories of regulation based on external forces: interest groups, legislatures, the executive branch, and the courts. The net neutrality regulatory arenas in the US and UK show few signs of regulatory "capture" insofar as it was originally conceived, with regulatory agencies producing regulation to the benefit of the industries they oversee (Stigler 1971). The FCC's actions were widely perceived to be against the interests of the broadband industry and Ofcom did not produce regulation for any industry sector. Moreover Ofcom's inaction was largely a result of its own instinctive dedication to competition principles – the agency itself originated the arguments used to rationalize its decision not to intervene. Thus this thesis adds empirical support to the sizable literature that questions the validity of the capture theory (Derthick and Quirk 1985; Noll and Owen 1983; Posner 1974; Wilson 1980).

The comparison between the US and the UK provides an important lesson in light of other theories that link particular interest group characteristics to regulatory outcomes. More fundamental than the groups' political salience, ability to efficiently organize, or breadth of represented interests (Becker 1983; Moe 1987; Peltzman 1976; Wilson 1975; Wilson 1980) is their impetus to act. If neither the public nor any sizable industry sector nor advocacy organizations are moved enough by a particular policy question to petition the regulator to take action, there is little reason for an agency to act on its own initiative. In Ofcom's case, this reticence was reinforced by a climate of limited budgets, skepticism about its policy role, and litigation-induced conservatism.

The two cases studied here also do not comport neatly with prominent theories of regulation based on political actors: the "bureaucratic view" that agencies act autonomously outside the bounds of oversight from other branches of government (Gilardi 2005; Majone 1999; Niskanen 1971; Thatcher 2002a), the theory of congressional dominance of agency behavior (McCubbins and Schwartz 1984; Weingast and Moran 1983), and president-centered theories of agency control (Moe 1985; Moe 1987b; Moe and Wilson 1994). The FCC and Ofcom both showed clear streaks of independence from the political system, but also influences from it. For example, Kevin Martin's aggressive response to the accusations against Comcast were met with consternation in both the legislative and executive branches, while Parliamentarians at times questioned Ofcom's corporate style of operation. At other times, however, the agencies' actions were in obvious alignment with the preferences of the President or Prime Minister: Chairman Genachowski's pursuit of industry-wide regulation and Ofcom's acceptance of its role as a technocratic policy implementer, for example. What consistently explains the agencies' actions over the time period studied is the role of reputation within each agency, not the ability of other political actors to control the agencies' actions.

Finally, the case of traffic management regulation shows the importance of a regulatory agency's life cycle stage in understanding the effects of external forces. Legal scholars have theorized that increases in judicial review cause agencies to become conservative, rulemaking processes to become ossified, and informal means of regulating to proliferate (McGarity 1992; Melnick 1992; Pierce 1995; Seidenfeld 1997). These effects appear much more prominently in the UK than in the US, although not for lack of litigation in the US. Instead, while the FCC's accustomization to litigation muted the extent to which litigation risk affected its decision-making, for Ofcom litigation risk was a newer and therefore highly influential phenomenon. There may indeed be a regulatory "life cycle" (Bernstein 1955), but in the case of traffic management, a later stage of the cycle implies an emboldened regulator, not the opposite. More broadly, the FCC's behavior reflects its greater de facto independence from all branches of government (Maggetti 2007). Although both agencies were endowed with significant formal independence by statute, the FCC's ability to take more political risks reflects the increased control it has accrued over its own agenda over a longer lifetime.

Although these findings provide little direct support for some of the most well-studied theories of regulation based on external forces, the relevant theoretical landscape is far broader, as discussed in Chapters 2 and 3. The next chapter generalizes the analysis of the findings in this thesis and interprets them in light of these broader theoretical foundations and the hypotheses offered at the outset of this study.

Chapter 10.
Conclusion

10.1 Introduction

Academic and policy debates about net neutrality and traffic management have been largely normative and theoretical, characterized by a wide diversity of predictions and opinions about the relationships between discrimination, competition, regulation, and innovation. This thesis makes a substantive, positive, empirical contribution to both scholarship and policy concerned with these issues. The inquiry was framed by two research questions:

> Question 1. Why do network operators take up discriminatory traffic management (or not)? In particular, how does competition in the market for broadband service influence network operators' traffic management decisions?
>
> Question 2: How does the institutional setting – the formal and informal constraints that comprise the regulatory environment – influence traffic management outcomes?

In answering these questions, this thesis explains the reasons that network operators do or do not take up discriminatory traffic management, why and how regulators choose to intervene or not in response, and the influence of competition on traffic management outcomes. These explanations are derived from a qualitative, comparative study consisting of elite interviews, participant observation, and documentary analysis. This chapter reviews these findings, explains how they contribute to scholarship concerning net neutrality and regulatory theory, and discusses their broader applicability. It concludes with a discussion of future directions for research and policy.

10.2 Research Findings and Contributions

10.2.1 Network Operator Decision-Making

Scholars and practitioners seeking to assess the motivations for network operators to take up discriminatory traffic management have to date primarily focused on operational drivers of a technical or business nature. The net neutrality literature points to three core types of rationales to explain the use of discriminatory traffic management – to control performance and/or cost, to segment the broadband market, and to disadvantage competing applications – and these motivations have been much debated by policy stakeholders.

This thesis demonstrates that controlling performance and/or cost is by far the most common operational concern that causes ISPs to adopt discriminatory traffic management solutions. The impact of peer-to-peer traffic on the performance of other applications was the central operational factor that drove adoption among US cable companies, supporting the suggestion in the literature that changing network usage patterns motivate the use of discriminatory solutions (Clarke 2009; Hahn, Litan, and Singer 2007; Marsden 2010; Shelanski 2007; van Schewick 2010). The engineering specifications of DSL and fiber networks lacked the same characteristics that created the performance problems for cable, allowing the US telcos to refrain from discriminatory traffic management without undue concern for performance degradation. In the UK, the wholesale market structure created a broadband landscape in which cost and performance were intimately intertwined and in which peer-to-peer management was viewed by many ISPs as a critical tool for controlling both. As price competition intensified after the expansion of LLU, the perception of that need only intensified. Together these findings support claims in the literature that bandwidth costs can drive adoption of discriminatory traffic management (Marsden 2010; van Schewick 2010), as can achieving a particular level of network performance within ISPs' budget constraints (Crocioni 2011; Hazlett and Wright 2011). However, the costs associated with the equipment used to conduct application-specific management can be just as important and can drive some operators away from application-specific solutions.

The UK provided examples of instances in which ISPs used discriminatory traffic management to segment their customer bases or as leverage in business negotiations, as envisioned by some net neutrality scholars (Litan and Singer 2007; Marcus 2008; Renda 2008; Valcke et al. 2009; Weisman 2010; Yoo 2004; Yoo 2005), but those were exceptions to the general trend. Assessing the extent of anti-competitive motivations would require more intimate knowledge of ISPs' decision processes than could be uncovered in this research.

Operational considerations are just one component of ISPs' overall decision-making, however. Operators' broader institutional settings – consisting of pressures and constraints from regulators, consumers, and other stakeholders – are fundamental factors that are determinative of whether they choose to pursue discriminatory traffic management. In the US, the regulatory histories of individual companies and their corresponding internal regulatory oversight structures circumscribed the set of traffic management choices that they conceptualized and selected. The telcos had battled for decades to have their regulatory requirements lifted, along the way creating corporate oversight structures whose ultimate effect was to prevent them from taking the kinds of regulatory risks that discriminatory traffic management entailed. By contrast, the cable companies' relative regulatory success and lack of internal policy oversight gave some cable marketing teams an opening to steer their businesses towards DPI-based traffic management solutions that they believed would provide them with customer insights while also solving performance problems. In the end, even in the absence of industry-wide regulation, the FCC's intervention in the *Comcast* proceeding at the urging of public interest advocates and Internet companies put both cable and telco ISPs on notice that the threat of regulation was real, and those that had adopted discriminatory traffic management reversed course.

The threat of regulation – or even significant consumer backlash – was nearly non-existent in the UK. At the time when ISPs were making early decisions about traffic management, Ofcom was focused on spurring local loop unbundling and promoting competition. Its initial position was that discriminatory traffic management was a technical necessity – an argument that even the ISPs themselves did not make as vociferously and that proved to be untrue in practice. The agency was not inclined to undermine its enduring dedication to competition by admitting that competition might not provide an adequate safeguard against discriminatory conduct. At the same time, UK users and consumer groups were loathe to complain visibly about discriminatory traffic management. Consumer groups adopted Ofcom's rhetoric while user complaints remained muted, in part out of users' fear of being associated with claims of illegal copyright infringement, as peer-to-peer users had been previously. UK ISPs seized the freedom that their institutional environment provided them to adopt discriminatory traffic management, often using blunt approaches that remained largely static over time.

Thus regulatory threat explains differences not just between the US and UK but also within the US. The threat of regulation was clear to US telecommunications providers with a long history of regulatory strife, less clear to US cable operators who made forays into discriminatory traffic management, and not present for UK providers. The evidence from these countries supports the hypothesis offered at the outset concerning the reasons for ISPs to adopt discriminatory traffic management:

> Hypothesis 1: The threat of regulation limiting how operators can manage traffic acts as an informal constraint on operator behavior. This threat is at least as important, if not more so, than the competitiveness of the market.

Where regulatory threat is present and internalized by ISPs, it fundamentally shapes traffic management, while its absence has an equally strong effect. This finding supports the arguments that scholars have made about the deterrent effect of regulatory threat (Felten 2006; Marsden 2007; Wu 2003; Wu 2004). Principles-based case-by-case enforcement, as was used in the *Comcast* proceeding, can clearly also have deterrent effects, as some scholars had envisioned (Atkinson and Weiser 2006; Bauer, Clark, and Lehr 2009; Greenstein 2007; Lessig 2006; Sluijs 2009; Weiser 2003; Weiser 2008). Regulators need not enact detailed, prescriptive rules in order to influence network operator behavior.

Overall, these findings provide an in-depth qualitative view of how technical, economic, and regulatory pressures feed into internal ISP decision-making about traffic management. They complement the quantitative work of Mueller and Asghari (2011) and Asghari, Van Eeten, and Mueller (2012) that draws links between discrimination, regulatory activity, and market characteristics using network traffic measurements and economic metrics. The qualitative insights herein also enhance the body of work that has sought to characterize ISP discrimination by analyzing network traffic and operator policies (Beverly, Bauer, and Berger

2007; Dischinger et al. 2008; Dischinger et al. 2010; Kreibich et al. 2010; Li and Losey 2009; Sidak 2006; Tariq et al. 2009; Weinsberg, Soule, and Massoulie 2011; Wu 2003; Wu 2007; Zhang, Mao, and Zhang 2009).

10.2.2 Influence of Competition on Discriminatory Traffic Management

With respect to competition, the UK case provides strong evidence that competition does not deter discrimination, and that instead it reinforces the drive to conduct discriminatory traffic management within an institutional setting that lacks countervailing pressures from regulators or consumers. Competition drove UK broadband prices down, making any cost-saving option attractive to ISPs, including peer-to-peer management. Most consumers did not understand traffic management or use it as a basis for switching even in a competitive marketplace with a handful of nondiscriminatory choices. Those consumers who were concerned about traffic management comprised a small enough group that ISPs felt safe in pursuing discriminatory strategies even if it meant losing a fraction of (high-volume) users.

The evidence from the UK contradicts the often-stated view in the net neutrality literature that competition reduces incentives to discriminate (Becker, Carlton, and Sider 2010; Cave and Crocioni 2007; Chirico, Haar, and Larouche 2007; Faulhaber and Farber 2010; Hahn, Litan, and Singer 2007; Nuechterlein 2009; Shelanski 2007). Even under competitive circumstances, ISPs can still be motivated to adopt discriminatory approaches for performance improvement, as envisioned by van Schewick (2010) and Wu (2003). The UK case demonstrates that the combination of information asymmetry (Lennett 2009; Marsden 2007; van Schewick 2007) and switching costs (Bar et al. 2000; Economides 2008; Krafft and Salies 2008; Wu 2007) can effectively dilute the market discipline that having multiple competing operators is intended to provide. Scholars have claimed that opening access networks to competition by regulatory mandate, as Ofcom did, would safeguard nondiscriminatory access to content and applications (Bar et al. 2000; Cooper 2003; Lemley and Lessig 2001). Open access can clearly have the opposite effect, driving prices down and bandwidth demand up, thus making discriminatory behavior a profitable strategy, as modeled by Hogendorn (2007). When a fraction of consumers are directly affected by discriminatory traffic management and when consumers generally do not internalize the costs to innovation of discriminatory practices – as proponents of regulation have argued (Lemley and Lessig 2001; Lessig 2001; van Schewick 2010; van Schewick 2012) – a competitive market with pervasive discriminatory traffic management can arise as it did in the UK. These findings support the formally modeled results of Reggiani and Valletti (2012) and Guo, Cheng, and Bandyopadhyay (2012) that show how discrimination can harm niche or fringe content or applications providers and their users while benefitting large, established providers.

The UK case also shed light on a question that was not at the center of this thesis, but that motivates many academic and policy discussions concerning net neutrality: does discrimination create barriers to application innovation? The findings contained herein demonstrate that discrimination can create costs for application developers whether their products are the targets of application-specific traffic management or not. That application-specific management requires network operators to observe and classify traffic means that some application developers must expend resources to avoid having their applications misclassified or risk performance degradation. Thus discrimination can lead to negotiations between application developers and network operators, as proponents of net neutrality regulation have feared. It can also lead to "arms races" between network operators seeking to identify applications and application developers aiming to avoid having their products' traffic managed (Lehr et al. 2006; Marsden 2010; Sandvig 2007). Operators in the UK were aware of these drawbacks but remained committed to application-specific management in many cases. There was no competitive, regulatory, or consumer force that compelled them to do otherwise.

10.2.3 Regulatory Decision-Making

The findings of this thesis reveal that telecommunications regulators operate within an "institutional matrix" (North 1990) that consists of a diversity of constraints that guide their regulatory decision-making concerning traffic management. The evidence presented here supports the claim that institutionalist scholars have made that "institutions matter" in determining regulatory behavior (Baldwin, Cave, and Lodge 2012; Black 1997). Traffic management regulation cannot be understood as a market good governed by the laws of supply and demand (Stigler 1971), but as the product of a multi-faceted institutional landscape (Galperin 2004).

Some of the constraints that telecommunications regulators face are formal, derived from law, regulation, or administrative practice, while others are informal, constructed based on conventions, norms, and perceptions (North 1990). These formal and informal factors, many of which are interrelated, fall across three of the four categories of regulatory theories identified in Chapter 3: institutional design, external forces, and internal characteristics, as shown in Table 2. The fourth category, nation-specific factors, is discussed further below.

	Formal constraints	**Informal constraints**
Institutional design	Governance structure	Perception of remit Endowment of policy-setting function
External forces	Standards for appeal	Perception of litigation risk Interest group participation
Internal characteristics		Reputation

Table 2. Constraints in the regulatory environment that influence traffic management.

Matrix of Institutional Influences

Analysis of the FCC and Ofcom demonstrates that certain formal constraints create the foundational structure in which informal constraints drive regulators' decisions about whether and how to engage with respect to traffic management. The institutional design consideration of agency governance structure circumscribes the kinds of reputational pressures that regulators face, which in turn shape their traffic management decisions. The FCC's commission structure vests substantial authority with the chairman, allowing him to dictate the agency's agenda and priorities. This gives the chairman significant freedom to pursue policy objectives according to his preferences, albeit within the confines of political considerations relating to Congress and the need to garner supporting votes from his fellow commissioners. As a result, the steps that the FCC takes are a strong reflection of each chairman's ideals and personal reputational aspirations. During the last decade, the agency had successive chairmen who, for different reasons, felt that regulatory intervention to deter discriminatory traffic management was necessary. Had there been different chairmen, these outcomes surely would not have been the same, although they still would have reflected the personal image that the chairmen intended to convey.

Ofcom's more corporate structure yielded a focus on regulatory activities that could build and sustain the agency's reputation in the long term. It was endowed from its creation with a strong vision of what it should be as an organization (regardless of who was leading it): a rigorous, evidence-based, industry-savvy regulator with a competition-centric viewpoint. As a policy issue, net neutrality presented challenges that carried reputational risk because determining a course to pursue required assessing ambiguity, conducting qualitative assessment, and possibly conceptualizing the issue outside the bounds of traditional competition analysis. As many corporations would have done (particularly in light of budgetary constraints, which the agency was facing), Ofcom proceeded with caution, declining to intervene to alter existing discriminatory practices. The combination of governance structure and reputation was in sharp contrast to that of the FCC, and produced a sharply divergent outcome.

An agency's perception of its own remit and the extent to which it is viewed as a policymaking body are also paramount in defining the space of traffic management approaches that it is willing to consider. The policy issues associated with ISP discriminatory conduct, including those related to innovation, go beyond the core competency that would be expected of a traditional competition regulator. Likewise, addressing those issues through regulatory intervention potentially requires the regulator to set a national Internet policy direction and decide technical and economic questions about the relative balance of power between different industries. To be able to fully consider policy questions concerning discriminatory traffic management, a telecommunications regulator must both perceive its remit broadly, as encompassing industrial and social policy, and be endowed with the ability to set such policy. The FCC is such a regulator and took advantage of the breadth of its remit and policymaking authority to act. By contrast, Ofcom viewed itself primarily as a competition regulator, was imbued with a technocratic culture, and had good political reasons for resisting the public appearance that it was engaging in setting an Internet policy direction for the UK. These were central factors contributing to its reticence to act.

It is important to recognize how powerful these factors are despite being legitimately informal. Ofcom for many years focused its telecommunications work on competition regulation, but this was based on its own choices and culture, not specific restrictions in its authorizing statute. The Communications Act 2003 tasks the agency with duties in relation to both citizens and consumers and its remit is not limited to competition matters (c. 21, s. 3(1)). Similarly, the extent to which Ofcom felt that it could engage in policymaking was clearly fluid and changed based on political circumstances, not amendments that were made to existing law or regulation. Recognizing this informality is crucial because it demonstrates the extent to which the factors that contribute to regulatory outcomes related to traffic management are self-determined by regulatory agencies, not bestowed by laws or legislatures. When the legal bounds of a regulator's authority are broad, as was the case for both the FCC and Ofcom, the regulator's conception of itself comes to hold tremendous sway over the path that the agency pursues. Narrow self-conceptions leave the innovation and social questions at the crux of traffic management policy out of scope.

In addition to institutional design factors, external forces related to judicial review present another instance in which formal constraints provide a foundation on which informal ones serve to influence traffic management outcomes. The legal structure that defines how and by whom a regulatory agency can be challenged in court sets the foundation for whether litigation risk factors into regulatory decisions about traffic management. In both the US and the UK, these legal standards were certainly taken into account; they took on much more prominence in the UK given that the Competition Appeal Tribunal has the power to decide Ofcom's rulings anew. But just as critical as the legal framework governing appeals is the regulatory agency's perception of its litigation risk and the effects that litigation may have on the agency. The FCC was so accustomed to being sued over its decisions that while it sought as best it could to avoid a lawsuit over the *Open Internet Order*, in the end the near-certainty of litigation was not enough to deter Chairman Genachowski and his fellow commissioners from enacting rules. The onset of constant litigation was much more recent and perceived much more potently at Ofcom, where it factored into the agency's decision to refrain from net neutrality intervention. Both countries' telecom industries were litigious, but the regulators' differing levels of attentiveness to this dynamic – resulting from their being at different stages in the regulatory agency life cycle – yielded differences in their willingness to risk litigation. These perceptions also tie back into reputation. Ofcom's caution was inspired by a desire to safeguard its organizational brand, while FCC chairmen could justify taking bold steps that would reflect well on them whether or not their decisions were later upheld in court.

The final informal constraint that clearly impacts traffic management regulatory decisions is the willingness of interest groups to engage in policy debate, particularly Internet companies and public interest advocates. Regulatory intervention does not arise without some constituency to request it and argue for it. That constituency was minimal in the UK, whereas in the US a powerful coalition of advocates and companies engaged in a high-profile public campaign that ultimately inspired regulatory intervention.

That traffic management regulation results from a matrix of institutional constraints provides support for regulatory theories that vest explanations for regulatory activity in a multi-faceted collection of factors, rather than seeking a single cause for regulatory outcomes (Galperin

2004). The "external signals" theory, which suggests that regulatory agencies maximize the positive feedback they receive from each branch of government, the press, and interest groups (Joskow 1974; Noll 1971; Noll 1985), is one example that takes account of the complexity of external influences on regulatory agencies. Noll's (1985) rejection of simple cause-effect models of regulation is supported by the case of traffic management regulation.

Even accounting for multiple external signals does not sufficiently explain regulatory outcomes in this case, however. Carpenter (2001; 2004; 2010) has emphasized the importance of reputation in explaining the behavior of regulators, and this thesis shows how reputation intersects with other factors to shape their decision-making. Supporting the findings of Gilad (2012), it is clear that understanding the unique reputational motivations that attach to particular regulatory roles helps to explain why regulators devote more attention to some external signals than others. The reputational differences between the FCC and Ofcom are vital to understanding their differing approaches to litigation and their willingness (or lack thereof) to engage in policy-setting activities.

In addition to external forces and reputational factors, an agency's own perception of itself, governance structure, and professional culture all contribute to defining its regulatory path. A complex regulatory arena begets a complexity of factors to explain it.

National Regulatory Styles

The analysis presented thus far does not speak to nation-specific factors, one of which was at the center of the second hypothesis set out at the beginning of the thesis:

> Hypothesis 2: National regulatory styles are a key determinant of traffic management outcomes. Consensual regulatory regimes are more likely to produce regulatory outcomes that do little to constrain network operator behavior; adversarial regulatory regimes are more likely to restrain network operators from discriminating for traffic management purposes.

This hypothesis was inspired by prior comparative studies of regulation in other industries that had emphasized the differences between American adversarialism and the more consensual styles of European regulators (Kagan 2003; Kagan and Axelrad 2000; Kelman

1981; Vogel 1986; Wilson 1989). While adversarial and consensual aspects of the US and UK telecommunications regulatory regimes influenced whether traffic management was placed on each regulator's agenda, the findings of this thesis do not provide strong evidence in support of theories concerning nation-specific regulatory styles as articulated in this literature. US telecommunications regulation certainly has adversarial characteristics: interest group participation creates cacophonous public debates and litigation is frequent. But it was not because of adversarialism that the FCC intervened to restrict discriminatory traffic management. The conflicts between interest groups over net neutrality inspired the interest of FCC chairmen, but they intervened as a means to accomplish their own policy goals and establish their own legacies rather than as a show of strength in the struggle with the industry they were tasked with overseeing. Chairman Genachowski in particular operated as a consensus-builder and expended tremendous resources in an attempt to find a regulatory compromise that the broadband industry would not challenge. All the while, litigation was viewed not as a weapon in the regulatory arsenal, but as a fact of life.

Without vocal demands from consumer groups and Internet companies, the regulatory environment in the UK did lack antagonism, but that was neither the sole cause nor a consistent explanation for Ofcom's actions. Some of the agency's initial positioning with respect to traffic management gelled with the arguments that network operators made about its utility and necessity, but Ofcom's overarching logic related more to the agency's own grounding in competition principles than to negotiation with or influence from the broadband industry. Ofcom's reticence to intervene with respect to traffic management resulted in part from the advent of adversarialism in the form of increased litigation, not from cooperation with broadband providers – a finding that runs counter to what would be expected from the national regulatory styles literature. Indeed, characterizing Ofcom's general approach to broadband regulation as consensual would be inconsistent with the agency's most significant and enduring broadband accomplishment: the functional separation of BT. This was nothing if not a direct challenge to the core of the business of the regulated industry. Thus the depiction of the UK regulatory environment as consensual is too sweeping and largely does not explain the permissive atmosphere for discriminatory traffic management that Ofcom helped to create.

There is clearly value in questioning how a particular instance of regulatory activity on a specific topic fits into larger national narratives concerning the relationship between industries and their regulatory overseers. In the case of traffic management, the adversarialism engendered by the participation of particular interest groups (or lack thereof) forms an important backdrop for regulatory activity, even if national regulatory styles more broadly lack explanatory power when compared to the institutional design factors, external forces, and internal characteristics discussed above.

In addition to national regulatory styles, scholarly attention has focused on how the relationship between potent ideas and nation-specific institutional structures produces specific policy outcomes in some nations (Derthick and Quirk 1985; Hall 1989; Hall 1993) and on how these ideas spread from one nation to another, particularly under harmonized European legal frameworks (Gilardi 2002; Gilardi 2005; Thatcher 2002a). This thesis demonstrates the power of ideas concerning competition in telecommunications, particularly that competitive discipline protects consumers from abuses more likely to arise in concentrated markets. This idea is common to most liberalized regulatory regimes and much regulatory rhetoric, but it has had exceptional influence over telecommunications regulation in the UK (as well as in other UK industries (Thatcher 2007)). Promoting competition was one of the key bases of Ofcom's creation; the formative imprint of the competitive ideal continues to shape how the agency conceptualizes regulatory problems and solutions. Ofcom's success in spurring a competitive broadband marketplace has been much revered by its counterparts in current and previous incarnations of the Body of European Regulators for Electronic Communications (BEREC), and that reverence reinforced some of the reputational drivers that deterred the agency from engaging in net neutrality policymaking. Thus the UK case provides evidence that ideas can be influential in state-specific ways and that the cross-pollination of regulatory activities in neighboring states can augment the power of ideas that are valued and promoted in multiple countries.

10.2.4 Broader Applicability of Findings

The findings discussed above reveal much about telecommunications market and regulatory behavior that applies beyond the context of traffic management.

When facing decisions about how to offer their products, Internet service providers are clearly concerned first and foremost with the combination of cost and performance. Broadband markets that are more competitive increase pressure on operators to cut costs while increasing capacity, which may influence a variety of the technical choices they face, such as how and where to interconnect with other networks or whether to install technology upgrades (IPv6 support or Domain Name System security extensions, for example). Operators in more concentrated markets may have more money to spend on these kinds of investments that users may not notice in the short term, but that can bolster the network's performance, accessibility, and security in the long term. That significant competition squeezes operators' margins has consequences not only for traffic management, but for any other choice an operator makes that involves a trade-off between cost and user experience, broadly defined.

Whether ISPs face significant competition or not, their behavior is deeply influenced by the regulatory environment, even in the absence of specific regulatory requirements or prohibitions. The threat of regulatory intervention can be enough to alter their behavior, as can a long history of regulatory confrontations that get reflected in internal corporate oversight structures. These informal regulatory influences can be sufficient to achieve policy goals without the need for regulators to articulate detailed rules. Those companies that have incorporated regulatory oversight into their internal decision-making processes can leverage those oversight structures as new regulatory issues arise.

These dynamics may apply to corporate choices related to data protection, content filtering, universal service, or other policy issues just as they do to traffic management. Examples of the role of regulatory threat abound in Internet policy. As a result of pressure – but not legislation or litigation – from the UK government, nearly every UK ISP has for years been conducting content filtering using the same third-party blacklist of sites said to be hosting illegal content (Clayton 2006). Meanwhile the industry of companies that find novel ways to exploit data about consumer web usage has flourished in the US over the last decade while the likelihood of consumer privacy regulation or legislation being adopted has been slight.

That competition can incentivize discrimination rather than discourage it may have important implications for other non-price-based values that policymakers expect competition to support. Competition in telecommunications policy is often viewed as a way to deliver better performance, improved customer service, or novel product features to consumers. In light of the evidence concerning nondiscrimination, it is worth questioning whether competition has delivered those benefits in practice and examining the relationship between the price declines that competition inspires and its other effects on how telecommunications products are offered.

This thesis also points to key barriers that prevent competition from functioning optimally: switching costs, lack of consumer understanding, and the inability of consumers and carriers to internalize spillovers associated with broadband Internet services. These limitations are not specific to traffic management and may influence the extent to which competition can have disciplining effects on network operators with respect to their product offerings, prices, or terms of service more broadly. Technological advances such as those discussed above related to IPv6 and DNS security provide examples of benefits that are unlikely to be fully internalized through the price system; relying on competition between providers to spur their adoption is unlikely to be sufficient.

Finally, the findings concerning the regulatory agencies are likewise generalizable to telecommunications more broadly, as it is the same agencies that make decisions about traffic management, spectrum allocation, media ownership, and many other issues. The FCC's chairman-centered structure, cavalier approach to litigation risk, and broad policy framework shape its actions on these issues just as Ofcom's more corporate, conservative, and competition-focused approach does. Furthermore, as Lunt and Livingstone (2012) have demonstrated, when communications regulation is conceived narrowly, framed primarily in economic terms, and limited to conclusions that can be drawn from quantitative analysis, its social and democratic implications can go unaddressed regardless of whether it pertains to the Internet, television, radio, or other communications mediums. The converse is just as true.

10.3 Future Directions

Net neutrality has been a contentious, high-profile telecommunications policy issue over the last decade and debates about discrimination on broadband networks show no signs of subsiding. In the US, the pending legal challenge to the FCC's *Open Internet Order* will leave both industry and policymakers to chart their courses through continuing legal uncertainty in the coming years. The UK government is developing plans for legislative reform related to communications, starting from a position that emphasizes the need for network operators to be transparent about traffic management but that is otherwise accepting of application-specific management (DCMS 2013). The European Commission is also crafting new telecommunications legislation that will address net neutrality and traffic management (COM(2013) 627 final 2013/0309 (COD)). National laws and co- or self-regulatory programs that limit discrimination continue to be debated and adopted in other countries, and neutrality issues continue to arise in global Internet governance discussions.

All of this ongoing policy activity means that net neutrality will continue to be ripe for scholarly attention. The key contribution that future academic scholarship can make in this area is to provide empirical evidence and analysis of the effects of different net neutrality policy approaches on network investment, application innovation, economic growth, and digital rights. The academic literature is replete with normative and theoretical arguments about these topics, but little concrete evidence about the effects of the particular choices that policymakers and network operators have made. This thesis elucidates the relationships between regulatory policy, discriminatory traffic management, and competition, but only minimally explores the broader effects of any particular policy or technical choice on innovation, consumers, or whole economies. Nor does it explore other forms of conduct that are just as central to net neutrality debates as fixed-line traffic management: paid prioritization, discrimination on mobile networks, or preferential interconnection agreements, for example. Because net neutrality policy choices can have a significant impact on the Internet's economic, social, and democratic benefits, the decisions that shape it should reflect rigorous analysis of available evidence about the consequences of existing policy choices.

Understanding the implications of institutional differences between different countries' regulatory settings will also require further scholarly attention. The FCC provides an example of a regulatory agency with a broad mandate that stretches beyond promoting competition and incorporates social and industrial policy objectives. The European telecommunications regulatory framework, in contrast, relies heavily on competition between network operators as the means for developing the optimal communications infrastructure to serve the public and the European economy. Ofcom has taken its competition responsibilities related to telecommunications to heart despite being tasked with broader duties. Does a national or regional telecommunications regulatory regime that relies primarily on competition between operators foster communications networks that meet policy goals related to network availability, performance, choice, and technological advancement? Or must regulation be conceptualized within a broader framework to ensure that those goals are met? This thesis provides answers to these questions with respect to the goal of nondiscrimination, but the other objectives of telecommunications regulation are equally important and deserve further study. As communications networks are becoming ever more intertwined with everyday life in advanced economies, it is critical to understand what sort of regulatory regime promotes networks that reflect the interests of their users.

The engineering choices that Internet service providers make are central in determining how the Internet is experienced by its users, which applications succeed or fail, and the overall potential for the Internet to support economic activity and human rights. These choices are driven by individual business needs, available technology, and ISPs' perceptions of the regulatory environment. Understanding why these decisions get made in particular ways and the influence of the regulatory context on the paths chosen is therefore vital to understanding how the future of the Internet might be affected by technological change, business developments, and regulatory activity. Policy stakeholders that move beyond normative debates and leverage the kinds of empirical insights contained in this thesis can more effectively shape that future to the benefit of the Internet and its users.

References

Albanesius, Chloe. 2007. “Comcast Admits Delaying, Not Blocking, P2P Traffic.” *PC Magazine*, October 22. http://www.pcmag.com/article2/0,2817,2204751,00.asp.

———. 2010. “Net Neutrality: The Heavy Hitters React.” *PC Magazine*, December 21. http://www.pcmag.com/article2/0,2817,2374620,00.asp.

Allison, Bill. 2009. “Fighting Net Neutrality, Telecom Companies, Outside Lobbyists, Cluster Contributions to Members of Congress”. Sunlight Foundation. http://reporting.sunlightfoundation.com/2009/fighting-net-neutrality-telecom-companies-outside-lobbyists-clu/.

Allot Communications. 2007. “Digging Deeper Into Deep Packet Inspection.” http://www.allot.com/Common/FilesBinaryWrite.aspx?id=3053.

Ammori, Marvin, Ben Scott, Andy Schwartzman, Harold Feld, and Parul Desai. 2007a. “Petition of Free Press et al. for Declaratory Ruling That Degrading an Internet Application Violates the FCC’s Internet Policy Statement and Does Not Meet an Exception for ‘Reasonable Network Management’.” http://www.freepress.net/files/fp_et_al_nn_declaratory_ruling.pdf.

———. 2007b. “Formal Complaint of Free Press and Public Knowledge Against Comcast Corporation For Secretly Degrading Peer-to-Peer Applications.” http://www.publicknowledge.org/pdf/fp_pk_comcast_complaint.pdf.

Analysys Mason. 2008. “Delivering high-quality video services online.” http://www.ofcom.org.uk/research/technology/research/emer_tech/hqvs/analysyshqvs.pdf.

———. 2010. “Europe’s Digital Deficit: Revitalising the Market in Electronic Communications.” http://www.ectaportal.com/en/upload/Press_Releases/2010/Europes_Digital_Deficit.pdf.

Asghari, Hadi, Michael van Eeten, and Milton Mueller. 2012. “Unravelling the Economic and Political Drivers of Deep Packet Inspection.” In *GigaNet 7th Annual Symposium*. Baku, Azerbaijan. http://measurementlab.net/sites/default/files/DPI_Drivers_HA_MvE_MM_2012.pdf.

AT&T. 2012. “AT&T High Speed Internet Terms of Use.” http://www.att.com/shop/internet/att-internet-terms-of-service.jsp.

Atherton, Martin. 2010. “BT Response to Ofcom Consultation on Traffic Management.” http://stakeholders.ofcom.org.uk/binaries/consultations/net-neutrality/responses/BT.pdf.

Atkinson, Robert D., and Philip J. Weiser. 2006. “A ‘Third Way’ on Network Neutrality”. The Information Technology and Innovation Foundation. http://www.itif.org/files/netneutrality.pdf.

Baldwin, Robert, Martin Cave, and Martin Lodge. 2012. *Understanding Regulation: Theory, Strategy, and Practice*. 2nd ed. Oxford: Oxford University Press.

Baldwin, Robert, and Christopher McCrudden. 1987. *Regulation and Public Law*. London: Weidenfeld and Nicolson.

Bangeman, Eric. 2007. “Comcast Traffic Blocking: Even More Apps, Groupware Clients Affected.” *Ars Technica*, October 22. http://arstechnica.com/uncategorized/2007/10/comcast-traffic-blocking-even-more-apps-groupware-clients-affected/.

Bar, François, Stephen Cohen, Peter Cowhey, Brad DeLong, Michael Kleeman, and John Zysman. 2000. “Access and innovation policy for the third-generation internet.” *Telecommunications Policy* 24 (6–7): 489–518.

Barnard, Chester I. 1938. *The Functions of the Executive*. Cambridge, MA: Harvard University Press.

Bartlett, David C., Jeffrey S. Lanning, John E. Benedict, and Brian Staihr. 2008. “Comments of Embarq In the Matter of Broadband Industry Practices; Petition of Free Press et al.

for Declaratory Ruling That Degrading an Internet Application Violates the FCC's Internet Policy Statement and Does Not Meet an Exception for 'Reasonable Network Management;' Vuze, Inc. Petition to Establish Rules Governing Network Management Practices by Broadband Network Operators, WC Docket No. 07-52." http://apps.fcc.gov/ecfs/document/view?id=6519841222.
Bauer, Steve, Dave Clark, and William Lehr. 2009. "The Evolution of Internet Congestion." In *37th Research Conference on Communication, Information and Internet Policy*. Arlington, VA. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1999830.
Baumgartner, Jeff. 2008. "Comcast Denied Set-Top Waiver (Again)." *Light Reading*, May 18. http://www.lightreading.com/set-top-boxes/comcast-denied-settop-waiver-again/240103032.
Baumol, William J., Martin E. Cave, Robert E. Litan, Peter Cramton, Robert W. Hahn, Thomas W. Hazlett, Paul L. Joskow, et al. 2007. "Economists' Statement on Network Neutrality Policy". Related Publication 07-08. AEI-Brookings Joint Center for Regulatory Studies. http://www.brookings.edu/views/papers/litan/200703jointcenter.pdf.
BBC. 2010. "BBC iPlayer Breaks New Record of 100 Million Requests." http://www.bbc.co.uk/pressoffice/pressreleases/stories/2010/01_january/14/iplayer.shtml.
Be. 2005. "'Be' Starts the Broadband Revolution: 24 Meg, for £24, 24 Hours Per Day." https://www.o2broadband.o2.co.uk/newsReleaseDetailed5.do.
Becker, Gary S. 1983. "A Theory of Competition Among Pressure Groups for Political Influence." *The Quarterly Journal of Economics* 98 (3) (August): 371–400.
Becker, Gary S., Dennis W. Carlton, and Hal S. Sider. 2010. "Net Neutrality and Consumer Welfare." *Journal of Competition Law and Economics* 6 (3): 497–519.
Bell, Gregory. 2003. "Failure to Thrive: QoS and the Culture of Operational Networking." In *Proceedings of the ACM SIGCOMM Workshop on Revisiting IP QoS: What Have We Learned, Why Do We Care*, 115–120. http://dl.acm.org/citation.cfm?id=944595.
Berkett, Neil. 2009. "Virgin Media Inc. at Goldman Sachs Communacopia Conference New York" presented at the Goldman Sachs Communacopia Conference New York, September 17, FD (Fair Disclosure) Wire.
———. 2010. "Virgin Media Inc. Analyst and Investors Day (New York)" presented at the Virgin Media Inc. Analyst and Investors Day (New York), December 3, FD (Fair Disclosure) Wire.
Bernstein, Marver H. 1955. *Regulating Business by Independent Commission*. Princeton, NJ: Princeton University Press.
Beverly, Robert, Steven Bauer, and Arthur Berger. 2007. "The Internet Is Not a Big Truck: Toward Quantifying Network Neutrality." In *Proceedings of the 8th International Conference on Passive and Active Network Measurement*, 135–144. PAM'07. Louvain-la-Neuve, Belgium. http://portal.acm.org/citation.cfm?id=1762888.1762907.
BIS. 2010. "Implementing the Revised EU Electronic Communications Framework: Overall Approach and Consultation on Specific Issues." https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/31567/10-1132-implementing-revised-electronic-communications-framework-consultation.pdf.
———. 2013. "Streamlining Regulatory and Competition Appeals: Consultation on Options for Reform." https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/207701/bis-13-876-regulatory-and-competition-appeals.pdf.
BIS and DCMS. 2009. "Digital Britain Final Report." http://www.official-documents.gov.uk/document/cm76/7650/7650.pdf.
Bishop, William. 1990. "A Theory of Administrative Law." *Journal of Legal Studies* 19 (1): 489.
"BitTorrent Manipulation in Selected Countries." 2012. Syracuse University School of Information Studies. http://dpi.ischool.syr.edu/countries.html.

Black, Julia. 1997. "New Institutionalism and Naturalism in Socio-Legal Analysis: Institutionalist Approaches to Regulatory Decision Making." *Law & Policy* 19 (1): 51–93.
Black, Julia, and Paul Walker, ed. 1998. *Commercial Regulation and Judicial Review*. Oxford: Hart Publishing.
Bourreau, Marc, Frago Kourandi, and Tommaso Valletti. 2013. "Net Neutrality with Competing Internet Platforms". Telecom ParisTech Department of Economics and Social Sciences Working Paper. http://www.iae.csic.es/investigatorsMaterial/a1314509555492591.pdf.
Bradshaw, Tim. 2009. "BT Seeks to End 'Free Ride' by Video Websites." *Financial Times*, June 11. http://www.ft.com/intl/cms/s/0/1c979154-5621-11de-ab7e-00144feabdc0.html#axzz2SRAL6sTQ.
Brennan, Timothy J. 2012. "Network Neutrality or Minimum Quality? Barking Up the Wrong Tree--and Finding the Right One." *CPI Antitrust Chronicle* Spring 2012, 3 (2) (March).
Breyer, Stephen. 1986. "Judicial Review of Questions of Law and Policy." *Administrative Law Review* 38 (4): 363–398.
Brianl. 2011. "Consolidated Latency Tracking and Reporting." February 4. http://us.battle.net/wow/en/forum/topic/1965992365?page=7#122.
Brill, Ed. 2009. "Hart v. Comcast, Proposed Class Action Settlement for Lotus Notes Throttling." http://www.edbrill.com/ebrill/edbrill.nsf/dx/hart-v.-comcast-proposed-class-action-settlement-for-lotus-notes-throttling.
Brill, Matthew A., and Jarrett S. Taubman. 2008. "Comments of Time Warner Cable Inc. In the Matter of Petition of Free Press et al. for Declaratory Ruling for Degrading an Internet Application Violates the FCC's Internet Policy Statement and Does Not Meet an Exception for Reasonable Network Management; Vuze, Inc. Petition to Establish Rules Governing Network Management Practices by Broadband Network Operators, WC Docket No. 07-52." http://apps.fcc.gov/ecfs/document/view?id=6519841176.
Brito, Jerry, Martin E. Cave, Robert W. Crandall, Larry F. Darby, Everett Ehrlich, Jeffrey A. Eisenach, Jerry Ellig, et al. 2010. "Net Neutrality Regulation: The Economic Evidence." http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1587058.
Broadband Stakeholder Group. 2011. "Broadband Providers Launch New Traffic Management Transparency Code: New Initiative Will Help Consumers and Policy Makers Make Informed Choices." March 15. http://www.broadbanduk.org/wp-content/uploads/2012/08/broadband_providers_launch_new_traffic_management_transparency_code_14_march_201111.pdf.
Brown, Craig J., and Timothy M. Boucher. 2007. "Comments of Qwest Communications International Inc. In the Matter of Broadband Industry Practices, WC Docket No. 07-52." http://apps.fcc.gov/ecfs/document/view?id=6519529463.
Brown, Maggie. 2010. "Media: Has Ofcom Changed Since the Tories Took Over?" *The Guardian*, October 4, sec. Media Pages.
Bryman, Alan. 2008. *Social Research Methods*. 3rd ed. Oxford: Oxford University Press.
BT. 2006. "BT Price List." http://stakeholders.ofcom.org.uk/binaries/telecoms/policy/bb/model.pdf.
Burgess, Robert G. 1995. *In the Field: An Introduction to Field Research*. London, UK: Routledge.
Buttery, Paul. 2010. "Virgin Media Investor Day" presented at the Virgin Media Investor Day, December. http://media.corporate-ir.net/media_files/IROL/13/135485/Investor_Day_slides_FINAL.pdf.
Cameron, David. 2009. "People Power - Reforming Quangos." http://www.conservatives.com/News/Speeches/2009/07/David_Cameron_People_Power_-_Reforming_Quangos.aspx.

Candeub, Adam. 2012. “An End to End-to-End? A Review Essay of Barbara van Schewick’s Internet Architecture and Innovation.” *Federal Communications Law Journal* 64 (3): 661–674.
Candeub, Adam, and Keith Brown. 2008. “What Do Bureaucrats Want? Estimating Regulator Preferences at the FCC.” In *Annual Meeting of the American Law and Economics Association*. New York. http://law.bepress.com/cgi/viewcontent.cgi?article=2493&context=alea.
Candeub, Adam, and Daniel John McCartney. 2010. “Network Transparency: Seeing the Neutral Network.” *Northwestern Journal of Technology and Intellectual Property* 8 (2): 228–246.
Cannon, Robert. 2001. “Where Internet Service Providers and Telephone Companies Compete: A Guide to the Computer Inquiries, Enhanced Service Providers and Information Service Providers.” *CommLaw Conspectus: Journal of Communications Law and Policy* 9: 49.
Cañon, Carlos. 2009. “Regulation Effects on Investment Decisions in Two-Sided Market Industries: The Net Neutrality Debate.” http://papers.ssrn.com/abstract=1374782.
Carney, Sean. 2006. “Bandwidth.com Names Julius Genachowski to Advisory Board; Investment Veteran Joins Leading Communications Provider to Grow Nationwide Business.” *PR Newswire*, September 1. http://www.prnewswire.com/news-releases/bandwidthcom-names-julius-genachowski-to-advisory-board-55807747.html.
Carpenter, Daniel P. 2001. *The Forging of Bureaucratic Autonomy: Reputations, Networks, and Policy Innovation in Executive Agencies, 1862-1928*. Princeton, NJ: Princeton University Press.
———. 2004. “Protection Without Capture: Product Approval by a Politically Responsive, Learning Regulator.” *American Political Science Review* 98 (4): 613–631.
———. 2010. *Reputation and Power: Organizational Image and Pharmaceutical Regulation at the FDA*. Princeton, NJ: Princeton University Press.
Carrigan, Christopher, and Cary Coglianese. 2011. “The Politics of Regulation: From New Institutionalism to New Governance.” *Annual Review of Political Science* 14: 107–129.
Casserly, James L., Ryan G. Wallach, Daniel K. Alvarez, Helgi C. Walker, Joseph W. Waz, Kathryn A. Zachem, Mary McManus, Thomas R. Nathan, and Gerard J. Lewis. 2008. “Reply Comments of Comcast Corporation In the Matter of Broadband Industry Practices, WC Docket No. 07-52.” http://apps.fcc.gov/ecfs/document/view?id=6519856208.
Casserly, James L., Ryan G. Wallach, Daniel K. Alvarez, Joseph W. Waz, Kathryn A. Zachem, Mary McManus, Thomas R. Nathan, and Gerard J. Lewis. 2008. “Comments of Comcast Corporation In the Matter of Broadband Industry Practices, WC Docket No. 07-52.” http://apps.fcc.gov/ecfs/document/view?id=6519840991.
Cave, Martin, Richard Collins, Nico van Eijk, Pierre Larouche, Luigi Prosperetti, Alexandre de Streel, and Tommaso Valletti. 2009. “Statement by European Academics on the Inappropriateness of Imposing Increased Internet Regulation in the EU.” http://cerre.eu/publications/statement-european-academics-inappropriateness-imposing-increased-internet-regulation-e.
Cave, Martin, and Pietro Crocioni. 2007. “Does Europe Need Network Neutrality Rules?” *International Journal of Communication* 1: 669–679.
Cave, Martin, Luigi Prosperetti, and Chris Doyle. 2006. “Where Are We Going? Technologies, Markets and Long-Range Public Policy Issues in European Communications.” *Information Economics and Policy* 18 (3): 242–255.
Cellan-Jones, Rory. 2009a. “BT Accused of iPlayer Throttling.” *BBC*, June 1, sec. Technology. http://news.bbc.co.uk/2/hi/technology/8077839.stm.
———. 2009b. “iPlayer: BBC v BT.” *BBC News Dot.life*. http://www.bbc.co.uk/blogs/technology/2009/06/iplayerbbc_v_bt.html.

———. 2009c. “BT and the iPlayer.” *BBC News Dot.life*. http://www.bbc.co.uk/blogs/technology/2009/06/bt_and_the_iplayer.html.
Chen, M. Keith, and Barry J. Nalebuff. 2006. “One-Way Essential Complements”. Cowles Foundation Discussion Paper 1588. Cowles Foundation for Research in Economics, Yale University. http://ideas.repec.org/p/cwl/cwldpp/1588.html.
Cheng, Hsing Kenneth, Subhajyoti Bandyopadhyay, and Hong Guo. 2011. “The Debate on Net Neutrality: A Policy Perspective.” *Information Systems Research* 22 (1) (March): 60–82.
Chirico, Filomena, Ilse M. Van Der Haar, and Pierre Larouche. 2007. “Network Neutrality in the EU.” *TILEC Discussion Paper No. 2007-030*.
Cho, Kenjiro, Kensuke Fukuda, Hiroshi Esaki, and Akira Kato. 2006. “The Impact and Implications of the Growth in Residential User-to-User Traffic.” In *Proceedings of the 2006 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, 207–218. Pisa, Italy: ACM. http://dl.acm.org/citation.cfm?id=1159938.
Choi, Jay Pil, and Byung-Cheol Kim. 2010. “Net Neutrality and Investment Incentives.” *RAND Journal of Economics* 41 (3): 446–471.
Cicconi, Jim. 2010. “A Few Thoughts on Today’s FCC Vote.” http://attpublicpolicy.com/government-policy/a-few-thoughts-on-todays-fcc-vote/.
Cisco. 2008. “Cisco Visual Networking Index 2007-2012.” June 16. http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481374_ns827_Networking_Solutions_White_Paper.html.
Clark, Nick. 2009. “Rivals Fail to Back BT Over Challenge to Online Video.” *The Independent*, June 12. http://www.independent.co.uk/news/business/news/rivals-fail-to-back-bt-over-challenge-to-online-video-1703244.html.
Clarke, Richard N. 2009. “Costs of Neutral/Unmanaged IP Networks.” *Review of Network Economics* 8 (1): 61–89.
Clayton, Richard. 2006. “Failures in a Hybrid Content Blocking System.” In *Privacy Enhancing Technologies*, edited by George Danezis and David Martin, 78–92. Lecture Notes in Computer Science 3856. Springer Berlin Heidelberg.
Clyburn, Mignon L. 2010. “Approving in Part, Concurring in Part Statement of Commissioner Mignon L. Clyburn Re: Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices, WC Docket No. 07-52.” http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-10-201A5.pdf.
Coalition of Broadband Users and Innovators. 2003. “Ex Parte Submission Re: CS Docket No. 02-52; CC Docket Nos. 02-33, 98-10 & 95-20; GN Docket No. 00-185”, July 17. http://apps.fcc.gov/ecfs/document/view?id=6514286199.
Coglianese, Cary. 2002. “Empirical Analysis and Administrative Law.” *University of Illinois Law Review* 2002 (4): 1111.
Cohen, David L. 2008. “Statement of David L. Cohen, Executive Vice President, Comcast Corporation; Federal Communications Commission Public En Banc Hearing on ‘Broadband Network Management Practices,’ WC Docket No. 07-52.” http://apps.fcc.gov/ecfs/document/view?id=6519844216.
Cooper, Alissa, Arnaud Jacquet, and Andrea Soppera. 2011. “Bandwidth Usage and Management: A UK Case Study.” In *39th Research Conference on Communications, Information and Internet Policy (TPRC), 2011*. Arlington, VA, USA. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1979754.
Cooper, Mark. 2003. “Open Communications Platforms: The Physical Infrastructure as the Bedrock of Innovation and Democratic Discourse in the Internet Age.” *Journal on Telecommunications & High Technology Law* 2 (1): 177.
Copps, Michael J. 2005. “Statement of Commissioner Michael J. Copps, Concurring Re: In the Matter of Appropriate Framework for Broadband Access to the Internet over Wireline Facilities (CC Docket No. 02-33) et al. Report and Order and Notice of Proposed Rulemaking.” http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-260433A4.pdf.

———. 2008. “Statement of Commissioner Michael J. Copps Re: Formal Complaint of Free Press and Public Knowledge Against Comcast Corporation for Secretly Degrading Peer-to-Peer Applications, File No. EB-08-IH-1518; Broadband Industry Practices, Petition of Free Press et al. for Declaratory Ruling That Degrading an Internet Application Violates the FCC’s Internet Policy Statement and Does Not Meet an Exception for ‘Reasonable Network Management,’ WC Docket No. 07-52.” http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-08-183A3.pdf.

Cowhey, Peter F., Jonathan D. Aronson, and John Richards. 2009. “Shaping the Architecture of the U.S. Information and Communication Technology Architecture: A Political Economic Analysis.” *Review of Policy Research* 26 (1-2): 105–125.

Crabtree, James. 2004. “Kip Meek; Can the Man from Ofcom Get More Competition Without Confusing the Consumer?” *New Statesman*, May 31.

Crawford, Susan P. 2007. “The Internet and the Project of Communications Law.” *UCLA Law Review* 55: 359.

———. 2008. “Transporting Communications.” *Boston University Law Review* 89: 871–937.

———. 2011. “The Communications Crisis in America.” *Harvard Law & Policy Review* 5-2.

———. 2013. *Captive Audience: The Telecom Industry and Monopoly Power in the New Gilded Age*. New Haven: Yale University Press.

Crocioni, Pietro. 2011. “Net Neutrality in Europe: Desperately Seeking a Market Failure.” *Telecommunications Policy* 35 (1) (February): 1–11.

Croley, Steven P. 2008. *Regulation and Public Interests: The Possibility of Good Regulatory Government*. Princeton, NJ: Princeton University Press.

Crowcroft, Jon. 2007. “Net Neutrality: The Technical Side of the Debate: A White Paper.” *SIGCOMM Comput. Commun. Rev.* 37 (1): 49–56.

Darroch, Jeremy. 2010. “Full Year 2009/2010 British Sky Broadcasting Group Plc Earnings Conference Call (US)” July 29, FD (Fair Disclosure) Wire.

Daver, Shernaz. 2005. “Truveo Adds IAC/InterActiveCorp Veteran to Board of Advisors.” *BusinessWire*, October 11. http://www.businesswire.com/news/home/20051011005416/en/Truveo-Adds-IACInterActiveCorp-Veteran-Board-Advisors.

Davidson, Paul. 2008. “FCC Chief Martin Hasn’t Lost Focus on Cable.” *USA Today*, March 11. http://usatoday30.usatoday.com/money/media/2008-03-10-fcc-martin_n.htm.

DCMS. 2011a. “Open Internet Roundtable Statement.” http://www.dcms.gov.uk/news/news_stories/7958.aspx.

———. 2011b. “Reforming the Appeals Regime for the Electronic Communications Sector Impact Assessment (IA).” https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/72915/Electronic_Communications_Sector_June2011_-_Appeals_IA.pdf.

———. 2011c. “Implementing the Revised EU Electronic Communications Framework - Appeals.” https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/72909/Electroninc_Communications_Sector_August2011_Appeals_condoc2.pdf.

———. 2013. “Connectivity, Content and Consumers: Britain’s Digital Platform for Growth.” https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/225783/Connectivity_Content_and_Consumers_2013.pdf.

Derthick, Martha, and Paul J. Quirk. 1985. *The Politics of Deregulation*. Washington: Brookings Institution.

DiIulio, John T. 1994. “The Cultural Bases of Behavior in Federal Government Bureaucracy.” *Journal of Public Administration Research & Theory* 4 (3): 277–318.

Dischinger, Marcel, Massimiliano Marcon, Saikat Guha, Krishna P. Gummadi, Ratul Mahajan, and Stefan Saroiu. 2010. “Glasnost: Enabling End Users to Detect Traffic Differentiation.” In *Proceedings of the 7th USENIX Conference on Networked Systems Design and Implementation*, 27–27. NSDI’10. Berkeley, CA, USA: USENIX Association. http://dl.acm.org/citation.cfm?id=1855711.1855738.

Dischinger, Marcel, Alan Mislove, Andreas Haeberlen, and Krishna P. Gummadi. 2008. "Detecting BitTorrent Blocking." In *Proceedings of the 8th ACM SIGCOMM Conference on Internet Measurement*, 3–8. Vouliagmeni, Greece: ACM. http://portal.acm.org/citation.cfm?id=1452520.1452523&coll=Portal&dl=GUIDE&CFID=75047289&CFTOKEN=91399611.

Dobbin, Frank. 1994. *Forging Industrial Policy: The United States, Britain, and France in the Railway Age*. Cambridge: Cambridge University Press.

Dobbin, Frank, and John R. Sutton. 1998. "The Strength of a Weak State: The Rights Revolution and the Rise of Human Resources Management Divisions." *American Journal of Sociology* 104 (2): 441–476.

Dutton, William H., and Grant Blank. 2011. "Next Generation Users: The Internet in Britain 2011". Oxford Internet Institute, University of Oxford. http://www.oii.ox.ac.uk/publications/oxis2011_report.pdf.

Dutton, William H., Corinna di Gennaro, and Andrea Millwood Hargrave. 2005. "The Internet in Britain: The Oxford Internet Survey (OxIS)". Oxford Internet Institute, University of Oxford. http://www.oii.ox.ac.uk/research/oxis/oxis2005_report.pdf.

Dyson, Kenneth. 1983. "The Cultural, Ideological and Structural Context." In *Industrial Crisis: A Comparative Study of the State and Industry*, edited by Kenneth Dyson and Stephen Wilks. Oxford: Basil Blackwell.

Eckersley, Peter, Fred von Lohmann, and Seth Schoen. 2007. "Packet Forgery By ISPs: A Report On The Comcast Affair." https://www.eff.org/files/eff_comcast_report.pdf.

Economides, Nicholas. 2008. "Net Neutrality, Non-Discrimination and Digital Distribution of Content Through the Internet." *I/S: A Journal of Law and Policy for the Information Society* 4: 209.

Economides, Nicholas, and Benjamin Hermalin. 2012. "The Economics of Network Neutrality." *Rand Journal of Economics* 43 (4): 602–629.

Economides, Nicholas, and Joacim Tåg. 2012. "Network Neutrality on the Internet: A Two-Sided Market Analysis." *Information Economics and Policy* 24 (2): 91–104.

ECTA. 2010. "Public Consultation on the Open Internet and Net Neutrality in Europe." http://ec.europa.eu/information_society/policy/ecomm/doc/library/public_consult/net_neutrality/comments/05internet_tech_industry_orgs/ecta.pdf.

Eggerton, John. 2009. "EXCLUSIVE: Kevin Martin – The Exit Interview." *Broadcasting & Cable*, January 4. http://www.broadcastingcable.com/article/161393-EXCLUSIVE_Kevin_Martin_The_Exit_Interview.php.

Erickson, Markham C. 2010. "Comments of the Open Internet Coalition In the Matter of Preserving the Open Internet, GN Docket 09-191; Broadband Industry Practices, WC Docket 07-52; Further Inquiry into Two Under-Developed Issues in the Open Internet Proceeding, DA 10-1667."

Erman, Jeffrey, Alexandre Gerber, Mohammad T. Hajiaghayi, Dan Pei, and Oliver Spatscheck. 2009. "Network-Aware Forward Caching." In *Proceedings of the 18th International Conference on the World Wide Web*, 291–300. WWW '09. Madrid, Spain: ACM. http://dl.acm.org/citation.cfm?id=1526749.

European Commission. 2010. "Questionnaire for the Public Consultation on the Open Internet and Net Neutrality in Europe." http://ec.europa.eu/information_society/policy/ecomm/doc/library/public_consult/net_neutrality/comments/nn_questionnaire.pdf.

European Competitive Telecommunications Association. 2009. "Regulatory Scorecard 2009." http://www.ectaportal.com/en/REPORTS/Regulatory-Scorecards/Regulatory-Scorecard-2009/.

Everything Everywhere. 2010. "Traffic Management and 'Net Neutrality:' Response of Everything Everywhere to the European Commission Consultation." http://ec.europa.eu/information_society/policy/ecomm/doc/library/public_consult/net_neutrality/comments/01operators_isps/everything_everywhere.pdf.

Faratin, Peyman, Clark, David, Steven Bauer, William Lehr, Patrick Gilmore, and Arthur Berger. 2008. “The Growing Complexity of Internet Interconnection.” *Communications & Strategies* 72 (4th quarter): 51–71.
Farrell, Joseph, and Philip J Weiser. 2003. “Modularity, Vertical Integration, and Open Access Policies: Towards a Convergence of Antitrust and Regulation in the Internt Age.” *Harvard Journal of Law & Technology* 17 (1): 85.
Faulhaber, Gerald R. 2011. “Economics of Network Neutrality: A Review.” *Communications & Convergence Review* 3 (1).
Faulhaber, Gerald R., and David J. Farber. 2010. “The Open Internet: A Customer-Centric Framework.” *International Journal of Communication* 4: 302–342.
FCC. 1968. “In the Matter of Use of the Carterfone Device in Message Toll Telephone Service; In the Matter of Thomas F. Carter and Carter Electronics Corp., Dallas, Tex. (Complainants), v. American Telephone and Telegraph Co., Associated Bell System Companies, Southwestern Bell Telephone Co., and General Telephone Co. of the Southwest (Defendants), Docket No. 16942, Docket No. 17073.”
———. 1971. “Computer I Final Decision and Order, 28 FCC 2d 267.”
———. 1980. “Computer II Final Decision, 77 FCC 2d 384.”
———. 1999. “Memorandum Opinion and Order in the Matter of Applications for Consent to the Transfer of Control of Licenses and Section 214 Authorizations from Tele-Communications, Inc., Transferor, to AT&T Corp., Transferee, CS Docket No. 98-178.”
———. 2000a. “Memorandum Opinion and Order in the Matter of Applications for Consent to the Transfer of Control of Licenses and Section 214 Authorizations from MediaOne Group, Inc., Transferor, To AT&T Corp., Transferee, CS Docket No. 99-251.”
———. 2000b. “Notice of Inquiry In the Matter of Inquiry Concerning High-Speed Access to the Internet Over Cable and Other Facilities, GN Docket No. 00-185.”
———. 2002. “Declaratory Ruling and Notice of Proposed Rulemaking In the Matter of Inquiry Concerning High-Speed Access to the Internet Over Cable and Other Facilities, GN Docket No. 00-185; Internet Over Cable Declaratory Ruling; Appropriate Regulatory Treatment for Broadband Access to Internet Over Cable Facilities, CS Docket No. 02-52.”
———. 2004. “Federal Communications Commission Releases Data on High-Speed Internet Access Services.” http://transition.fcc.gov/Bureaus/Common_Carrier/Reports/FCC-State_Link/IAD/hspd1204.pdf.
———. 2005a. “Order in the Matter of Madison River Communications, LLC and Affiliated Companies, File No. EB-05-IH-0110, Acct. No. 200532080126, FRN: 0004334082.”
———. 2005b. “Report and Order and Notice of Proposed Rulemaking In the Matters of Appropriate Framework for Broadband Access to the Internet Over Wireline Facilities, CC Docket 02-33, et al.”
———. 2005c. “Policy Statement In the Matters of Appropriate Framework for Broadband Access to the Internet Over Wireline Facilities, CC Docket 02-33, et al.”
———. 2005d. “Memorandum Opinion and Order in the Matter of Verizon Communications Inc. and MCI, Inc. Applications for Approval of Transfer of Control, WC Docket No. 05-75.”
———. 2005e. “Memorandum Opinion and Order in the Matter of SBC Communications Inc. and AT&T Corp. Applications for Approval of Transfer of Control, WC Docket No. 05-65.”
———. 2006. “Memorandum Opinion and Order In the Matter of AT&T Inc. and BellSouth Corporation Application for Transfer of Control, WC Docket 06-74.”
———. 2007a. “Notice of Inquiry In the Matter of Broadband Industry Practices, WC Docket No. 07-52.”
———. 2007b. “Second Report and Order in the Matter of Service Rules for the 698-746, 747-762 and 777-792 MHz Bands, WT Docket No. 06-150.”

———. 2008. “Memorandum Opinion and Order in the Matters of Formal Complaint of Free Press and Public Knowledge Against Comcast Corporation for Secretly Degrading Peer-to-Peer Applications, File No. EB-08-IH-1518; Broadband Industry Practices Petition of Free Press et al. for Declaratory Ruling That Degrading an Internet Application Violates the FCC’s Internet Policy Statement and Does Not Meet an Exception for ‘Reasonable Network Management,’ WC Docket No. 07-52.”

———. 2009. “Notice of Proposed Rulemaking In the Matter of Preserving the Open Internet, GN Docket 09-191; Broadband Industry Practices, WC Docket No. 07-52.”

———. 2010a. “Notice of Inquiry In the Matter of Framework for Broadband Internet Service, GN Docket No. 10-127.”

———. 2010b. “Report and Order In the Matter of Preserving the Open Internet, GN Docket 09-1913; Broadband Industry Practices, WC Docket 07-52.”

———. 2012. “Internet Access Services: Status as of June 30, 2011.” http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-314630A1.pdf.

———. 2013. “Commissioners from 1934 to Present.” http://www.fcc.gov/leadership/commissioners-1934-present.

Felten, Edward W. 2006. “Nuts and Bolts of Network Neutrality.” *AEI-Brookings Joint Center for Regulatory Studies* (Related Publication 06-23). http://www.reg-markets.org/admin/authorpdfs/redirect-safely.php?fname=../pdffiles/php9e.pdf.

Ferguson, Andrew. 2006. “Update from Nildram on Traffic Management.” http://www.thinkbroadband.com/news/2847-update-from-nildram-on-traffic-management.html.

———. 2009a. “BT Total Excess Usage Charge to Rise to £1 Per GB.” http://www.thinkbroadband.com/news/3878-bt-total-excess-usage-charge-to-rise-to-1-per-gb.html.

———. 2009b. “O2 Off-Net Customers Now Have Traffic Management.” http://www.thinkbroadband.com/news/4064-o2-off-net-customers-now-have-traffic-mangement.html.

———. 2010. “O2 Publishes Traffic Management Information for Products.” http://www.thinkbroadband.com/news/4373-o2-publishes-traffic-management-information-for-products.html.

Fiorina, Morris P. 1982. “Legislative Choice of Regulatory Forms: Legal Process or Administrative Process?” *Public Choice* 39 (1): 33–66.

Fluendy, Simon. 2007. “Pay Up If You Jam the Web, BBC Warned.” *Daily Mail*, August 12. http://www.dailymail.co.uk/money/news/article-1612982/Pay-up-if-you-jam-the-web-BBC-warned.html.

Frieden, Robert M. 2006. “Network Neutrality or Bias?–Handicapping the Odds for a Tiered and Branded Internet.” In *Bepress Legal Series*. Paper 1755.

Frischmann, Brett M. 2005. “An Economic Theory of Commons Management.” *Minnesota Law Review* 89 (4): 917–1030.

Frischmann, Brett M., and Mark A. Lemley. 2006. “Spillovers.” *American Law & Economics Annual Meetings* (Paper 27).

Frischmann, Brett M., and Barbara van Schewick. 2007. “Network Neutrality and the Economics of an Information Superhighway: A Reply to Professor Yoo.” *Jurimetrics* 47: 383–428.

FTC. 2000. “Agreement Containing Consent Orders In the Matter of America Online, Inc., a Corporation, and Time Warner, Inc., a Corporation, File No. 001 0105.”

———. 2007. “Broadband Connectivity Competition Policy FTC Staff Report.” http://www.ftc.gov/reports/broadband/v070000report.pdf.

Galperin, Hernan. 2004. “Beyond Interests, Ideas, and Technology: An Institutional Approach to Communication and Information Policy.” *The Information Society* 20 (3): 159–168.

Ganley, Paul, and Ben Allgrove. 2006. “Net Neutrality: A User’s Guide.” *Computer Law & Security Review* 22 (6): 454–463. http://www.sciencedirect.com/science/article/pii/S0267364906000902.

Geller, Henry. 1974. “A Modest Proposal for Modest Reform of the Federal Communications Commission.” *Georgetown Law Journal* 63: 705.
Genachowski, Julius. 2010a. “The Third Way: A Narrowly Tailored Broadband Framework.” http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-297944A1.pdf.
———. 2010b. “Statement of Chairman Julius Genachowski Re: Framework for Broadband Internet Service, GN Docket No. 10-127.” http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-10-114A2.pdf.
———. 2010c. “Statement of Chairman Julius Genachowski Re: Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices, WC Docket No. 07-52.” http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-10-201A2.pdf.
Gibbons, Thomas. 2005. “Competition Policy and Regulatory Style -- Issues for OFCOM.” *Info* 7 (5): 42–51.
Gilad, Sharon. 2012. “Attention and Reputation: Linking Regulators’ Internal and External Worlds.” In *Executive Politics in Times of Crisis*, edited by Martin Lodge and Kai Wegrich. New York: Palgrave Macmillan.
Gilardi, Fabrizio. 2002. “Policy Credibility and Delegation to Independent Regulatory Agencies: A Comparative Empirical Analysis.” *Journal of European Public Policy* 9 (6) (January): 873–893.
———. 2005. “The Institutional Foundations of Regulatory Capitalism: The Diffusion of Independent Regulatory Agencies in Western Europe.” *Annals of the American Academy of Political and Social Science* 598: 84–101.
Glover, Michael, Edward Shakin, William H. Johnson, and John T. Scott. 2008. “Comments of Verizon and Verizon Wireless In the Matter of Petition for Rulemaking to Establish Rules Governing Network Management Practices by Broadband Network Operators, Petition for Declaratory Ruling Regarding Internet Management Policies, Broadband Industry Practices, WC Docket No. 07-52.” http://apps.fcc.gov/ecfs/document/view?id=6519841190.
Glover, Michael, Edward Shakin, William H. Johnson, John T. Scott, William D. Wallace, Walter E. Dellinger, Irv Gornstein, and Henry Weissman. 2010. “Comments of Verizon and Verizon Wireless In the Matter of Preserving the Open Internet, GN Docket 09-191; Broadband Industry Practices, WC Docket No. 07-52.” http://apps.fcc.gov/ecfs/document/view?id=7020378541.
Gold, Raymond L. 1957. “Roles in Sociological Field Observations.” *Social Forces* 36: 217.
Goldberg, Henry, Devendra T. Kumar, and Jay Monahan. 2007. “Petition for Rulemaking In the Matter of of Vuze, Inc. Petition to Establish Rules Governing Network Management Practices by Broadband Network Operators; Broadband Industry Practices, WC Docket No. 07-52.” http://fjallfoss.fcc.gov/ecfs/document/view?id=6519811711.
Golden, Marissa Martino. 2000. *What Motivates Bureaucrats?: Politics and Administration During the Reagan Years*. New York: Columbia University Press.
Government Accountability Office. 2006. “Broadband Deployment Is Extensive Throughout the United States, but It Is Difficult to Assess the Extent of Deployment Gaps in Rural Areas (GAO-06-426).” http://www.gao.gov/new.items/d06426.pdf.
Greenstein, Shane. 2007. “Economic Experiments and Neutrality in Internet Access.” *Innovation Policy and the Economy* 8: 59–109.
Greenstein, Shane, and Ryan McDevitt. 2011. “Evidence of a Modest Price Decline in US Broadband Services.” *Information Economics and Policy* 23 (2): 200–211.
Guo, Hong, Hsing Kenneth Cheng, and Subhajyoti Bandyopadhyay. 2012. “Net Neutrality, Broadband Market Coverage, and Innovation at the Edge.” *Decision Sciences* 43 (1): 141–172.
———. 2013. “Broadband Network Management and the Net Neutrality Debate.” *Production and Operations Management*. http://onlinelibrary.wiley.com/doi/10.1111/poms.12020/full.

Gustin, Sam. 2013. “FCC Chairman Julius Genachowski Stepping Down After Contentious Term.” *Time*, March 22. http://business.time.com/2013/03/22/fcc-chairman-julius-genachowski-stepping-down-after-contentious-term-reports/.

Hahn, Robert W., Robert E. Litan, and Hal J. Singer. 2007. “The Economics of Wireless ‘Net Neutrality’.” *Journal of Competition Law and Economics* 3 (3): 399–451.

Hall, Peter A. 1989. *The Political Power of Economic Ideas: Keynesianism Across Nations*. Princeton, NJ: Princeton University Press.

———. 1993. “Policy Paradigms, Social Learning, and the State: The Case of Economic Policymaking in Britain.” *Comparative Politics* 25 (3): 275–296.

Halperin, Morton H. 1974. *Bureaucratic Politics and Foreign Policy*. Washington: Brookings Institution.

Hancher, Leigh, and Michael Moran. 1989. “Organizing Regulatory Space.” In *Capitalism, Culture, and Economic Regulation*, edited by Leigh Hancher and Michael Moran. Oxford: Oxford University Press.

Hancock, Donald. 1983. “Comparative Public Policy: An Assessment.” In *Political Science: The State of the Discipline*, edited by Ada W. Finifiter. Washington: American Political Science Association.

Hanretty, Chris, and Christel Koop. 2012. “Measuring the Formal Independence of Regulatory Agencies.” *Journal of European Public Policy* 19 (2): 198–216.

Hazlett, Thomas W., and Dennis L. Weisman. 2009. “Market Power in U.S. Broadband Services.” George Mason University Law and Economics Research Paper Series. George Mason University. http://www.law.gmu.edu/assets/files/publications/working_papers/0969MarketPowerinUSBroadband.pdf.

Hazlett, Thomas W., and Joshua D. Wright. 2011. “The Law and Economics of Network Neutrality.” George Mason University Law and Economics Research Paper Series. George Mason University. http://www.law.gmu.edu/assets/files/publications/working_papers/1136Law&EconomicsofNetworkNeutrality.pdf.

Heclo, Hugh. 1977. *A Government of Strangers: Executive Politics in Washington*. Washington: Brookings Institution Press.

Heidenheimer, Arnold J., Hugh Heclo, and Carolyn Teich Adams. 1990. *Comparative Public Policy: The Politics of Social Choice in America, Europe, and Japan*. New York: St. Martin’s Press.

Heilemann, John. 2007. “Money Chooses Sides.” *New York Magazine*, April 23. http://nymag.com/news/politics/30634/.

Hemphill, C. Scott. 2008. “Network Neutrality and the False Promise of Zero-Price Regulation.” *Yale Journal on Regulation* 25 (2): 135.

Hemphill, C. Scott, and Tim Wu. 2013. “Parallel Exclusion.” *Yale Law Journal* 122 (5) (March): 1182.

Hermalin, Benjamin E., and Michael L. Katz. 2007. “The Economics of Product-Line Restrictions with an Application to the Network Neutrality Debate.” *Information Economics and Policy* 19 (2): 215–248.

Herman, Bill D. 2006. “Opening Bottlenecks: On Behalf of Mandated Network Neutrality.” *Federal Communications Law Journal* 59 (1): 103–156.

Hermes Project. 2011. “Sky Starts Marketing to Non TV Customers.” http://www.thehermesproject.com/2011/03/sky-starts-marketing-to-non-tv.html.

Highfield, Ashley. 2007a. “The Net At Breaking Point? Pt 1.” http://www.bbc.co.uk/blogs/bbcinternet/2007/11/the_net_at_breaking_point_part_1.html.

———. 2007b. “The Net At Breaking Point? Pt 2.” http://www.bbc.co.uk/blogs/bbcinternet/2007/11/the_net_at_breaking_point_pt_2.html.

———. 2008a. "'Hidden Costs' Of Watching TV Online?" http://www.bbc.co.uk/blogs/bbcinternet/2008/04/hidden_costs_of_watching_tv_on.html.
———. 2008b. "'The BBC Should Not Have to Subsidize the ISPs' Ashley Highfield, the BBC's Director of Future Media and Technology, Explains the Corporation's Position." *The Daily Telegraph*, April 10.
Hilton, George W. 1972. "The Basic Behavior of Regulatory Commissions." *American Economic Review* 62 (2): 47–54.
Hirsch, Paul M. 1995. "Tales from the Field." In *Studying Elites Using Qualitative Methods*, edited by Rosanna Hertz and Jonathan B. Imber. Thousand Oaks, CA, USA: Sage Publications.
Ho, Daniel E. 2007. "Congressional Agency Control: The Impact of Statutory Partisan Requirements on Regulation." *American Law and Economics Association Working Paper* (73). http://dho.stanford.edu/research/partisan.pdf.
Hogendorn, Christiaan. 2007. "Broadband Internet: Net Neutrality Versus Open Access." *International Economics & Economic Policy* 4 (2): 185–208.
———. 2012. "Spillovers and Network Neutrality." In *Regulation and the Performance of Communication and Information Networks*, edited by Gerald R. Faulhaber, Gary Madden, and Jeffrey Petchey. Cheltenham: Edward Elgar.
Horn, Murray J. 1995. *The Political Economy of Public Administration: Institutional Choice in the Public Sector*. Cambridge, UK: Cambridge University Press.
Horwitz, Robert Britt. 1989. *The Irony of Regulatory Reform: The Deregulation of American Telecommunications*. New York: Oxford University Press.
House Energy and Commerce Committee. 2008. "Deception and Distrust: The Federal Communications Commission Under Chairman Kevin J. Martin."
House of Commons. 2007. "Ofcom Annual Plan 2007-8: Oral Evidence Taken Before the Culture, Media and Sport Committee and the Trade and Industry Committee." http://www.publications.parliament.uk/pa/cm200607/cmselect/cmtrdind/459/7041701.htm.
———. 2009. "Culture, Media and Sport Committee - Minutes of Evidence: Ofcom Annual Plan 2009-10." http://www.publications.parliament.uk/pa/cm200910/cmselect/cmcumeds/77/7701.htm.
———. 2011a. "Committee of Public Accounts - Ofcom: The Effectiveness of Converged Regulation." http://www.publications.parliament.uk/pa/cm201011/cmselect/cmpubacc/688/688.pdf.
———. 2011b. "Culture, Media and Sport Committee - Minutes of Evidence: The Work of Ofcom." http://www.publications.parliament.uk/pa/cm201012/cmselect/cmcumeds/956/11050301.htm.
House of Lords. 2012. "Unrevised Transcript of Evidence Taken Before the Select Committee on Communications Inquiry on Superfast Broadband." http://www.parliament.uk/documents/lords-committees/communications/Superfastbroadband/ucCOMMS190612ev9.pdf.
Hundt, Reed H., and Gregory L. Rosston. 2006. "Communications Policy for 2006 and Beyond." *Federal Communications Law Journal* 58 (1).
Imtech Telecom. 2005. "Imtech Telecom UK Case Study: Plusnet Selects Ellacoya Networks." http://www.imtechtele.com/content/references/_doc/UK/uk_0005.pdf.
Ingram, Peter. 2006. "Net Neutrality: Battle for a New World Order or Much Ado About Nothing?" presented at the Communications Research Networks (CRN) Workshop, December 14. http://media.ofcom.org.uk/2006/12/14/communications-research-networks-crn-workshop/.
Jackson, Mark. 2002. "FairADSL ISP on Service Problems." October. http://www.ispreview.co.uk/ispnews/archives2002/arc6.shtml.

———. 2011. “UPD UK ISP TalkTalk Uses Traffic Management to Block Users from Playing OnLive.” September 30. http://www.ispreview.co.uk/story/2011/09/30/uk-isp-talktalk-uses-traffic-management-to-block-customers-from-playing-onlive.html.
Jamison, Mark A., and Janice Alane Hauge. 2008. “Getting What You Pay For: Analyzing the Net Neutrality Debate.” http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1081690.
Jitsuzumi, Toshiya. 2011. “Discussion on Network Neutrality: Japan’s Perspective.” *Communications & Convergence Review* 3 (1): 71–89.
Jordan, Scott. 2007. “A Layered Network Approach to Network Neutrality.” *International Journal of Communication* 1: 427–460.
———. 2009. “Implications of Internet Architecture on Net Neutrality.” *ACM Trans. Internet Technol.* 9 (2) (May): 5:1–5:28.
Jordan, Scott, and Arijit Ghosh. 2010. “A Framework for Classification of Traffic Management Practices as Reasonable or Unreasonable.” *ACM Transactions on Internet Technology* 10 (3): 1–23.
Joskow, Paul L. 1974. “Inflation and Environmental Concern: Structural Change in the Process of Public Utility Price Regulation.” *Journal of Law and Economics* 17 (2): 291–327.
“Julius Genachowski.” 2012. *The Washington Post*, July 25, sec. Politics. http://www.washingtonpost.com/politics/julius-genachowski/gIQAHAht9O_topic.html.
Kagan, Elena. 2001. “Presidential Administration.” *Harvard Law Review* 114 (8): 2245.
Kagan, Robert A. 2003. *Adversarial Legalism: The American Way of Law*. Cambridge: Harvard University Press.
Kagan, Robert A., and Lee Axelrad. 2000. *Regulatory Encounters: Multinational Corporations and American Adversarial Legalism*. Berkeley: University of California Press.
Kantar Media. 2012. “OCI Tracker Benchmark Study Q3 2012.” http://stakeholders.ofcom.org.uk/binaries/research/telecoms-research/online-copyright/Kantar-Media.pdf.
Kanuparthy, P., and C. Dovrolis. 2010. “DiffProbe: Detecting ISP Service Discrimination.” In *2010 Proceedings IEEE INFOCOM*, 1–9. IEEE. http://dl.acm.org/citation.cfm?id=1833745.
Katzmann, Robert A. 1980. *Regulatory Bureaucracy: The Federal Trade Commission and Antitrust Policy*. Cambridge, MA: MIT Press.
Kaufman, Herbert. 1960. *The Forest Ranger: A Study in Administrative Behavior*. Baltimore: Resources for the Future by Johns Hopkins Press.
Kelman, Steven. 1981. *Regulating America, Regulating Sweden: A Comparative Study Of Occupational Safety and Health Policy*. Cambridge: MIT Press.
Kiddoo, Jean L. 2010. “Re: Ex Parte Notice; WC Docket No. 09-191 and 07-52”, May 7. http://apps.fcc.gov/ecfs/document/view?id=7020450131.
Kiedrowski, Tom. 2007. “Joint CEPS and Progress for Freedom Conference” presented at the Joint CEPS and Progress for Freedom conference, February 26. http://media.ofcom.org.uk/2007/02/26/joint-ceps-and-progress-for-freedom-conference/.
Kisielowska-Lipman, Marzena. 2012. “Lost on the Broadband Super Highway: Consumer Understanding of Information on Traffic Management.” http://www.consumerfocus.org.uk/files/2012/11/Lost-on-the-broadband-super-highway.pdf.
Knieps, Gunter, and Patrick Zenhausern. 2008. “The Fallacies of Network Neutrality Regulation.” *Competition and Regulation in Network Industries* 9 (2): 119.
Krafft, Jackie, and Evens Salies. 2008. “The Diffusion of ADSL and Costs of Switching Internet Providers in the Broadband Industry: Evidence from the French Case.” *Research Policy* 37 (4): 706–719.

Krämer, Jan, and Lukas Wiewiorra. 2012. “Network Neutrality and Congestion Sensitive Content Providers: Implications for Content Variety, Broadband Investment, and Regulation.” *Information Systems Research* (May 7). http://isr.journal.informs.org/cgi/doi/10.1287/isre.1120.0420.
———. 2013. “Net Neutrality and Fairness: An Empirical Investigation of Consumers’ Appreciation of Quality of Service in Internet Access”. Karlsruhe Institute of Technology Working Paper. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2031034.
Krämer, Jan, Lukas Wiewiorra, and Christof Weinhardt. 2012. “Net Neutrality: A Progress Report.” *Telecommunications Policy*. doi:doi:10.1016/j.telpol.2012.08.005.
Kreibich, Christian, Nicholas Weaver, Boris Nechaev, and Vern Paxson. 2010. “Netalyzr: Illuminating the Edge Network.” In *Proceedings of the 10th Annual Conference on Internet Measurement*, 246–259. IMC ’10. Melbourne, Australia: ACM. http://conferences.sigcomm.org/imc/2010/papers/p246.pdf.
Kumar, Vishesh. 2008. “Cable Prepares an Answer to FiOS.” *Wall Street Journal*, February 14. http://online.wsj.com/article/SB120295689385867313.html.
Labaton, Stephen. 2007. “F.C.C. Planning Rules to Open Cable Market.” *The New York Times*, November 10, sec. Washington. http://www.nytimes.com/2007/11/10/washington/10cable.html.
Lahtinen, Sebastien. 2003. “Nildram Introduces ‘Traffic Shaping’...” http://www.thinkbroadband.com/news/951-nildram-introduces-traffic-shaping.html.
Lasar, Matthew. 2008. “Interview: Laying It on the Line with FCC Chair Kevin Martin.” *Ars Technica*, October 6. http://arstechnica.com/features/2008/10/fcc-interview-kevin-martin/.
Leaver, Clare. 2009. “Bureaucratic Minimal Squawk Behavior: Theory and Evidence from Regulatory Agencies.” *American Economic Review* 99 (3) (May): 572–607.
Lee, Robin S., and Tim Wu. 2009. “Subsidizing Creativity Through Network Design: Zero-Pricing and Net Neutrality.” *Journal of Economic Perspectives* 23 (3): 61–76.
Lehr, William E., Sharon E. Gillett, Marvin A. Sirbu, and Jon M. Peha. 2006. “Scenarios for the Network Neutrality Arms Race.” In *34th Research Conference on Communication, Information, and Internet Policy (TPRC)*. Arlington, VA, USA. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2104379.
Lemley, Mark A, and Lawrence Lessig. 2001. “The End of End-to-End: Preserving the Architecture of the Internet in the Broadband Era.” *UCLA Law Review* 48: 925.
Lennett, Benjamin. 2009. “Dis-Empowering Users vs. Maintaining Internet Freedom: Network Management and Quality of Service (QoS).” *CommLaw Conspectus: Journal of Communications Law and Policy* 18 (1): 97.
Lessig, Lawrence. 1998. “The New Chicago School.” *The Journal of Legal Studies* 27 (2): 661–691.
———. 1999. *Code: And Other Laws of Cyberspace*. New York: Basic Books.
———. 2001. *The Future of Ideas: The Fate of the Commons in a Connected World*. Random House Digital, Inc.
———. 2006. “Testimony Before the United States Senate, Committee on Commerce, Science, and Transportation, at Its Hearing on: Network Neutrality”. 109th Cong., 2nd Session. http://www.lessig.org/blog/archives/Lessig_Testimony_2.pdf.
Levinson, David, and Andrew Odlyzko. 2008. “Too Expensive to Meter: The Influence of Transaction Costs in Transportation and Communication.” *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 366 (1872) (June 13): 2033–2046.
Li, Chiehyu, and James Losey. 2009. “Bandwidth Caps for Residential High-Speed Internet in the U.S. and Japan”. New America Foundation. http://www.newamerica.net/files/Bandwidth%20Caps%20for%20High-Speed%20Internet%20in%20the%20U.S.%20and%20Japan.pdf.
Lijphart, Arend. 1971. “Comparative Politics and the Comparative Method.” *The American Political Science Review* 65 (3): 682–693.

Litan, Robert E., and Hal J. Singer. 2007. “Unintended Consequences of Net Neutrality Regulation.” *Journal on Telecommunications & High Technology Law* 5 (3): 533.
Livingstone, Sonia. 2008. “What is the Citizen’s Interest in Communication Regulation? Ofcom’s Agenda for ‘Citizens, Communications and Convergence’.” In *Media@lse Fifth Anniversary Conference*. London: LSE Research Online. http://eprints.lse.ac.uk/21561/.
London Economics. 2011. “Consumer Information on Broadband Speed and Net Neutrality Experiment.” http://stakeholders.ofcom.org.uk/binaries/consultations/net-neutrality/statement/Consumer_information1.pdf.
Lowi, Theodore J. 1979. *The End of Liberalism: The Second Republic of the United States*. 2nd ed. New York: W. W. Norton & Company.
Lunt, Peter, and Sonia Livingstone. 2012. *Media Regulation: Governance and the Interests of Citizens and Consumers*. London: SAGE Publications Ltd.
Magat, Wesley A., Alan J. Krupnick, and Winston Harrington. 1986. *Rules in the Making: A Statistical Analysis of Regulatory Agency Behavior*. Washington: Resources for the Future.
Maggetti, Martino. 2007. “De Facto Independence After Delegation: A Fuzzy-Set Analysis.” *Regulation & Governance* 1 (4): 271–294.
Majone, Giandomenico. 1994. “The Rise of the Regulatory State in Europe.” *West European Politics* 17 (3): 77–101.
———. 1999. “The Regulatory State and its Legitimacy Problems.” *West European Politics* 22 (1) (January): 1–24.
Marcus, J. Scott. 2008. “Network Neutrality: The Roots of the Debate in the United States.” *Intereconomics* 43 (1): 30–37.
Marcus, J. Scott, and Juan Rendon Schneir. 2010. “Drivers and Effects of the Size and Composition of Telecoms Regulatory Agencies.” In *European Regional ITS Conference*. Copenhagen, Denmark. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1675705.
Marsden, Christopher T. 2007. “Net Neutrality and Consumer Access to Content.” *SCRIPT-ed* 4 (4): 407–435.
———. 2008. “Net Neutrality: The European Debate.” *Journal of Internet Law* 12 (2): 1–16.
———. 2010. *Net Neutrality: Towards a Co-Regulatory Solution*. London: Bloomsbury Academic.
———. 2013. “Network Neutrality: A Research Guide.” In *Research Handbook on Governance of the Internet*, edited by Ian Brown. Cheltenham: Edward Elgar.
Marsden, Christopher T., and Jonathan Cave. 2007. “Beyond the ‘Net Neutrality’ Debate: Price and Quality Discrimination in Next Generation Internet Access.” In *35th Research Conference on Communication, Information, and Internet Policy (TPRC)*. Alexandria, VA. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2103636.
Marshall, Catherine, and Gretchen B. Rossman. 2010. *Designing Qualitative Research*. Thousand Oaks, CA, USA: SAGE.
Martin, James, and James Westall. 2007. “Assessing the Impact of BitTorrent on DOCSIS Networks.” In *Proceedings of the 2007 IEEE Broadnets*. Raleigh, North Carolina, USA. http://www.cs.clemson.edu/~jmarty/papers/bittorrentBroadnets.pdf.
Martin, Kevin J. 2005a. “Comments on Commission Policy Statement.” http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-260435A2.pdf.
———. 2005b. “Chairman Kevin J. Martin Comments on Adoption of Wireline Broadband Internet Access Order.” http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-260433A2.pdf.
———. 2005c. “Chairman Kevin J. Martin Comments on Adoption of SBC/AT&T and Verizon/MCI Merger Orders.” http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-261936A2.pdf.
———. 2006. “Written Statement Of The Honorable Kevin J. Martin Chairman Federal Communications Commission Before the Committee on Commerce, Science &

Transportation U.S. Senate". 109th Cong., 2nd Session. http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-267390A1.pdf.
———. 2007a. "Written Statement of The Honorable Kevin J. Martin Chairman Federal Communications Commission Before the Committee on Energy and Commerce U.S. House of Representatives". 110th Cong., 1st Session. http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-271486A1.pdf.
———. 2007b. "Statement of Chairman Kevin J. Martin Re: Broadband Industry Practices, Notice of Inquiry, WC Docket No. 07-52." http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-07-31A2.pdf.
———. 2008a. "Statement of Chairman Kevin J. Martin, En Banc Hearing of the Federal Communications Commission." http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-280446A1.pdf.
———. 2008b. "Statement by FCC Chairman Kevin J. Martin on Announcement by Comcast and BitTorrent." http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-281165A1.pdf.
———. 2008c. "Written Statement Of The Honorable Kevin J. Martin Chairman Federal Communications Commission Before the United States Senate Committee on Commerce, Science and Transportation". 110th Cong., 2nd Session. http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-281690A1.pdf.
———. 2008d. "Statement of Kevin J. Martin Re: Formal Complaint of Free Press and Public Knowledge Against Comcast Corporation for Secretly Degrading Peer-to-Peer Applications; Broadband Industry Practices, File No. EB-08-IH-1518, WC Docket No. 07-52." http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-08-183A2.pdf.
Martin, Kevin J., and Deborah Taylor Tate. 2006. "Joint Statement of Chairman Kevin J. Martin and Commissioner Deborah Taylor Tate Re: AT&T Inc. and BellSouth Corporation Application for Transfer of Control, WC Docket No. 06-74." http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-06-189A2.pdf.
Mashaw, Jerry L., and David L. Harfst. 1986. "Regulation and Legal Culture: The Case of Motor Vehicle Safety." *Yale Journal on Regulation* 4 (2): 257.
McCubbins, Mathew D., Roger G. Noll, and Barry R. Weingast. 1987. "Administrative Procedures as Instruments of Political Control." *Journal of Law, Economics, & Organization* 3 (2): 243–277.
McCubbins, Mathew D., and Thomas Schwartz. 1984. "Congressional Oversight Overlooked: Police Patrols Versus Fire Alarms." *American Journal of Political Science* 28 (1) (February): 165.
McDowell, Robert M. 2007. "Statement of Commissioner Robert M. McDowell Regarding Broadband Industry Practices, Notice of Inquiry, WC Docket No. 07-52 (FCC 07-31)." http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-07-31A1.pdf.
———. 2013. "Statement of Commissioner Robert M. McDowell on the Announced Departure of Chairman Julius Genachowski." http://transition.fcc.gov/Daily_Releases/Daily_Business/2013/db0322/DOC-319732A1.pdf.
McGarity, Thomas O. 1992. "Some Thoughts on Deossifying the Rulemaking Process." *Duke Law Journal* 41 (6): 1385.
Meinrath, Sascha D., and Victor W. Pickard. 2008. "The New Network Neutrality: Criteria for Internet Freedom." *International Journal of Communications Law and Policy* (12): 225–243.
Melnick, R. Shep. 1983. *Regulation and the Courts: The Case of the Clean Air Act*. Washington: Brookings Institution Press.
———. 1992. "Administrative Law and Bureaucratic Reality." *Administrative Law Review* 44 (2): 245.
Mitnick, Barry M. 1980. *The Political Economy of Regulation*. New York: Columbia University Press.
Moe, Terry M. 1985. "Control and Feedback in Economic Regulation: The Case of the NLRB." *The American Political Science Review* 79 (4): 1094–1116.

———. 1987a. "An Assessment of the Positive Theory of 'Congressional Dominance'." *Legislative Studies Quarterly* 12 (4): 475–520.
———. 1987b. "Interests, Institutions, and Positive Theory: The Politics of the NLRB." *Studies in American Political Development* 2 (-1): 236–299.
Moe, Terry M., and Scott A. Wilson. 1994. "Presidents and the Politics of Structure." *Law and Contemporary Problems* 57 (2): 1.
Mohammed, Arshad. 2006. "Verizon Executive Calls for End to Google's 'Free Lunch'." *The Washington Post*, February 7, sec. Technology. http://www.washingtonpost.com/wp-dyn/content/article/2006/02/06/AR2006020601624.html.
Mooyaart, Xavier. 2012. "Traffic Management and 'Net Neutrality': Response of Everything Everywhere to the Ofcom Discussion Document." http://stakeholders.ofcom.org.uk/binaries/consultations/net-neutrality/responses/Everything_Everywhere_.pdf.
Moran, Michael. 2003. *The British Regulatory State*. USA: Oxford University Press.
MoveOn.org. 2013. "Top Ten Examples of Grassroots Mobilization in the Net Neutrality Fight." Accessed July 5. http://civic.moveon.org/save_the_internet/littleguy.html.
Mueller, Milton L. 2007. "Net Neutrality as Global Principle for Internet Governance". Internet Governance Project. http://www.internetgovernance.org/wordpress/wp-content/uploads/NetNeutralityGlobalPrinciple.pdf.
———. 2011. "DPI Technology from the Standpoint of Internet Governance Studies: An Introduction". Syracuse University School of Information. http://dpi.ischool.syr.edu/Technology_files/WhatisDPI-2.pdf.
Mueller, Milton L., and Hadi Asghari. 2011. "Deep Packet Inspection and Bandwidth Management: Battles over BitTorrent in Canada and the United States." In *39th Research Conference on Communication, Information, and Internet Policy (TPRC)*. Arlington, VA. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1989644.
Murray-Watson, Andrew. 2007. "Internet Groups Warn BBC over iPlayer Plans." *The Independent*, August 12, sec. Business & Money.
Musacchio, John, Galina Schwartz, and Jean Walrand. 2009. "A Two-Sided Market Analysis of Provider Investment Incentives with an Application to the Net-Neutrality Issue." *Review of Network Economics* 8 (1): 22–39.
Napoli, Philip M. 1998. "Government Assessment of FCC Performance: Recurring Patterns and Implications for Recent Reform Efforts." *Telecommunications Policy* 22 (4–5) (May): 409–418.
Nardotto, Mattia, Tommaso Valletti, and Frank Verboven. 2012. "Unbundling the Incumbent: Evidence from UK Broadband." *Centre for Economic Policy Research Discussion Paper Series* No. 9194 (October). http://www.aueb.gr/conferences/Crete2013/papers/Valletti.PDF.
National Audit Office. 2010. "The Effectiveness of Converged Regulation." http://www.nao.org.uk/publications/1011/ofcom.aspx.
National Cable and Telecommunications Association. 2003. "Ex Parte Re: CS Docket No. 02-52", September 8. http://apps.fcc.gov/ecfs/document/view?id=6514882243.
———. 2010. "Statement Regarding FCC Approval of the Open Internet Item." http://www.ncta.com/news-and-events/media-room/article/2312.
"Net Firm Warns on Web Video Costs." 2007. *BBC News*, August 13, sec. Technology. http://news.bbc.co.uk/2/hi/technology/6944176.stm.
Niskanen, William A. 1971. *Bureaucracy and Representative Government*. New York: Aldine Atherton.
———. 1975. "Bureaucrats and Politicians." *Journal of Law and Economics* 18 (3): 617–643.
Noam, Eli M. 2009. *Media Ownership and Concentration in America*. Oxford: Oxford University Press.
Noll, Roger G. 1971. "The Behavior of Regulatory Agencies." *Review of Social Economy* 29 (1): 15–19.

———. 1985. “Government Regulatory Behavior: A Multidisciplinary Survey and Synthesis.” In *Regulatory Policy and the Social Sciences*, edited by Roger G. Noll. Berkeley: University of California Press.
———. 1989. “Economic Perspectives on the Politics of Regulation.” In *Handbook of Industrial Organization*, edited by Richard Schmalensee and Robert Willig, 2:1253–1287. Amsterdam: Elsevier.
Noll, Roger G., and Bruce M. Owen. 1983. *The Political Economy of Deregulation: Interest Groups in the Regulatory Process*. Washington: American Enterprise Institute for Public Policy Research.
North, Douglass C. 1990. *Institutions, Institutional Change and Economic Performance*. Cambridge: Cambridge University Press.
———. 2005. “Institutions and the Performance of Economies Over Time.” In *Handbook of New Institutional Economics*, edited by Mary M. Menard and Mary M. Shirley. Berlin/Heidelberg: Springer-Verlag.
Nuechterlein, Jonathan E. 2009. “Antitrust Oversight of an Antitrust Dispute: An Institutional Perspective on the Net Neutrality Debate.” *Journal on Telecommunications & High Technology Law* 7 (1): 19.
Nuechterlein, Jonathan E., and Philip J. Weiser. 2005. *Digital Crossroads: American Telecommunications Policy in the Internet Age*. USA: The MIT Press.
O’Connell, Patricia. 2005. “Online Extra: At SBC, It’s All About ‘Scale and Scope’.” *BusinessWeek: Online Magazine*, November 7. http://www.businessweek.com/magazine/content/05_45/b3958092.htm.
Obama for America. 2007. “Barack Obama: Connecting and Empowering All Americans Through Technology and Innovation.”
Odlyzko, Andrew. 2009. “Network Neutrality, Search Neutrality, and the Never-Ending Conflict Between Efficiency and Fairness in Markets.” *Review of Network Economics* 8 (1): 40–60.
OECD. 2006. “Broadband Penetration, Historic, G7 Countries.” http://www.oecd.org/sti/broadbandandtelecom/38449381.xls.
———. 2011. “OECD Communications Outlook 2011.” http://www.oecd.org/document/44/0,3746,en_2649_34225_43435308_1_1_1_1,00.html.
Ofcom. 2004a. “Annual Report 2003/4.” http://www.ofcom.org.uk/about/annual-reports-and-plans/annual-reports/ofcom-annual-report-2003-04/.
———. 2004b. “Strategic Review of Telecommunications Phase 1 Consultation.” http://stakeholders.ofcom.org.uk/binaries/consultations/telecoms_review1/summary/tcoms_rev.pdf.
———. 2004c. “Review of the Wholesale Broadband Access Markets.” http://stakeholders.ofcom.org.uk%2Fbinaries%2Fconsultations%2Fwbamp%2Fsummary%2Fbroadbandaccessreview.pdf&ei=aU1XULbXBqL10gGUzoC4AQ&usg=AFQjCNFvIiNRIQEu-3KV0BN1-ZrU7RDv0Q.
———. 2004d. “Ofcom’s Broadband Framework.” http://media.ofcom.org.uk/2004/05/13/ofcoms-broadband-framework/.
———. 2004e. “Strategic Review of Telecommunications Phase 2 Consultation.” http://stakeholders.ofcom.org.uk/binaries/consultations/telecoms_p2/summary/maincondoc.pdf.
———. 2005a. “Undertakings Which Have Been Offered by British Telecommunications Plc (‘BT’): a Notice Under Section 155(1) of the Enterprise Act 2002:” http://stakeholders.ofcom.org.uk/consultations/sec155/.
———. 2005b. “Broadband Regulation.” http://stakeholders.ofcom.org.uk%2Fbinaries%2Fconsultations%2Frwlam%2Fstatement%2Fbbr.pdf&ei=hE9XUPvhH-KA0AGWoYDwCw&usg=AFQjCNFLSg744jPnZL3_wbY-3qbF8V_MBw.

———. 2006a. “Annual Report 2005/6.” http://www.ofcom.org.uk/about/annual-reports-and-plans/annual-reports/ofcom-annual-report-2005-06/h/files/2010/06/annual_rpt_0506.pdf.
———. 2006b. “Ofcom Response to the European Commission Consultation on Content Online in the Single Market.” http://ec.europa.eu/avpolicy/docs/other_actions/contributions/ofcom_col_en.pdf.
———. 2006c. “The Consumer Experience Research Report.” http://stakeholders.ofcom.org.uk/binaries/research/consumer-experience/research1.pdf.
———. 2007a. “Annual Report 2006/7 – Section A: Introduction.” http://www.ofcom.org.uk/files/2010/06/sectiona2.pdf.
———. 2007b. “The Communications Market 2007.” http://stakeholders.ofcom.org.uk/binaries/research/cmr/cm07_3.pdf.
———. 2008a. “Annual Report 2007/8.” http://www.ofcom.org.uk/files/2010/06/annrep0708.pdf.
———. 2008b. “Annual Plan 2008/9.” http://www.ofcom.org.uk/files/2010/06/annplan0809.pdf.
———. 2008c. “Voluntary Code of Practice: Broadband Speeds.” http://stakeholders.ofcom.org.uk/telecoms/codes-of-practice/broadband-speeds-cop/voluntary-codes-of-practice/.
———. 2009a. “UK Broadband Speeds 2009.” http://stakeholders.ofcom.org.uk/binaries/research/telecoms-research/broadbandspeeds.pdf.
———. 2009b. “The Communications Market 2009.” http://stakeholders.ofcom.org.uk/binaries/research/cmr/cmr09.pdf.
———. 2009c. “The Consumer Experience 2009.” http://stakeholders.ofcom.org.uk/binaries/research/consumer-experience/research09.pdf.
———. 2010a. “Annual Plan 2010/11.” http://www.ofcom.org.uk/files/2010/06/annplan1011.pdf.
———. 2010b. “Review of the Wholesale Broadband Access Markets: Consultation on Market Definition, Market Power Determinations and Remedies.” http://stakeholders.ofcom.org.uk/binaries/consultations/wba/summary/wbacondoc.pdf.
———. 2010c. “Traffic Management and ‘Net Neutrality’: A Discussion Document.” http://stakeholders.ofcom.org.uk/consultations/net-neutrality/.
———. 2010d. “UK Broadband Speeds, May 2010.” http://stakeholders.ofcom.org.uk/binaries/research/telecoms-research/bbspeeds2010/bbspeeds2010.pdf.
———. 2010e. “Review of the Wholesale Broadband Access Markets: Second Consultation on Market Definition, Market Power Determinations and Remedies.” http://stakeholders.ofcom.org.uk/consultations/wholesale-broadband-markets/.
———. 2010f. “Traffic Management and ‘Net Neutrality’.” http://stakeholders.ofcom.org.uk/consultations/net-neutrality/?showResponses=true.
———. 2010g. “Strategic Review of Consumer Switching: A Consultation on Switching Processes in the UK Communications Sector.” http://stakeholders.ofcom.org.uk/binaries/consultations/consumer-switching/summary/switching.pdf.
———. 2011a. “The Communications Market 2011.” http://stakeholders.ofcom.org.uk/binaries/research/cmr/cmr11/UK_CMR_2011_FINAL.pdf.
———. 2011b. “Approach to Net Neutrality.” http://stakeholders.ofcom.org.uk/binaries/consultations/net-neutrality/statement/statement.pdf.

———. 2013a. “Board Procedures.” http://www.ofcom.org.uk/about/how-ofcom-is-run/ofcom-board-2/board-procedures/.
———. 2013b. “The Consumer Experience 2012.” http://stakeholders.ofcom.org.uk/binaries/research/consumer-experience/tce-12/Consumer_Experience_Researc1.pdf.
Ofcom Consumer Panel. 2007. “Consumer Panel Asks Ofcom to Produce a Code of Practice to Address Broadband Speed Concerns.” http://www.communicationsconsumerpanel.org.uk/downloads/NewsandMedia/Press-Releases/2007/19%20Dec%202007%20Consumer%20Panel%20asks%20Ofcom%20to%20produce%20a%20code%20of%20practice.pdf.
Ogus, Anthony I. 1994. *Regulation: Legal Form and Economic Theory*. Portland, Oregon: Hart Publishing.
Olson, Mancur. 1965. *The Logic of Collective Action: Public Goods and the Theory of Groups*. Cambridge, MA: Harvard University Press.
Olson, Mary K. 1995. “Regulatory Agency Discretion Among Competing Industries: Inside the FDA.” *Journal of Law, Economics and Organization* 11 (2): 379.
———. 1996. “Substitution in Regulatory Agencies: FDA Enforcement Alternatives.” *Journal of Law, Economics and Organization* 12 (2): 376.
———. 1997. “Firm Characteristics and the Speed of FDA Approval.” *Journal of Economics & Management Strategy* 6 (2) (June): 377–401.
Open Rights Group. 2010. “Net Neutrality.” August 19. http://wiki.openrightsgroup.org/wiki/Net_Neutrality.
OpenSecrets.org. 2013. “OpenSecrets.org Lobbying Database.” http://www.opensecrets.org/lobby/index.php.
Orlowski, Andrew. 2008. “BBC Tells ISPs to Get Stuffed.” *The Register*, April 3. http://www.theregister.co.uk/2008/04/03/bbc_highfield_isp_threat/.
Ostrander, Susan A. 1995. “‘Surely You’re Not in This Just to Be Helpful’: Access, Rapport, and Interviews in Three Studies of Elites.” In *Studying Elites Using Qualitative Methods*, edited by Rosanna Hertz and Jonathan B. Imber. Thousand Oaks, CA, USA: Sage Publications.
Owen, Bruce M. 2007. “Antecedents to Net Neutrality.” *Regulation* (Fall).
———. 2011. “Antitrust and Vertical Integration in ‘New Economy’ Industries with Application to Broadband Access.” *Review of Industrial Organization* 38 (4): 363–386.
Owen, Bruce M., and Gregory Rosston. 2006. “Local Broadband Access: Primum Non Nocere or Primum Processi? A Property Rights Approach.” In *Net Neutrality or Net Neutering: Should Broadband Internet Services Be Regulated?*, edited by Thomas M. Lenard and Randolph J. May, 163–194. Springer US.
Palmer, Maija. 2007. “ISPs Warn BBC over New iPlayer Service”, August 12. http://www.ft.com/intl/cms/s/0/f3428cd4-48fb-11dc-b326-0000779fd2ac.html#axzz2SRAL6sTQ.
Peha, Jon M. 2007. “The Benefits and Risks of Mandating Network Neutrality, and the Quest for a Balanced Policy.” *International Economics and Economic Policy* 1: 644–668.
Peltzman, Sam. 1976. “Toward a More General Theory of Regulation.” *Journal of Law and Economics* 19 (2): 211–240.
Pierce, Richard J. 1988. “Two Problems in Administrative Law: Political Polarity on the District of Columbia Circuit and Judicial Deterrence of Agency Rulemaking.” *Duke Law Journal* 1988 (2-3): 300.
———. 1995. “Seven Ways to Deossify Agency Rulemaking.” *Administrative Law Review* 47 (1): 59–96.
Plissoneau, Louis, Jean-Laurent Costeaux, and Patrick Brown. 2005. “Analysis of Peer-to-Peer Traffic on ADSL.” In *Proceedings of the 6th International Workshop on Passive and Active Network Measurement (PAM)*. Boston, MA, USA: Springer. http://www.pamconf.net/2005/PDF/34310070.pdf.

Plusnet. 2012a. “Product Archive: Broadband Plus.” http://www.plus.net/support/broadband/products/archive/broadband_plus_archive.shtml.
———. 2012b. “Broadband Product Archive.” http://www.plus.net/support/broadband/products/archive/index.shtml.
———. 2012c. “Broadband.” http://www.plus.net/support/broadband/products/archive/bbyw_archive.shtml#trafficPriority.
Posner, Richard A. 1971. “Taxation by Regulation.” *Bell Journal of Economics* 2 (1): 22–50.
———. 1974. “Theories of Economic Regulation.” *National Bureau of Economic Research Working Paper Series* No. 41 (May).
Powell, Alison, and Alissa Cooper. 2011. “Net Neutrality Discourses: Comparing Advocacy and Regulatory Arguments in the United States and the United Kingdom.” *The Information Society* 27 (5) (October): 311–325.
Powell, Michael K. 2004. “Preserving Internet Freedom: Guiding Principles for the Industry.” http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-243556A1.pdf.
Prince, Jeffrey, and Shane Greenstein. 2013. “Does Service Bundling Reduce Churn?” *Journal of Economics and Management Strategy,* forthcoming.
Prosser, Tony. 1997. *Law and the Regulators*. Oxford: Clarendon Press.
———. 1999. “Theorising Utility Regulation.” *The Modern Law Review* 62 (2): 196–217.
———. 2004. “Regulation, Markets, and Legitimacy.” In *The Changing Constitution*, edited by Jeffrey Erman and Dawn Oliver. Oxford: Oxford University Press.
———. 2010. *The Regulatory Enterprise*. Oxford: Oxford University Press.
Prüfer, Jens, and Eric Jahn. 2007. “Dark Clouds Over the Internet?” *Telecommunications Policy* 31 (3–4): 144–154.
Ragin, Charles. 1994. *Constructing Social Research*. Thousand Oaks, CA, USA: Pine Forge Press.
“Re: BT Throttling P2P Traffic!!!!!!!!!!!!!!!!!!!!!” 2010. October. http://community.bt.com/t5/BB-Speed-Connection-Issues/BT-Throttling-P2P-Traffic/m-p/56929.
Read, Darren. 2012. “Net Neutrality and the EU Electronic Communications Regulatory Framework.” *International Journal of Law and Information Technology* 20 (1) (March 1): 48–72.
Reardon, Marguerite. 2006. “Qwest CEO Supports Tiered Internet.” *CNET*, March 15. http://news.cnet.com/2100-1034_3-6050109.html.
Reggiani, Carlo, and Tommaso Valletti. 2012. “Net Neutrality and Innovation at the Core and at the Edge.” *University of Manchester School of Economics Discussion Paper Series 1202* (January). http://www.socialsciences.manchester.ac.uk/disciplines/economics/research/discussionpapers/pdf/EDP-1202.pdf.
Renda, Andrea. 2008. “I Own the Pipes, You Call the Tune? The Network Neutrality Debate and Its (Ir)relevance for Europe”. Centre for European Policy Studies. http://www.ceps.eu/node/1579.
Richards, Ed. 2008. “Institution of Engineering and Technology Speech, Broadband Britain – Towards the Next Generation” April 16. http://media.ofcom.org.uk/2008/04/16/institution-of-engineering-and-technology-speech-broadband-britain-towards-the-next-generation-wednesday-16-april-2008/.
———. 2010. “Competition Law and the Communications Sector” presented at the UCL Jevons Institute for Competition Law and Economics Annual Colloquium, July 13. http://media.ofcom.org.uk/2010/07/13/competition-law-and-the-communications-sector/.
Richardson, Jeremy John, ed. 1982. *Policy Styles in Western Europe*. London: Allen & Unwin.

Richardson, Jeremy John, Gunnel Gustafsson, and Grant Jordan. 1982. "The Concept of Policy Style." In *Policy Styles in Western Europe*, edited by Jeremy Richardson. London: Allen & Unwin.
Richardson, Tim. 2003. "NTL Seeks to Clarify 1GB/day Broadband Cap." *The Register*, February 10. http://www.theregister.co.uk/2003/02/10/ntl_seeks_to_clarify_1gb/.
Richmond, Shane. 2011. "Superfast Broadband 'for 90 Per Cent of Britain by 2015'." *The Telegraph*, May 12. http://www.telegraph.co.uk/technology/broadband/8510062/Superfast-broadband-for-90-per-cent-of-Britain-by-2015.html.
Rose, Richard. 1973. "Comparing Public Policy: An Overview." *European Journal of Political Research* 1 (1) (April): 67–94.
Rosen, Andrea. 2012. "File 545613, Internet Traffic Management Practice ('ITMP'), Section 36 of the Telecommunications Act, S.C.1993, C.38, as Amended ('Act'), and Paragraphs 126 and 127 of Telecom Regulatory Policy CRTC 2009-657 ('TRP CRTC 2009-657')", February 29. http://www.crtc.gc.ca/eng/archive/2012/lt120229.htm.
Rosston, Gregory L., and Michael D. Topper. 2009. "An Antitrust Analysis of the Case for Wireless Network Neutrality". Stanford Institute for Economic Policy Research. http://businessinnovation.berkeley.edu/Mobile_Impact/Rosston_Topper_Wireless_Net_Neutrality_11_02_09.pdf.
Rothstein, Henry F. 2004. "Precautionary Bans or Sacrificial Lambs? Participative Risk Regulation and the Reform of the UK Food Safety Regime." *Public Administration* 82 (4): 857–881.
Rubin, Herbert J., and Irene S. Rubin. 2004. *Qualitative Interviewing: The Art of Hearing Data*. 2nd ed. Thousand Oaks, CA, USA: Sage Publications, Inc.
Russell, Milton, and Robert B. Shelton. 1974. "A Model of Regulatory Agency Behavior." *Public Choice* 20 (1): 47–62.
Ryan, Gery W, and H. Russell Bernard. 2003. "Techniques to Identify Themes." *Field Methods* 15 (1) (February 1): 85–109.
Ryan, John M. 2011. "Re: Preserving the Open Internet, GN Docket No. 09-191", February 17. http://apps.fcc.gov/ecfs/document/view?id=7021030031.
Saltzer, J. H., D. P. Reed, and D. D. Clark. 1984. "End-to-end Arguments in System Design." *ACM Trans. Comput. Syst.* 2 (4): 277–288.
Sandvig, Christian. 2007. "Network neutrality is the new common carriage." *Info : the Journal of Policy, Regulation and Strategy for Telecommunications, Information and Media* 9 (2/3): 136–147.
Sandvine. 2010a. "2009 Global Broadband Phenomena." http://www.sandvine.com/downloads/documents/2009%20Global%20Broadband%20Phenomena%20-%20Full%20Report.pdf.
———. 2010b. "Comments of Sandvine Incorporated In the Matter of Preserving the Open Internet, GN Docket 09-191; Broadband Industry Practices, WC Docket 07-52." http://apps.fcc.gov/ecfs/document/view?id=7020370020.
Schlick, Austin. 2010. "A Third-Way Legal Framework for Addressing the Comcast Dilemma." http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-297945A1.pdf.
Schuett, Florian. 2010. "Network Neutrality: A Survey of the Economic Literature." *Review of Network Economics* 9 (2): 1–13.
Schwartz, Morris S., and Charlotte Green Schwartz. 1955. "Problems in Participant Observation." *American Journal of Sociology* 60 (4) (January 1): 343–353.
Scott, Colin. 1998. "The Proceduralization of Telecommunications Law." *Telecommunications Policy* 22 (3): 243–254.
Scott, Douglas. 2007. "Regulating the Internet: Speech at Westminster eForum" presented at the Westminster eForum, March 21. http://media.ofcom.org.uk/2007/03/21/regulating-the-internet-speech-at-westminster-eforum/.

Seidenfeld, Mark. 1997. “Demystifying Deossification: Rethinking Recent Proposals to Modify Judicial Review of Notice and Comment Rulemaking.” *Texas Law Review* 75 (3): 483.

Shaffer, Gwen Lisa, and Scott Jordan. 2013. “Classic Conditioning: The FCC’s Use of Merger Conditions to Advance Policy Goals.” *Media, Culture & Society* 35 (3) (April 1): 392–403.

Shelanski, Howard A. 2007. “Adjusting Regulation to Competition: Toward a New Model for U.S. Telecommunications Policy.” *Yale Journal on Regulation* 24 (1): 55.

Sherwood, Chris. 2010. “Yahoo! Response to the European Commission’s Public Consultation on the Open Internet and Net Neutrality in Europe.” http://ec.europa.eu/information_society/policy/ecomm/doc/library/public_consult/net_neutrality/comments/07internet_application_content_providers/yahoo.pdf.

Sidak, J. Gregory. 2006. “A Consumer-Welfare Approach to Network Neutrality Regulation of the Internet.” *Journal of Competition Law and Economics* 2 (3): 349.

Sidak, J. Gregory, and David J. Teece. 2010. “Innovation Spillovers and the ‘Dirt Road’ Fallacy: The Intellectual Bankruptcy of Banning Optional Transactions for Enhanced Delivery Over the Internet.” *Journal of Competition Law and Economics* 6 (3): 521–594.

Simon, Herbert A. 1955. “A Behavioral Model of Rational Choice.” *The Quarterly Journal of Economics* 69 (1) (February): 99–118.

Singer, Hal Jr. 2007. “Net Neutrality: A Radical Form of Non-Discrimination.” *Regulation* 30 (2): 36–41.

Skocpol, Theda. 1979. *States and Social Revolutions: A Comparative Analysis of France, Russia and China*. Cambridge, UK: Cambridge University Press.

Skocpol, Theda, and Margaret Somers. 1980. “The Uses of Comparative History in Macrosocial Inquiry.” *Comparative Studies in Society and History* 22 (02): 174.

Sky. 2010. “Response by British Sky Broadcasting Group PLC to Ofcom’s Discussion Document ‘Traffic Management and Net Neutrality’.” http://stakeholders.ofcom.org.uk/binaries/consultations/net-neutrality/responses/Sky.pdf.

Skype. 2010. “Skype Comments on Ofcom Discussion Paper ‘Traffic Management and Net Neutrality’.” http://stakeholders.ofcom.org.uk/binaries/consultations/net-neutrality/responses/Skype.pdf.

Sluijs, Jasper P. 2009. “Network Neutrality Between False Positives and False Negatives: Introducing a European Approach to American Broadband Markets.” *Federal Communications Law Journal* 62 (1): 77–117.

Sluijs, Jasper P., Florian Schuett, and Bastian Henze. 2011. “Transparency Regulation in Broadband Markets: Lessons from Experimental Research.” *Telecommunications Policy* 35 (7): 592–602.

Smith, James K. 2008. “AT&T Ex Parte Statement: Broadband Industry Practices (WC Docket 07-52)”, March 12. http://apps.fcc.gov/ecfs/document/view?id=6519866328.

Smith, Paul. 2006. “The Politics of UK Television Policy: The Making of Ofcom.” *Media, Culture & Society* 28 (6) (November 1): 929–940. http://mcs.sagepub.com/content/28/6/929.

Sohn, Gigi, Jef Pearlman, Marvin Ammori, Adam Lynn, Ben Scott, Andy Schwartzman, Harold Feld, and Parul Desai. 2008. “Reply Comments of Free Press et al. In the Matter of Petition of Free Press et al. for Declaratory Ruling for Degrading an Internet Application Violates the FCC’s Internet Policy Statement and Does Not Meet an Exception for Reasonable Network Management; Vuze, Inc. Petition to Establish Rules Governing Network Management Practices by Broadband Network Operators, WC Docket No. 07-52; Commercial Availability of Navigation Devices, CS Docket No. 97-80.” http://apps.fcc.gov/ecfs/document/view?id=6519856406.

Spence, David B. 1997. “Administrative Law and Agency Policy-Making: Rethinking the Positive Theory of Political Control.” *Yale Journal on Regulation* 14 (2): 407.

———. 1999. “Managing Delegation Ex Ante: Using Law to Steer Administrative Agencies.” *Journal of Legal Studies* 28 (2): 413.
Speta, James B. 2000a. “Handicapping the Race for the Last Mile: A Critique of Open Access Rules for Broadband Platforms.” *Yale Journal on Regulation* 17: 39.
———. 2000b. “The Vertical Dimension of Cable Open Access.” *University of Colorado Law Review* 71 (4): 975–1010.
———. 2002a. “A Vision of Internet Openness by Government Fiat.” *Northwestern University Law Review* 96 (4): 1553–1578.
———. 2002b. “A Common Carrier Approach to Internet Interconnection.” *Federal Communications Law Journal* 54 (2) (March): 225–279.
———. 2011. “Supervising Managed Services.” *Duke Law Journal* 60 (8): 1715–1759.
Stewart, Richard B. 1975. “The Reformation of American Administrative Law.” *Harvard Law Review* 88 (8): 1667.
Stigler, George J. 1971. “The Theory of Economic Regulation.” *Bell Journal of Economics* 2 (1): 3–21.
Stigler, George J., and Claire Friedland. 1962. “What Can Regulators Regulate? The Case of Electricity.” *Journal of Law and Economics* 5: 1–16.
Svensson, Peter. 2007. “Comcast Actively Hinders Subscribers’ File-Sharing Traffic, AP Testing Shows.” *Associated Press*, October 19.
TalkTalk Group. 2010a. “TalkTalk Group Response, Traffic Management and ‘Net Neutrality’: A Discussion Document.” http://stakeholders.ofcom.org.uk/binaries/consultations/net-neutrality/responses/TTG.pdf.
———. 2010b. “European Commission Consultation / Questionnaire on Open Internet and Net Neutrality in Europe: TalkTalk Group Response.” http://ec.europa.eu/information_society/policy/ecomm/doc/library/public_consult/net_neutrality/comments/01operators_isps/talktalk.pdf.
Tambini, Damian. 2012. “Consumer Representation in UK Communications Policy and Regulation.” *Info* 14 (2) (March 2): 3–16.
Tariq, Mukarram Bin, Murtaza Motiwala, Nick Feamster, and Mostafa Ammar. 2009. “Detecting Network Neutrality Violations with Causal Inference.” In *Proceedings of the 5th International Conference on Emerging Networking Experiments and Technologies*, 289–300. CoNEXT ’09. Rome, Italy: ACM. http://dl.acm.org/citation.cfm?id=1658972.
Taylor, Russ. 2006. “Rethinking Reform of the FCC: A Reply to Randolph May.” *Federal Communications Law Journal* 58: 263.
Telefónica. 2010. “Telefónica O2 UK Response to: ‘Traffic Management and “Net Neutrality”’: A Discussion Document.” http://stakeholders.ofcom.org.uk/binaries/consultations/net-neutrality/responses/Telef_nica_O2_UK.pdf.
Telegeography. 2011. “Global Internet Geography Executive Summary.” http://www.telegeography.com/page_attachments/products/website/research-services/global-internet-geography/0003/1871/GIG_Executive_Summary.pdf.
Telewest. 2004. “Blueyonder Acceptable Use Policy.” *Archive from April 4, 2004*. April 4. http://web.archive.org/web/20040404225138/http://www.blueyonder.co.uk/blueyonder/getContent.jspx?page=h_services_aup.
Teplitz, Steven N., Terri B. Natoli, Matthew A. Brill, Brian W. Murray, Marc Lawrence-Apfelbuam, and Jeffrey Zimmerman. 2010. “Comments of Time Warner Cable Inc. In the Matter of Preserving the Open Internet, GN Docket 09-191; Broadband Industry Practices, WC Docket 07-52.” http://apps.fcc.gov/ecfs/document/view?id=7020376160.
Thaler, Richard H., and Cass R. Sunstein. 2009. *Nudge*. New York, NY: Penguin Books.
Thatcher, Mark. 2002a. “Delegation to Independent Regulatory Agencies: Pressures, Functions and Contextual Mediation.” *West European Politics* 25 (1): 125–147.

———. 2002b. “Regulation After Delegation: Independent Regulatory Agencies in Europe.” *Journal of European Public Policy* 9 (6): 954–972.
———. 2007. “Regulatory Agencies, the State and Markets: A Franco-British Comparison.” *Journal of European Public Policy* 14 (7): 1028–1047.
Thomas, Robert J. 1995. “Interviewing Important People in Big Companies.” In *Studying Elites Using Qualitative Methods*, edited by Rosanna Hertz and Jonathan B. Imber. Thousand Oaks, CA, USA: Sage Publications.
Thompson, Ken. 2011. “Re: Complaint Regarding Rogers’ Internet Traffic Management Practices; Your File # 522253 and File # 517209”, March 22. http://aegir.openmedia.ca/sites/openmedia.ca/files/Rogers%20letter%20to%20CRTC%20re%20ITMPs%20Mar22-11.pdf.
Tomlinson, Dave. 2008. “How UK ISPs Are Charged for Broadband – The Cost of IPStream.” http://community.plus.net/blog/2008/02/28/how-uk-isps-are-charged-for-broadband-the-cost-of-ipstream/.
Topolski, Robb. 2007. “Comcast Is Using Sandvine to Manage P2P Connections - DSLReports Forums.” May 12. http://www.dslreports.com/forum/r18323368-Comcast-is-using-Sandvine-to-manage-P2P-Connections.
Towny, Chris. 2006. “Cable Forum: ntl:Telewest Admits Traffic Shaping on its Network.” http://www.cableforum.co.uk/article/292/ntltelewest-admits-traffic-shaping-on-its-network.
Trebing, Harry M. 1981. “Equity, Efficiency, and the Viability of Public Utility Regulation.” In *Applications of Economic Principles in Public Utility Industries*, edited by Werner Sichel and Thomas Gies. Ann Arbor: University of Michigan.
Tropina, Tatiana, Jason Whalley, and Peter Curwen. 2010. “Functional Separation Within the European Union: Debates and Challenges.” *Telematics and Informatics* 27 (3) (August): 231–241.
Tullock, Gordon. 1965. *The Politics of Bureaucracy*. Washington: Public Affairs Press.
U.S. Office of Personnel Management. 2013. “FedScope: Federal Human Resources Data, Federal Communications Commission.” http://www.fedscope.opm.gov/.
Ulaby, Neda. 2008. “Kevin Martin’s Contentious Turn at Helm of FCC.” *All Things Considered*. NPR. http://www.npr.org/templates/story/story.php?storyId=18711487.
Useem, Michael. 1995. “Reaching Corporate Executives.” In *Studying Elites Using Qualitative Methods*, edited by Rosanna Hertz and Jonathan B. Imber. Thousand Oaks, CA, USA: Sage Publications.
Valcke, Peggy, Liyang Hou, David Stevens, and Eleni Kosta. 2009. “Guardian Knight or Hands Off? The European Response to Network Neutrality: Legal Considerations on the Electronic Communications Reform.” *Communications & Strategies* 72 (April 8): 89–113.
van Schewick, Barbara. 2007. “Towards an Economic Framework for Network Neutrality Regulation.” *Journal on Telecommunications & High Technology Law* 5 (2): 329–391.
———. 2010. *Internet Architecture and Innovation*. Cambridge: MIT Press.
———. 2012. “Network Neutrality and Quality of Service: What a Non-Discrimination Rule Should Look Like”. Stanford Center for Internet and Society. http://cyberlaw.stanford.edu/downloads/20120611-NetworkNeutrality.pdf.
Verizon. 2012. “Verizon Online Network Management Guide for Broadband Internet Access Services.” http://www22.verizon.com/terms/networkmanagementguide.
Viscusi, W. Kip, Joseph E. Harrington, and John M. Vernon. 2005. *Economics of Regulation and Antitrust*. 4th ed. Cambridge: The MIT Press.
Vogel, David. 1978. “Why Businessmen Distrust Their State: The Political Consciousness of American Corporate Executives.” *British Journal of Political Science* 8 (1): 45–78.
———. 1986. *National Styles of Regulation: Environmental Policy in Great Britain and the United States*. Ithaca: Cornell University Press.

Waddington, David. 1994. "Participant Observation." In *Qualitative Methods in Organizational Research*, edited by Catherine Cassell and Gillian Symon. London: Sage Publications.
Wade, William, and Christopher Forsyth. 2004. *Administrative Law*. USA: Oxford University Press.
Walden, Ian. 2012. *Telecommunications Law and Regulation*. Oxford: Oxford University Press.
Wallsten, Scott, and Stephanie Hausladen. 2009. "Net Neutrality, Unbundling, and Their Effects on International Investment in Next-Generation Networks." *Review of Network Economics* 8 (1) (March): 90–112.
Watal, Aparna. 2011. "A Co-Regulatory Approach to Reasonable Network Management." *Journal of Information Policy* 1: 155–173.
Watson, Neil. 2009. "ISPs vs BBC iPlayer: Missing the Point?" *The Register*, June 26. http://www.theregister.co.uk/2009/06/26/isps_against_bbc_iplayer_comment/.
Wein, Josh. 2007. "Comcast Accused Of Limiting BitTorrent Traffic." *Communications Daily*, August 21.
Weingast, Barry R., and Mark J. Moran. 1983. "Bureaucratic Discretion or Congressional Control? Regulatory Policymaking by the Federal Trade Commission." *Journal of Political Economy* 91 (5): 765–800.
Weinsberg, Udi, Augustin Soule, and Laurent Massoulie. 2011. "Inferring Traffic Shaping and Policy Parameters Using End Host Measurements." In *2011 Proceedings IEEE INFOCOM*, 151–155. Shanghai, China. http://www.bibsonomy.org/bibtex/2d2eb46e975a059a50578734a6256c409/dblp.
Weiser, Philip J. 2003. "Toward a Next Generation Regulatory Strategy." *Loyola University Chicago Law Journal* 35 (1): 41–86.
———. 2008. "The Next Frontier for Network Neutrality." *Administrative Law Review* 60 (2): 273.
———. 2009. "Institutional Design, FCC Reform, and the Hidden Side of the Administrative State." *Administrative Law Review* 61 (4): 675.
Weisman, Dennis L. 2010. "A 'Principled' Approach to the Design of Telecommunications Policy." *Journal of Competition Law and Economics* 6 (4): 927–956.
Weisman, Dennis L., and Robert B. Kulick. 2010. "Price Discrimination, Two-Sided Markets, and Net Neutrality Regulation." *Tulane Journal of Technology and Intellectual Property* 13: 81.
Wellman, Stephen. 2007. "Comcast Is Blocking More Than BitTorrent, including Lotus Notes." *InformationWeek*, October 22. http://www.informationweek.com/mobility/business/comcast-is-blocking-more-than-bittorrent/229214621.
Werbach, Kevin. 2005. "Breaking the Ice: Rethinking Telecommunications Law for the Digital Age." *Journal on Telecommunications & High Technology Law* 4 (1): 59.
Whitt, Richard S. 2004. "A Horizontal Leap Forward: Formulating a New Communications Public Policy Framework Based on the Network Layers Model." *Federal Communications Law Journal* 56 (3): 587–672.
Wiewiorra, Lukas. 2012. "The Economics of Net Neutrality: Implications of Priority Pricing in Access Networks". Karlsruhe Institute of Technology. http://digbib.ubka.uni-karlsruhe.de/volltexte/1000028857.
Wilkin, Mark. 2011. "World of Warcraft Issues *Updated 14/02/11*." February 14. http://community.virginmedia.com/t5/Broadband-down-your-phone-line/World-of-Warcraft-Issues-Updated-14-02-11/td-p/161857.
Williams, Christopher. 2007. "ISPs Hijack BBC in Tiered Services Push." *The Register*, August 13. http://www.theregister.co.uk/2007/08/13/isp_bbc_iplayer_neutrality/.
———. 2008. "BBC vs ISPs: Bandwidth Row Escalates as Tiscali Wades In." *The Register*, April 9. http://www.theregister.co.uk/2008/04/09/bbc_tiscali_iplayer/.
Wilson, Alexandra M., Lauren M. Van Wazer, Grace Koh, John P. Spalding, and Alysia M. Long. 2010. "Comments of Cox Communications, Inc. In the Matter of Preserving

the Open Internet, GN Docket 09-191; Broadband Industry Practices, WC Docket 07-52." http://apps.fcc.gov/ecfs/document/view?id=7020378743.
Wilson, James Q. 1975. "The Politics of Regulation." In *Social Responsibility and the Business Predicament*, edited by James W. McKie. Washington: Brookings Institution.
———. 1980. *The Politics of Regulation*. New York: Basic Books.
———. 1989. *Bureaucracy: What Government Agencies Do And Why They Do It*. New York: Basic Books.
"World of Warcraft Latency and Lag Spikes." 2011. February. http://community.virginmedia.com/t5/Up-to-30Mb-Setup-Equipment/World-of-Warcraft-latency-and-lag-spikes/m-p/349083/highlight/true#M151086.
Wu, Tim. 2003. "Network Neutrality, Broadband Discrimination." *Journal on Telecommunications & High Technology Law* 2 (1): 141.
———. 2004. "The Broadband Debate, A User's Guide." *Journal on Telecommunications & High Technology Law* 3 (1): 69–95.
———. 2007. "Wireless Carterfone." *International Journal of Communication* 1: 389–426.
Wu, Tim, and Lawrence Lessig. 2003. "Ex Parte Submission in CS Docket No. 02-52", August 22. http://apps.fcc.gov/ecfs/document/view?id=6514683884.
Wyse, Alistair. 2008. "Plusnet Broadband: Arbor customer summit" June 2. http://www.arbornetworks.com/images/Webcasts/plusnet/.
Yackee, Jason Webb, and Susan Webb Yackee. 2010. "Administrative Procedures and Bureaucratic Performance: Is Federal Rule-Making '"Ossified"'?" *Journal of Public Administration Research & Theory* 20 (2): 261–282.
Yoo, Christopher S. 2004. "Would Mandating Broadband Network Neutrality Help or Hurt Competition? A Comment on the End-to-End Debate." *Journal on Telecommunications & High Technology Law* 3 (1): 23.
———. 2005. "Beyond Network Neutrality." *Harvard Journal of Law & Technology* 19: 1.
———. 2006. "Network Neutrality and the Economics of Congestion." *Georgetown Law Journal* 94 (6): 1847–1908.
———. 2007. "What Can Antitrust Contribute to the Network Neutrality Debate?" *International Journal of Communication* 1: 493–530.
Zachem, Kathryn. 2008. "Comcast Ex Parte Re: In the Matter of Formal Complaint of Free Press and Public Knowledge Against Comcast Corporation for Secretly Degrading Peer-to-Peer Applications, File No. EB-08-IH-1518; In the Matter of Broadband Industry Practices; Petition of Free Press et al. for Declaratory Ruling That Degrading an Internet Application Violates the FCC's Internet Policy Statement and Does Not Meet an Exception for 'Reasonable Network Management,' WC Docket No. 07-52", September 19. http://apps.fcc.gov/ecfs/document/view?id=6520169715.
Zhang, Ying, Zhuoqing Morley Mao, and Ming Zhang. 2009. "Detecting Traffic Differentiation in Backbone ISPs with NetPolice." In *Proceedings of the 9th ACM SIGCOMM Internet Measurement Conference*, 103–115. IMC '09. Chicago, Illinois, USA: ACM. http://dl.acm.org/citation.cfm?id=1644893.1644905.
Zinman, Jack S., Gary L. Phillips, Paul K. Mancini, Jonathan E. Nuechterlein, and Lynn R. Charytan. 2007. "Repy Comments of AT&T Inc. In the Matter of Broadband Industry Practices, WC Docket 07-52." http://apps.fcc.gov/ecfs/document/view?id=6519558101.