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.