# 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.