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

7.1 Introduction

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

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

7.2 Rationales for Traffic Management

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

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

7.2.1 Performance and Cost

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

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

Differences Between Network Technologies and Markets

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

The US cable operators that deployed peer-to-peer management solutions were also motivated by the need to improve performance, but the nature of both their problems and chosen solutions was different. Most cable operators were focused on reducing upstream congestion, and they therefore chose peer-to-peer management solutions that limited the number of upstream connections. They were less concerned about controlling the overall growth of their backhaul bandwidth than with managing localized congestion problems in their access networks, which were the most expensive portions of the network to be constantly upgrading (one partial exception was RCN, which cited transit costs as one of

several rationales for its peer-to-peer management). The difference in approach between the two countries' operators was related not only to the physical differences between cable and DSL technologies, but to market dynamics: while both cable and DSL operators in the UK justified their decisions on the basis of downstream bandwidth costs, by and large neither cable nor DSL operators in the US did so.

Costs of Traffic Management Equipment

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

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

Some US cable operators that adopted the TCP reset solution viewed it as "very cheap." Compared to the hundreds of millions of dollars being spent to upgrade to DOCSIS 3.0, the DPI equipment cost one company "[on] the order of single millions of dollars to deploy across [the] whole network." Some ISPs realized savings by purchasing lower-end DPI

equipment with simpler processing technology than what might be found in higher-end switches or routers; one engineer described the equipment as "just a PC with 2 Gbit NICs [network interface cards] in it."

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

Beyond network size and architecture, ISPs differed as to their fundamental outlook about the long-term value of application-specific management. In contrast to the idea that DPI and bandwidth investment were interchangeable solutions to the same problem, ISPs that abandoned or never deployed application-specific traffic management felt that the value to be gained from a DPI deployment – whatever its cost – could not be rationalized compared to

the kinds of things that ISPs were accustomed to spending money on, including network expansion. One US engineer aptly summarized this idea:

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

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

Costs of Video Streaming

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

Tension Between the ISPs and the BBC

Well in advance of the launch of the BBC's iPlayer video service in 2007, BBC executives and staff had been meeting regularly with executives from the large ISPs to discuss the potential impact of iPlayer traffic on their networks. While those meetings were usually productive and collegial, at times the tension between the ISPs and the BBC over iPlayer

traffic growth spilled into the press. Several months prior to the iPlayer's official launch in 2007, a number of press reports cited ISPs (most notably Tiscali) threatening to place restrictions on iPlayer traffic unless the BBC agreed to pay for at least some of the bandwidth costs associated with the delivery of that traffic ("Net Firm Warns on Web Video Costs" 2007; Fluendy 2007; Murray-Watson 2007; Palmer 2007; Williams 2007). Although BT was mentioned in some of those reports, the company denied having any such complaints about iPlayer at the time (Williams 2007).

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

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

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

BT's Decision

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

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

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

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

7.2.2 Market Segmentation and Anti-Competitive Motivations

Beyond performance and cost control, the net neutrality literature points to two other potential rationales for discriminatory traffic management: segmenting broadband markets and undermining competing services. Some commentators have suggested that by using different traffic management strategies for different broadband products, ISPs can offer a

menu of products that may suit the needs of Internet users better than a suite of product choices that are all offered on a nondiscriminatory basis (Litan and Singer 2007; Marcus 2008; Renda 2008; Valcke et al. 2009; Weisman 2010; Yoo 2004; Yoo 2005). The market segmentation rationale implies that an ISP either offers some products with discriminatory traffic management and others without it, or that discriminatory traffic management is applied differently to different products, such that the applications that receive priority treatment vary between products, for example. Market segmentation also requires that these differences be advertised to potential customers – otherwise they would not have the information needed to select the broadband product most appropriate to their needs.

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

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

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

Market Segmentation in the UK

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

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

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

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

Peer-to-Peer and Video Streaming Management

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

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

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

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

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

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

7.2.3 Summary

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

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

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

7.3 Relationship Between Competition and Discrimination

The second key debate concerning net neutrality relates to whether competition can deter broadband providers from discriminating against applications. One broadly held view is that because discrimination can impair application performance, ISPs in competitive markets will be reluctant to take up discriminatory strategies for fear of losing customers (Becker, Carlton,

and Sider 2010; Cave and Crocioni 2007; Chirico, Haar, and Larouche 2007; Faulhaber and Farber 2010; Hahn, Litan, and Singer 2007; Nuechterlein 2009; Shelanski 2007). This premise is fundamental to both the EU telecommunications regulatory framework (OJL 337/11, 2009) and Ofcom's regulatory approach (Ofcom 2011b), and it is what drives regulatory interest in increased transparency about broadband operators' practices. Unsurprisingly, many ISPs have argued publicly that competition serves to deter discrimination (Atherton 2010; C. J. Brown and Boucher 2007; Casserly, Wallach, Alvarez, Walker, et al. 2008; ECTA 2010; Sky 2010; TalkTalk Group 2010a; Telefónica 2010).

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

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

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

7.3.1 Lack of Consumer Awareness and Understanding

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

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

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

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

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

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

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

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

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

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

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

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

7.3.2 Secondary Role of Traffic Management in Switching Decisions

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

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

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

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

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

With consumers reasoning this way, competition among broadband providers might have safeguarded applications that consumers already valued very highly – applications whose performance was satisfactory enough for consumers to enjoy. But for applications that have yet to obtain deal-breaker status, or for new applications that consumers have yet to even discover, discrimination against them would not necessarily be eradicated through market discipline, and could therefore serve to prevent those applications from ever performing well enough – and becoming so important to consumers – that they could trigger a switching decision. This is a specific manifestation of the argument put forth by Lessig and van Schewick that consumer preferences would not adequately reflect the importance of

application innovation: it is not that consumers do not care at all about less popular or established applications, but that they do not care enough to switch their ISP over them. These findings also comport with a number of formal economic models that show that discriminatory regimes can harm niche applications and their users while benefitting large, established providers and their users (Bourreau, Kourandi, and Valletti 2013; Guo, Cheng, and Bandyopadhyay 2012; Reggiani and Valletti 2012).

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

7.3.3 Competitive Discipline When Unwanted Customers Are Affected

Despite most consumers not understanding traffic management or caring about it enough to switch ISP, a small minority of UK Internet users did highly value having a broadband connection that does not discriminate against specific applications. Because peer-to-peer traffic management has been so common, heavy peer-to-peer users provide a useful minority group to study. Estimates of how many UK broadband users are consistent users of peer-to-

peer applications vary; on the network where participant observation was conducted, for example, significantly less than 1% of customers were considered heavy peer-to-peer users. Data collected for an Ofcom copyright infringement study in 2012 indicated that at least 4% of UK broadband users (about 800,000 people) had downloaded or shared files via peer-to-peer applications in the three months prior to the study, although the study did not differentiate between casual and consistent peer-to-peer users (Kantar Media 2012).

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

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

For some, it was "a question mark." One policy executive's musings about his company's heavy peer-to-peer customers revealed the inherent tension between wanting to maximize the size of the customer base and minimize expense: "To be fair, we've got quite a few actually. Because our policy is quite nice, actually. It's not that stringent. So unfortunately, we've got quite a few. But, okay, fair enough." For the ISP in question it was difficult to say whether

these customers were desired or not, but clearly they came to the ISP for having known about the lax traffic management policy. Other operators were more willing to cast these customers off. As one engineer explained, "you want to manage [peer-to-peer traffic] and if anything the attitude has been, well, if they get a bit annoyed and they don't like that, I mean, okay . . . they can go somewhere else."

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

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

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

This logic is supported experimentally in the finding of Sluijs et al. (2011) that if only a subset of consumers are well informed about the quality of their own broadband connections,

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

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

7.3.4 Barriers to Switching

Finally, although this study did not include data collection concerning specific switching costs, it is worth analyzing the extensive consumer switching data collected by Ofcom. Indeed, judging based only on the amount of time and resources that Ofcom has spent studying consumer switching behavior and improving relevant regulations, it is reasonable to

conclude that significant barriers to switching existed when the agency first took up the issue in 2006 and that barriers persisted in the years that followed. Between 2006 and 2012, Ofcom conducted three public consultations, ran two separate working groups with industry participation, hosted two workshops, and conducted or procured at least 15 separate studies with the involvement of seven different outside research firms on the topic of consumer switching.

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

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

7.3.5 Summary

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

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

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

7.4 Impact of Discrimination on Application Development

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

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

7.4.1 Difficulties with Traffic Management in Practice

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

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

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

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

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

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

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

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

Some of the attitudes expressed within the network operator community about these kinds of problems reveal a fundamental disconnect between the way that operators view the Internet as a platform for application development and the way that application developers and others might view it. If providers assume that "HTTPS is normally just used for banking," that "the problem is that some gaming companies use peer-to-peer signatures," or that "we have problems when . . . [gaming providers] start using the same protocols that the BitTorrents of the world use," as interviewees claimed, it puts the blame for the miscategorization on application developers rather than on ISPs for choosing to discriminate based on guesses

about the applications associated with particular traffic. This is exactly the opposite view of application providers (and net neutrality advocates), who assume that they can use any open, non-proprietary application protocol and have their applications perform in a predictable way. The root of the problem is not their choice of application protocol, but the fact that ISPs have chosen to interfere with particular application protocols for traffic management purposes. Some ISPs, having made significant investments in DPI platforms, came to view the application behavior as abnormal rather than viewing their own interference with application traffic that way.

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

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

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

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

No technology is perfect, and it would be unreasonable to expect that any traffic management system, whether application-specific or application-agnostic, could be deployed without operators later discovering bugs in need of fixing. But what differentiates application-specific management is the fundamental premise that network operators (and the DPI vendors they work with) can distinguish applications targeted for management from all others accurately enough to achieve some desired traffic management effect without prior coordination with

application providers. Taken together, the classification errors of the sort described in this section represent not just transient software bugs, but an underlying weakness in the notion that any network-based classification engine can stay ahead of the multitude of ongoing, uncoordinated innovations and application changes taking place at the edges of the network. Furthermore, discriminatory traffic management turns network operators into arbiters of application-specific network performance. When the arbiters make mistakes, applications suffer the consequences differentially. As the next section demonstrates, operators have sought to correct these errors by involving application providers in their decision-making – creating a cost to application innovation that would not exist were it not for application-specific management.

7.4.2 Engaging Developers to Resolve Traffic Management Problems

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

As perceived by network operators, the willingness of application developers to help ISPs sort out these problems appears to have varied, both over time and from developer to developer. While some companies were described as being "happy to speak" to individual ISPs, others were "quite difficult to deal with." One interviewee noted that some developers seemed to react to user complaints with the "kind of attitude where 'it's not our problem, go

and talk to your ISP'" because they were "fundamentally anti-traffic-management as a principle." Upon recognizing the latest in a series of problems related to the classification of World of Warcraft traffic, a Virgin Media support forum manager exposed the reticence of some gaming companies:

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

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

different technologies and vendors. Small developers simply might not have had the resources to spend on that level of engagement, as one ISP policy executive readily admitted:

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

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

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

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

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

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

Such are the visions of a discriminatory future: where a multitude of application developers are in constant contact with ISPs as their applications get updated and as new applications get introduced; where application developers expend resources to understand how DPI classifiers work so that they can help ISPs and vendors craft application signatures; and where developers engage directly with the technology companies whose equipment is used to degrade application performance. While these visions may never be fully realized, the seeds of this future were at least partially sown when application developers began to engage with ISPs to resolve misclassification issues, incurring the kinds of costs to application

development and innovation that some academic and policy stakeholders had predicted. For application developers, particularly small upstarts with few resources to spare, it was a lose-lose situation with barriers in both directions: they could either sink time and resources into solving the problems that ISPs created for them for the sake of restoring their own products' performance, or decline to engage in the hope that ISPs would find their own solutions or shift to nondiscriminatory strategies before their product performance became a casualty of imperfect DPI technology.

7.4.3 Commitment to Discrimination Despite its Impact on Innovation

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

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

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

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

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

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

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

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

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

7.4.4 Summary

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

Application developers have choices about how to resolve these problems, but all of them raise barriers to further application development and innovation. Some operators have sought developers' ongoing assistance in improving traffic classification, and some developers have complied. The alternative is for developers to see the performance of their applications suffer on particular networks. Operators are aware that this puts application developers in a bind, but the benefits they derive from perpetuating application-specific management outweigh

their concerns for how it may be affecting the environment for application development. They are willing to take certain steps to minimize the impact of application-specific management on developers, but not so willing that they would rethink their overall approach to traffic management.

7.5 Conclusion

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

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

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

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

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