

**From Network Neutrality to Data Neutrality:
A Techno-Economic Framework and Research Agenda¹**

Robert Easley*, Hong Guo*, Jan Krämer**

*University of Notre Dame, Notre Dame, IN, USA, {reasley, hguo}@nd.edu

**University of Passau, Passau, Germany, jan.kraemer@uni-passau.de

Forthcoming: Information Systems Research

Abstract

The Internet has assumed a central role in the global economy facilitating commerce and communication, and is thus central to many areas of IS research. In particular, IS researchers played a critical role in the academic discourse on net neutrality, which has recently informed new regulatory frameworks in the US and Europe. We discuss and categorize the various issues and key trade-offs that are still being debated in the context of net neutrality, and identify open research questions in this domain. Based on these insights, we argue that net neutrality, which is concerned with a gatekeeper at the infrastructure level, may just be part of a larger debate on data neutrality, where the gatekeeper may rather control a software platform. We provide several examples of potential data neutrality issues and generalize the key trade-offs in the context of a proposed four-step framework for identifying and organizing promising areas of IS research on data neutrality.

Keywords: Net neutrality, public policy, Internet governance, network management, data management, data neutrality, platforms, gatekeeping

¹ All the authors contributed equally and are listed in the alphabetical order.

1. Introduction

Net neutrality (or network neutrality) issues have been at the center of a worldwide Internet policy debate during the past decade. In general, net neutrality refers to a network design principle stating that all data packets should be treated equally regardless of their content, sites, and platforms (Wu 2003). The focus of the net neutrality debate is concerned with the various network management practices that last-mile network providers (NPs) should be allowed to pursue, being the central gatekeepers between consumers and content providers (CPs) at the broadband infrastructure level. Both in the US and in Europe recently new regulatory frameworks on net neutrality were adopted, which limit the NPs freedom in choosing their network management practices, but also remain vague with respect to the fine line that delineates neutral from non-neutral practices. To this end, we present a classification of the many network management practices that can affect the consumers' Quality of Service (QoS) on the Internet, articulating whether they would be considered network neutral, while focusing on the particular question of who pays for any incremental improvements. We thus map the landscape of the current debate, which enables us to highlight where IS research has made a significant contribution to our understanding, and which important questions remain unaddressed.

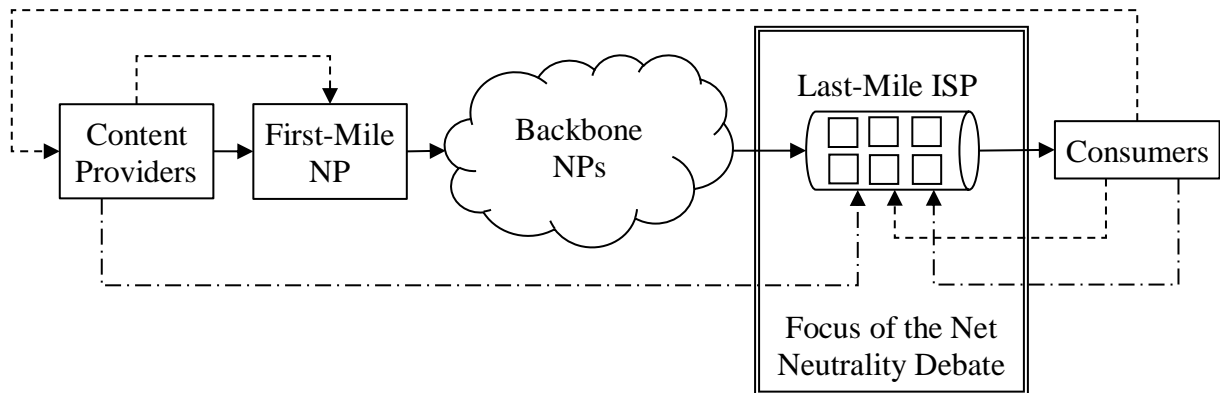
Based on these insights, we then broaden the perspective by highlighting that similar gatekeepers exist at the software level, which also control the flow of information between consumers and CPs. This allows us to identify other domains in which a related set of issues arise, which we refer to as “data neutrality”, and which we deem a promising area of research for IS scholars. We extend our analysis of net neutrality issues to this general case by proposing a four-step research framework to provide IS researchers guidance in organizing research

programs on data neutrality. We then illustrate its application to several emerging issues in this broader data neutrality area.

2. Net Neutrality

2.1. Focus of the net neutrality debate

The first thing one should understand about the net neutrality debate is that it applies only to the “last-mile” NP, i.e., an Internet Service Provider (ISP) paid by a consumer for the connection to their Internet-connected device(s), as shown in Figure 1. Confusion can arise here since a typical Internet consumer also uploads content to the Internet using the same ISP, so a useful further clarification is that the domain of the debate focuses on the downloading of content. A further complication involves digital convergence, either of devices, such as smart phones that merge telephony and data consumption; of networks, such as cable, telephone or wireless networks which are all interconnected and may all serve as part of the Internet; or of firms that may play multiple roles, for example as both NPs and CPs.



Notes: Solid lines represent data flows, whereas dashed lines represent payment flows. Short-dash lines represent current payment flows, and long dash-dot lines represent potential additional payment flows that are examined in the net neutrality debate.

Figure 1: The Internet Ecosystem

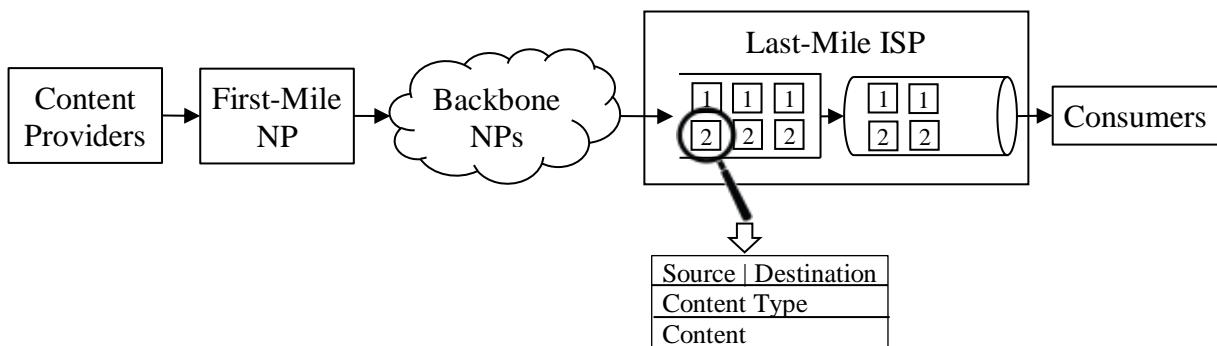
Figure 1 presents a stylized abstraction of the Internet, designed to isolate the focus of the net neutrality debate. It shows data and payment flows in the Internet ecosystem, encompassing consumers, who request data from the CPs; first-mile and backbone NPs, which receive and relay the data of the CP; and the last-mile ISP, which is the focus of the net neutrality debate because it controls the final and crucial part of the transmission system to consumers (depicted as a “pipe”) through which the data packets (depicted as squares) coming from various CPs are sent. CPs pay a first-mile NP to deliver their content to the Internet backbone, and consumers pay a last-mile ISP for delivery of requested content, as represented with short-dash lines.

The consumer may pay the CP as well, either directly via a subscription fee (e.g., as in the case of Netflix) or indirectly through viewing and clicking advertisements (e.g., as in the case of Google). However, the potential payment of last-mile ISPs by CPs is a contentious aspect of the net neutrality debate because the payment would be for prioritization or to prevent blocking of their data packets, which would require packet discrimination and would thus violate net neutrality. Many last-mile ISPs have argued that the high volume of delivery from a particular CP alone justifies a demand for payment, regardless of prioritization, and some have even succeeded (see Ray and Leach 2013 for the case of Orange and Google), but generally CPs argue they already pay for delivery via a chain of peering and transit arrangements between the various NPs involved in transmitting the data packets through the Internet backbone. Similarly, ISPs could implement network management practices that would allow them to collect extra payments from consumers, e.g., for prioritizing their data or for lifting an imposed data cap, as represented with long dash-dot lines. In this context the net neutrality principle has become known as a *zero-price rule*, which bans the last-mile ISP from leveraging extra payments from either the CP- or the consumer-side (Schuett 2010, Krämer et al. 2013). However, on a more technical level, the

net neutrality principle demands a *no-discrimination rule*, which bans the use of certain network management practices within the last-mile network.

2.2. Network management practices

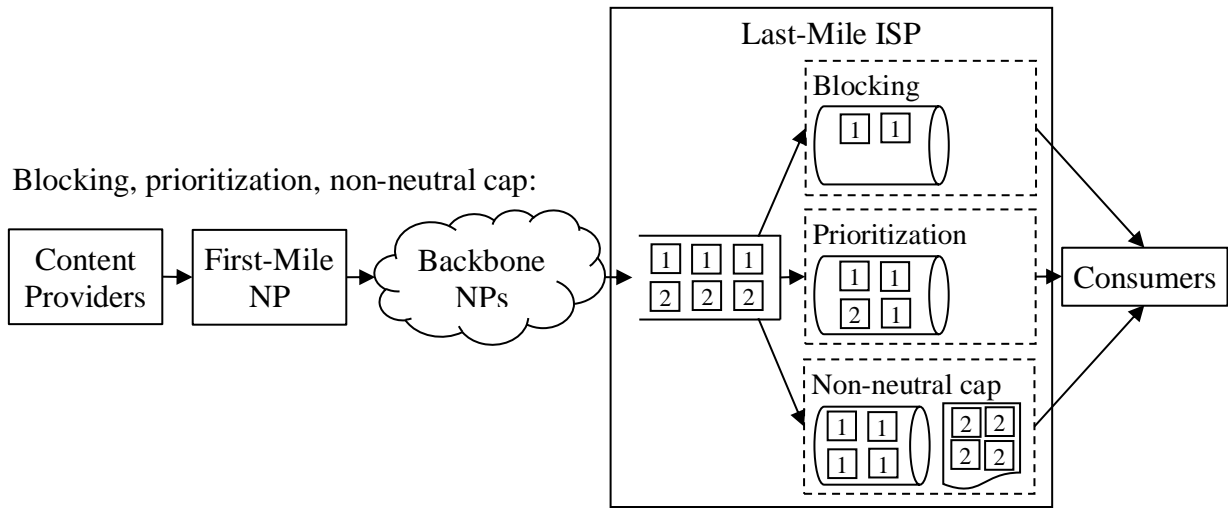
Figure 2 expands the representation of the last-mile ISP to illustrate the arrival of data packets, which can potentially be inspected for specific information regarding source, destination, content type, and content. However, when the net neutrality principle applies, last-mile ISPs must not act upon this information and must transmit packets through their network in a first-come-first-served manner, using only the destination information needed to pass them through in the order received. This is shown in Figure 2, where packets may arrive from different CPs (e.g., 1 and 2 as Google and Microsoft), be requested by different consumers (e.g., 1 and 2 as heavy and light users), contain different content types (e.g., 1 and 2 as email and Voice-over-IP, or VoIP), or contain different content (e.g., 1 and 2 as text with and without forbidden words), but are nonetheless being passed on in a net-neutral manner. Thus, the no-discrimination rule applies and, consequently, the zero-price rule also applies, as the last-mile ISP has no means to leverage additional payments from either CPs or consumers in this way.



Notes: Packet numbers 1 and 2 can be interpreted to represent a different source (i.e., CP), destination (i.e., consumer), content type, or content.




Figure 2: The Net Neutrality Principle

Figure 3 illustrates various observed network management practices that affect QoS for consumers, separating those that violate the net neutrality principle (i.e., the no-discrimination rule) from those that do not. The net neutrality violations illustrated include *blocking*, in which data packets are not delivered to the consumer based on the source, the destination, the content type, or the content itself. For example, a mobile phone provider could block content from a competing service (source); a capacity-constrained ISP could block data packets requested by heavy users (destination) or bandwidth-consuming HD video traffic (type); or an ISP could block content of a politically sensitive nature (content).



Network management practices that violate the net neutrality principle

Network management practices that do not violate the net neutrality principle



Notes: Packet numbers 1 and 2 can be interpreted as representing different source, destination, content type or content. Practices that violate net neutrality principles are illustrated as blocking 2, prioritizing 1 over 2, and counting 2 (but not 1) toward a cap. Practices that do not violate these principles include counting all packets toward a cap, and achieving faster delivery to the ISP of 1 (relative to 2) via a CDN or paid peering arrangements.

Figure 3: Network Management Practices

An ISP could also choose to *prioritize* delivery of certain data packets over others, here illustrated with prioritizing data packets of type 1. This is possible as data packets are generally first stored in router queues before they are transmitted. Prioritization in the technical sense then means that some packets are moved ahead in the router's queue, i.e., packets are not handled on a first-come-first-served basis anymore. Prioritization can result in a near-blocking of content, because some types of content are effectively blocked if they are slowed to the point of impairing normal consumption. Note that although some forms of packet discrimination may be desired by the consumer, for example prioritizing the delivery of Internet Protocol television (IPTV) over email packets, such discrimination remains nonetheless a technical violation of net neutrality.

ISPs may also set usage caps for their consumers, such that when total usage reaches a threshold, consumers or CPs need to pay extra for additional data. *Non-neutral caps* result when caps are applied to certain data packets (e.g., coming from certain CPs) but not others, and thus violate the net neutrality principle. We illustrate the case where data packets of type 2 are counted against the consumer's data allowance cap, but packets of type 1 are not. Thus, as the cap for type 2 is reached, those packets are not forwarded anymore. This practice is also known as *zero-rating*, *toll-free data*, or *sponsored data* (of CP 1). The first non-violating network management practice also involves usage caps, but *neutral caps* refer to usage caps that apply to all data packets equally. In contrast with the non-neutral caps (i.e., data caps with zero-rating), here the data packets of all CPs are counted against the cap, and thus, as the cap for that consumer is reached, no further data packets are forwarded to the consumer.

Other network management practices involve bypassing some or all of the Internet backbone NPs to achieve what we refer to as *packet precedence*, as distinct from prioritization, because it does not involve discrimination in the data transmission within the last-mile ISP's

network, but simply depends on arranging for packets to arrive at the last-mile ISP ahead of other packets. Independent companies like Akamai, Limelight or Level 3 operate *content delivery networks* (CDNs), charging CPs for hosting their content on servers closer to the last-mile ISPs, sometimes residing even within the last-mile ISP's network. Finally, some CPs may directly pay and connect to last-mile ISPs through *paid peering*. For example, Netflix's paid peering arrangement with Comcast improves the video streaming experience of its own customers (Yu 2014). In Figure 3, the CDN and paid-peering packets arrive with a higher QoS to the last-mile ISP, and thus also to the consumer, as they bypass potential congestion on the way, but because the ISP treats the packets the same once they arrive this does not violate the net neutrality principle.

Based on the previous discussion, we now propose a classification that characterizes network management practices based on two dimensions – the specific network traffic management mechanisms discussed above, and different payment structures. These dimensions correspond to the notion of net neutrality as a *no-discrimination rule* and as a *zero-price rule*. Some network management practices are triggered by the ISPs themselves, i.e., they involve no *additional* payment by CPs or consumers to the last-mile ISP (beyond the standard Internet access fees) and thus conform with the zero-price rule. Other network management practices are paid for by CPs and yet others are selected and paid for by the consumers.

As shown in Column 2 of Table 1 and detailed above, network management practices range from packet discrimination based on content provider, content type, content, or destination; to data caps with and without zero-rating; and then on to content delivery networks, and peering. Whenever a network management practice involves packet discrimination *within* the last-mile ISP's transmission system (as exemplified in the top part of Figure 3), the no-

discrimination rule is violated and consequently this network management practice is classified as non-neutral. When a network management practice occurs *outside* of the transmission system of the last-mile ISP (as exemplified in the lower part of Figure 3), it does not violate the no-discrimination rule and thus the net neutrality principle is generally maintained.

Table 1: Network Management Practices

<i>Net neutrality violation?</i>	<i>Network management practices</i>	<i>Neither side pays extra to ISP</i>	<i>Content provider side pays extra to ISP</i>	<i>Consumer side pays extra to ISP</i>
<i>Non-neutral:</i> Network management WITHIN the transmission system of the last-mile ISP	<i>Packet discrimination based on</i> <ul style="list-style-type: none"> • content provider or source • content type • content • consumer or destination 	<i>ISP-driven packet discrimination, e.g.,</i> <ul style="list-style-type: none"> • Blocking web sites • Throttling or blocking certain protocols (e.g., P2P, VoIP) • Prioritizing certain protocols (e.g., IPTV) 	<i>CP-driven packet discrimination, e.g.,</i> <ul style="list-style-type: none"> • CP pay termination fee to make content available at ISP • Pay-for-priority arrangements between CP and ISP 	<i>Consumer-driven packet discrimination, e.g.,</i> <ul style="list-style-type: none"> • VoIP option in mobile broadband networks • Pay-for-priority arrangements between consumer and ISP
	<i>Data cap with managed service exemption</i>	ISP exempts only its own content offerings	Certain CPs pay for exemption as a managed service (zero-rating)	Consumers pay for certain CPs to be exempt from cap
<i>Neutral:</i> Network management OUTSIDE the transmission system of the last-mile ISP	<i>Data cap without exemption</i>	Data cap cannot be lifted	One or many CPs pay to lift consumers' data cap for all CPs	Consumers buy additional data allowance
	<i>Packet precedence achieved using</i> <ul style="list-style-type: none"> • CDNs • Peering 	<ul style="list-style-type: none"> • Use of independent CDN by ISP • Peering 	<ul style="list-style-type: none"> • CP pays ISP's in-house CDNs • Paid Peering 	<ul style="list-style-type: none"> • Consumers subsidize ISP's in-house CDN

However, as these neutral network management practices may affect QoS, they may still have an effect on the contractual relationship between the last-mile ISP and CPs or consumers such that the zero-price rule may in fact not apply. We adopt the common viewpoint and denote the practices that obey the no-discrimination rule as neutral network management practices, though they violate the zero-price rule. Note that our classification therefore demonstrates that the no-discrimination and zero-price dimensions of net neutrality may not always be aligned.

To exemplify the applicability of this classification, we briefly mention and locate some concrete examples of well-known network management practices and highlight whether they would be considered net-neutral or not.

2.2.1. Packet discrimination

These network management practices include all the types of packet discrimination that occur inside the last-mile ISP's transmission system. Depending on whether additional payments from its consumers, the CPs or neither side influence the ISP's decision concerning packet discrimination, we speak of consumer-driven, content-provider-driven or ISP-driven packet discrimination. All three of these network management practices violate the no-discrimination rule and hence, are denoted as non-neutral:

Consumer-driven packet discrimination: In this case consumers pay extra to the ISP in order to gain access to or prioritize content from their preferred CPs. For example, to date many mobile broadband ISPs prohibit their customers from using unaffiliated VoIP services unless they pay extra to the ISP. Consumers may also be given the option of paying to prioritize certain types of content (e.g., VoIP data) or certain content (e.g., videos of their favorite sports events) or any of their data flows relative to other concurrent users of the network.

Content-provider-driven packet discrimination: These non-neutral packet discrimination practices are implemented by the ISP to provide a better QoS to those CPs that pay additional fees. In extreme cases, CPs may need to pay a "termination fee" simply to make their content available to consumers. For example, Google pays France Telecom-Orange for the traffic that it sends to their network, because Orange threatened to cut Google off the African market, where it enjoys significant market power in Internet subscriptions (Ray and Leach 2013). CPs may also be given an option to pay for prioritization, i.e., preferential delivery of their content. This type

of non-neutral network management practice has been a key focus of discussion by the FCC and other regulatory agencies around the world (Kang 2014).

ISP-driven packet discrimination: These non-neutral network management practices are undertaken by the ISP without specific financial compensation from either CPs or consumers. Some cases may involve a general network management effort for congestion-sensitive content, such as live television (IPTV) which may affect many customers. Some may actually be mandated for technical or legal reasons, such as prioritizing emergency calls over IP or blocking illegal or harmful content. Other prominent examples include Telus, a Canadian telecommunications company, which in July 2005 blocked its Internet subscribers from accessing the CP Voices for Change, a website run by striking union members (CBC 2005). Similarly some ISPs (e.g., Comcast, Bell Canada, etc.) have reportedly throttled peer-to-peer (P2P) traffic (Mueller and Asghari 2012) and others (e.g., Madison River) have interfered with VoIP traffic (FCC 2011).

2.2.2. Data-cap-based network management

These network management practices are based on controlling consumers' total usage through implementing data caps, and may or may not violate the net neutrality principle. On the one hand, data-cap-based practices are non-neutral if the traffic to and from certain CPs or types are not counted towards the cap. For example, Comcast exempts some of its Xfinity on-demand videos through Xbox 360 from its 250GB broadband cap (Orland 2012). Certain CPs may then be given the option of paying extra to the ISP in order to become a zero-rated "managed service" which is exempt from the cap, as is the case, for example, with Airtel Zero in India (Russell 2015). Of course, it is also possible that consumers would pay extra to assure that certain CPs are exempt from the cap. On the other hand, data-cap-based practices are neutral if all packets,

independent of their source, content or type, are considered the same and thus are counted towards cap. This holds even if some consumers or CPs may make additional payments to exceed or extend their data allowance, i.e., although the zero-price rule is violated.

2.2.3. Packet precedence

In order to improve their consumers' QoS, CPs may make special arrangements to get their packets to arrive sooner and with higher reliability at the ISP's last-mile transmission system. All of these practices are considered neutral although some may involve a violation of the zero-price rule. CDNs and peering are two such arrangements (McMillan 2014). Some big CPs operate their own CDNs. Google Global Cache, Google's own content delivery platform, hosts Google's popular content such as YouTube at the edge of last-mile networks. Netflix also has its own CDN built on capacity leased from Cogent (a tier 1 ISP). In these cases, CPs do not pay extra to the ISP, but possibly to some third-party, in order to improve the QoS for consumers. However, the last-mile ISP itself may offer an in-house CDN service, as Comcast has done (Brodkin 2014), so that CPs willing to pay extra can have their content hosted there. In addition, CPs may get direct network connections to last-mile ISPs through peering in order to achieve higher precedence for their data packets. Although peering can be based on mutual agreement without direct financial compensation, recently last-mile ISPs and CPs have engaged in paid peering, where CPs make additional payments to the ISP for peering arrangements. For example, Netflix recently struck paid peering deals with Comcast and Verizon (Gustin 2014, Rogowsky 2014). Finally, it is also conceivable that consumers, or consumer groups in the same community, may pay to subsidize CDNs or peering for certain content.

2.3. Key trade-offs and insights for net neutrality regulation

Research on net neutrality has grown considerably over the past decade, with a strong focus on using economic models to examine the impact of various forms of CP-driven packet discrimination under a variety of assumptions. We do not attempt a full review of the still evolving net neutrality literature here, and refer interested readers to excellent reviews elsewhere (Schuett 2010, Faulhaber 2011, Krämer et al. 2013, Greenstein et al. 2016). Instead, we focus on summarizing the key trade-offs that have been identified in IS and economics papers that have advanced the debate thus far. These trade-offs also correspond to the key issues that policymakers need to consider when deciding to what extent net neutrality regulation is warranted. In addition, for the papers mentioned below, in Table A1 in the appendix we also offer a more detailed summary of the main assumptions (i.e., the market conditions considered) and findings, and how these relate to the key trade-offs that we discuss.

2.3.1. Trade-off 1 – Affiliation: What are the incentives of the vertically integrated ISPs to interfere with unaffiliated content?

One of the most pertinent issues in the public debate on net neutrality is the interference of a vertically integrated ISP (i.e., an ISP integrated with a CP) with unaffiliated content. The vertically integrated ISP has the ability to interfere either directly by blocking the unaffiliated content or slowing it down, or indirectly by prioritizing its affiliated content (recall the examples from ISP-driven packet discrimination). However, the incentives for ISPs to degrade unaffiliated content are not so clear. On the one hand, the integrated ISP can internalize the value created by the affiliated CP but not that created by unaffiliated CPs. On the other hand, content (both affiliated and unaffiliated) is complementary to the ISPs' Internet access services, with the availability of more and better content increasing the value of its access services. This trade-off

was analyzed in a number of papers (see, e.g., Guo et al. 2010, Broos and Gautier 2016, Dewenter and Rösch 2016), that characterize the conditions under which the ISP has an incentive to discriminate against or even block rival content. As can be seen from Table A1, these papers find that a vertically integrated ISP will not necessarily block or degrade unaffiliated content. On the one hand, the ISP will generally accommodate unaffiliated content if it is deemed valuable to Internet customers (Broos and Gautier 2016, Dewenter and Rösch 2016), because such content serves as a one-way essential complement (Chen and Nalebuff 2006) that will drive up the value of the Internet access, which drives up profits for the vertically integrated ISP. In extreme cases, the ISP may even prioritize unaffiliated content over its own (Guo et al. 2010). On the other hand, if the unaffiliated CP offers content that is similar to that of the vertically integrated ISP or if it is likely to diminish network effects for the integrated ISP's content, then blocking of the rival CP may occur (Dewenter and Rösch 2016).

2.3.2. Trade-off 2 – Compatibility: Does net neutrality lead to more or less incompatibility of content or Internet fragmentation?

A related but distinct trade-off in the net neutrality debate concerns the ISPs' incentives to fragment the Internet, such that not all content can be reached from each ISP. In other words, a CP may choose to be incompatible with other ISPs. For example, absent net neutrality, this may occur because a CP does not want to pay the termination fee at each ISP and therefore chooses not to be available at that ISP. However, fragmentation due to incompatibility of CPs and ISPs may also occur in agreement with (and possibly compensated by) an ISP, who seeks to attain a competitive advantage over other ISPs by offering its customers exclusive access to a CP (Lee and Wu 2009). It has been argued that net neutrality regulation would prevent such fragmentation and consequently the conditions under which CPs and ISPs choose to be

compatible or not have been analyzed in detail by Kourandi et al. (2015) as well as D'Annunzio and Russo (2015) in the context of the net neutrality debate. They show under various market conditions (see Table A1) that net neutrality regulation is neither a necessary, nor a sufficient condition to prevent Internet fragmentation. However, Internet fragmentation is, everything else being equal, less likely to occur under net neutrality regulation.

2.3.3. Trade-off 3 – Innovation: Does net neutrality lead to more or less innovation or variety of content?

Perhaps the most intricate issue of the net neutrality debate is which network management regime will stimulate content innovation and thus increase content variety in the long run. On the one hand, different services have different requirements for QoS and thus packet discrimination may allow innovation of content that relies on QoS. On the other hand, allowing paid prioritization may favor financially strong firms that are not necessarily the most innovative. Although innovation is difficult to model formally and there is certainly no standard way in doing so, several papers have analyzed this trade-off in more detail (see, e.g., Hermalin and Katz 2007, Krämer and Wiewiorra 2012, Bourreau et al. 2015, Reggiani and Valletti 2016, Guo and Easley 2016). Different approaches yield conflicting results (see Table A1 for a comparison of the market conditions that these papers consider as well as their key findings). Guo and Easley (2016), for example, identify cases in which net neutrality regulation is likely to lead to more content innovation or variety than packet discrimination. Other studies find that while net neutrality allows a level playing field among CPs with equal requirements for QoS, it puts more congestion-sensitive CPs with higher QoS requirements at a disadvantage. Under packet discrimination, the ISP can balance the QoS requirements of the CPs better, which encourages

entry of congestion-sensitive CPs that would not have been able to operate profitably under a net neutrality regime.

2.3.4. Trade-off 4 – Investment: Does net neutrality foster or stifle investment incentives for the ISPs?

Investments in network infrastructure are not only believed to spur economic growth in general (see, e.g., Röller and Waverman 2001, Czernich et al. 2011), but also to have a positive effect on the Internet economy in particular. Proponents of net neutrality have argued that packet discrimination practices would entail less investments by the ISP, because the ISP would have an incentive to keep the network capacity scarce in order to maintain a high value for prioritization (Krämer et al. 2013). Conversely, net neutrality opponents have argued that non-zero pricing under packet discrimination would shift some of the CPs' profits to the ISP, who would then have the means and the incentive to expand network capacity in order to accommodate more congestion-sensitive CPs. Formal analysis of this issue shows that in some cases (see Table A1 for a comparison) infrastructure investments are indeed likely to be lower under net neutrality (Krämer and Wiewiorra 2012, Economides and Hermalin 2012, Bourreau et al. 2015), while in others (see, e.g., Choi and Kim 2010, Cheng et al. 2011), the ISP's incentive to keep network capacity scarce outweighs its incentive to reduce network congestion through infrastructure investment.

2.3.5. Trade-off 5 – Welfare: Who are the winners and losers of net neutrality? Is the economy as a whole better or worse off?

Finally, factoring in all possible trade-offs, particularly including those mentioned above, the main question for policymakers is whether net neutrality would increase or decrease social welfare and, more precisely, who (CP, ISP, consumers) will likely gain and who will likely lose

from this network regime. While each of the papers mentioned above only considers a subset of the described trade-offs, many times focusing on just one of them, almost all papers explicitly evaluate the considered trade-offs by their effect on social welfare. From Table A1 in the appendix, which highlights the trade-offs and market conditions considered in the key papers, it is apparent that in many papers the welfare results are mixed, identifying cases in which either net neutrality or packet discrimination yield higher social welfare. Some papers, in particular those that explicitly take congestion or long-term infrastructure investments into account, come to a negative conclusion on the impact of net neutrality regulation. This is due to the fact that net neutrality regulation can stifle infrastructure investments (see Trade-off 4) and may lead to an inefficient allocation or inefficient traffic inflation (see, e.g., Cheng et al. 2011, Peitz and Schuett 2016). Others identify specific scenarios in which net neutrality may improve social welfare (see, e.g., Economides and Tåg 2012, D'Annunzio and Russo 2015). On a more nuanced note, ISPs tend to be better off when packet discrimination is allowed, and CPs tend to be worse off (see, e.g., Krämer and Wiewiorra 2012), because non-zero pricing tends to shift welfare from CPs to ISPs. Consumers may be better off in the absence of net neutrality, when the ISP is able to offer lower Internet subscription fees since, in the absence of the zero-price rule, it must not reap all surplus from the consumer side (see, e.g., Economides and Tåg 2012). This can be especially important in developing countries to increase Internet take up. These welfare results generally also hold when both CPs and ISPs are in competition (Guo et al. 2016).

2.4. Prospects for future IS research on net neutrality

Strikingly, almost all papers mentioned above and listed in Table A1 in the Appendix focus on either ISP-driven or CP-driven packet discrimination scenarios, whereas research is scarce (with exceptions as noted) on user-driven packet discrimination (see Krämer and Wiewiorra 2015),

data caps (see Economides and Hermalin 2014) and QoS practices that occur outside the ISP's transmission system, i.e., paid peering (see Coucheney et al. 2014) and CDNs (see Chiang and Jhang-Li 2014). These network management scenarios are also not explicitly addressed by the existing legislation in the US and Europe, which in itself may encourage further research in this area. Where legislation has been adopted it also poses some interesting implementation and design challenges, such as how permitted network management practices should best be presented to consumers in order to achieve transparency, which is explicitly required by the regulations.

The conclusions reached with respect to the evaluations of the trade-offs are often dependent on whether and how network congestion is explicitly modeled. Usually a M/M/1 queuing model is assumed, but the realism of this assumption with respect to actual Internet traffic today, as well as the impact of alternative assumptions on a model's outcome, are yet unexplored, although IS researchers have considered related issues. For example, Johar et al. (2011) analyze how congestion impacts consumers' incentives to use an Internet service (in their context a peer-to-peer sharing service) and Masuda and Whang (2006) study optimal tariffs for telecommunications services to examine their impact on congestion.

Another limitation is that all existing papers consider QoS, which involves technical differences (e.g., reliability, congestion, etc.) in traffic flows coming from different CPs. It is crucially important to be able to distinguish QoS from QoE (quality of experience, i.e., the behavioral perception of transmission quality) in this context. QoE, however, also depends on the type of content delivered (e.g., for the same level of QoS, email is likely to have a higher QoE than real-time video), and users' preferences and expectations. It would be very valuable to have a measure that enables researchers to compare QoE across applications, such that the winners

and losers of net neutrality regulation could be more clearly identified. However, such a measure does not yet exist², although IS researchers have vast experience in measuring user's perceptions when interacting with IT, as well as the impact of IT usage on the systems' performance. For example, McKinney et al. (2002) as well as Chen and Hitt (2002) link customer's perception of "system quality" (a measure akin to QoE) to consumer satisfaction and retention.

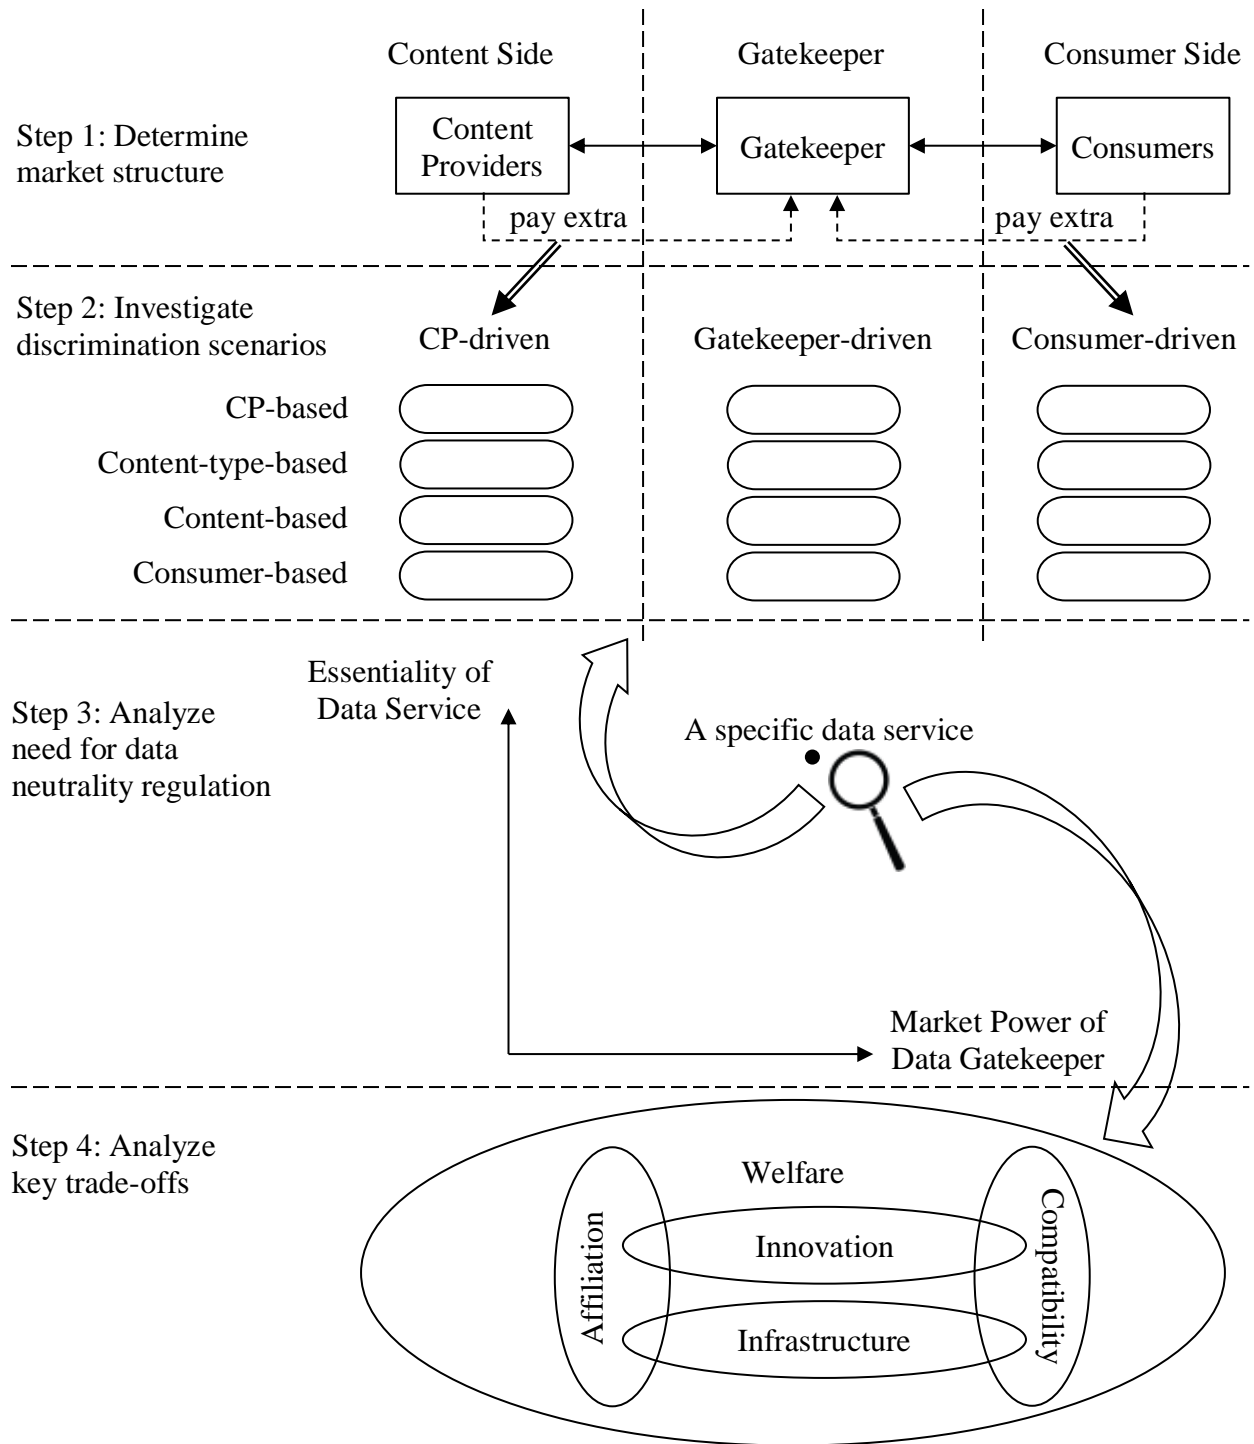
Existing research on net neutrality has methodologically almost exclusively relied on analytical economic modeling. Now, that both the US and Europe have adopted net neutrality legislation, whereas other countries with advanced broadband infrastructure have not (e.g. Japan or South Korea), comparative empirical research has become feasible. Such research could corroborate the (sometimes contradictory) theoretical results and yield significant new insights for the evaluation of the key trade-offs, e.g., on consumer behavior and CP innovation. Empirical IS research can provide significant new insights from investigation of ISPs' reactions to net neutrality regulation, e.g., with respect to tariffs (e.g., zero-rated plans) and infrastructure investments. Here, net neutrality is linked to issues of the first order digital divide through its potential impact on the affordability of access to the Internet. This has been a perennial issue for IS research (see for example, Dewan and Riggins 2005, Hsieh et al. 2008, Dewan et al. 2010, Hsieh et al. 2011), but currently that research is not linked to the net neutrality debate. For example, if packet discrimination leads to more infrastructure investments (see Trade-off 4) and lower prices for consumers (see Trade-off 5), developing countries may have lower incentives to adopt net neutrality regulation than developed countries.

² In fact, a recent workshop jointly sponsored by NSF and FCC calls for more collaborative research for QoE (Bustamante et al. 2015).

3. Towards an Agenda for Research on Data Neutrality

3.1. A techno-economic framework for research on data neutrality

The net neutrality debate focuses on neutrality with respect to how ISPs, as *gatekeepers* of an essential platform that is able to control the data flow between CPs and consumers, may handle the data that passes through their last-mile networks, i.e., whether they prioritize, throttle, cap or block it. However, the debate on net neutrality, although multifaceted in itself, may just be the onset of a larger debate on *data neutrality*, i.e., how other platform gatekeepers in the Internet that connect CPs with consumers may exert control over the data services that run over their platform. Generalizing from the structure implemented above for our discussion of net neutrality, we present in Figure 4 a techno-economic framework for structuring research on data neutrality issues.



Notes: In Step 1, solid lines represent data flows, whereas dashed lines represent payment flows.

Figure 4: A Techno-Economic Framework for Research on Data Neutrality

The framework encompasses four important steps which will be outlined briefly here, and then explored in more depth in the context of several emerging research domains. In Step 1, the critical issue is to determine the market structure, i.e., to identify the three players – the gatekeeper and the two parties that it connects – and trace the flow of payments, as we did in Section 2.1 for net neutrality. Step 2 involves investigation and classification of different drivers and types of discrimination practices in play as seen in Section 2.2 in the context of net neutrality. Step 3 concerns the need for neutrality regulation. This step was not directly addressed with net neutrality above since the importance of net neutrality regulation can be considered self-evident, at least in the US and Europe, from the extensive history of legal conflict between industry and regulators. In principle it is possible to apply the framework we develop above to domains that are simply not significant, so it is critical for researchers to be able to establish, if not predict, the importance of a new domain of investigation. We propose that this would involve measures of what we term the essentiality of the data service provided, and the market power of the gatekeeper. Step 4 involves analysis of (a subset of) the five key trade-offs identified above – for innovation and variety of content, infrastructure investment, potential integration of the gatekeeper and CP, compatibility constraints that may be imposed by the gatekeeper in collaboration with a CP, and overall welfare effects – as discussed in Section 2.3 for net neutrality regulation.

3.2. Step 1: Determine market structure

In Step 1 we are concerned primarily with clearly identifying the gatekeeper and the two sides of the market, and with making some preliminary observations about the neutrality issues that arise. We now present an initial analysis of the market structure in several emerging areas of IS research.

3.2.1. Search neutrality

The first emergence of data neutrality as an issue over and beyond net neutrality occurred with respect to *search neutrality* (Odlyzko 2009). Search engines are the gatekeepers between consumers and CPs and the main asserted non-neutral conduct here was that Google would rank those search results that related to its own or affiliated services higher in the organic (unpaid) search results (FTC 2013). This is analogous to the issue of ISP-driven packet discrimination in the context of the net neutrality debate. However, unaffiliated CPs can also pay the search engine directly in order to be displayed as sponsored search results above the organic search results. This is clearly analogous to CP-driven packet discrimination, yet has not drawn regulatory attention. Moreover, search neutrality is not constrained to pure search engines such as Google. Potential search neutrality issues also apply to providers with search functions, such as searching for products on Amazon and searching for people on Facebook. For example, recently both Apple and Google have announced that sponsored search results are forthcoming for the AppStore (Perez 2016) and PlayStore (Siliski 2015), respectively.

3.2.2. Operating system neutrality

Data non-neutrality practices are particularly striking in the context of operating systems, which are the gatekeepers between the hardware and the software (apps) that run on this hardware. Accordingly, here non-neutral practices may involve interference with respect to which software can be installed, and with respect to which hardware may be used with the operating system.

There is also another related neutrality issue with respect to the app store, which usually is an integral part of the operation system. Apps stores can implicitly control the types of apps that are available on a device by choosing not to make certain apps available (e.g., some apps may not pass Apple's approval process) or by imposing restrictions on an apps functionality. For

example, Apple's iOS operating system comes with several pre-installed apps (e.g., iMessage), which may affect users' choices of alternative apps (e.g., WhatsApp) or services (e.g., SMS). While this may be done with good intentions to ensure quality and security (akin to "reasonable" network management practices of ISPs), there is currently no neutrality regulation in place that governs such pre-configurations and restrictions, or that demands transparency in the app store's approval process.

3.2.3. Browser neutrality

Our third example of data neutrality concerns web browsers, including their associated plug-ins, as well as other pieces of software that may interfere with how web content is displayed, i.e., how CPs may interact with consumers. For example, the company Eyeo offers a very popular browser plug-in called Adblock Plus that can effectively block advertisements on web sites from being shown. Evidently major CPs like Google, Amazon and Microsoft have paid Adblock Plus in order not to have the ads on their sites blocked (Cookson 2015). Clearly, this puts these CPs at an advantage in the competition for advertisers over other CPs that did not, or could not, pay Adblock Plus for preferential treatment. The relationship to the net neutrality debate is evident, as the business model of many CPs relies on advertisement, although in this specific example blocking occurs with respect to content (ads) that typically creates a disutility for consumers.

3.3. Step 2: Investigate discrimination scenarios

In this step we aim to investigate different potential discrimination scenarios. We believe that the general forms of discrimination practices that arise in the context of data neutrality are similar to those in the net neutrality debate. To see this more clearly, in Table 2 we have organized various non-neutral data management practices in the same logical framework as previously for net neutrality in Table 1.

Table 2: Data Management Practices

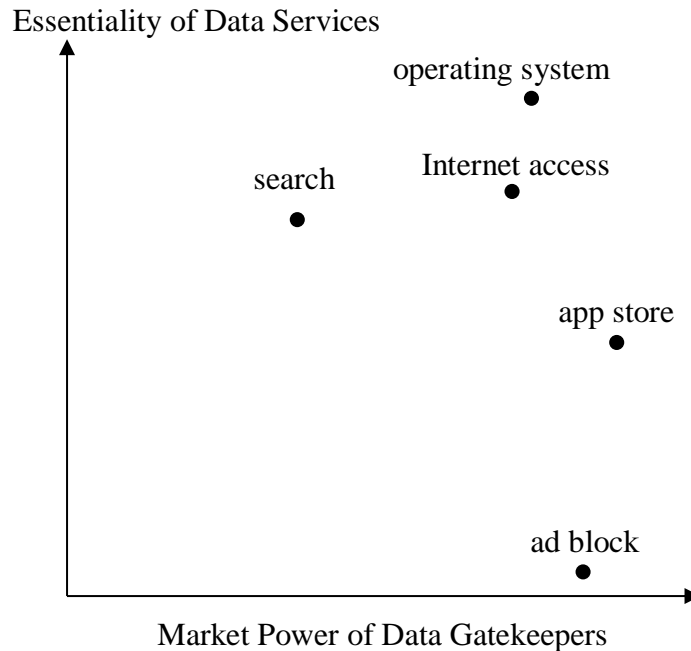
<i>Data services</i>	<i>Related agents (CP, Gatekeeper, Consumer)</i>	<i>Data management practices</i>	<i>Neither side pays extra to gatekeeper</i>	<i>CP side pays extra to gatekeeper</i>	<i>Consumer side pays extra to gatekeeper</i>
<i>Search neutrality</i>	The search service (e.g., Google, Facebook, Amazon) is the gatekeeper. Consumers search for content or products of interest.	Managerial and other preferences are reflected in search ranking algorithm.	Search provider boosts ranking of sites in which it has a financial interest or which are affiliated.	CPs pay search service to achieve a higher ranking.	Consumer pays (e.g., in data) to achieve a better fitting, personalized search result.
<i>Operating system neutrality</i>	The device manufacturer or OS is the gatekeeper. Consumers get content or services through the use of device or OS.	Discrimination of functionality based on CP or app provider, consumer device, or NP.	Tightly controlling software/hardware ecosystem (e.g., by disallowing jailbreak), pre-installation of software to displace competitive apps, or interference with functionality of alternative apps.	CP or manufacturer pays to make app available or for priority placement (e.g., on home screen).	Consumer pays to unlock certain features of the OS or device.
<i>App store neutrality</i>	The app store is the gatekeeper (usually associated with the OS), app developers and consumers are the other two parties.	App prioritization or blocking based on the developer, the app or app type, or the consumer.	App promotion tools are provided for discriminant app discovery. App store provides customized featured apps for consumers.	Developers pay to be featured or to be listed	Consumers pay to have certain apps listed or featured.
<i>Ad block neutrality</i>	Ad block provider is the gatekeeper. Consumers may use ad block software to block ads from publishers/advertisers.	Ad blocking based on the ad publisher, network, type, content; or on the advertiser or consumer.	Making it free for small websites and blogs to be listed in the whitelist of acceptable ads.	Advertisers pay not to be blocked.	Consumers pay to undo the whitelisting of certain ads.

3.4. Step 3: Analyze the need for neutrality regulation

Although the previously discussed examples show close parallels to the net neutrality debate, there are several subtle differences that require investigation. For example, in the more general context of data neutrality, payments for preferential treatment may not only be made in money, but also in “data”. Moreover, all of the examples in the context of data neutrality above entail

software platforms as the gatekeepers, for which our conclusions for possible policy interventions may be quite different than for infrastructure platforms, such as a last-mile networks, e.g., because congestion in access to the service is not a central issue here, or because the competitive dynamics may be significantly different for software vs. infrastructure platforms. So one of the underlying research questions that must be addressed is under which conditions we should actually care about data neutrality regulation in general, for example, in the same way that we care about net neutrality regulation. Informed by the specific debate on net neutrality, we suggest that the initial screening for the need for data neutrality regulation in general should be based on two key factors: First, how much *market power* does the data gatekeeper have? And, second, how *essential* is the data service for (the interaction with) consumers? As we will discuss next, both questions pose significant research problems on their own.

Figure 5 provides an illustrative outcome of the envisaged assessment of market power and essentiality for different data gatekeepers that might be under scrutiny in the context of data neutrality. The closer a data service is located to the top right corner, the more likely it is to be scrutinized by researchers and policymakers. It is worth noting that the need for data neutrality regulation for various data services may change over time. For example, as consumers become more and more dependent on their mobile devices, the essentiality of mobile devices increases. As a result, the need for data neutrality regulation may become stronger over time.



Notes: The essentiality of Internet access, coupled with gatekeeper ISPs having significant market power, together underlie the net neutrality debate. In the case of browser neutrality (ad block) the essentiality of the service is not yet at a level to merit scrutiny, while app store neutrality may soon be. Operating system neutrality was attracting regulatory attention even in 2001 for Microsoft’s bundling of its IE browser (U.S. Court of Appeals 2001), and continues to arise with platforms like Android. Search neutrality is an area attracting increased attention.

Figure 5: Analyze the Need of Neutrality Regulation for Different Data Services

Market power is an abstract economic and legal term that denotes that a firm can behave relatively free of competitive constraints. Traditionally, in economic and legal analysis, market power is assessed by market share (e.g., by the Herfindahl-Hirschman index) or the ability to raise prices above marginal costs (e.g., by the Lerner index). However, in the context of data services, such measures are often inappropriate. Many firms in the digital economy enjoy very large market shares (due to network effects), but this does not necessarily mean that they enjoy market power. For example, a search engine could have dominant market share when offering a free service to consumers, but its demand may be very elastic such that, once it decides to demand a price, many consumers are likely to switch to another search engine. Moreover, in the

digital economy marginal costs are usually close to or even at zero, whereas fixed costs can be significant. In addition, generally platforms operate as a two-sided market, where even in a monopolistic market prices on one market side may be below or at marginal costs (see Parker and Van Alstyne 2005). This means that prices on either side of the market are not a good signal for market power, because low profit margins on one market side (reflecting low market power) are usually coupled with high profit margins on the other market side (reflecting high market power). Generally, the extent of market power is determined by the substitutability of the firm's offering with another which is already available (demand substitutability) or may readily become available (supply substitutability) once the firm in question tries to exploit its market power. However, the determination of such substitutability patterns is likely to be very complex for data services and can be highly application specific, as it requires intimate knowledge of user adoption decisions as well as technology.

Thus, the development of a market power metric is a perfect research area for IS scholars. For example: to what extent is Facebook substitutable by other social networking platforms (e.g., Twitter or LinkedIn) or other CPs that offer consumers individually curated news content (e.g., Kite)? How large are the network effects and demand-side economies of scale, which play a key role in establishing barriers to entry for new competitors. The answers will also depend, at least in part, on technology characteristics (e.g., functionality, reliability), users' expectations, preferences and adoption patterns, and other firms' abilities to offer similar products. Moreover, and closely related, it is crucial to determine the extent of switching costs that arise ex-post, i.e., after a user has opted for a specific service. For example, ISPs are believed to have market power because they enjoy a termination monopoly over their subscribers. That is, even though ISPs may be in competition for consumers ex-ante, once a consumer has subscribed to the ISP, there

exist significant switching costs (e.g., due to a fixed contract length) that grant the ISP a termination monopoly. The same type of switching costs can arise for data services, although they are usually much harder to estimate because they are typically constituted by network effects (i.e., lock-in), and personalization (i.e., past interaction with the service). This presents another promising and relevant area of empirical IS research, which would be very valuable for the determination of market power, and thus the need for regulatory intervention.

Yet, market power alone is not sufficient to warrant a regulatory intervention. A policy intervention always has to address a societal problem of significant magnitude. In other words, the platform or data service under consideration must also be “essential” to consumers or for the interaction with consumers. As opposed to “market power”, the term “essentiality” does not currently represent a well-developed economic or legal concept. By “essential” we mean how important the service is in terms of the economic benefit (or loss) that it creates for society as a whole. Note that essentiality and market power represent two distinct characteristics of data services, as a firm may enjoy market power over a service that is not essential and vice versa. For example, it may well be that Eyeo, the provider of the ad-blocking browser plug-in Adblock-Plus has more market power than AT&T with respect to the specific data service that it provides. However, we would probably argue that Internet access is more essential to society than any given browser plug-in. Nevertheless, one should not underestimate the potential disruptive effect that ad-blocking software can have on the Internet economy, which relies heavily on advertisement revenues (Scott 2017). We currently do not have a broadly accepted measure of “essentiality” across different data services, which could serve as a yardstick for policy intervention. We believe that IS researchers are well-positioned to develop such a measure for

essentiality and therefore we highlight this as another promising research area for theoretical and empirical IS research.

3.5. Step 4: Analyze key trade-offs

In this step we seek to identify the specific trade-offs that should be researched in the context of data neutrality. Just as net neutrality can be seen as one distinct issue in the broader concept of data neutrality, we believe that it is useful to organize the discussion of prospective research areas in data neutrality by generalizing the five key trade-offs that have been proven central in the net neutrality debate. The first two trade-offs (affiliation and compatibility) are static trade-offs that may lead to allocation inefficiency. The third and fourth trade-offs (innovation and investment) are dynamic trade-offs that may affect CPs' ongoing incentive to innovate and gatekeepers' ongoing incentive to invest respectively. Finally, welfare analysis takes both the static and dynamic trade-offs into consideration.

Similar to the body of literature that has considered the net neutrality debate, researchers will usually want to address only a subset of these trade-offs in a given research project, rather than to address all trade-offs at the same time. However, there are also likely to be interdependencies between the trade-offs. This is, for example, self-evident for welfare analysis (Trade-off 5), for which all of the other trade-offs have to be taken into account.

3.5.1. Trade-off 1 – Affiliation: What are the incentives of vertically integrated gatekeepers to interfere with unaffiliated content?

Research issues concerning this trade-off occur whenever the data gatekeeper is also providing content for which it acts as a gatekeeper. For example, this trade-off is considered central in the case of Google being accused of ranking affiliated services higher in the (organic) search results. Similar issues can also arise in the context of app stores, e.g., when Apple ranks its own apps

higher, or in the context of operating system neutrality when the device manufacturer preinstalls its own apps on the devices, both of which can have a tremendous impact on the success of an app.

There are also reasons not to engage in favoring of affiliated content, because it may deteriorate the quality of the gatekeeper's matching service and thus, particularly if the gatekeeper is not in a strong position of market power, endanger its core business model. For example, if Google distorts search results too much in its own favor, then users might switch to another search engine, such as Microsoft's Bing. This process may be irreversible (unless Bing, in turn, jeopardizes this position) and ultimately constitute Bing as the new gatekeeper. From a policy perspective, the key research question is to find, both theoretically and empirically, the general conditions under which gatekeepers have an incentive to favor affiliated content, and under which conditions this is harmful, e.g., with respect to other trade-offs (compatibility, innovation, infrastructure, and welfare). If favoring affiliated content is harmful and the gatekeeper has incentives to do so, then data neutrality regulation may be warranted.

Data sharing, a weaker form of vertical integration, also raises interesting research issues. For example, a search engine may announce a plan to rank those content sites higher that provide an API through which the search engine can analyze user behavior on the content site *after* the referral. On the one hand, providing incentives for data sharing can be considered to be welfare improving because it increases the information base on which algorithms are optimized. On the other hand, there may be concerns of a further manifestation of market power or undermining of privacy. In this regard, data neutrality, research on privacy preserving mechanisms, and data analytics, which is a hot topic for IS research in itself (Agarwal and Dhar 2014), are tightly related.

3.5.2. Trade-off 2 – Compatibility: What are the incentives of gatekeepers (CPs) to make their platforms (content) compatible with and without neutrality regulation?

Gatekeepers often seek incompatibility as a way to offer exclusive content and secure their gatekeeper position. Research on the compatibility of the products offered by different firms has been a perennial issue for IS scholars in the context of software platforms (see, e.g., Adner et al. 2016). However, it would be valuable to extend this research by viewing it under the lens of data neutrality: For example, the new General Data Protection Regulation that comes into effect in the European Union in 2018 requires that users must be empowered to transfer their data from one CP to another, i.e., that firms must establish data compatibility. It is not yet clear, however, how to achieve this in practice and how this will affect the way in which data is collected and stored by the CPs. Moreover, no research has yet examined how such neutrality regulation with respect to data compatibility affects the gatekeepers' incentives to provide compatible platforms in general. On the one hand, data compatibility should facilitate the emergence of compatible platforms (indeed this was the rationale behind the regulation). On the other hand, the opposite effect is also feasible since forced compatibility with respect to data sharing may provide a stronger incentive for incompatibility of the service itself. Similarly, in another context one could ask whether data neutrality regulation of an app store that would prohibit favoring certain apps would lead to more or less compatibility between apps. For example, currently WhatsApp is not compatible with other messenger applications, nor is it compatible with all devices, such as tablets or Wi-Fi only devices. Incompatibility is usually detrimental to consumers' welfare because it limits their choices, so research examining these and related questions would be valuable to assess the impact of data neutrality regulation.

3.5.3. Trade-off 3 – Innovation: Does neutrality lead to more or less innovation or variety of content?

This trade-off has been central in the context of net neutrality regulation since the Internet is considered to be an extraordinary source of innovation, where it has been argued that innovations occur so frequently precisely because there is no gatekeeper that controls the innovation process. For the same reasons this trade-off is also central in the larger context of data neutrality, where gatekeepers can also be software platforms. For example, IS researchers have been trying to assess the demand for apps (i.e., a software innovation) based on observable characteristics (see, e.g., Ghose and Han 2014), and particularly based on their ranking position (see, e.g., Garg and Telang 2014). So, should we not be concerned with search engine's or app store's business model to sell a prioritized listing to CPs or apps, in the same way as we are concerned with the ISP's desire to sell prioritized data delivery to CPs? Viewing the above IS research in the context of data neutrality would give an indication how much app developers might be willing to pay to be listed higher in the app store's ranking. Alternatively, if such pay-for-priority arrangements were prohibited, previous IS research has shown how such rule changes can affect the quality of apps (e.g., Claussen et al. 2013) as well as the drivers for intra-platform competition (Tiwana 2015).

Another approach to this issue would be to estimate how much less (venture) capital would be required to compete in a data neutral vs. non-neutral app store and thus how many more innovative apps might receive funding above the critical threshold required to pursue the software project. Likewise, to assess the impact of data neutrality regulation in the context of browser neutrality, it would be worthwhile to investigate how the (selective) blocking of ads interferes with CP business models, and therefore their entry and exit decisions. This would

clearly have an impact on which (innovative) business models are sustainable and how much content will be available.

3.5.4. Trade-off 4 – Investment: Does data neutrality foster or stifle investment incentives for the platform operators?

Software platforms arguably create value for society, and thus investments in the development of such platforms should generally be desirable in the same sense that investment in broadband infrastructure is considered desirable. But how do the incentives to build a software platform change if the platform is subject to data neutrality regulation once it becomes successful (with high levels of essentiality and market power)? For example, in 2005 Google chose to compete with Apple's iOS by investing in its own mobile operating system Android. After Android became a success, Google has pursued non-neutral data practices on the Android platform. Although Android's core is open source, Google has imposed requirements on device manufacturers that wanted to install *any* of Google's mobile applications (such as the Play Store, Google Search, Maps or YouTube) that tightly integrate with the Android operating system and, if not installed, would significantly diminish its value. These obligations include that all of Google's applications be pre-installed and that they be placed prominently on the user's device, e.g., on the home screen (Edelman 2015). Clearly, as both the space on the home screen as well as the space to store apps on mobile devices is limited, this conduct is non-neutral with respect to alternative mobile applications. Should we then not have the same concerns as in the net neutrality debate about app innovation and fair competition among apps?

It could be argued that Google's incentives to innovate and invest in Android may have been driven by the fact that discrimination (ex-post) was possible. Although, data neutrality regulation may be desirable ex-post, i.e., after the investment has been made, it may not be

desirable ex-ante, because it diminishes the initial incentive for the investment. Eventually this may lead to forgone investments and thus a forgone opportunity of competition between platforms (iOS and Android in this case). Ironically, competition would likely limit the market power of either platform and therefore the need for data neutrality regulation in the first place. Hence, more empirical and theoretical research is needed to better understand the trade-off inherent to data neutrality regulation between fostering investments in the long run and fostering fair and equal competition in the short run.

3.5.5. Trade-off 5 – Welfare: Who are the winners and losers of data neutrality? Is the economy as a whole better or worse off?

With respect to policy implications, it is paramount to identify the welfare consequences of data neutrality regulation for the different stakeholders involved, both in the short run as well as in the long run. It is apparent from the preceding discussion that welfare consequences may be very different for different stakeholders, and that the outcomes of such analyses are likely to be very case specific. Therefore, researchers will typically base their evaluation on a subset of the trade-offs mentioned above, and a full picture emerges from the host of complementary analysis. Nevertheless, the ultimate goal should be to derive patterns from a richer set of research results that prescribe more general conditions under which neutrality regulation is warranted.

In this overarching context, it is also interesting to study the relationship between device manufacturers, operation system developers, network operators and content providers, each of which control a different part of the value chain that they seek to monetize through discriminatory practices. As exemplified above, content providers that decry non-neutral network management practices by ISPs may exercise similar practices themselves at different levels of

the value chain. The combined effect of this complex interplay of non-neutral practices itself provides an interesting area of future research.

4. Conclusions

The public policy debate concerning net neutrality regulation and its potential impact has been long and arduous, in part at least because of the inherent complexity of the issues involved. We present a multi-step framework that is designed to help both IS researchers and public policymakers integrate the many issues involved in a comprehensive way while still maintaining focus on the key players, and the essential trade-offs in the debate. We exemplify this framework by reviewing the issues and research on net neutrality. To this end, we first provide a conceptual abstraction of the Internet that allows us to clearly define the central gatekeeper – the last-mile ISP – while still representing the many other parties involved in Internet transactions. We trace the existing payment flows as well as potential and proposed payment flows and use this to motivate the second step, in which we clearly delineate which are neutral and which non-neutral, providing a kind of road map to the myriad issues that may arise. Given the maturity of the field, we are then able to summarize numerous research results as they relate to a set of key trade-offs that we identify with respect to impact on incentives for affiliation, compatibility, innovation, infrastructure, and overall welfare results.

Recognizing that this multi-step process provides a framework that helps to conceptually organize the many substantial research contributions to the net neutrality debate, we then turn our attention to emerging domains where similar concerns may arise. In looking to the future, it becomes critical to have a method of determining whether a domain where neutrality issues may arise will turn out to be important enough to merit our collective research attention. We thus propose that some research attention be devoted to the development of metrics for the market

power of gatekeepers, beyond those currently available, and for the essentiality of data services. In cases where a service becomes essential, such as is arguably the case with search services, and where market power is concentrated, as may be argued for certain mobile platforms, there is potential for public policy intervention to come to be viewed as necessary, as has happened, for example, with the classification of the Internet as a utility service in the US.

Combining this analysis with the structure of our net neutrality analysis, we develop a general framework that can be used to help guide future IS research on important emerging domains of what we term data neutrality. We provide three examples of general domains of data neutrality: search, operating systems, and browsers, and in some cases specific examples within those domains, such as ad block add-ins to browsers, or app stores associated with specific operating systems. We mention recent research in these emerging areas, but our primary focus is on demonstrating how our framework, and in particular the consideration of the trade-offs, can help define specific research questions and guide future IS research.

References

- Adner, R., J. Chen, F. Zhu. 2016. Frenemies in platform markets: The case of Apple's iPad vs. Amazon's Kindle. Working paper.
- Agarwal, R., V. Dhar. 2014. Editorial - big data, data science, and analytics: The opportunity and challenge for IS research. *Information Systems Research*, 25 (3), 443-448.
- Bourreau, M., F. Kourandi, T. Valletti. 2015. Net neutrality with competing internet platforms. *Journal of Industrial Economics*, 63 (1), 30-73.
- Brodkin, J. 2014. It's not a "fast lane" but Comcast built a CDN to charge for video delivery. *Ars Technica*, May 19.
- Broos, S., A. Gautier. 2016. Competing one-way essential complements: The forgotten side of net neutrality. Working paper.
- Bustamante, F.E., D. Clark, and N. Feamster. 2015. *Tracking quality of experience in the internet*. NSF/FCC Workshop, October 21-22, Available at: <http://aqualab.cs.northwestern.edu/NSFWorkshop-InternetQoE>.
- CBC. 2005. Telus cuts subscriber access to pro-union website. *CBC News*, July 24.
- Chen, M.K., B.J. Nalebuff. 2006. One-way essential complements. Working paper.

- Chen, P., L. M. Hitt. 2002. Measuring switching costs and the determinants of customer retention in internet-enabled businesses: A study of the online brokerage industry. *Information Systems Research*, 13 (3), 255-274.
- Cheng, H. K., S. Bandyopadhyay, H. Guo. 2011. The debate on net neutrality: A policy perspective. *Information Systems Research*, 22 (1), 60-82.
- Chiang, I. R., J. Jhang-Li. 2014. Delivery consolidation and service competition among internet service providers. *Journal of Management Information Systems*, 31 (3), 254-286.
- Choi, J. P., B. Kim. 2010. Net neutrality and investment incentives. *RAND Journal of Economics*, 41 (3), 446-471.
- Claussen, J., T. Kretschmer, P. Mayrhofer. 2013. The effects of rewarding user engagement: The case of facebook apps. *Information Systems Research*, 24 (1), 186-200.
- Cookson, R. 2015. Google, Microsoft and Amazon pay to get around ad blocking tool. *Financial Times*, February 1.
- Coucheney, P., P. Maillé, B. Tuffin. 2014. Network neutrality debate and ISP inter-relations: Traffic exchange, revenue sharing, and disconnection threat. *Netnomics*, 15 (3), 155-182.
- Czernich, N., O. Falck, T. Kretschmer, L. Woessmann. 2011. Broadband infrastructure and economic growth. *Economic Journal*, 121 (552), 505-532.
- D'Annunzio, A., A. Russo. 2015. Net neutrality and internet fragmentation: The role of online advertising. *International Journal of Industrial Organization*, 43, 30-47.
- Dewan, S., F. J. Riggins. 2005. The digital divide: Current and future research directions. *Journal of the Association for Information Systems*, 6 (12), 298-337.
- Dewan, S., D. Ganley, K. L. Kraemer. 2010. Complementarities in the diffusion of personal computers and the internet: Implications for the global digital divide. *Information Systems Research*, 21 (4), 925-940.
- Dewenter, R., J. Rösch. 2016. Net neutrality and the incentives (not) to exclude competitors. *Review of Economics*, 67 (2), 209-229.
- Economides, N., B. E. Hermalin. 2014. The strategic use of download limits by a monopoly platform. *RAND Journal of Economics*.
- Economides, N., B. Hermalin. 2012. The economics of network neutrality. *RAND Journal of Economics*, 43 (4), 602-629.
- Economides, N., J. Tåg. 2012. Net neutrality on the internet: A two-sided market analysis. *Information Economics and Policy*, 24 (2), 91-104.
- Edelman, B. 2015. Does Google leverage market power through tying and bundling? *Journal of Competition Law and Economics*, 11 (2), 365-400.
- Faulhaber, G. R. 2011. Economics of net neutrality: A review. *Communications & Convergence Review*, 3 (1), 7-25.
- FCC. 2011. *In the matter of Madison River Communications, LLC and affiliated companies*. Federal Communications Commission, Available at: https://apps.fcc.gov/edocs_public/attachmatch/DA-05-543A2.pdf.
- FTC. 2013. *Statement of the Federal Trade Commission regarding Google's search practices*. Federal Trade Commission, January 3, Available at: http://www.ftc.gov/sites/default/files/documents/public_statements/statement-commission-regarding-googles-search-practices/130103brillgooglesearchstmt.pdf.
- Garg, R., R. Telang. 2014. Inferring app demand from publicly available data. *MIS Quarterly*, 37 (4), 1253-1264.

- Ghose, A., S. P. Han. 2014. Estimating demand for mobile applications in the new economy. *Management Science*, 60 (6), 1470-1488.
- Greenstein, S., M. Peitz, T. Valletti. 2016. Net neutrality: A fast lane to understanding the trade-offs. *Journal of Economic Perspectives*, 30 (2), 127-150.
- Guo, H., H. K. Cheng, S. Bandyopadhyay. 2013. Broadband network management and the net neutrality debate. *Production and Operations Management*, 22 (5), 1287-1298.
- . 2012. Net neutrality, broadband market coverage and innovations at the edge. *Decision Sciences*, 43 (1), 141-172.
- Guo, H., S. Bandyopadhyay, A. Lim, Y. Yang, H. K. Cheng. 2016. Effects of competition among internet service providers and content providers on the net neutrality debate. *MIS Quarterly*, forthcoming.
- Guo, H., R. F. Easley. 2016. Network neutrality versus paid prioritization: Analyzing the impact on content innovation. *Production and Operations Management*, 25 (7), 1261-1273.
- Guo, H., S. Bandyopadhyay, H. K. Cheng, Y. Yang. 2010. Net neutrality and vertical integration of content and broadband services. *Journal of Management Information Systems*, 27 (2), 243-275.
- Gustin, S. 2014. Netflix pays Verizon in streaming deal, following Comcast pact. *Time*, April 28.
- Hermalin, B. E., M. L. Katz. 2007. The economics of product-line restrictions with an application to the network neutrality debate. *Information Economics and Policy*, 19, 215-248.
- Hsieh, J. J. P., A. Rai, M. Keil. 2011. Addressing digital inequality for the socioeconomically disadvantaged through government initiatives: Forms of capital that affect ICT utilization. *Information Systems Research*, 22 (2), 233-253.
- . 2008. Understanding digital inequality: Comparing continued use behavioral models of the socio-economically advantaged and disadvantaged. *MIS Quarterly*, 32 (1), 97-126.
- Johar, M., S. Menon, V. Mookerjee. 2011. Analyzing sharing in peer-to-peer networks under various congestion measures. *Information Systems Research*, 22 (2), 325-345.
- Kang, C. 2014. FCC approves plan to consider paid priority on internet. *The Washington Post*, May 15.
- Kourandi, F., J. Krämer, T. Valletti. 2015. Net neutrality, exclusivity contracts and internet fragmentation. *Information Systems Research*, 26 (2), 320-338.
- Krämer, J., L. Wiewiorra. 2012. Network neutrality and congestion sensitive content providers: Implications for content variety, broadband investment and regulation. *Information Systems Research*, 23 (4), 1303-1321.
- Krämer, J., L. Wiewiorra. 2015. When 'Just' is just not enough. *Business & Information Systems Engineering*, 57 (5), 325-338.
- Krämer, J., L. Wiewiorra, C. Weinhardt. 2013. Net neutrality: A progress report. *Telecommunications Policy*, 37 (9), 794-813.
- Lee, R. S., T. Wu. 2009. Subsidizing creativity through network design: Zero-pricing and net neutrality. *Journal of Economic Perspectives*, 23 (3), 61-76.
- Masuda, Y., S. Whang. 2006. On the optimality of fixed-up-to tariff for telecommunications service. *Information Systems Research*, 17 (3), 247-253.
- McKinney, V., K. Yoon, F. Zahedi. 2002. The measurement of web-customer satisfaction: An expectation and disconfirmation approach. *Information Systems Research*, 13 (3), 296-315.
- Mcmillan, R. 2014. What everyone gets wrong in the debate over net neutrality. *Wired*, June 23.

- Mueller, M. L., H. Asghari. 2012. Deep packet inspection and bandwidth management: Battles over BitTorrent in Canada and the United States. *Telecommunications Policy*, 36 (6), 462-475.
- Odlyzko, A. 2009. Network neutrality, search neutrality, and the never-ending conflict between efficiency and fairness in markets. *Review of Network Economics*, 8 (1), 40-60.
- Orland, K. 2012. Comcast: Xbox 360 on demand streams won't count against data caps. *Ars Technica*, March 26.
- Parker, G. G., M. W. Van Alstyne. 2005. Two-sided network effects: A theory of information product design. *Management Science*, 51 (10), 1494-1504.
- Peitz, M., F. Schuett. 2016. Net neutrality and inflation of traffic. *International Journal of Industrial Organization*, 46, 16-62.
- Perez, S. 2016. Apple to introduce search ads on app store along with changes to app review, discovery and splits. *TechCrunch*, June 8.
- Ray, B., A. Leach. 2013. Orange boasts: We made google PAY US for traffic. *The Register*, January 17.
- Reggiani, C., T. Valletti. 2016. Net neutrality and innovation at the core and at the edge. *International Journal of Industrial Organization*, 45, 16-27.
- Rogowsky, M. 2014. Comcast-Netflix didn't violate net neutrality but it wasn't benevolent, it was business. *Forbes*, February 24.
- Röllner, L., L. Waverman. 2001. Telecommunications infrastructure and economic development: A simultaneous approach. *American Economic Review*, 91 (4), 909-923.
- Russell, J. 2015. Bharti Airtel fuels net neutrality fire in India with free access to partner apps. *TechCrunch*, April 6.
- Schuett, F. 2010. Network neutrality: A survey of the economic literature. *Review of Network Economics*, 9 (2), 1-15.
- Scott, M. 2017. Use of ad-blocking software rises by 30% worldwide. *New York Times*, January 31.
- Siliski, M. 2015. *A new way to promote your app on Google Play*. Android Developers Blog, February 26, Available at: <http://android-developers.blogspot.ca/2015/02/a-new-way-to-promote-your-app-on-google.html>.
- Tiwana, A. 2015. Evolutionary competition in platform ecosystems. *Information Systems Research*, 26 (2), 266-281.
- U.S. Court of Appeals. 2001. United States of America, appellee v. Microsoft corporation, appellant, 253 F.3d 34 (D.C. cir. 2001). *U.S. Court of Appeals for the District of Columbia Circuit - 253 F.3d 34 (D.C. Cir. 2001)*, June 28.
- Wu, T. 2003. Network neutrality, broadband discrimination. *Journal of Telecommunications and High Technology Law*, 2, 141-179.
- Yu, R. 2014. Netflix cuts deal with Comcast to speed service. *USA Today*, February 26.

Appendix

Table A1 reviews the main stream of the literature, which has considered the impact of net neutrality in comparison to packet discrimination. Thereby we note the main assumptions made along with the findings for the net neutrality case.

Table A1: Literature Review

Papers	Modeling Assumptions / Market Conditions Considered							Trade-offs considered: Findings for Net Neutrality Case
	Market Structure	Broadband Market Coverage	Consumer Heterogeneity	CP Heterogeneity	CP Entry	Multi-homing MH / Single-homing SH	Network Congestion	
Hermalin and Katz (2007)	ISP: Monopoly or Duopoly CP: No Competition	Full	Homogeneous	Heterogeneous in consumers' valuation	Continuum of CPs with entry	Consumers MH CPs	Not modeled	<ul style="list-style-type: none"> • Social Welfare: likely lower • Content Innovation: lower
Choi and Kim (2010)	ISP: Monopoly CP: Duopoly	Full	Heterogeneous in content preference	Heterogeneous in markup	Two CPs with no entry	Consumers SH CPs	M/M/1 queuing system	<ul style="list-style-type: none"> • Social Welfare: mixed • Infrastructure Investment: higher • Content Innovation: mixed
Guo et al. (2010)	ISP: Monopoly CP: Duopoly with ISP-CP integration	Full	Heterogeneous in content preference	Heterogeneous in revenue rate	Two CPs with no entry	Consumers SH CPs	M/M/1 queuing system	<ul style="list-style-type: none"> • Social Welfare: mixed when the ISP integrates with a CP • Affiliation: Vertically integrated ISP may degrade or even prioritize, competing content.
Cheng et al. (2011)	ISP: Monopoly CP: Duopoly	Full	Heterogeneous in content preference	Heterogeneous in revenue rate	Two CPs with no entry	Consumers SH CPs	M/M/1 queuing system	<ul style="list-style-type: none"> • Social Welfare: lower • Infrastructure Investment: higher
Economides and Hermalin (2012)	ISP: Monopoly CP: No Competition	Full	Homogeneous	Heterogeneous in congestion sensitivity	Continuum of CPs with no entry	Consumers MH CPs	Congestion modeled based on bandwidth division	<ul style="list-style-type: none"> • Social Welfare: higher or lower, depending on the total traffic • Infrastructure Investment: lower
Economides and Tåg (2012)	ISP: Monopoly or Duopoly CP: No Competition	Partial	Heterogeneous in ISP preference	Heterogeneous in entry cost	Continuum of CPs with entry	Consumers MH CPs; with duopoly ISPs, CPs MH ISPs and consumers SH ISPs	Not modeled	<ul style="list-style-type: none"> • Social Welfare: higher, for most parameter values.
Guo et al. (2012)	ISP: Monopoly CP: Duopoly	Partial	Heterogeneous in content preference	Heterogeneous in revenue rate	Two CPs with no entry	Consumers SH CPs	M/M/1 queuing system	<ul style="list-style-type: none"> • Social Welfare: lower • Content Innovation: higher because the less effective CP may be driven out of the market without net neutrality.
Krämer and Wiewiorra (2012)	ISP: Monopoly CP: No Competition	Full & Partial	Homogeneous	Heterogeneous in congestion sensitivity	Continuum of CPs with entry	Consumers MH CPs	M/M/1 queuing system	<ul style="list-style-type: none"> • Social Welfare: lower • Infrastructure Investment: lower • Content Innovation: lower (in the long run)

Papers	Modeling Assumptions / Market Conditions Considered							Trade-offs considered: Findings for Net Neutrality Case
	Market Structure	Broadband Market Coverage	Consumer Heterogeneity	CP Heterogeneity	CP Entry	Multi-homing MH / Single-homing SH	Network Congestion	
Guo et al. (2013)	ISP: Monopoly CP: Duopoly	Full	Heterogeneous in content preference	Heterogeneous in revenue rate	Two CPs with no entry	Consumers SH CPs	M/M/1 queuing system	<ul style="list-style-type: none"> • Social Welfare: lower
Bourreau et al. (2015)	ISP: Duopoly	Full	Homogeneous	Heterogeneous in congestion sensitivity	Continuum of CPs with entry	Consumers MH CPs and SH ISPs; CPs MH ISPs	M/M/1 queuing system	<ul style="list-style-type: none"> • Social Welfare: lower • Infrastructure Investment: lower • Content Innovation: lower
D'Annunzio and Russo (2015)	ISP: Duopoly CP: Duopoly (incl. case with integrated ISP-CPs)	Full	Heterogeneous in ISP preference	Homogenous (ex-ante) or heterogeneous in consumer attention	Two CPs with no entry	Consumers MH CPs and SH or MH ISPs; CPs MH or SH ISPs	Not modeled	<ul style="list-style-type: none"> • Social Welfare: equal or higher, depending on degree of CP competition • Compatibility: equal or higher; NN ensures compatibility, but only if CPs are not vertically integrated
Kourandi et al. (2015)	ISP: Duopoly CP: Duopoly	Full	Heterogeneous in ISP preference	Heterogeneous in consumer valuation	Two CPs with no entry	Consumers MH CPs and SH or MH ISPs; CPs MH or SH ISPs	Not modeled	<ul style="list-style-type: none"> • Social Welfare: higher, equal or lower, depending on degree of CP competition • Compatibility: higher, but NN does not ensure compatibility
Broos and Gautier (2016)	ISP: Monopoly and Duopoly CP: Duopoly	Full (for welfare analysis)	Heterogeneous in content preference	Heterogeneous in content and quality	Two CPs with no entry	Consumers SH CPs	Not modeled	<ul style="list-style-type: none"> • Social Welfare: higher or lower • Affiliation: Vertically integrated ISP does not block rival CP if it is valuable to consumers
Dewenter and Rösch (2016)	ISP: Monopoly CP: Duopoly with ISP-CP integration	Partial	Heterogeneous in valuation for content	Heterogeneous in content	Two CPs with no entry	Consumers MH CPs	Not modeled	<ul style="list-style-type: none"> • Social Welfare: not considered • Affiliation: Vertically integrated ISP does not block rival CP if it offers differentiated content and if the loss in own network effects is not too strong
Guo and Easley (2016)	ISP: Monopoly CP: No Competition	Partial	Heterogeneous in content valuation	Heterogeneous in congestion sensitivity	Continuum of CPs with entry	Consumers MH CPs	Congestion modeled based on bandwidth division	<ul style="list-style-type: none"> • Social Welfare: lower • Content Innovation: higher
Guo et al. (2016)	ISP: Duopoly CP: Duopoly	Full	Heterogeneous in content and ISP preference	Heterogeneous in revenue rate	Two CPs with no entry	Consumers SH CPs	M/M/1 queuing system	<ul style="list-style-type: none"> • Social Welfare: lower • Content Innovation: the dominant CP may be worse off

Papers	Modeling Assumptions / Market Conditions Considered							Trade-offs considered: Findings for Net Neutrality Case
	Market Structure	Broadband Market Coverage	Consumer Heterogeneity	CP Heterogeneity	CP Entry	Multi-homing MH / Single-homing SH	Network Congestion	
Peitz and Schuett (2016)	ISP: Monopoly CP: No Competition	Full	Homogeneous	Heterogeneous in congestion sensitivity	Continuum of CPs with no entry	Consumers MH CPs	Congestion modeled as probability of on time delivery of content	• Social Welfare: lower
Reggiani and Valletti (2016)	ISP: Monopoly One large CP and many small CPs	Full	Homogenous	Heterogeneous in number of content (large vs small) and development costs for content	Continuum of CPs with entry	Consumers MH CPs	M/M/1 queuing system	<ul style="list-style-type: none"> • Social Welfare: lower or higher, depending on ad revenue rate • Content Innovation: higher for small CPs, lower or higher for large CP. • Infrastructure Investment: lower or higher, depending on innovation outcome

Notes: The modeling assumptions and thus the considered market conditions of the papers listed here vary with respect to several dimensions as reflected in the columns above:

- First, different market structures are assumed that consider either monopolistic or duopolistic ISPs and/or CPs, as well as vertically integrated or non-integrated ISPs. If vertically integrated ISPs are considered, the focus of the analysis is on studying the ISPs incentives to sabotage rival CPs' content (Trade-off 1).
- Second, many papers assume full market coverage, in which case consumer entry cannot be considered.
- Third, some papers assume that consumers are homogeneous, whereas others assume that consumers are heterogeneous. In the latter case, it is with regard to their preferences for either ISP or content.
- Fourth, CPs may either be considered as homogeneous or are treated as heterogeneous in one (and only one) respect, which varies from markups, revenues, and entry costs, to congestion sensitivity and consumer valuation. Thus, the impacts of traffic management mechanisms on different CPs vary. Also note that when modeling heterogeneous CPs or consumers, all papers in the literature consider CPs or consumers to be heterogeneous along one dimension and the same along all other dimensions.
- Fifth, some papers take the number of active CPs in the market to be fixed, whereas others allow for entry and exit of CPs in response to changing market conditions.
- Sixth, depending on the market structure, consumers (CPs) may multi-home or single-home CPs (ISPs). In cases where CPs form a duopoly, CP entry is not considered, and consumers usually single-home CPs, meaning they acquire content from just one CP. Where there is a continuum of CPs, the possibility of CP entry is sometimes accounted for, and consumers multi-home CPs. At the same time, consumers usually single-home ISPs, but sometimes consumers are also allowed to multi-home ISPs. Likewise, CPs are usually assumed to multi-home ISPs, but sometimes CPs may also single-home ISPs, which then leads to Internet fragmentation as not all content is available at every ISP.
- Finally, congestion is an important issue as a major motivation for packet prioritization and some papers incorporate it, typically with a standard M/M/1 queuing model. Recent efforts using a bandwidth division approach to model prioritized "fast lanes" enable further examination of congestion externalities that captures consumers' strategic response to congestion, i.e., consumer demand is reduced in response to increased congestion.