International Journal of Industrial Organization 45 (2016) 16–27

Contents lists available at ScienceDirect

International Journal of Industrial Organization

journal homepage: www.elsevier.com/locate/ijio

Net neutrality and innovation at the core and at the edge☆

Carlo Reggiani [a], Tommaso Valletti [b,*]

[a] School of Social Sciences, University of Manchester, UK
[b] Imperial College London, University of Rome II and CEPR. Address: Imperial College Business School, South Kensington Campus, London SW7, 2AZ, UK

ARTICLE INFO

Article history:
Received 23 December 2014
Received in revised form 31 December 2015
Accepted 31 December 2015
Available online 11 January 2016

JEL classification:
D4
L12
L43
L51
L52

Keywords:
Internet
Net neutrality
Congestion
Innovation

ABSTRACT

How would abandoning Internet net neutrality affect content providers that have different sizes? We model an Internet broadband provider that can offer a different quality of service (priority) to heterogeneous content providers. Internet users can potentially access all content, although they browse and click ads with different probabilities. Net neutrality regulation effectively protects innovation done at the edge by small content providers. Prioritization, instead, increases both infrastructure core investment and welfare only if it sufficiently stimulates innovation from the large content provider.

1. Introduction

The Internet has probably been the fastest developing industry of the last two decades. From the early development as an experimental network linking a limited number of computers, it has now become one of the key priorities for policy makers around the world, as it is seen as an engine to economic growth (Czernich et al., 2011; Mayo and Wallsten, 2011). The Internet is delivered by broadband providers who can use their infrastructure to set particular terms for access to applications and content (e.g., websites, services, protocols). These access terms are discussed under the heading of "net neutrality" (henceforth, NN), generating one of the most hotly debated issues in communications policy in the US, the EU and elsewhere.

NN is commonly defined as the principle for which all the traffic on the Internet should be treated equally and it has often been linked to the "end to end" principle.[1] These principles are thought to have guaranteed openness and free access to the Internet; their operation, however, has been questioned by the establishment of broadband as the standard delivering technology.[2] The US Federal Communications Commission (FCC) published in April 2015 the final rule on its new regulations on "Promoting and Protecting the Open Internet".[3] The NN debate received massive attention in the US: a record of 4 million comments were submitted to the FCC. Even President Obama made very strong comments in favor of NN in November 2014.[4]

From an economic viewpoint, the issue related to NN is that broadband allows for web traffic management techniques. These techniques can be used, e.g., for quality discrimination of data packets or use of termination charges for data traffic. From this angle, NN is mainly a data treatment (and pricing) issue with possible redistributive consequences. While the debate is complex, the following schematization is useful. On the one side stand proposers of a regulation that bans discrimination of data packets and guarantees open and equal access to the net (or "openists", according to Wu, 2004); on the other side it is believed that

☆ We thank the editor, two anonymous referees, as well as Bruno Jullien, Fabio Manenti, Joacim Tag, Greg Taylor, Joao Vareda, and seminar participants at Bern, Bologna, Manchester, Northwestern University, Rhodes (CRESSE), Stockholm (EARIE), Telecom ParisTech, and Valencia (JEI).

* Corresponding author.
E-mail addresses: carlo.reggiani@manchester.ac.uk (C. Reggiani), t.valletti@imperial.ac.uk (T. Valletti).

[1] The principle was that the transmission and routing of Internet traffic should be "dumb", not interfering with information packets sent between sender and receiver (Saltzer et al., 1984).

[2] Broadband adopts the TCP/IP protocols that allow discrimination between data packages.

[3] https://www.fcc.gov/document/protecting-and-promoting-open-internet-nprm.

[4] http://www.nytimes.com/2014/11/11/technology/obama-net-neutrality-fcc.html? _r = 0. In Europe, the European Commission introduced highly debated rules in October 2015.

http://dx.doi.org/10.1016/j.ijindorg.2015.12.005
0167-7187/

the Internet needs no regulation and will develop better by letting the market forces operate freely (or "deregulationists").

Valid arguments have been proposed by both sides. One of the main stances of "openists" is that NN is needed to protect the innovation of small start up content providers (henceforth, CPs), where among those there may be tomorrow's giants like Facebook or Google. Innovation at the "edge" of the network is one of the key attributes of the Internet and discrimination constitutes a potential harm to it (Lessig, 2001; Lee and Wu, 2009). On the other hand, the main counter argument of "deregulationists" is based on the need of Internet service providers (henceforth, ISPs) to get an appropriate remuneration for the use of the infrastructure, which is seen as the best way to guarantee investment for maintenance and expansion of the capacity of the network (the "core" of the Internet). This concern is becoming more prominent due to the increasing diffusion of bandwidth-intensive applications (Yoo, 2005; Van Schewick, 2006; Becker et al., 2010).

In this paper we develop a two-sided model of the Internet to analyze the possible tensions at the"core" versus the "edge". We focus on the two polar cases of NN regulation and priority pricing (henceforth, PP). In the model, a monopolist ISP allows CPs to reach final users of the Internet. Importantly, the model captures one of the defining features of the Internet, that is, the heterogeneous size of CPs. In particular, we assume there is one large, established CP and also a number of small CPs that constitute a fringe. The main contribution of our paper is to investigate how a NN regulation would affect the division of resources between different players and their incentives to innovate. To the best of our knowledge, the heterogeneous size of the CPs has not been explicitly addressed in the formal literature on NN. In our model, CPs fund themselves through advertising revenues that are related to the clicks received on the content supplied. The likelihood of being clicked, and consequently the resources available through advertising, can be affected by the priority regime. Final users, on their side, desire to potentially access as much content and applications as possible. The ISP owns the infrastructure to connect the CPs to end users. Its incentives to invest in maintaining and extending the network, which affects congestion, is also ultimately related to the regime adopted.

The analysis of the model provides important results, including some "rules of thumb" for policy makers. First, NN or PP do not per se affect congestion: for a given level of traffic, average congestion is unaffected by the priority regime. The priority regime, however, alters the incentives of the ISP to build capacity. In particular, PP leads to more investment at the core only if it sufficiently boosts innovation from the established large CP: we show that this is a necessary but not sufficient condition. Second, NN is the appropriate regime to protect innovation at the "edge", as it guarantees more supply of content and apps from small CPs. On the other hand, there are circumstances where the overall content supplied is higher under PP as, in that case, the large provider expands its supply to more than compensate for loss of content from the fringe of small CPs. This result is related to the crucial role played in our model by the advertising revenues per click. When the revenues per click are not too high, the larger CP reacts positively to the decrease in the content supplied by the fringe, implied by PP: this may lead to higher overall content. However, as the ad revenues increase, the strategic response of the large provider to reduced content from the fringe is to reduce its content supply as well: this avoids a "cannibalization" effect and a reduction of the clicks received overall. The size of advertising revenues is thus critical in our results and is related to the effectiveness of ads in online platforms: the more effective adverts, the less likely that PP can be welfare enhancing. Thus we also study the extent to which private incentives of the ISP to adopt a particular regime are not aligned with those of a social planner. There is only one limiting case, when content is "king" (that is, content from any CP is highly valued), in which the private and social choice of the regime always coincide.

While the model is developed having in mind the specificities of the NN debate that centres around CPs and ISPs, we note that it can also be reinterpreted to deal with other platform environments. For instance, think of Facebook that invests in a platform that hosts different user-generated content as well as advertising banners. Facebook can use data analytics to allow for more targeted advertising. The regime equivalent to prioritization would be one where Facebook is allowed to price discriminate among advertisers based on data analytics. Instead, a neutral regime is one where all advertisers must be treated equally. Of course there are differences, as Facebook users do not pay (contrary to the ISP's subscribers), and even uniform advertising rates would be positive (while, under NN, termination fees are set as zero). Still, the analogy lies in an economic environment with a multi-sided platform and where some policy regulation affects prices and transactions on one side, having deep consequences over the entire ecosystem.

The rest of the paper is structured as follows. Section 2 briefly reviews the most closely related work to better locate the contribution of the paper. Section 3 introduces the basic model. Section 4 presents the results of the analysis. Section 5 discusses the robustness of the findings. Section 6 concludes with policy implications.

2. Related work and contribution

Two facts have been long recognized about work on NN: first, as the debate is fierce, much has been written about it;[5] second, there is a disproportion between the law, policy and advocacy papers and the formal economic analysis. In recent years, however, the economics literature has developed at a fast pace and in different directions.[6]

The structure of the Internet industry naturally invokes a two-sided market approach: ISPs are the platforms that connect CPs to final users. Economides and Tag (2012) present a static model of charges imposed by the ISP to CPs for "traffic termination" to consumers. Their paper shares with ours the approach but key problems of the Internet such as traffic congestion and bandwidth allocation and ISP's investment are not addressed.

Cheng et al. (2011), Choi and Kim (2010) and Kramer and Wiewiorra (2012) use, as we do, the M/M/1 approach: borrowed from queuing theory, it is considered a good proxy for actual congestion on the Internet. Cheng et al. (2011) and Choi and Kim (2010) consider similar models in which users access exclusively only one of two content providers;[7] total supply of content is fixed so priority only affects the market shares. In Cheng et al. (2011) both CPs can get priority. This leads to a prisoners' dilemma and similar CPs both buy priority: the result is no effect on congestion and more surplus extracted by the ISP. Choi and Kim (2010) consider the case in which CPs bargain with the ISP to obtain exclusive priority for their traffic; CPs are charged a fee only if they opt for priority. The impact of NN on investment crucially depends on the inelastic content supply of the CPs: as more capacity means less value for priority, the ISP has less incentive to invest when NN is abandoned. Kramer and Wiewiorra (2012) consider a continuum of CPs differently sensitive to congestion. In the long run, the welfare superior regime is the one leading

[5] In November 2015, a casual SSRN search returned 450 papers with "network neutrality" in the title or abstracts. A similar Google search provided over 11 million hits.

[6] For example, recently, Njoroge et al. (2013), Bourreau et al. (2015) and Choi et al. (2015) have looked at the role of competition between ISPs in the context of NN, while Kourandi et al. (2015) and D'Annunzio and Russo (2015) consider competition among CPs. Peitz and Schuett (2015) and Jullien and Sand-Zantman (2015) study traffic inflation in presence of congestion sensitive data. Gans (2015) considers direct payments from consumers to CPs, which are ignored in other models. Schuett (2010), Kramer et al. (2013) and Greenstein et al. (2016) provide overarching surveys on NN.

[7] While this might be a characterization of particular situations where content providers are substitutes between each other (e.g., a subscriber will typically want to use only one search engine, and will decide, for instance, between either Google or Bing), it cannot capture the fact that most of the Internet content has a different nature, that is, subscribers want to see (and do see) both Google and Facebook, which cannot be modelled as mutually exclusive choices.