Neutral and Non-Neutral Countries in a Global Internet: What Does it Imply?

Patrick Maillé, Bruno Tuffin

► To cite this version:

Patrick Maillé, Bruno Tuffin. Neutral and Non-Neutral Countries in a Global Internet: What Does it Imply?. GECON 2019: 16th International Conference on the Economics of Grids, Clouds, Systems, and Services, Sep 2019, Leeds, United Kingdom. pp.111-123, 10.1007/978-3-030-36027-6_10 . hal-02128437v2

HAL Id: hal-02128437

<https://hal.inria.fr/hal-02128437v2>

Submitted on 4 Jul 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Neutral and Non-Neutral Countries in a Global Internet: What Does it Imply?

Patrick Maillé¹ and Bruno Tuffin²

¹ IMT Atlantique, IRISA, UMR CNRS 6074, F-35700 Rennes, France

`patrick.maille@imt.fr`

² Inria, Univ Rennes, CNRS, IRISA, Rennes, France

`bruno.tuffin@inria.fr`

Abstract. Network neutrality is being discussed worldwide, with different countries applying different policies, some imposing it, others acting against regulation or even repealing it as recently in the USA. The goal of this paper is to model and analyze the interactions of users, content providers, and Internet service providers (ISPs) located in countries with different rules.

To do so, we build a simple two-regions game-theoretic model and focus on two scenarios of net neutrality relaxation in one region while it remains enforced in the other one. In a first scenario, from an initial situation where both regions offer the same basic quality, one region allows ISPs to offer fast lanes for a premium while still guaranteeing the basic service; in a second scenario the ISPs in both regions play a game on quality, with only one possible quality in the neutral region, and two in the non-neutral one but with a regulated quality ratio between those.

Our numerical experiments lead to very different outcomes, with the first scenario benefiting to all actors (especially the ones in the relaxed-neutrality region) and the second one mainly benefiting mostly to ISPs while Content Providers are worse off, suggesting that regulation should be carefully designed.

Keywords: Net neutrality · Service differentiation · Game theory · Regulation.

1 Introduction

The network neutrality debate has been raging for close to two decades and is still a very sensitive issue worldwide. Basically, network neutrality is “the principle that traffic should be treated equally, without discrimination, restriction or interference, independent of the sender, receiver, type, content, device, service or application”. This type of definition was introduced by the Federal Communications Commission, the regulator in the US in 2005, and by the European Union in 2014, among others. It is part of the principles for an open Internet according to which resources available on the Internet should be easily accessible to all entities. The debate has been highlighted in 2005 with Ed Whitacre, CEO of the Internet service provider (ISP) AT&T, complaining that distant content

providers (CPs) were using his network without financially participating to its infrastructure maintenance and upgrade, while at the same time the proportion of telecommunications economy coming from advertisement and going to CPs was increasing. The threat to differentiate traffic or even block some services raised a lot of protests. Since, numerous attempts to discriminate traffic have been observed, such as ISP Madison River Communications fined in 2005 for preventing its clients from using VoIP in competition with its own “voice” offer, Comcast blocking in 2007 BitTorrent (P2P) traffic, or the recent exclusion of some traffic from data caps in wireless subscription offers (the so-called zero rating). For more on Net neutrality and its history, the reader is advised to look at [5, 7–9, 11, 13, 14] and references therein.

An important issue barely addressed in the literature is that while neutrality principles are imposed in many countries, it is not the case everywhere. As of March 2019, we can define two, and even three sets of countries regarding net neutrality³:

- Countries having passed laws to protect neutrality; this includes all European Union, Canada, most of South America, Japan, India, etc. Remark though that rules are more or less strict depending on the country; for example some authorize sponsored data (that is, the possibility for content providers to pay for their traffic and exclude it from the users data cap) while other don’t.
- Countries against neutrality as recently the USA, or other big countries such as Russia, China, etc. Claimed reasons are not always the same: economic efficiency in the USA, control of content/traffic by deep packet inspection in China, or congestion control in Russia.
- Countries still in the process of deciding, such as Australia or Uruguay.

As a consequence, in a global Internet, some ISPs are allowed to differentiate service for their users but also of traffic of CPs originated or going through their country. This difference of rules could end up with differentiated services for users located in a foreign and neutral country, even if neutrality is imposed and applied there. This issue is particularly exacerbated by the recent decision, in 2017, of the authorities in the USA to repeal neutrality. The USA being the origin or intermediate of an important part of worldwide traffic, it seems to us that studying the relations between countries applying different neutrality policies is becoming particularly important.

Our goal is therefore to investigate what it implies on all actors (ISPs, CPs, users) to have both neutral and non-neutral countries. To start, we will limit ourselves in this paper to two interacting countries. The actors we consider in each country are CPs, potentially deciding between classes of service for their traffic, ISPs deciding connection prices for given qualities of service (QoS) offered to CPs, and users with demand level depending on the QoS they experience. The interactions between countries come from users requesting traffic to CPs in

³ See https://en.wikipedia.org/wiki/Net_neutrality_by_country for an exhaustive list or https://www.reddit.com/r/MapPorn/comments/7k9wus/status_of_net_neutrality_around_the_world2060x1400/ for an instructive map

their own and also in foreign countries, and from traffic having to go through ISPs in the two countries, potentially applying different neutrality policies. Non-cooperative game theory [4, 12] is used to analyse the interactions of selfish actors. The question we would like to answer is: Is there a “winner” with such a heterogeneous situation? We wish to compare the output with the situation of a fully neutral Internet.

The literature on modeling and analysis of network neutrality through game theory has been extensive (see among others [1–3, 6, 9, 10] and references therein), trying to answer various questions, but to our knowledge no work has been dealing with the impact of interactions between countries applying different rules. Again, given the current tendency of countries to evolve in different political directions, this is becoming an issue of primary interest.

The results provided in this paper highlight how the specifics of net neutrality regulation relaxation can affect all actors, including those in the region where neutrality remains enforced. Our numerical results show in particular that if the non-neutral ISP still has to offer the same quality as the neutral ISP (in addition to an improved quality), then all actors are likely to benefit from the relaxation, while only the non-neutral ISP would benefit if it is given more freedom (even its hosted CPs would prefer the all-neutral situation).

The remainder of the paper is organized as follows. Section 2 introduces our two-zones, several-CP model with the available strategies and utility functions for all actors. We then investigate in Section 3 the case when the non-neutral ISP can only offer a high-quality service in addition to the basic one (the one in the neutral zone). In Section 4 we analyze numerically, on the same instance, what happens when the regulator imposes a given ratio between the low-quality and high-quality services offered by the non-neutral ISP, while the neutral ISP fixed its quality to maximize its revenue. We provide conclusions and suggest directions for future work in Section 5.

2 Model

2.1 Topology and actors

The topology we study in this paper is described in Figure 1. We consider two geographic areas, with a single ISP and several CPs in each one; the set of CPs in the neutral (resp., non-neutral) area is denoted by $\mathcal{L}_N$ (resp., $\mathcal{L}_D$). We assume a peering relationship between the two ISPs. We will consider that net neutrality is enforced in one area, while neutrality constraints will be relaxed in the other one. The neutral area may correspond to the European Union and the non-neutral one to the United States, after net neutrality regulations have been repealed.

2.2 Service qualities and prices

Let q_N be the performance level offered by the neutral ISP, and assume that the non-neutral ISP can offer two different quality levels, namely q_L and q_H

Fig. 1. Topology

(with $q_L > q_H$). We assume $q_N, q_L, q_H \in [0, 1]$. Without loss of generality, the higher the value of q , the better the performance. The per-unit-of-volume prices (paid by CPs to their host ISP) are p_N (in the neutral area), and p_L and p_H for performance levels q_L and q_H , respectively, in the non-neutral area.

Prices will be assumed fixed, imposed by the regulator(s) or decided through competition: we only focus on quality levels as decision variables for ISPs.

2.3 CP demands and utilities

CPs also have a per-unit-of-volume gain (from advertisement) a_ℓ for CP ℓ . Each CP has a volume demand coming from its own area, plus another one coming from the other area, both depending on the offered performance q (defined just after). Let $D_i^j(q)$ be the demand of content of CP j from customers in area i . We will for example later consider the standard linear expressions $D_i^j(q) = \beta_i^j + \alpha_i^j q$. Users in area i and looking at CP j located in that same area just need to use the local ISP. Thus $q = q_i$, the performance at ISP i chosen (if a choice is available) by that CP. On the other hand, when accessing the CP j from the other area, traffic goes through both ISPs, and overall performance is assumed to be the product of the two ISP qualities $q = q_N \cdot q_D$, so that a null quality at a network leads to a null quality along the path, and a perfect quality at a node (that is $q_N = 1$ or $q_D = 1$) reduces the overall quality to the quality of the other node along the path.

The utility of CP ℓ in the non-neutral area is the difference between its advertising revenue and payments to its ISP, both of which are proportional to its total demand:

ℓ in non-neutral area $\Rightarrow$

$$U_\ell^C = (a_\ell - p_{k_\ell}) (D_D^\ell(q_{k_\ell}) + D_N^\ell(q_{k_\ell} q_N)),$$

where $k_\ell \in \{L, H\}$ is the performance level chosen by the CP.

Similarly, the utility of a CP ℓ in the neutral area is given by

ℓ in neutral area $\Rightarrow$

$$U_\ell^C = (a_\ell - p_N) D_N^\ell(q_N) + (a_\ell - (p_N + p_{k_\ell} - p_L)) D_D^\ell(q_N q_{k_\ell}),$$

with decision $k_\ell \in \{L, H\}$ representing the quality that the CP has selected for its flows reaching the non-neutral zone. Note that we assume that if the CP selects the high quality, it then has to pay an extra $p_H - p_L$ per demand unit to the non-neutral ISP.

2.4 ISP utilities

The revenues of ISPs are made of the gain from the volume of data flowing through it, minus the cost for maintaining a network with the given performance level. We therefore get, for the ISP in the neutral area,

$$U_N^I = p_N \sum_{\ell \in \mathcal{L}_N} (D_N^\ell(q_N) + D_D^\ell(q_N q_{k_\ell})) - f_N,$$

and for the ISP in the non-neutral area we obtain

$$U_U^I = \sum_{\ell \in \mathcal{L}_N} (p_{k_\ell} - p_L) D_D^\ell(q_N q_{k_\ell}) + \sum_{\ell \in \mathcal{L}_D} p_{k_\ell} (D_D^\ell(q_{k_\ell}) + D_N^\ell(q_{k_\ell} q_N)) - f_D,$$

where f_N and f_D are the cost functions. We assume that the cost borne by an ISP is made of the demand level at the quality times the (unit) cost $c_i(q_i)$ ($i \in \{N, H, L\}$) to provide this level:

$$\begin{aligned} f_N &= \left(\sum_{\ell \in \mathcal{L}_N} (D_N^\ell(q_N) + D_D^\ell(q_N q_{k_\ell})) + \sum_{\ell \in \mathcal{L}_D} D_N^\ell(q_{k_\ell} q_N) \right) c_N(q_N) \\ f_D &= \left(\sum_{\ell \in \mathcal{L}_D: k_\ell=L} [D_D^\ell(q_L) + D_N^\ell(q_L q_N)] + \sum_{\ell \in \mathcal{L}_N: k_\ell=L} D_D^\ell(q_N q_L) \right) c_D(q_L) \\ &\quad + \left(\sum_{\ell \in \mathcal{L}_D: k_\ell=H} [D_D^\ell(q_H) + D_N^\ell(q_H q_N)] + \sum_{\ell \in \mathcal{L}_N: k_\ell=L} D_D^\ell(q_N q_H) \right) c_D(q_H). \end{aligned}$$

Remark that total demand through a network is composed of the total demand of subscribers to this network, but also of the demand of users of the other network to CPs in this local network.

The cost functions c_N and c_D will be assumed increasing convex, with value 0 at 0 and $+\infty$ at 1. The typical example we will use is $c_N(q) = c_D(q) = \alpha \frac{q}{1-q}$ with some conversion rate α .

2.5 Analysis of the interactions

Each actor takes its own decision, but not at the same time scale. We end up with a Stackelberg game, where:

1. ISPs play a game on performance levels, the one in the neutral area choosing q_N , the one in the differentiated area choosing q_L and q_H (prices being fixed, as described previously);

2. The CPs decide the class of service in the differentiated zone;
3. Demand is computed depending on those strategies.

The game is played by backward induction, i.e., each decision maker is assumed to be able to anticipate the outcome of the later stages when selecting an action.

2.6 Different ways of relaxing neutrality

We take as a reference basis the situation where both areas are neutral, and investigate what happens when one region allows some non-neutral behavior from its ISP. To avoid extremely unfair behaviors, we assume that there remains some (relaxed) regulation regarding the non-neutral ISP's actions. For example, we may assume that the ISP may only offer an *improved* service in addition to that of the reference situation (our first scenario), or that the regulator imposes a *fixed quality difference* between the high- and low-quality services. Of course, other rules can be imagined, but we think those two are sufficiently simple and realistic to be worth considering.

3 A first scenario: opening to non-neutrality in a zone from an all-neutral situation

In this section, we investigate some possible outcomes when, from a situation where both zones are neutral and provide the same quality q_N , the possibility of creating fast-lanes is opened in one zone while the minimum quality should remain q_N . This could for example correspond to the situation in the United States, where previous neutrality recommendations have been repealed.

With our previous notations, this corresponds to q_N being fixed, and the new decision variable being q_H .

The other parameters are also assumed fixed, with values given below.

$$\begin{array}{cccc} p_L & p_H & p_N & \alpha \\ \hline 0.5 & 0.6 & 0.5 & 0.02 \end{array}$$

Those values have been chosen quite arbitrarily, but so that the players' strategies have an impact. We nevertheless think they can be realistic.

The sets of CPs are $\mathcal{L}_N = \{1_N, 2_N, 3_N\}$ and $\mathcal{L}_D = \{1_D, 2_D, 3_D\}$. We consider linear demand functions for CPs, given in Table 1. In particular, we are considering pairs of CPs, i.e., CPs that have the same demand functions in both regions and only differ by being attached to the neutral or the non-neutral ISP. This will help analyze whether a given CP prefers being hosted by a neutral or a non-neutral ISP. We also consider three kinds of CPs, with one kind (CPs 1_N and 1_D) equally of interest in both areas, and each of the two other types mainly of interest for users in one area: CPs 2_N and 2_D rather target users in the non-neutral area, while CPs 3_N and 3_D produce content that interests mostly users in the neutral area.

CP index	CP location	CP adv. rev. a	Demand in neutral region	Demand in non-neutral region
1 _N	Neutral ISP	1	$0.1+1.0q$	$0.1+1q$
2 _N	Neutral ISP	1	$0.05+0.5q$	$0.2+2q$
3 _N	Neutral ISP	1	$0.1+1.0q$	$0.036+0.36q$
1 _D	Non-Neutral ISP	1	$0.1+1.0q$	$0.1+1q$
2 _D	Non-Neutral ISP	1	$0.05+0.5q$	$0.2+2q$
3 _D	Non-Neutral ISP	1	$0.1+1.0q$	$0.036+0.36q$

Table 1. CP demand functions and advertisement revenue factors. **Fig. 2.** CP utilities versus q_H , for $q_L = q_N = 0.5$. The situation when $q_H = 0.5$ corresponds to the all-neutral case. The marks highlight the values after which a CP selects q_H rather than q_L .**Fig. 3.** ISP revenues versus q_H when $q_L = q_N$.

We display in Figure 2 the utilities of all CPs when q_H varies, with $q_N = q_L$ fixed. We first remark that, as expected, when $q_H = q_N$ each CP is better off being hosted in the region where most of its demand lies:

- being equally attractive in both regions, CPs 1_N and 1_D get the same utility, suggesting that a “virtual CP1” considering where to locate its service would be indifferent;
- similarly, a “virtual CP2” (resp., CP3) asking the same question would compare the utility of our CPs 2_N and 2_D (resp., CPs 3_N and 3_D) and prefer to be in the situation of CP 2_D (resp., 3_N). This was to be expected, since CPs will seek to minimize the multiplicative impact of qualities on the demand from the alien area.

As q_H increases, for our parameter values each CP opts for the improved quality after a certain point (indicated in the figure). Note that the values for this point conform to the intuition: a CP will switch sooner to high-quality if the impact on its demand is larger, i.e., if its target market is in the non-neutral zone and if the CP is itself hosted in the non-neutral zone. For the same reason, the improved quality q_H favors more the CPs located in the non-neutral area than their neutral-area-located counterparts, due to the multiplicative term in the demand from the area remote to the CP. When q_H becomes high, all CPs – including the one mostly targeting users in the neutral area, our “virtual CP3” – prefer being hosted in the non-neutral area than the neutral one.

The net revenues (utilities) of both ISPs are also plotted in Figure 3, illustrating how the non-neutral ISP could choose its high-quality level q_H . We notice some discontinuities, corresponding to CPs switching from low-quality q_L to high-quality q_H . Those switches first have no noticeable impact on the neutral ISP, but then a positive impact since demand increases with q_H ; however we do not consider here what would happen if CPs decide to switch regions (as we pointed out above, the non-neutral ISP becomes more attractive to CPs as q_H increases).

Finally, in Figures 4 to 6 we vary the common value of q_H and q_L , and consider that the non-neutral ISP selects the quality q_H maximizing its net revenue: that best-response quality q_H is plotted in Figure 4, while Figures 5 and 6 show the utilities of the actors. As could be expected, when the “low-quality” levels $q_N = q_L$ is high enough, a high-quality level does not make a significant enough difference, and is therefore not implemented by the non-neutral ISP. With our parameter values, this happens when q_N exceeds 0.71. Note also that before that point, the optimal q_L is not monotone in q_N : it first slowly decreases for low values of q_N , but for q_N high enough there seems to be a need to maintain q_H a certain level above q_N to keep attracting CPs, i.e., justifying the price difference, as illustrated in Figure 3 when $q_N = 0.65$.

The impact on each individual CP is shown in Figure 5: when the neutral (basic) quality is low, all CPs prefer being hosted by the non-neutral ISP since they can benefit from its high-quality level and the corresponding demands. However, when that basic quality increases, the gain from the high-level quality q_H is less significant, and after some threshold (around 0.35 for our example)

Fig. 4. Best-response $q_H(q_N)$ when $q_L = q_N$ (no value is shown when offering a high-quality service is not beneficial, i.e., when the non-neutral ISP prefers offering only q_N).

Fig. 5. CP utilities (revenues) versus $q_N = q_L$, with best-response q_H from the non-neutral ISP.

the CP mainly targeting users in the neutral zone prefers to be hosted in that zone. The CP targeting the non-neutral area always prefers being hosted by the non-neutral ISP, while the CP equally demanded in both regions prefers the non-neutral one until the non-neutral ISP non longer offers service differentiation.

We also display in Figure 6 the revenues of both ISPs, and the cumulative revenues of the CPs in each zone. This allows to see what an advantage it is for

Fig. 6. ISP and cumulative CP utilities (revenues) versus q_N , with best-response q_H .

an ISP to be allowed to offer an improved service; we also see the two regimes (when it is beneficial for the non-neutral ISP to offer $q_H > q_N$, and when it is not).

Finally, in Figures 7 and 8 we compare the all-neutral situation (when the same quality q_N is offered in both areas) to the partially neutral one studied here, where a high-quality q_H can be offered in the non-neutral region. We observe that for our parameters, relaxing the neutrality constraint benefits to all actors, except (for a small range of values of q_N) the neutral ISP. All CPs benefit from

that relaxation, in particular those hosted by the non-neutral ISP. Note however, again, that we did not consider CP mobility in our model (CPs that would switch ISPs).

Fig. 7. CP utilities (revenues) differences when switching from an all-neutral situation to a non-neutral ISP optimizing q_H , versus $q_N = q_L$.

Fig. 8. ISP utilities (revenues) differences when switching from an all-neutral situation to a non-neutral ISP optimizing q_H , versus $q_N = q_L$.

4 A second scenario: quality game between ISPs, with a fixed q_H/q_L ratio in the non-neutral region

We now consider a different scenario, where both ISPs play a non-cooperative game on their qualities—the neutral ISP chooses q_N and the non-neutral ISP chooses q_L and q_H —but a fixed ratio between q_H and q_L is imposed to the non-neutral ISP, to avoid excessive differences between basic and improved services. For the numerical investigations, we consider the same CPs with their specificities (location and demand functions) as in the previous scenario. They are given in Table 1.

When the quality q_N is fixed and the non-neutral ISP varies q_H (and q_L to maintain the quality ratio), the utilities of the CPs are plotted in Figure 9. We observe the same trend as with the previous scenario: the CPs hosted in the non-neutral region prefer the “higher quality” q_H over q_L sooner (i.e., for smaller values of q_H) than their counterparts hosted in the neutral region. Also, CPs whose main demand lies in the non-neutral region switch sooner to q_H . Another thing worth noting is that for all CPs, the switching point is below the quality value q_N in the neutral network: here even if q_H cannot be said to be high quality, CPs still choose it to avoid the worse quality q_L .

The net revenues (utilities) of both ISPs are also plotted in Figure 10, illustrating how the non-neutral ISP could choose its high-quality level q_H . For our

Fig. 9. CP utilities versus q_H , for $q_N = 0.5$ and the ratio q_H/q_L fixed to 2. The marks show after which value of q_H the CPs prefer q_H over q_L .

Fig. 10. ISP revenues versus q_H , for $q_N = 0.5$ and the ratio q_H/q_L fixed to 2.

numerical values, the non-neutral ISP should choose a high quality around 0.8, and therefore a basic quality 0.4.

However, in this scenario we do not consider q_N fixed but rather determined through a non-cooperative game played between ISPs. To analyze that game, we plot the ISP best-responses (q_H versus q_N , since q_L is directly determined by q_H) in Figure 11, which exhibits continuous best-response that intersect at only one point, a (stable) Nash equilibrium. Note that the best-response q_H is above

Fig. 11. ISP best-responses when the ratio q_H/q_L is fixed to 2.

Fig. 12. Equilibrium qualities when the ratio q_H/q_L varies.

q_N , and seems to equal q_N when q_N is large enough. Here at the equilibrium we have $q_H > q_N$.

We then vary in Figure 12 the regulated ratio q_H/q_L , and plot the equilibrium values of the qualities. The resulting revenues of ISPs and CPs are displayed in Figures 13 and 14. The figures highlight two regimes:

Fig. 13. ISP net revenues at equilibrium when the ratio q_H/q_L varies.

Fig. 14. CP net revenues at equilibrium when the ratio q_H/q_L varies.

- when the ratio q_H/q_L is low, the non-neutral ISP cannot set a significantly higher quality q_H to justify the price difference, and all CPs prefer the basic quality q_L . The results therefore do not depend on the specific value of the regulated ratio, provided it is low enough. Note that the equilibrium quality q_L , the one all CPs choose, is slightly below the quality of the neutral ISP, even though their price is the same. That regime is actually equivalent to the all-neutral equilibrium (when the ratio is 1, i.e., no differentiation is allowed); we have the same qualities (except q_H , that is not chosen by any CP) and the same utilities for stakeholders. Therefore, such a limited relaxation of neutrality constraints has no impact.
- When the regulated ratio q_H/q_L is large enough, the price difference is justified by the quality difference, hence all CPs select that high-quality. Again, as a result the equilibrium qualities do not depend on the regulated ratio in that regime. Note that in this regime, the chosen quality in the non-neutral zone q_H is strictly above the one in the neutral zone. It is also worth noting that both ISPs—even the neutral one—prefer this regime over the all-neutral situation, while it is the opposite for CPs: all of them were better off in the all-neutral setting than in this relaxed scenario, the most affected being the CPs hosted in the non-neutral zone.

Between those two regimes, there is a small range of values for the regulated ratio, where some CPs would select q_L and others q_H . The impacts on the actors in that limited range are less clear.

5 Conclusions

Changing the net neutrality rules in a part of the Internet may affect actors in other parts; in this paper we have focused on a simple scenario with two domains having different regulations, and have investigated the resulting decisions from ISPs (in terms of offered qualities) and the consequences on content providers.

The specifics of the new rules can lead to very different outcomes. With the same parameter values, our numerical analysis has for example shown that, when a non-neutral ISP can offer two qualities q_H and q_L and a neutral ISP only one quality q_N :

- if we impose $q_L = q_N$, i.e., non-neutrality is only allowed to offer improved service, then all stakeholders—in particular the non-neutral ISP and its hosted CPs, but also the other ones to a lesser extent—would prefer a relaxation of neutrality rules;
- if a given ratio q_H/q_L is imposed, then either nothing is changed with respect to the neutral situation, or all CPs are worse off.

Those results suggest that regulatory decisions about relaxing net neutrality rules should be made with great care, taking into account the whole ecosystem that is affected.

The setting described in this paper opens several perspectives for future work. A first direction would be to develop the analytical study of the model proposed here rather than investigate numerical examples. Also, one can imagine other versions of regulation to constrain the non-neutral ISP in other ways, with the objective to favor user welfare, CP innovation, or fairness among users or among CPs.

References

1. E. Altman, A. Legout, and Y. Xu. Network non-neutrality debate: An economic analysis. In *Proc. of Networking*, pages 68–81, Barcelona, Spain, May 2011. LNCS 6641, Springer Verlag.
2. P. Coucheney, P. Maillé, and B. Tuffin. Impact of Reputation-Sensitive Users and Competition Between ISPs on the Net Neutrality Debate. *IEEE Transactions on Network and Service Management*, 2013.
3. P. Coucheney, P. Maillé, and B. Tuffin. Network neutrality debate and ISP inter-relations: Traffic exchange, revenue sharing, and disconnection threat. *Netnomics*, 1(3):155–182, 2014.
4. D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, Cambridge, Massachusetts, 1991.
5. T.M. Lenard and R.J. (Eds.) May. *Net Neutrality or Net Neutering: Should Broad-band Internet Services be Regulated*. Springer, 2006.
6. T.T.B. Ma, D.M. Chiu, J.C.S. Lui, V. Misra, and D. Rubenstein. On cooperative settlement between content, transit, and eyeball internet service providers. *IEEE/ACM Transactions on Networking*, 19(3):802–815, June 2011.

7. P. Maillé, P. Reichl, and B. Tuffin. Internet governance and economics of network neutrality. In A. Hadjiantonis and B. Stiller, editors, *Telecommunications Economics - Selected Results of the COST Action IS0605 EconTel*, pages 108–116. Lecture Notes in Computer Science 7216, Springer Verlag, 2012.
8. P. Maillé, G. Simon, and B. Tuffin. Toward a net neutrality debate that conforms to the 2010s. *IEEE Communications Magazine*, 54(3):94–99, 2016.
9. P. Maillé and B. Tuffin. *Telecommunication Network Economics: From Theory to Applications*. Cambridge University Press, 2014.
10. P. Njoroge, A. Ozdaglar, N. Stier-Moses, and G. Weintraub. Investment in two sided markets and the net neutrality debate. Technical Report DRO-2010-05, Columbia University, Decision, Risk and Operations Working Papers Series, 2010.
11. A. Odlyzko. Network neutrality, search neutrality, and the never-ending conflict between efficiency and fairness in markets. *Review of Network Economics*, 8(1):40–60, 2009.
12. M. Osborne and A. Rubinstein. *A Course in Game theory*. MIT Press, 1994.
13. H. Schulzrinne. Network neutrality is about money, not packets. *IEEE Internet Computing*, 22(6):8–17, 2018.
14. T. Wu. Network neutrality, broadband discrimination. *Journal of Telecommunications and High Technology*, 2003.