



HAL
open science

Impact of Heterogeneous Neutrality Rules with Competitive Content Providers

Khushboo Agarwal, Patrick Maillé, Bruno Tuffin

► **To cite this version:**

Khushboo Agarwal, Patrick Maillé, Bruno Tuffin. Impact of Heterogeneous Neutrality Rules with Competitive Content Providers. IFIP/IEEE International Symposium on Integrated Network Management, May 2021, Bordeaux, France. hal-02463485

HAL Id: hal-02463485

<https://inria.hal.science/hal-02463485v1>

Submitted on 1 Feb 2020

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.

Impact of Heterogeneous Neutrality Rules with Competitive Content Providers

Khushboo Agarwal

IIT Bombay

Mumbai, India

agarwal.khushboo@iitb.ac.in

Patrick Maillé

IMT Atlantique

Rennes, France

patrick.maille@imt.fr

Bruno Tuffin

Inria, Univ Rennes, CNRS, IRISA

Rennes, France

bruno.tuffin@inria.fr

February 1, 2020

Abstract

Network neutrality has recently been repealed in the United States, leading to a worldwide Internet with different imposed policies. We build and analyze in this paper a game-theoretic model representing the interactions between users, network providers and content providers in this heterogeneous regulation context, and investigate the impact of two neutrality relaxation policies in a part of the world on all actors, compared with a fully-neutral network. Our results show that neutrality repeal may *only* favor the ISP in the differentiation-authorized zone, but no other actor, and that it can be worse off for everybody if the regulation procedures are very strict in the neutral area.

1 Introduction

The *network neutrality* debate is still active, more than 15 years after Tim Wu first coined the term [15]. Indeed, while a consensus was slowly seeming to emerge among regulators toward protecting neutrality (in the USA from the FCC in 2015 [5] and in Europe with the BEREC 2016 Net Neutrality guidelines [2]), major countries like the United States have recently taken steps to repeal those protections in order to favor market freedom and stimulate economic growth, by “remov[ing] unnecessary regulations to promote broadband investment” [6].

Detailed technical interpretations slightly differ among countries and stakeholders, but network neutrality can be summarized as the principle according to which “all

traffic should be treated equally, without discrimination, restriction or interference, independent of the sender, receiver, type, content, device, service or application” [2]. In particular, in neutrality-enforced networks, network service providers (or ISPs, Internet Service Providers) are not allowed to charge their customers (users, or content providers (CPs)) extra for an improved service. For more on net neutrality, the arguments of both sides, and the history of the debate, the reader is advised to look at [7, 9, 11, 10, 13, 14, 15] and references therein.

With the different regulation stances among countries, we end up with a global network (the internet) where some parts are subject to neutrality rules—which can be seen as constraints for ISPs and protections for users and CPs—, while others are not and ISPs have the flexibility to create differentiated services. But the network being global, those two worlds interact: users in neutral areas can access content hosted in non-neutral areas, and vice-versa, so that the decisions made in an area (e.g., service qualities) can affect the whole internet ecosystem.

The goal in this paper is to investigate the impact of those regulatory disparities on all the actors involved: ISPs, users, and CPs. In a preliminary work in that direction [12], we considered two areas, one neutral and one non-neutral (that can be seen as, e.g., the European Union and the United States after net neutrality regulations have been repealed, respectively), with some CPs based in each area. The CP demands (and hence, revenues) were assumed to depend only on the quality their users experienced, and CPs then had to select the most favorable quality for that purpose, when available. Our conclusions suggested that the type of regulatory relaxation (e.g., allowing the non-neutral ISPs to only offer quality improvement over a reference situation, or imposing a ratio between the highest and lowest quality levels) can have dramatically different consequences for stakeholders, ranging from benefitting all CPs and users to benefitting only ISPs at the expense of CPs and users.

An aspect that was missing in the model of [12], and that we develop here, is the price competition among CPs: in addition to playing on the quality (a discrete choice in the non-neutral zone and no choice in the neutral one), we now allow CPs to decide the price to charge for their service, and we introduce a specific demand function to model the interactions among CPs. The objective of that new model is to get closer to reality, with prices strongly affecting demands, so as to anticipate who the winners and losers of neutrality rules relaxation will be.

The comparison of the outputs in the three scenarios (fully neutral, partially neutral, and totally non-neutral in a zone) illustrate that repealing neutrality is always worse for CPs, users, and the ISP in the neutral area, whatever the neutrality level imposed by the regulator in the neutral zone; only the non-neutral ISP may benefit from it. Even worse, it may lead to some content being inaccessible for some users, and no winner at all if neutrality rules are strict for CPs in the neutral area.

The remainder of this paper is organized as follows. Our main assumptions in terms of network topology, types of actors, CP demand functions, and order of decisions, are detailed in Section 2. We then develop in Section 3, as a reference situation, the case when the whole network is fully neutral. We then treat a first scenario in Section 4, where the CPs located in the neutral area cannot choose the quality for their

flows crossing the non-neutral area. Another scenario, where that possibility is added, is analyzed numerically in Section 5. Each time, we find numerically the outcomes (equilibria) that can be expected and observe their implications for the stakeholders. We finally draw some conclusions and propose directions for future work in Section 6.

2 Model

2.1 Topology: two peering ISPs

We consider the network topology depicted in Figure 1 and first described in [12]. We limit ourselves to two countries/areas in interactions with potentially different neutrality rules for sake of simplicity, even if more countries can be considered. In




Figure 1: Topology

each country, we consider a single ISP (all ISPs can be aggregated otherwise). With respect to [12], we here deal with CPs in competition, introducing an additional level of game; we thus focus on the impact of this competition by considering a CP denoted by CP_D in the area potentially applying differentiation and another CP_N in the neutral area. We also introduce the price charged by CPs to users. We assume a peering link between areas, with no charge between ISPs.

2.2 Service qualities

Service can be differentiated by ISPs thanks to the provided quality of service (QoS). Let q_N be the performance level offered by the neutral ISP. For the differentiated ISP, we will consider two scenarios in order to make comparisons: first the case where it is also neutral in the sense that it offers a single class of service with QoS q_D , and then the case where it is non-neutral, with two offered qualities q_L and q_H . All qualities are assumed to be in $[0, 1]$. When users in Area $i \in \{N, D\}$ download/upload content to the CP in the same area, traffic only goes through the ISP i and the experienced quality is the one selected (no choice in case of neutrality, but a choice may be possible in the differentiated area) by CP i in its host zone. On the other hand, when accessing 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 $q = q_N \cdot q_{C_j}$, so that a null

quality at a network leads to a null quality along the path, and a perfect quality on one network means the perceived overall quality is the QoS of the other network. This is illustrated in Figure 2. A typical example of such a product form is when QoS




Figure 2: Qualities perceived by users, depending on their host ISP and the host ISP of the CP they use, in the non-neutral scenarios (in the neutral scenario, $q_L = q_H = q_D$).

is the probability of successful transmission and losses occur independently on both networks.

2.3 Monetary exchanges

This paper focuses on the determination of qualities by ISPs, but prices are also involved in the model (and decided through another game or by a regulator; on the other hand prices by CPs are decided in this paper). 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 if differentiation is implemented. If no differentiation is implemented in Area D , the per-unit-of-volume price is fixed to p_D . CPs on the other hand get revenue from two different sources:

- A per-unit-of-volume gain from advertisement a_i for CP i , with $i \in \{N, D\}$. Indeed advertisement revenue can reasonably be assumed proportional to the amount of data “sold”;
- Revenue from users, assumed in this paper proportional to the consumed volume of data. We denote by r_i , for $i \in \{N, D\}$ the per-unit-of-volume price for data of CP i .

All those monetary flows are illustrated in Figure 3.

2.4 User demands

We finally denote by D_i^j the demand for content of CP j from customers in area i . Demand also depends on all prices and qualities (as can be seen below), but we remove this dependence in the notation for sake of clarity and conciseness.

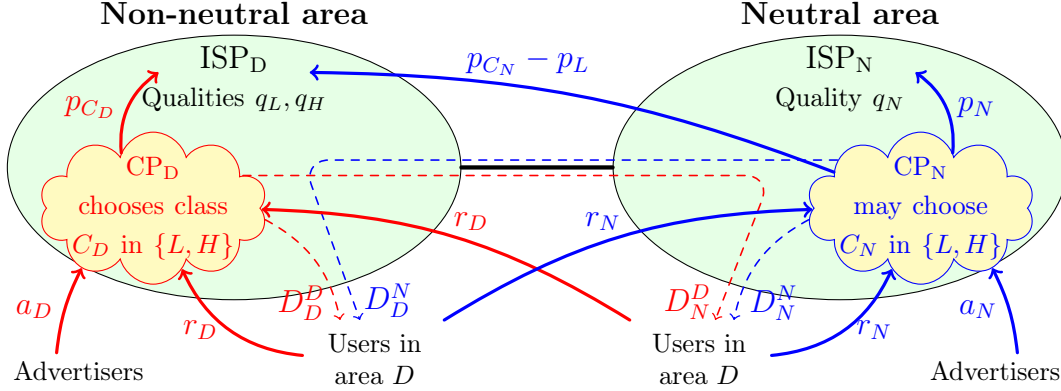


Figure 3: Monetary flows (per unit of volume) in the model (solid lines), in the non-neutral scenarios (in the neutral scenario, $q_L = q_H = q_D$). The dashed lines represent the data flow volumes. Note that CPs only pay their host ISP, for all their outgoing flows, except when CP_N selects the quality q_H for its flows crossing ISP_D (in that case that CP pays ISP_D for the extra quality).

Following the ideas developed in [3, 4], user CP selection is from a discrete choice model. Using a Logit model, the valuation, or level of satisfaction, of a user u in Area j for a CP i is random and given by

$$V_{u,i,j} = v_{i,j} + \kappa_{u,i,j}$$

where

- $v_{i,j} = \beta \log \left(\frac{q_{i,j}}{r_i} \right)$ with β a sensitivity parameter. A logarithmic form of utility is justified by the Weber-Fechner Law, a key principle in psychophysics describing the general relationship between the magnitude of a physical stimulus and its perceived intensity within the human sensory system [8]. The fraction in the log represents the quality per monetary unit.
- The $\kappa_{u,i,j}$ are i.i.d. random variables following a Gumbel distribution with mean 0, which leads to simple analytical formulas for the user repartition among choices.

We also assume that there is a fictitious quality per unit of price, $q'_0 > 0$, corresponding to the minimal quality under which users prefer to subscribe to no CP, that choice being considered an outside option.

This leads (see [1] to get how this is derived) to the proportion (or demand, assuming a total volume 1) of users D_N^N (resp. D_D^D) of the neutral area choosing CP N (resp. CP D), since users in Area N have the choice between no CP (with

quality per unit price q'_0), CP N (with q_N/r_N) and CP D (with $q_N q_{C_D}/r_D$):

$$D_N^N = \frac{(q_N/r_N)^\beta}{(q_N q_{C_D}/r_D)^\beta + (q_N/r_N)^\beta + (q'_0)^\beta} \quad (1)$$

$$D_N^D = \frac{(q_N q_{C_D}/r_D)^\beta}{(q_N q_{C_D}/r_D)^\beta + (q_N/r_N)^\beta + (q'_0)^\beta} \quad (2)$$

(the remaining part chooses the outside option—*not subscribing*—, estimating that the offered quality per unit of money is not sufficient). Without loss of generality, we assume a mass 1 of potential demand in the neutral area.

Similarly, the demand of content from users in Area D is D_D^N (resp. D_D^D) for CP N (resp. CP D), given by, assuming a total potential demand mass m_D ,

$$D_D^N = m_D \frac{(q_N q_{C_N}/r_N)^\beta}{(q_{C_D}/r_D)^\beta + (q_N q_{C_N}/r_N)^\beta + (q'_0)^\beta} \quad (3)$$

$$D_D^D = m_D \frac{(q_{C_D}/r_D)^\beta}{(q_{C_D}/r_D)^\beta + (q_N q_{C_N}/r_N)^\beta + (q'_0)^\beta}. \quad (4)$$

It is interesting to note that, with respect to our first paper on heterogeneous neutrality rules [12], the demand function is totally different to take into account the competition between CPs; this could change the output and adds a level of game in the model.

2.5 Utilities

Utilities of CPs and ISPs are represented by their revenues. For CPs, it is their demand times the unit gain (advertisement plus subscriptions), to which the transmission cost to the host ISP has to be subtracted. This gives U_N^{CP} and U_D^{CP} for the CPs in respectively Area N and D :

$$\begin{aligned} U_D^{CP} &= (a_D + r_D - p_{C_D})(D_N^D + D_D^D) \\ U_N^{CP} &= (a_N + r_N - p_N)(D_N^N + D_D^N) + (p_L - p_{C_N})D_D^N. \end{aligned}$$

The difference for the CP in the neutral area comes from the fact that it may not only have to pay its per-unit-of-volume price p_N to ISP N , but also the price difference to ISP D to be in the upper class, if chosen, for the traffic of users in Area D : If CP N chooses the low quality for its flows on ISP D , i.e., $C_N = L$ then there is no extra fee and all its flows are charged p_N . This latter case corresponds to the current situation where CPs are charged only by their host ISP and not the distant one, which is specifically what triggered the net neutrality debate.

The utilities of ISPs are made of their revenue minus the cost of the architecture for providing a given quality. Following [12], the cost f_i borne by an ISP $i \in \{N, D\}$ is made of the demand level at the quality times the (unit) cost $c_i(q)$ to provide this level:

$$\begin{aligned} f_N &= (D_N^D + D_N^N + D_D^N) c_N(q_N) \\ f_D &= (D_N^D + D_D^D) c_D(q_{C_D}) + D_D^N c_D(q_{C_N}), \end{aligned}$$

where we use the classes chosen by the CPs. It gives the utilities:

$$\begin{aligned} U_N^{ISP} &= p_N(D_N^N + D_D^N) - f_N \\ U_D^{ISP} &= p_{C_D}(D_D^D + D_N^D) + (p_{C_N} - p_L)D_D^N - f_D. \end{aligned}$$

2.6 Order of decisions

Decisions are taken at different time scales, with some players choosing first their strategy. More specifically:

1. ISPs play the game on the QoS levels to be offered to the CPs: the ISP in the neutral area chooses the value q_N and the ISP in the non-neutral area chooses the optimal values of q_L, q_H (or q_D in the fully neutral scenario). Note that the strategy of an ISP impacts the other in terms of demand.
2. CPs choose the class of service to be provided to the users in the non-neutral region and the price to be charged to their consumers, acknowledging the competition that they will face.
3. Users select their preferred CP.

The game is played by backward induction, meaning that players at a given level anticipate the decisions of actors playing next when optimizing their strategy.

For comparison purposes, we will consider three different scenarios:

1. When neutrality is imposed in both areas.
2. When neutrality is relaxed in one of the areas. In this situation, we further split our study into two scenarios:
 - (a) The case where the CP in the neutral area is always served with the low quality in the differentiated ISP network (that CP flows are always treated with low quality q_L in the non-neutral area).
 - (b) The case where both CPs can freely select the class of service for their flows in the differentiated area.

3 Benchmark scenario: a fully-neutral network

We first look at the benchmark case where both ISPs are neutral, for comparison purposes. The game is then reduced to the following set of ordered decisions and is (again) analyzed by backward induction:

1. ISPs play a game on the QoS to be offered to the CPs, i.e., the ISP in the neutral area chooses q_N and the ISP in the “non-neutral area” (that is actually neutral in this scenario) chooses q_D .
2. CPs play a game over the prices to charge their users, r_N and r_D ,
3. Users select their CP based on prices and qualities.

For a self-contained section, we rewrite the demand and utility functions with the full-neutrality restriction:

$$\begin{aligned}
D_N^N &= \frac{(q_N/r_N)^\beta}{(q_N q_D/r_D)^\beta + (q_N/r_N)^\beta + (q'_0)^\beta} \\
D_N^D &= \frac{(q_N q_D/r_D)^\beta}{(q_N q_D/r_D)^\beta + (q_N/r_N)^\beta + (q'_0)^\beta} \\
D_D^N &= m_D \frac{(q_N q_D/r_N)^\beta}{(q_D/r_D)^\beta + (q_N q_D/r_N)^\beta + (q'_0)^\beta} \\
D_D^D &= m_D \frac{(q_D/r_D)^\beta}{(q_D/r_D)^\beta + (q_N q_D/r_N)^\beta + (q'_0)^\beta} \\
U_D^{CP} &= (a_D + r_D - p_D)(D_N^D + D_D^D) \\
U_N^{CP} &= (a_N + r_N - p_N)(D_N^N + D_D^N) \\
U_N^{ISP} &= p_N(D_N^N + D_D^N) - (D_N^D + D_N^N + D_D^N)c(q_N) \\
U_D^{ISP} &= p_D(D_D^D + D_N^D) - (D_N^D + D_D^D + D_D^N)c(q_D).
\end{aligned}$$

Throughout the paper, we will use the following parameter values, unless mentioned otherwise: $p_N = 0.5, p_D = 0.5, q'_0 = 0.8, \alpha = 0.02, m_D = 1, \beta = 1.5, a_N = 0.1, a_D = 0.1, q_N = 0.6, q_D = 0.5$.

Figure 4 displays CP N revenue in terms of its decision variable (price) r_N for several choices of prices r_D of the opponent while the other parameters are fixed. It




Figure 4: CP N utility in terms of r_N , with $q_D = 0.5, q_N = 0.6, a_D = 0.1, a_N = 0.1, p_D = 0.5, p_N = 0.5$, starting at $r_N = p_N - a_N$ to ensure a non-negative revenue per customer.

can be seen that there is an optimal best response r_N^* maximizing CP N revenue and that this best response seems to increase with r_D . It is also interesting to note that the revenue curves are not monotonous in r_D : for small values of r_N CP N gets more revenue when r_D is small, but it is the opposite when r_N is large.

Figure 5 displays the best responses of both CPs in terms of the price of the opponent on the same figure. A Nash equilibrium, that is a profile of strategies from which no player has an interest to deviate unilaterally, is then an intersection point

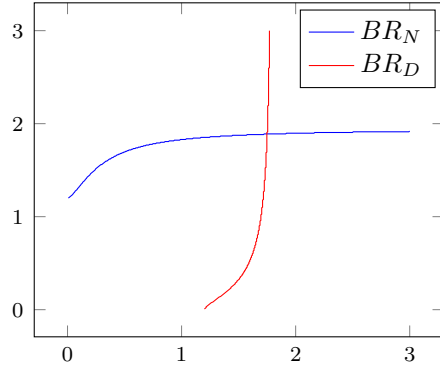


Figure 5: CP best responses, with $q_D = 0.5, q_N = 0.6, a_D = 0.1, a_N = 0.1, p_D = 0.5, p_N = 0.5$

of both curves. We can see that such an equilibrium exists and is unique. It can also be checked that best responses are increasing.

ISPs play first, but anticipating the strategies of CPs. We therefore use for each profile of qualities chosen by ISPs the corresponding Nash equilibrium in the pricing game between CPs. Figure 6 displays the revenue of ISP N in terms of q_N for several values of q_D (again, using the Nash equilibrium for the pricing game between CPs). A




Figure 6: ISP N utility for several values of q_D , with $a_D = 0.1, a_N = 0.1, p_D = 0.5, p_N = 0.5$

maximum value is readily observed. The larger the quality of the opponent, the larger the revenue of the ISP: it comes from a positive externality on users more interested in network usage, yielding an impact in all zones. We can also see that even if the curves are distinct, the values at which the maximum is obtained close.

This can also be observed on the best-response curves displayed in Figure 7. Best responses seem to evolve linearly in terms of the quality chosen by the other ISP, and to decrease but at a very small rate. The best strategy therefore only slightly depends on the strategy of the opponent. It can be seen that there is for this upper-level quality game a single Nash equilibrium.

For the parameters values provided at the beginning of this section, we give in




Figure 7: ISP best responses, with $a_D = 0.1, a_N = 0.1, p_D = 0.5, p_N = 0.5$

Table 1 the output of the game described previously. This will be used as a benchmark for comparison with the non-neutral scenarios of the next sections.

q_D^*	q_N^*	r_D^*	r_N^*	$(U_D^{CP})^*$	$(U_N^{CP})^*$	$(U_D^{ISP})^*$	$(U_N^{ISP})^*$	$(D_D^D)^*$	$(D_D^N)^*$	$(D_N^D)^*$	$(D_N^N)^*$
0.71	0.72	2.022	2.033	0.517	0.523	0.1379	0.1374	0.198	0.121	0.120	0.200

Table 1: Output of the game in the fully-neutral scenario

4 Scenario: CP in the neutral area not free to pay for quality in the non-neutral zone

We now consider the situation where the ISP in Zone D is allowed to offer a differentiated service while the one in Zone N is not allowed to (US versus Europe-like situation). To start, we assume in this section that the CP in Zone N is not allowed by its local rules to even benefit from a differentiated service for its traffic through the ISP in Zone D (to reach customers in that zone).

For a self-contained section, we rewrite the demand and utility functions as follows:

$$\begin{aligned}
D_N^N &= \frac{(q_N/r_N)^\beta}{(q_N q_{C_D}/r_D)^\beta + (q_N/r_N)^\beta + (q'_0)^\beta} \\
D_N^D &= \frac{(q_N q_{C_D}/r_D)^\beta}{(q_N q_{C_D}/r_D)^\beta + (q_N/r_N)^\beta + (q'_0)^\beta} \\
D_D^N &= m_D \frac{(q_N q_L/r_N)^\beta}{(q_{C_D}/r_D)^\beta + (q_N q_L/r_N)^\beta + (q'_0)^\beta} \\
D_D^D &= m_D \frac{(q_{C_D}/r_D)^\beta}{(q_{C_D}/r_D)^\beta + (q_N q_L/r_N)^\beta + (q'_0)^\beta} \\
U_D^{CP} &= (a_D + r_D - p_{C_D})(D_N^D + D_D^D) \\
U_N^{CP} &= (a_N + r_N - p_N)(D_N^N + D_D^N) \\
U_N^{ISP} &= p_N(D_N^N + D_D^N) - (D_N^D + D_D^N + D_D^D)c(q_N) \\
U_D^{ISP} &= p_D(D_D^D + D_N^D) - (D_N^D + D_D^D)c_D(q_{C_D} - D_D^N c_D(q_L)).
\end{aligned}$$

The game is played in the following order, with the described decision variables:

1. ISPs play the game on the QoS to be offered to the CPs (q_N for ISP N , and two values q_L, q_H for ISP D , with $0 < q_L < q_H < 1$).
2. The CP in the non-neutral area chooses at the same time the quality class (L or H) its traffic will be conveyed on, and the price r_D charged to its users, while the CP in the neutral area decides only the price r_N charged to its users according to Equations (1) to (4).
3. Users select their preferred CP.

Decisions being taken anticipatively, we first look at the game between CPs. CP N only chooses its price r_N , but CP D 's decision is made of two components:

- one continuous (its price r_D),
- and one discrete (the class L or H , depending on the quality levels q_L and q_H to which its traffic will be oriented in Zone D).

This makes the game more complicated to analyze, in particular computing best-responses is a bit harder, as explained now. For CP N , for each possible choice (r_D, C) with $C \in \{L, H\}$, we can easily compute the price $r_N = \text{BR}_N(r_D, C)$ maximizing CP N revenue, exactly like in the previous section. To compute the best response of CP D in terms of a choice r_N of CP N , we just compute the optimal price r_D^L and revenue when Class L is selected, then the optimal price r_D^H and revenue when Class H is selected, and choose the class (and associated price) giving the larger revenue, that is

$$(r_D^*, C) = \arg \max_{(r,c) \in \{(r_D^L, L), (r_D^H, H)\}} U_D^{CP}(r, c)$$

where we abusively give as only inputs of the utility the price and class of CP D , omitting the price and quality level in the neutral area.

Figure 8 displays as an illustration CP D revenue in terms of its decision variable (price) r_D for several choices of prices r_N and the two class choices; it illustrates that

a maximum exists, and that the quality chosen by the CP may vary with the price set by the opponent.




Figure 8: CP D utility in terms of r_D , with $q_L = 0.5, q_H = 0.685, q_N = 0.6, a_D = 0.1, p_L = 0.5, p_H = 1.4$

To graphically determine if there is a Nash equilibrium in this game with two components for CP D , we can just draw:

- two best-response curves for CP N in terms of r_D : one corresponding to the selection of Class L by CP D and the other to the selection of Class H , drawn in two different colors;
- the best response value r_D^* in terms of r_N for CP D with a color (the same than above) used when the optimal class is L and the other color when it is H .

A Nash equilibrium is then an intersection point between a best-response curve of CP N with a best-response curve of CP D with the same color.

Figure 9 illustrates this. A Nash equilibrium is an intersection between blue




Figure 9: CP best responses, with $q_L = 0.5, q_H = 0.685, q_N = 0.6, a_D = 0.1, a_N = 0.1, p_L = 0.5, p_H = 1.4, p_N = 0.5$

curves or between red curves/ It can be readily checked that there is here a single Nash equilibrium, where Class H is selected by CP D .

It is then interesting to investigate CP D 's chosen service class at Nash equilibrium in terms of the quality levels, while the other parameters are kept fixed. Figure 10 illustrates with different colors the class chosen by the CP in the non-neutral area, for each couple (q_L, q_H) (with $q_H \geq q_L$). The result is expected: when q_H is close to q_L ,

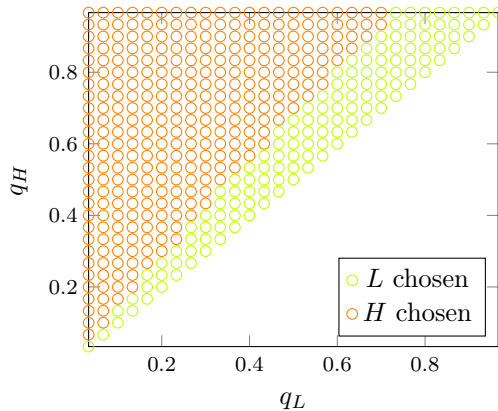


Figure 10: Chosen quality class C_D in terms of (q_L, q_H) , when $q_N = 0.6, a_D = 0.1, a_N = 0.1, p_L = 0.5, p_H = 1.4, p_N = 0.5$

it is not worthwhile to “upgrade” to Class H , but there there is a threshold for each value of q_L over which it becomes valuable. Moreover, the larger q_L the larger the difference of quality is required to opt for Class H (translated as an enlarged band as q_L increases).

Let us now consider the game on qualities, played by ISPs. Figure 11 displays the revenue of ISP N in terms of q_N for several values of q_H when q_L is fixed (using the Nash equilibrium for the pricing game between CPs). Here too, the max value (that




Figure 11: ISP N utility for several values of q_H , with $a_D = 0.1, a_N = 0.1, p_L = 0.5, p_H = 1.4, p_N = 0.5$

is, a best response) can be observed.

Figure 12 displays in 3D the best responses for the quality game with ISP D playing a two-dimensional strategy (optimizing over q_L and q_H). A Nash equilibrium




Figure 12: ISP best responses, with $a_D = 0.1, a_N = 0.1, p_L = 0.5, p_H = 1.4, p_N = 0.5$

can be observed, actually located on $q_L = 0$.

We give in Table 2 the output of the game with partial neutrality, using the above Nash equilibrium for the quality game. Comparisons will be made in Section 6, but it

q_L^*	q_H^*	q_N^*	r_D^*	r_N^*	$(U_D^{CP})^*$	$(U_N^{CP})^*$	$(U_D^{ISP})^*$	$(U_N^{ISP})^*$	$(D_D^D)^*$	$(D_D^N)^*$	$(D_N^D)^*$	$(D_N^N)^*$
0.0	0.25	0.73	4.044	2.3025	0.0868	0.3758	0.044077	0.08752	0.0210	0.0	0.0106	0.19755

Table 2: Output of the game in the partially-neutral scenario

has to be noted that $q_L = 0$ is different from just defining a single quality class since it impacts the foreign customers. It “forces” the local CP to choose the high class but gives no service to the customers abroad. From this point of view it is a bad signal if we want to ensure access to all content.

5 Scenario: All CPs can differentiate service in the non-neutral area

We now tackle the situation when the CP in the neutral area (CP N) can also ask for service differentiation in Zone D , that is, the model fully described in Section 2. Again, except specified otherwise, the parameter values will be the same as in previous sections for comparison purposes.

In the CP game, since here both CPs have two decision variables, the price (continuous value) and the class L or H to be associated with, the Nash equilibrium is computed a bit differently:

- We draw two curves of best response of CP N in terms of r_D : one corresponding to the selection of Class L by CP D and the other to the selection of Class H , drawn in two different colors. Each time, we differentiate whether CP N selects

L or H by using a dashed or plain line. As before, the trade-off is between setting a low price with low (cheap) quality, and high price with a higher quality service.

- Similarly, we draw two curves of best response of CP D in terms of r_N .

A Nash equilibrium is then an intersection point between best-response curves *with corresponding quality selections*, that is,

$$\begin{cases} (r_N, q_{C_N}) = \text{BR}_N(r_D, q_{C_D}) \text{ with class } C_N \text{ selected} \\ (r_D, q_{C_D}) = \text{BR}_D(r_N, q_{C_N}) \text{ with class } C_D \text{ selected,} \end{cases}$$

with $C_i \in \{L, H\}$ the quality selected by CP i . Figure 13 illustrates this. It can be




Figure 13: CP best responses, with $q_L = 0.5, q_H = 0.685, q_N = 0.6, a_D = 0.1, a_N = 0.1, p_L = 0.5, p_H = 1.4, p_N = 0.5$

readily checked that there is here also a single Nash equilibrium. Note that the two curves $\text{BR}_N(r_D, q_L)$ and $\text{BR}_N(r_D, q_H)$ are very close to each other with the set of parameters, Class H is never selected by CP N .

Figure 14 displays in terms of (q_L, q_H) (with other parameters fixed) the zones over which qualities L or H are selected by CPs: (I, J) with $I, J \in \{L, H\}$ means CP D chooses I and CP N chooses J . Again, when the difference between qualities is low, both CPs prefer to stay with Class L , while when it is high, they both prefer H , with the required difference increasing with the quality q_L . In between, there is a band of values where CP D only prefers H , because the class choice impacts all its traffic and has therefore more importance.

Figure 15 displays in 3D the best responses for the quality game with ISP D playing a two-dimensional strategy. Here again, we end up with $q_L = 0$ for the Nash equilibrium, but with slightly different values of q_N and q_H .

We give in Table 3 the output of the game where CP N can also differentiate service for traffic in Zone D . Note that allowing the CP in the neutral zone to choose

q_L^*	q_H^*	q_N^*	r_D^*	r_N^*	$(U_D^{CP})^*$	$(U_N^{CP})^*$	$(U_D^{ISP})^*$	$(U_N^{ISP})^*$	$(D_D^D)^*$	$(D_D^N)^*$	$(D_N^D)^*$	$(D_N^N)^*$
0.0	0.5667	0.74667	4.35	2.7805	0.2826	0.4529	0.1737	0.09074	0.05683	0.07817	0.03428	0.15723

Table 3: Output of the game in the non-neutral scenario

its class of service in the non-neutral zone ends up with all content accessible, in




Figure 14: Chosen qualities classes (C_D, C_N) in terms of (q_L, q_H) , when $q_N = 0.6, a_D = 0.1, a_N = 0.1, p_L = 0.5, p_H = 1.4, p_N = 0.5$




Figure 15: ISP best responses, with $a_D = 0.1, a_N = 0.1, p_L = 0.5, p_H = 1.4, p_N = 0.5$

contrast to the previous case where content from CP N was only accessible in Zone N .

6 Comparisons and conclusions

Table 4 summarizes the main metrics at equilibrium in the three scenarios. Several

q_L^*	q_H^*	q_N^*	r_D^*	r_N^*	$(U_D^{CP})^*$	$(U_N^{CP})^*$	$(U_D^{ISP})^*$	$(U_N^{ISP})^*$	$(D_D^D)^*$	$(D_D^N)^*$	$(D_N^D)^*$	$(D_N^N)^*$
$q_D^* = 0.71$		0.72	2.022	2.033	0.517	0.523	0.1379	0.1374	0.198	0.121	0.120	0.200
0.0	0.25	0.73	4.044	2.3025	0.0868	0.3758	0.044077	0.08752	0.0210	0.0	0.0106	0.19755
0.0	0.5667	0.74667	4.35	2.7805	0.2826	0.4529	0.1737	0.09074	0.05683	0.07817	0.03428	0.15723

Table 4: Comparison of equilibrium outcomes for the fully-neutral case (*first line of results*), the partially-neutral case (*second line*) and the non-neutral case (*last line*).

points are worth noticing.

We first observe the quality levels. As previously mentioned, relaxing neutrality constraints leads to $q_L = 0$, i.e., forcing to select the “premium” service to have some service at all. But this premium service is actually less good than in the fully neutral situation, and for the partially-neutral scenario is not even accessible to CPs hosted in the neutral area, leading to those CPs not being accessible by users in the non-neutral region. In summary, in terms of quality:

- the fully-neutral scenario should be preferred,
- but if neutrality is not imposed in zone D (e.g., in the United States), then the regulator of the neutral area would prefer the non-neutral situation to the partially-neutral one, and allow its CPs to choose their quality abroad and attract some demand.

Now let us focus on the prices r_D and r_N charged by CPs to users. We note that as the constraints are relaxed (on neutrality for ISPs, on quality choices for CPs), those prices tend to increase in both areas.

As a result (of qualities decreasing and prices increasing), demands decrease when leaving the fully-neutral situation. In the neutral zone, we remark that the total demand of the CP ($D_N^N + D_D^N$) is slightly larger in the non-neutral case than in the fully-neutral case, even if demand from users within its zone decreases. For CP D , the demand from both zones is larger in the fully-neutral than in the partially-neutral case. Note however that the cumulated demand from users in the neutral zone ($D_N^N + D_N^D$) is slightly larger in the partially-neutral situation.

In terms of CP utilities, in accordance with their position in the debate, a fully-neutral situation is better. But a non-neutral scenario is preferred to a partially-neutral one, for both CPs.

As regards ISP utilities, note that the partially-neutral case is the worst for both, while as expected the ISP in the non-neutral (resp., neutral) zone prefers the non-neutral (resp., neutral) scenario. This can be a justification for a direct transition between the two extremes. Remark that in a cooperative world looking at the sum of ISP utilities, the neutral scenario is the best.

In summary, our results suggest that:

- only the non-neutral ISP benefits from relaxing neutrality rules;
- if non-neutrality is allowed in a zone, then it is better to avoid the partially-neutral situation: CPs in the neutral zone should be allowed to select quality for their flows crossing the non-neutral zone.

References

- [1] M. Ben-Akiva and S. Lerman. *Discrete choice analysis*. The MIT Press, Cambridge, Massachusetts, 1985.
- [2] BEREC. Berec guidelines on the implementation by national regulators of european net neutrality rules. Document number: BoR (16) 127, Aug 2016. available at https://berec.europa.eu/eng/document_register/subject_matter/berec/regulatory_best_practices/guidelines/6160-berec-guidelines-on-the-implementation-by-national-regulators-of-european-net-neu
- [3] 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.
- [4] 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.
- [5] Federal Communications Commission. Protecting and promoting the open internet, April 2015. available at <https://www.federalregister.gov/documents/2015/04/13/2015-07841/protecting-and-promoting-the-open-internet>.
- [6] Federal Communications Commission. Restoring internet freedom, Jan 2018. available at <https://docs.fcc.gov/public/attachments/FCC-17-166A1.pdf>.
- [7] T.M. Lenard and R.J. (Eds.) May. *Net Neutrality or Net Neutering: Should Broadband Internet Services be Regulated*. Springer, 2006.
- [8] P. Maillé, P. Reichl, and B. Tuffin. Economics of Quality of Experience. In Antoniadis, Antonis, Stiller, and Burkhard, editors, *Telecommunications Economics - Selected Results of the COST Action IS605 EconTel*, pages 158–166. Lecture Notes in Computer Science 7216, Springer Verlag, 2012.
- [9] 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.
- [10] 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.
- [11] P. Maillé and B. Tuffin. *Telecommunication Network Economics: From Theory to Applications*. Cambridge University Press, 2014.

- [12] P. Maillé and B. Tuffin. Neutral and Non-Neutral Countries in a Global Internet: What Does it Imply? In *Proc. of GECON - 16th International Conference on the Economics of Grids, Clouds, Systems, and Services*, Leeds, United Kingdom, September 2019. Springer.
- [13] 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.
- [14] H. Schulzrinne. Network neutrality is about money, not packets. *IEEE Internet Computing*, 22(6):8–17, 2018.
- [15] T. Wu. Network neutrality, broadband discrimination. *Journal of Telecommunications and High Technology*, 2003.