Impact of Competition Between ISPs on the Net Neutrality Debate

Pierre Coucheney, Patrick Maillé, Bruno Tuffin

► To cite this version:

Pierre Coucheney, Patrick Maillé, Bruno Tuffin. Impact of Competition Between ISPs on the Net Neutrality Debate. IEEE Transactions on Network and Service Management, 2013, 10 (4), pp.425-432. hal-00681402v2

HAL Id: hal-00681402

<https://hal.inria.fr/hal-00681402v2>

Submitted on 25 Nov 2013

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 Competition Between ISPs on the Net Neutrality Debate

Pierre Coucheney, Patrick Maillé, Bruno Tuffin

Abstract—Network neutrality is the topic of a vivid and very sensitive debate, in both the telecommunication and political worlds, because of its potential impact in everyday life. That debate has been raised by Internet Service Providers (ISPs), complaining that content providers (CPs) congest the network with insufficient monetary compensation, and threatening to impose side payments to CPs in order to support their infrastructure costs. While there have been many studies discussing the advantages and drawbacks of neutrality, there is no game-theoretical work dealing with the observable situation of competitive ISPs in front of a (quasi-)monopolistic CP. Though, this is a typical situation that is condemned by ISPs, and, according to them, another reason of the non-neutrality need.

We develop and analyze here a model describing the relations between two competitive ISPs and a single CP, played as a three-level game corresponding to three different time scales. At the largest time scale, side payments (if any) are determined. At a smaller time scale, ISPs decide their (flat-rate) subscription fee (toward users), then the CP chooses the (flat-rate) price to charge users. Users finally select their ISP (if any) using a price-based discrete choice model, and decide whether to also subscribe to the CP service. The game is analyzed by backward induction. As a conclusion, we obtain among other things that non-neutrality may be beneficial to the CP, and not necessarily to ISPs, unless the side payments are decided by ISPs.

Index Terms—Network neutrality, Game theory, Pricing

I. INTRODUCTION

There has recently been a strong debate around the so-called *network neutrality*. The debate has been ignited by the increasing traffic asymmetry between Internet Service Providers (ISPs), mainly due to some prominent and resource consuming content providers (CPs) which are usually connected to a single ISP. The typical example is YouTube (owned by Google), accessed by all users while hosted by a single Tier 1 ISP, and whose traffic now constitutes a non-negligible part of the whole Internet traffic. Another example is the subscription-based video service Netflix, that is in the US the most bandwidth-consuming source of traffic, representing 23.3% of the total Internet traffic in late 2011 [1], while having commercial relationships with only one ISP. For those reasons, there has been a surge of protest among ISPs, complaining that the current Internet business model where ISPs charge both end-users and content providers directly connected to them, and have public peering or transit agreements with other ISPs, is not relevant anymore. The main solution proposed is that

ISPs should also charge content providers that are associated with other ISPs [2], as first advocated by Ed Whitacre (CEO of AT&T) at the end of 2005 [3].

The underlying concern is that investment is made by ISPs but content providers get an important part of the dividends. The revenue arising from on-line advertising (meaning showing graphical ads on regular web pages) is estimated at approximately a \$24 billion in 2009 [4], while textual ads on search pages has led to a combined revenue of \$8.5 billion in 2007 [5], those figures increasing every year. Meanwhile, transit prices - which constitute the main source of revenues for transit ISPs - are decreasing. ISPs argue that there is no sufficient incentive for them to continue to invest on the network infrastructures if most benefits go to content providers. The threat is to lower the quality of service of CPs that do not pay any fee to them, or even to block their traffic. This possibility has led to protests from CPs and user associations, complaining that this might impact the network development and is an impingement of freedom of speech [3]. The debate was thus launched, essentially at the law and policy makers level, to decide whether the Internet should be *neutral*, i.e., all packets should receive equal treatments in terms of price and service. In the US, the Federal Trade Commission (FTC) released in 2007 a report not supporting neutrality constraints, increasing the debate at the political level. This debate is also active in Europe and in France, as illustrated by the open consultation on network neutrality launched in 2010. For instance, the French regulation authority, ARCEP, has published in its response a proposal intending to define how net neutrality could be implemented [6], [7].

There has been an increasing attention in the literature on providing a mathematical analysis of the advantages and drawbacks of network neutrality. The idea is to investigate the output of the interactions between selfish actors that are end users, CPs and ISPs, using the framework of non-cooperative game theory [8], [9]. Let us briefly describe here, non exhaustively, some important existing works in that direction. In [10], [11], the authors propose to share the revenue among providers using the Shapley value, the only mechanism that satisfies a set of axioms representing a sense of fairness; in this case CPs participate to the network access cost. The work in [12] analyzes how neutrality or non-neutrality affects provider investment incentives, network quality, and user prices. A similar comparison is made in [13] between a two-sided pricing scheme where ISPs are allowed to charge CPs, and one-sided pricing where such side-payments are not allowed. In each case, at the equilibrium of the game, the levels of investment in content and architecture are determined. The paper gives

Pierre Coucheney and Bruno Tuffin are with INRIA Rennes Bretagne-Atlantique, Campus de Beaulieu, 35042 RENNES Cedex - France

Patrick Maillé is with Télécom Bretagne, 2, rue de la Chataigneraie, 35576 Cesson Sévigné Cedex - France

Manuscript received July 24, 2012; revised January 29, 2013; associate editor: Dr. Marwan Krunz.

conditions on the ratio between parameters characterizing advertising rates and end-user price sensitivity, under which a non-neutral network outperforms a neutral one in terms of social welfare. On the other hand, [14] investigates the case where ISPs negotiate joint investment contracts with a CP in order to enhance the quality of service and increase industry profits. It is found that an unregulated regime leads to higher quality investments, but that ISPs have an incentive to degrade content quality. The paper [15] studies the implications of non-neutral behaviors, taking into account advertising revenues and considering both cooperative and non-cooperative scenarios. In [16], we analyze and compare thanks to game-theoretic tools three different situations of interactions *between* ISPs: the case of peering between the ISPs, the case where ISPs do not share their traffic (exclusivity arrangements), and the case where they fix a transfer price per unit of volume. The paper supports the transit price scenario and suggests a limited regulation (enforcing global connectivity) to prevent incumbent ISPs from having a dominant position in the bargaining. Finally, in [17], a game-theoretic model is considered with a single CP, a single ISP, a (consumers') demand function that depends on price and quality of service, and involving advertisement and network investment components.

In those works, there is in general a single ISP, and one or several CPs. Though, in practice, we often have ISPs in competition for customers, while for many services, the CPs are in a quasi monopoly, a characteristic ISPs complain about. (Typical examples are YouTube for non-copyrighted videos, and Netflix for movies and TV shows in the US.) We propose to specifically address this issue in this paper. Remark that in addition to [13], considering competitive ISPs has been proposed in [18], but with competition over consumers, quality and prices for heterogeneous CPs: none of those works consider a monopolistic CP as can be encountered for some applications. We already addressed this type of problem in [19], whose model was inspired by [15], users were assumed to always go with the cheapest provider. As a consequence, we ended up with a price war (a classical Bertrand competition) such that ISPs always decrease their subscription price in order to attract all demand.

We consider here a more realistic user association model such that users make their choice still based on the price of ISPs, but also on other unknown considerations, hence the use of a classical discrete choice model as in [16]. This requires a derivation of results totally different from [19]. To analyze that situation, we propose a multi-level game where decisions are taken at different time scales. The solutions of the games at the largest time scales, played first, are determined using *backward induction*, meaning that players *anticipate* the impact on, and the resulting solution of, the games played later on at smaller time scales.

The paper is organized as follows. Section II presents the basic assumptions of the model we are going to consider, the different levels of game, and the mathematical description of the investigated comparison between the neutral and the non-neutral regimes. It also describes how users select their ISP (if any), and how the aggregated demand at the CP is determined. The next sections present the various game levels

Fig. 1. Charging interactions between stakeholders. Prices p_1 , p_A and p_B are positive flat rates, whereas q_A and q_B are positive per volume unit prices.

for providers' decisions: we describe in Section III how the CP, anticipating the decisions of end users, chooses the content price. At a higher level, by backward induction, ISPs play a game on the access charge for end users; this competition is described in Section IV. We then describe the game at the highest level, on the economic relationships between the ISPs and the CP, by determining the side payments of the CP to the ISPs in Section V. We address the case when those prices are fixed by ISPs, based on a game; we also look at the case when they are decided by the CP, or by a regulator (maximizing the supply chain value for instance). Section VI concludes by discussing the impact and relevance of side payments on the providers' revenues, highlighting that it is not always in the interest of the ISPs (but could be), while an appropriate choice of side payments may increase the CP revenue. We also give in that section directions for future research.

II. MODEL

A. Topology and Pricing Structure

We consider a single CP, whose parameters will be indexed by 1, and two ISPs, named (and indexed by) A and B . The access prices charged to users are *flat rate* subscription fees, denoted by p_1 , p_A and p_B for respectively the CP, ISP A and ISP B . In order to study non-neutrality, we also introduce side payments q_A and q_B representing the per unit of volume prices that the CP has to pay to A and B , respectively. All prices are assumed to be positive. Finally, the set of end users is considered continuous and (without loss of generality) of mass one, so that we will indifferently refer to "mass" and "proportion" of users. The charges imposed by actors to other players are summarized in Figure 1, the arrows indicating the cash flows.

B. User Demand

Users have to pay both the ISP and the CP to access the content. Users first choose their ISP, to which they pay a flat-rate subscription fee, and then subscribe to the CP too if its (flat) fee is not too high. Since users need an Internet access, not only to reach the content of the CP, but also for other purposes (e-mail, web browsing, ...), we decouple the ISP choice from the (individual) decision to subscribe to the CP, hence the independence from p_1 . Among those access subscribers, the proportion who also subscribe to the CP depends on p_1 on the other hand, but not on the

subscription price to the ISPs as a first-order model, assuming as a consequence that users do not relate their consumption of content to the previously paid ISP subscription fee.

Let us focus first on the ISP selection by users. In a previous work [19], we have considered users simply selecting the cheapest ISP (or choosing it randomly if price equality holds). However, this does not take into account the phenomenon of *stickiness* or *loyalty* of the users, highlighted in [20]. In the model considered here, user choices are influenced by the ISP subscription prices, but also by other considerations (reputation or preferences) that can be modeled as an additive noise to the main criteria determining the choice. We consider here a *discrete choice model*, a standard paradigm in economy, transportation, etc., and surprisingly not often applied in telecommunications. It is a common way to model how decision makers choose among a set of alternatives. Mathematically, we assume that a user has a valuation of the form $v_i = \beta \log(1/p_i) + \kappa_i$ for ISP i , and selects the highest-valued option (see [16] for details). The term κ_i is an individual-specific random term, taking into account unknown aspects and assumed to follow a Gumbel distribution as in standard discrete choice models [21]. We additionally assume that there is a fictitious price p_0 , assumed to be strictly positive, and representing the cost of the outside option, i.e., the perceived cost of not having access to the Internet. Thus if the (random) valuation associated with that outside option is larger than the ones associated with each ISP, the user prefers not to join the network. The parameter $\beta > 0$ represents the *user sensitivity* to the subscription prices: values of β close to zero lead to an uniform choice over the three alternatives (connecting to one of the two ISPs, or not having access to the Internet) regardless of the prices set, whereas large values of β make the users choose the least expensive option. The term $\log(1/p_i)$ expresses the dissatisfaction for higher prices, the logarithm being used to represent the fact that the same variation of price is felt smaller at high prices than at low prices: users are sensitive to relative price variations rather than absolute ones. Finally, the case $p_i = 0$ leads to the maximal possible valuation independently of the random terms, meaning that ISP i attracts all users, or (say, from symmetry) half of them if the other ISP chooses a null price as well.

At a macroscopic level, by discrete choice analysis, the proportion (or equivalently, the mass) σ_i of users selecting ISP i (with also $j \in \{A, B\}$; $j \neq i$), can be shown to equal from our expression of v_i (see [21] for a general derivation)

$$\sigma_i = \begin{cases} \frac{p_i^{-\beta}}{p_A^{-\beta} + p_B^{-\beta} + p_0^{-\beta}} & \text{if } p_A > 0 \text{ and } p_B > 0 \\ 1 & \text{if } p_i = 0 \text{ and } p_j > 0 \\ 1/2 & \text{if } p_A = 0 \text{ and } p_B = 0 \\ 0 & \text{if } p_i > 0 \text{ and } p_j = 0. \end{cases} \quad (1)$$

This repartition function expresses the fact that all users select an ISP if at least one of the subscription prices is null ($p_A = 0$ or $p_B = 0$). Of course, the higher the price, the fewer subscribers; this effect increasing with the user sensitivity to price β .

In this paper, we propose a new aggregated user demand in terms of data volume on ISP i , where users having selected an ISP then decide whether to use the content offered by the CP, depending on the *flat-rate* price p_1 . We consider that users' willingness-to-pay for the CP service (i.e., the access to all the content offered by the CP) follows an exponential distribution with mean value $1/\alpha > 0$ over the population, independently of the ISP choice. Therefore, a proportion $e^{-\alpha p_1}$ of each ISP's subscribers also subscribes to the CP, hence a number $\sigma_i e^{-\alpha p_1}$ subscribing to both the CP and ISP i . Note that the consequence of our assumption is an exponentially decreasing demand function (in terms of the CP price), a usual setting in economic theory.

We assume that there is an average volume D_0 of data that a user downloads from the CP if subscribing to it. The value will be helpful to compute the *volume-based* transit costs for the CP to the ISPs. This volume is assumed to be independent of the ISP choice. Instead of considering an average volume, we could say that it is the same value for all users without changing the expressions; an average value just allows to take into account the potential variations between users that are averaged when summing over all subscribers of an ISP. D_0 is additionally assumed independent of p_1 , meaning that users just get what they "need" if subscribing. As a result, the data volume for users subscribing simultaneously to the CP and to ISP i is given by:

$$D_i = D_0 \sigma_i e^{-\alpha p_1}. \quad (2)$$

The parameter $\alpha > 0$ can be interpreted as the sensitivity of users to the CP price: the global demand (sum of demands on all ISPs) is a decreasing function of α .

Notice that demand does not directly depend on the side payments q_A and q_B . But the introduction of side payments will induce a reaction on the prices p_A , p_B and p_1 set by ISPs and the CP at equilibrium, which, in turn, indirectly affects demand. Finally, the global volume demand for CP data $D_A + D_B$ equals $(\sigma_A + \sigma_B) D_0 e^{-\alpha p_1}$.

C. Utility and revenue functions

Among the proportion $\sigma_A + \sigma_B$ of users having accepted to pay for an access to the network, and then paying a flat-rate price p_1 to the CP, some would have accepted to pay more to benefit from the content of the CP. The surplus of users that would have accepted to pay p is $p - p_1$, while the proportion of users willing to pay more than p is $e^{-\alpha p_1}$.

We can then compute the *user welfare* associated with the existence of the CP, as the sum over all users of the benefit they make accessing the content of the CP. Note that this does not include the benefit that users make by selecting an ISP, which is associated with other (free) on-line services. From our assumption of an exponential distribution of users' willingness-to-pay for the content (yielding a density $\alpha e^{-\alpha x}$),

the user welfare due to the CP can be computed as:

$$\begin{aligned} \text{UW}_{\text{CP}} &= (\sigma_A + \sigma_B) \int_{p_1}^{\infty} \alpha e^{-\alpha x} (x - p_1) dx \\ &= (\sigma_A + \sigma_B) \frac{e^{-\alpha p_1}}{\alpha} \\ &= \frac{D_A + D_B}{\alpha D_0}. \end{aligned} \quad (3)$$

The utilities (revenues) of the ISPs come from the end users subscription fee, and from the CP through the possible side payment. The first one depends on the mass of users with the ISP, and the second one on the total amount of volume downloaded by users. Hence, for ISP i ($i \in \{A, B\}$), the revenue is

$$U_i = p_i \sigma_i + q_i D_i. \quad (4)$$

We can remark here that the revenue is always positive since we do not consider the cost of the network.

The utility of the CP in this model is the sum of revenues gained by users subscribing through A and through B . Those gains come from the flat-rate subscriptions by users through each ISP i , $p_1 \sigma_i e^{-\alpha p_1}$, but the volume-based side payments $q_i D_i$ paid to i also have to be taken into account. The CP net benefit (utility) is thus given by

$$\begin{aligned} U_1 &= (p_1 \sigma_A e^{-\alpha p_1} - q_A D_A) + (p_1 \sigma_B e^{-\alpha p_1} - q_B D_B) \\ &= (p_1/D_0 - q_A) D_A + (p_1/D_0 - q_B) D_B. \end{aligned} \quad (5)$$

Since p_1 is decided after q_A and q_B , the CP can also ensure a positive revenue by setting $p_1/D_0 \geq \max(q_A, q_B)$.

D. Multi-stage Decision Problem

The decision variables are the prices p_1, p_A, p_B, q_A, q_B , impacting end users (demand), as well as revenues of providers. Those variables are decided at different time scales or levels, that can be described as follows.

- 1) At the largest time scale, the side payments q_A and q_B are decided. In the neutral case, they are either fixed to 0, or determined as a common value. They can be different in the non-neutral case, and can be determined either by the ISPs (in a game), the CP, or a regulator. All those options will be investigated. Those determinations will be obtained anticipating the solution of the games below whatever the values of q_A and q_B (the so-called *backward induction*).
- 2) At a smaller time scale, for fixed values of q_A and q_B , the ISPs fix their prices p_A and p_B during a non-cooperative game to attract customers and maximize their revenues. Here again, the decisions are made anticipating the solutions at lower levels.
- 3) At an even smaller time scale, the CP sets the price p_1 .

Finally, for those fixed values of p_1, p_A, p_B, q_A, q_B , users choose their ISP (if not too expensive), and decide whether to use the service offered by the CP, as described by formulas (1) and (2).

All those interacting levels are now solved by backward induction, from the smallest time scale to the largest one.

It is possible to perform the same analysis with a different hierarchy in the three levels of game, assuming for example that subscription prices are decided before the side payments. Though, we think it is reasonable to consider that the decisions on side-payments to take place on a large time scale, given the difficulty, span, and costs of such decision processes. Similarly, regarding the time scale difference between prices set by the CP and by the ISPs, our approach is consistent with the (commonly spread) vision that CPs can adapt faster than ISPs, in part because of the different contract durations binding users to providers (larger durations with ISPs than with CPs), the various difficulties to switch ISPs, etc.; see also [22].

III. CONTENT PROVIDER PRICE DETERMINATION

The CP aims at maximizing his revenue U_1 , for fixed values of p_A, p_B, q_A, q_B , making use of what the total user demand $D_A + D_B$, with D_i given by (2), will be. For convenience, we define $P_i := p_i^\beta$.

Proposition 1. *Given the side payments q_A and q_B and the prices p_A and p_B decided by the ISPs, the price of the CP maximizing its revenue (5) is*

$$p_1^* = \begin{cases} \frac{P_A}{P_A + P_B} (D_0 q_B + \frac{1}{\alpha}) + \frac{P_B}{P_A + P_B} (D_0 q_A + \frac{1}{\alpha}) & \text{if } p_A > 0 \text{ or } p_B > 0, \\ D_0 \frac{q_A + q_B}{2} + \frac{1}{\alpha} & \text{if } p_A = 0 \text{ and } p_B = 0 \end{cases} \quad (6)$$

Proof: We first consider the case $p_A > 0$ and $p_B > 0$. The derivative $\frac{\partial U_1}{\partial p_1}$ of the CP revenue is then equal to

$$P_0 e^{-\alpha p_1} \frac{P_A (\alpha q_B + (1 - \alpha p_1)/D_0) + P_B (\alpha q_A + (1 - \alpha p_1)/D_0)}{P_0 P_A + P_0 P_B + P_A P_B}$$

which is strictly positive until p_1 achieves the value given in the first equation of (6), and strictly negative after. Hence the result.

If $p_A = 0$ (and then $P_A = 0$) and $p_B > 0$ (the opposite case is symmetric and then omitted), the CP revenue is $e^{-\alpha p_1} (p_1 - D_0 q_A)$, whose derivative is $e^{-\alpha p_1} (1 - \alpha(p_1 - D_0 q_A))$.

Finally, if $p_A = p_B = 0$, then the CP revenue is $\frac{1}{2} e^{-\alpha p_1} (2p_1 - D_0(q_A + q_B))$, and its derivative is $e^{-\alpha p_1} (1 - \alpha p_1 + \frac{\alpha}{2} D_0(q_A + q_B))$. ■

Notice that the optimal price does not depend on the outside option valuation p_0 . One can also check that it increases with the price p_i of ISP i that has the biggest side payment q_i , and decreases with the other price. In the limit (that can be interpreted as *neutral*) case $q_A = q_B = q$, the optimal pricing for the CP is $q D_0 + 1/\alpha$ whatever the value of p_A and p_B . In that case, as an important remark, the CP's revenue is $\frac{D_A + D_B}{D_0 \alpha}$, which corresponds to the CP-related user welfare: the interest of users and that of the CP coincide here. Finally, we can remark that the optimal price for the CP is always greater than $D_0 \min(q_A, q_B) + 1/\alpha$, because it is a convex combination of $D_0 q_A + 1/\alpha$ and $D_0 q_B + 1/\alpha$. In particular, it is greater than the inverse of the user price sensitivity α .

IV. PRICING GAME BETWEEN ISPS

Before the users decide which ISP to join and the CP chooses p_1 , the ISPs play a pricing game, making use of what the CP and users decisions would be. In this section, we determine the Nash equilibrium solutions of this pricing game in an analytical way when there are no side payments, and numerically (because intractable) in the general case. Recall (see [9]) that a Nash equilibrium would be a price profile (p_A, p_B) such that no ISP can improve (strictly) his utility by unilaterally changing his price. The best-response curves are defined as (by expliciting the dependence of U_A and U_B on p_A, p_B)

$$\begin{aligned} \text{BR}_A(p_B) &= \arg \max_{p_A \geq 0} U_A(p_A, p_B) \text{ and} \\ \text{BR}_B(p_A) &= \arg \max_{p_B \geq 0} U_B(p_A, p_B). \end{aligned}$$

With those notations, a Nash equilibrium is a point $(p_A^{\text{NE}}, p_B^{\text{NE}})$ for which $\text{BR}_A(p_B^{\text{NE}}) = p_A^{\text{NE}}$ and $\text{BR}_B(p_A^{\text{NE}}) = p_B^{\text{NE}}$. Graphically, if we draw the two best-response curves on the same figure, the set of Nash equilibria is then the (possibly empty) set of intersection points of those curves.

A. No side payments

In the case where no side payments are established, $q_A = q_B = 0$, we get a simple formulation for the revenue of ISPs. From the previous section, the optimal CP pricing is $1/\alpha$. Using the notation $P_i := p_i^\beta$, the revenue of ISP A is then (the revenue of ISP B being symmetrical)

$$U_A = \begin{cases} \frac{P_0 P_B P_A}{P_0 P_A + P_0 P_B + P_A P_B} & \text{if } p_A > 0 \text{ and } p_B > 0 \\ 0 & \text{if } p_A = 0 \text{ or } p_B = 0 \end{cases} \quad (7)$$

We first stress that $p_A = p_B = 0$ is a Nash equilibrium since no player can strictly increase his revenue by unilaterally changing his action: the revenue always remains equal to zero. But setting one's price to zero is a dominated strategy, that is strictly dominated as soon as the adversary price is not zero: it always yields no revenue whereas a strictly positive revenue can be guaranteed with any other choice. Therefore it is not likely to be chosen by ISPs if another equilibrium exists.

Proposition 2. *Assuming that there are no side payments, i.e. $q_A = q_B = 0$, then*

- if $\beta \leq 1$, there is no Nash equilibrium different from $(0, 0)$ with finite prices: the only other alternative is both ISPs setting infinitely large prices ($P_A^{\text{NE}} = P_B^{\text{NE}} = \infty$),
- if $1 < \beta < 2$, there is a unique Nash equilibrium different from $(0, 0)$ with $P_A^{\text{NE}} = P_B^{\text{NE}} = \frac{2-\beta}{\beta-1} P_0$,
- if $\beta \geq 2$, $(0, 0)$ is the unique Nash equilibrium, yielding no revenue for the ISPs.

The proof relies on the following general result about symmetric games.

Lemma 1. *If the best response functions BR_A and BR_B are*

- equal: $\text{BR}_A = \text{BR}_B = \text{BR}$,
- single-valued,

- strictly increasing,

then (p_A, p_B) is a Nash equilibrium if and only if $p_A = p_B = p$ with p a fixed point of the best-response function: $p = \text{BR}(p)$.

Proof: (Lemma) The pair of prices (p_A, p_B) is a Nash equilibrium if and only if $p_A = \text{BR}(p_B)$ and $p_B = \text{BR}(p_A)$. Let us suppose that $p_A \neq p_B$, for instance $p_A > p_B$ (should the indexes be permuted). Then:

$$p_B = \text{BR}(p_A) > \text{BR}(p_B) = p_A,$$

where the inequality comes from the strict increasingness of BR , hence a contradiction. At Nash equilibrium, $p_A = p_B$ is then a necessary condition, and from the definition of such an equilibrium, it is necessary and sufficient to have $p = \text{BR}(p)$. ■

Proof: (Proposition) Assuming that $p_A > 0$ and $p_B > 0$, the derivative of ISP A revenue (7) is

$$\frac{P_B P_0}{(P_A P_B + P_B P_0 + P_A P_0)^2} (P_A(1-\beta)(P_B + P_0) + P_B P_0).$$

Hence the derivative has the same sign as the affine function of P_A : $P_A(1-\beta)(P_B + P_0) + P_B P_0$. When $\beta \leq 1$, that derivative is always strictly positive for $P_A \geq 0$, thus (by symmetry) each ISP should set infinitely large prices.

When $\beta > 1$, the derivative of ISP A revenue is strictly positive while P_A is smaller than the unique root of the affine function above, and negative afterwards. Given $p_B > 0$, the best-response of ISP A is then

$$\text{BR}_A(P_B) = \frac{P_B P_0}{(\beta-1)(P_B + P_0)}.$$

For the case $\beta > 1$, notice that the best response is the same function for ISP B due to symmetry, that equals $\text{BR}(P) = \frac{1}{(\beta-1)(1/P_0+1/P)}$, and is a strictly increasing function of P . Hence it follows from the previous lemma that every Nash equilibrium is symmetric, which results here in the necessary and sufficient condition at Nash equilibrium $P_B^{\text{NE}} = P_A^{\text{NE}} = \text{BR}_A(P_A^{\text{NE}})$. The last equation has a unique strictly positive solution in the case $1 < \beta < 2$, the one given in the proposition, and no solution otherwise. ■

This proposition shows in particular that, when the price sensitivity of users is high ($\beta \geq 2$), we are led to the same price war as in the model of Bertrand competition studied in [19]. But for smaller levels of price sensitivity, this does no longer happen: the price set by ISPs at equilibrium is strictly positive, hence providing some revenue from users to both ISPs.

B. Positive side payments

In the general case, the computation of the Nash equilibrium or even the best response function is not analytically tractable. We are then led to study numerically the price competition between ISPs. From here, we take $\alpha = 1$, $p_0 = 1$, $D_0 = 1$, and $\beta = 1.5$. We remark that the choices of p_0 and D_0 are made without loss of generality (they correspond to unit changes of ISP prices and data volumes). The choice of α also corresponds to a unit change in the CP prices, and is therefore linked to that of p_0 (i.e. for a given user sensitivity to the

Fig. 2. Set $\mathcal{P}_w$ of side payments (q_A, q_B) for which the Nash equilibrium is such that $p_A^{\text{NE}} = p_B^{\text{NE}} = 0$, i.e. price war holds.

CP price, the value of α is a consequence of the monetary unit choice, and thus is determined by the choice for p_0). Therefore our choice of α and of β are subject to discussion. We however stress here that we have also run experiments with other values (in the range $(1, 2)$ for parameter β) and observed similar results.

Let us first remark that if the price set by an ISP (say ISP A) is equal to zero, then the other ISP (ISP B), at a Nash equilibrium, sets his price to zero as well. This is because ISP B does not attract any users if his price is not zero, and then his revenue is null, whereas the revenue is $\frac{1}{2}q_B D_0 e^{-\frac{\alpha}{2}D_0(q_A+q_B)-1}$ otherwise, hence strictly positive.

Numerical computations show that there is a set of side payments q_A and q_B for which the price war phenomenon between ISP happens. That set is shown in Figure 2. In the following, we will denote by $\mathcal{P}_w$ the set of side payments for which a price war between ISPs takes place, leading to null subscription prices.

Figures 3 to 6 show the prices and revenues at Nash equilibrium (recall that this is for q_A and q_B fixed). Numerical computations point out the fact that the revenue of the CP, and the user welfare he creates, are always equal at equilibrium, the reason why we do not plot user welfare here. While this equality is clear when side payments are the same, our numerical results suggest that it remains true in the general case.

Figure 3 represents the price p_A^{NE} set by ISP A at Nash equilibrium, when the side payment q_A varies, and for several ISP B side payments q_B . This reveals that the price at equilibrium first decreases with the side payment set by the ISP, and then increases. For some value of the opponent ISP side payment, it goes to zero when a threshold is reached. That threshold corresponds to the case where the side payment revenue that ISPs get by setting their prices to zero, and then attracting the whole set of users, becomes larger than the one they get on a limited market share with both the subscription fees and the side payments. This is the price war situation: equilibrium subscription prices of both ISPs fall down to 0 as the side payments enter the “price war zone” $\mathcal{P}_w$ highlighted in Figure 2, and each ISP finally only attracts half of the user

Fig. 3. ISP A user price p_A^{NE} at equilibrium as a function of the side payments q_A with $q_B \in \{0, 1, 2, 3\}$.

Fig. 4. CP optimal price at equilibrium as a function of the side payments q_A with $q_B \in \{0, 1, 2, 3\}$.

set, hence the discontinuity. Finally, we observe that there is no monotonicity in the opponent side payment.

Figure 4 displays the optimal price p_1^* of the CP in terms of q_A for different values of q_B . We can notice the discontinuity due to the price war thresholds (for the cases $q_B = 1$ and $q_B = 2$, since there are no such thresholds for the two other cases from Figure 2). However, this discontinuity is barely visible, which is quite surprising when compared to that observed at the ISP level. It can also be remarked that the optimal price increases with q_A (and q_B) both before and after the thresholds, but not in general. One can also check here the general property that $p_1^* \geq D_0 \min(q_A, q_B) + 1/\alpha$, which, in particular, ensures a strictly positive revenue to the CP.

Figure 5 shows that the revenue of ISPs is not monotonic with the side payment. Again, some discontinuities occur when side payments enter and leave the price war zone $\mathcal{P}_w$. Moreover, depending on the ISP B side payment, the maximal revenue of ISP A is reached either for a null side payment (e.g., when $q_B = 3.0$) or for a strictly positive value (e.g., when $q_B = 0.0$). This illustrates that predicting the effect of the various parameters in the output of the game(s) is a difficult task and game-theoretic tools are helpful here to get a result: a side payment increase can indeed lead to a content

Fig. 5. ISP A revenue at equilibrium as a function of the side payment q_A with $q_B \in \{0, 1, 2, 3\}$.

Fig. 6. CP revenue at equilibrium as a function of the side payment q_A , with $q_B \in \{0, 1, 2, 3\}$.

price increase by the CP to compensate for the loss, and a reaction of the other ISP; this also has an effect on demand and depending on the variations, this may or may not result in an ISP revenue increase.

On the other hand, the CP revenue, plotted in Figure 6, has a tendency to decrease with side payments, even if it is not strictly the case, as can be seen for $q_B = 1$ and small values of q_A . When a discontinuity occurs, the CP revenue is maximized for the smallest side payment leading to the price war. The non-monotonicity is due to the threshold for the price war between ISPs, inducing a discontinuity in the ISP prices. Because of those prices becoming null, the revenue jumps up.

Finally, over the price war set $\mathcal{P}_w$, the CP subscription price is $p_1 = D_0 \frac{q_A + q_B}{2} + \frac{1}{\alpha}$ and the revenues are

$$U_i = \frac{1}{2} q_i D_0 e^{-\frac{\alpha D_0}{2} (q_A + q_B) - 1} \quad (8)$$

$$U_1 = \text{UW}_{\text{CP}} = \frac{1}{\alpha} e^{-\frac{\alpha D_0}{2} (q_A + q_B) - 1}. \quad (9)$$

V. SIDE PAYMENTS DETERMINATION

We consider at the highest level three possibilities for the choice of the side payments q_A and q_B . We first look at the case when they are determined by the CP (even if unlikely in practice), then the case when they result from a game played

Fig. 7. CP revenue at equilibrium as a function of the side payment q_A , with $q_B \in \{0.2, 0.3, 0.4\}$. The maximal revenue is reached when $q_A = q_B = 0.3$.

between ISPs, and finally the case when they are determined by a regulator (e.g., to maximize social welfare).

Since we don't have the analytical expression for the ISPs price at Nash equilibrium, we provide numerical results, where we take $\alpha = 1$, $\beta = 1.5$, $p_0 = 1$ and $D_0 = 1$.

A. Determined by the CP

The revenue of the CP is maximized when the side payments are $q_A = q_B = 0.3$ as illustrated in Figure 7 (instead of plotting a hard-to-read 3D-curve of CP revenue in terms of q_A and q_B , we have preferred to draw 2D-curves in terms of one of the parameters for various values of the second parameter close to optimal). It is interesting to notice that, for such a set of side payments, there is a price war on the user prices, i.e., $p_A^{\text{NE}} = p_B^{\text{NE}} = 0$ here. In fact it corresponds to the symmetric ($q_A = q_B$) point of the price war set $\mathcal{P}_w$ described in Figure 2 for which the sum of side payments is minimized. Indeed, if $p_A^{\text{NE}} = p_B^{\text{NE}} = 0$, then the revenue of the CP can be rewritten as $U_1 = \frac{1}{\alpha} e^{-\frac{\alpha D_0}{2} (q_A + q_B) - 1}$, and then is maximized when $q_A + q_B$ is minimal.

At this point the revenues of the stakeholders are $U_A = U_B \approx 0.04$, and $\text{UW}_{\text{CP}} = U_1 \approx 0.27$. Hence the revenues of ISPs are much smaller than the one of the CP. Note that the situation is quite counter-intuitive, since the CP gains to introduce side payments. This is because those payments exacerbate the competition between ISPs, which is beneficial to end users, and finally to the CP who can reach more customers.

B. Determined by the ISPs, through a game

If we instead assume that the side payments are non-cooperatively determined by the ISPs, we are led to study the best response of each ISP to the other ISP side payment. As shown in Figure 8, the best response is first increasing in the other ISP side payments, and then falls to zero above a threshold, which is approximately 2.80. Since the best-response to a null price is 2.80, it follows that $(0, 2.80)$ and the symmetric point $(2.80, 0)$ are Nash equilibria, and they are the only ones. It is interesting to notice that the resulting side

Fig. 8. The optimal side payment of each ISP as a function of the opponent ISP side payment. There is a threshold $q_i \approx 2.80$ beyond which the best response falls to zero. There are two Nash equilibria (NE) where one ISP sets its side payment to zero.

payments are not symmetric at Nash equilibrium, so are the revenues equalling 0.42 for the ISP with side payment 2.80 and 0.34 for the other ISP.

Now let us compare that outcome to the case without any side payments. From Subsection IV-A, with $\beta = 1.5$, the revenue of ISPs is 0.33. Hence the ISPs global revenue increases by about 15%, which goes in the direction of ISPs arguments about side payments improving their revenue. On the other side, the CP revenue decreases from 0.25 to 0.06, hence losing nearly 75% of its value. The benefit of ISPs is then at the expense of the CP and consequently of the user welfare.

C. Determined by a regulator

A regulator can either decide to maximize the revenue of the supply chain (sum of utilities of the ISPs plus the CP), the user welfare (end-users surplus), or the social welfare (including user welfare and all providers utilities).

The total value of the *supply chain* is the total revenue got from the users, i.e., $U_1 + U_A + U_B$.

User welfare can be decomposed into two components: the user welfare due to the existence of the CP -that is computed in (3)-, and the user welfare due to the presence of the ISPs. Let us focus on the latter part: we have assumed that users not connected to the Internet perceive a cost p_0 (thus p_0 can be seen as the value of the connectivity service). When a user decides to subscribe to ISP i and pays the corresponding price p_i , its benefit is then $p_0 - p_i$ with respect to the no-ISP situation: the user does not bear anymore the cost p_0 of not having Internet access, and instead perceives the monetary cost p_i . Aggregating over the whole population, the user welfare that is due to the presence of the ISPs (with their prices p_A and p_B) equals

$$UW_{\text{ISPs}} = \sigma_A(p_0 - p_A) + \sigma_B(p_0 - p_B).$$

The global user welfare generated by the system (ISPs and CP) is therefore

$$UW = UW_{\text{CP}} + UW_{\text{ISPs}} \quad (10)$$

Finally, *social welfare* is defined as the overall value of the service for the society. It therefore includes the surpluses of all actors, and equals $SW = U_1 + U_A + U_B + UW$. Note that Social Welfare also corresponds to the total value that the service has for subscribers, without considering any monetary exchanges because they stay within the society. We indeed obtain, simplifying the sum of the terms in SW:

$$SW = (\sigma_A + \sigma_B) \left(p_0 + \left(p_1 + \frac{1}{\alpha} \right) e^{-\alpha p_1} \right),$$

where the term $(\sigma_A + \sigma_B)p_0$ is the value of the connectivity service for ISPs' subscribers, and the other term is the value of the CP service for CP subscribers, computed as

$$(\sigma_A + \sigma_B) \int_{p_1}^{\infty} \alpha e^{-\alpha x} x \, dx.$$

1) *Side payments to maximize User Welfare*: Since CP revenue and user welfare are equal at Nash equilibrium, it follows that user welfare is maximized when the CP revenue is maximized. This case has already been treated in Subsection V-A.

2) *Side payments to maximize Social Welfare*: We have obtained numerically that social welfare is maximized for the same side payments as the ones maximizing the CP revenue and the user welfare.

3) *Side payments to maximize the supply chain value*: The supply chain value is maximized when the side payments are both null, which has been studied in Subsection IV-A. In this neutral case the revenue of ISPs is approximately 0.33 whereas that of the CP (and the induced user welfare) is 0.25. We can remark that among the three alternatives considered in this subsection, this one leads to the fairest revenue sharing between stakeholders.

VI. DISCUSSION, CONCLUSIONS AND FUTURE WORK

We have provided in this paper a model describing the interactions between two ISPs in competition, a CP, and end users connecting to the network. With respect to the literature, we believe that considering competitive ISPs and a single CP is a more realistic representation of the current network where we have a quasi-monopoly for some applications (for instance YouTube or Netflix), while several ISPs are in competition (an argument of ISPs). The goal is to study the impact of side payments on providers' revenues, and conclude whether they can help reduce the unfairness of the current revenue sharing among all actors, as claimed by ISPs in the current network neutrality debate. We restricted the analysis to two ISPs because in many cases, competition is limited to a very small number of ISPs, and very often only two ISPs are available to a given user [22]; but a numerical analysis can be extended to more than two.

In this paper, we have presented a three-level game where (from the largest to the shortest time scale) the side payments are first determined, then a pricing game is played between ISPs, followed by the content provider price, and finally, knowing all those prices, end users choose their ISP (or none if too expensive) and possibly decide to subscribe also to the CP service. All those levels are played by backward induction,

meaning that players anticipate the solutions of the later games when choosing their strategies.

Our results have highlighted the fact that side payments, *unless decided by ISPs*, have little chance to address the concern from ISPs regarding the fairness of the revenue sharing associated with users accessing content through their infrastructures. This is due, to a great extent, to the competition played among ISPs on the access prices that drives their revenues to low values. On the other hand, the CP being in a monopolist situation, always obtains significant revenues. An interesting paradox we have highlighted is that side payments may be beneficial to the CP. Nevertheless, when side payments are decided by the ISPs (non-cooperatively), it can be beneficial to them, but at the expense of both the CP and the users. Remarkably, the side payments maximizing social and user welfare are the same than those maximizing the CP revenue. But looking at the whole supply chain, in order to avoid too big disparities between revenues of providers, the neutral case is the most suitable. If the side payments are decided non-cooperatively by ISPs, in our experiment, one (only) is a big winner, while the other ISP gains a bit more than in the neutral case. This asymmetry may be a problem and can create complicated tensions and negotiations.

As future research, we would like to go into several directions: first to include several CPs with different contents, but such that some end users are targeting only a subset of them, for all possible subsets. ISPs may also charge each other to let the CPs not connected to them reach their own customers (transit pricing). Other extensions to our model could include architecture investment and content innovation characteristics, for the ISPs and the CP respectively.

ACKNOWLEDGMENTS

This work has been supported by INRIA's ARC project MENEUR and EuroNF NoE's project INNIS on network neutrality, and by the French ANR VERSO project CAPTURES on competition between providers.

REFERENCES

- [1] Sandvine, "Global Internet phenomena report: Fall 2011," available at www.sandvine.com, 2011.
- [2] T. Lenard and R. E. May, *Net Neutrality or Net Neutering: Should Broadband Internet Services be Regulated*. Springer, 2006.
- [3] A. Odlyzko, "Network neutrality, search neutrality, and the never-ending conflict between efficiency and fairness in markets," *Review of Network Economics*, vol. 8, no. 1, pp. 40–60, 2009.
- [4] A. Ghosh and A. Sayedi, "Expressive auctions for externalities in online advertising," in *Proceedings of the 19th World Wide Web conference (WWW)*, Raleigh, NC, USA, April 2010.
- [5] D. Liu, J. Chen, and A. Whinston, "Competing keyword auctions," in *Proceedings of 4th Workshop on Ad Auctions*, Chicago, IL, USA, July 2008.
- [6] ARCEP, "Neutralité de l'internet et des réseaux : propositions et orientations," 2010, http://www.arcep.fr/uploads/tx_gspublication/net-neutralite-orientations-sept2010.pdf.
- [7] S. Wong, R.-M. J., and E. Altman, "Public consultations on Net Neutrality 2010: USA, EU and France," SSRN, Tech. Rep., 2010, http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1684086.
- [8] D. Fudenberg and J. Tirole, *Game Theory*. MIT Press, Cambridge, Massachusetts, 1991.
- [9] M. Osborne and A. Rubinstein, *A Course in Game theory*. MIT Press, 1994.

- [10] R. Ma, D.-M. Chiu, J. Lui, V. Misra, and D. Rubenstein, "Interconnecting eyeballs to content: A Shapley value perspective on ISP peering and settlement," in *Proc. of International Workshop on Economics of Networked Systems (NetEcon)*, 2008, pp. 61–66.
- [11] —, "On cooperative settlement between content, transit and eyeball internet service providers," in *Proc. of ACM International Conference on Emerging Networking EXperiments and Technologies (CoNEXT)*, 2008.
- [12] J. Musacchio, J. Walrand, and G. Schwartz, "Network neutrality and provider investment incentives," in *Proc. of 41st Asilomar Conference on Signals, Systems and Computers (ACSSC)*, Nov 2007, pp. 1437 – 1444.
- [13] J. Musacchio, G. Schwartz, and J. Walrand, "A two-sided market analysis of provider investment incentives with an application to the net-neutrality issue," *Review of Network Economics*, vol. 8, no. 1, 2009.
- [14] C. Saavedra, "Bargaining power and the net neutrality debate," 2009, sites.google.com/site/claudiasaavedra/attachments/bargaining_power.pdf.
- [15] E. Altman, P. Bernhard, S. Caron, G. Kesidis, J. Rojas-Mora, and S. Wong, "A study of non-neutral networks with usage-based prices," in *Proc. of 3rd ETM Workshop*, Amsterdam, The Netherlands, Sept 2010.
- [16] P. Coucheney, P. Maillé, and B. Tuffin, "Network Neutrality Debate and ISP Inter-Relations: Traffic Exchange, Revenue Sharing, and Disconnection Threat," Inria, Rapport de recherche, Jan. 2012. [Online]. Available: <http://hal.inria.fr/hal-00659278>
- [17] E. Altman, A. Legout, and Y. Xu, "Network non-neutrality debate: An economic analysis," in *Proceedings of Networking*. Barcelona, Spain: LNCS 6641, Springer Verlag, May 2011, pp. 68–81.
- [18] P. Njoroge, A. Ozdaglar, N. Stier-Moses, and G. Weintraub, "Investment in two sided markets and the net neutrality debate," Columbia University, Decision, Risk and Operations Working Papers Series, Tech. Rep. DRO-2010-05, 2010.
- [19] F. Boussion, P. Maillé, and B. Tuffin, "Net neutrality debate: Impact of competition among ISPs," in *Proceedings of the Fourth International Conference on Communication Systems and NETWORKS (COMSNETS)*, Bangalore, India, 2012.
- [20] S. Caron, G. Kesidis, and E. Altman, "Application neutrality and a paradox of side payments," in *Proceedings of the Re-Architecting the Internet Workshop*. ACM, 2010.
- [21] M. Ben-Akiva and S. Lerman, *Discrete choice analysis*. The MIT Press, Cambridge, Massachusetts, 1985.
- [22] N. Economides, "The economics of the internet backbone," in *Handbook of Telecommunications*, S. Majumdar et al., Ed. Amsterdam: Elsevier Publishers, 2005, vol. 2, pp. 373–412.

Pierre Coucheney received his Master degree in applied mathematics from universit Paris 6, and his Ph.D. in computer science from university of Grenoble. He is currently in a post-doctoral position at Inria-Rennes, working on the network neutrality topic. His research interests include game theory and its applications to the optimization and economic modelling of networks.

Patrick Maill graduated from Ecole polytechnique and Telecom ParisTech, France, in 2000 and 2002, respectively. He has been an assistant professor at the Networks, Security, Multimedia department of Telecom Bretagne since 2002, where he obtained his PhD in applied mathematics in 2005, followed by a 6-month visit to Columbia University in 2006. His research interests are on game theory and economic concepts applied to telecommunication ecosystems: resource pricing, routing, consequences of user selfishness on network performance.

Bruno Tuffin received his PhD degree in applied mathematics from the University of Rennes 1 (France) in 1997. Since then, he has been with INRIA in Rennes. He spent eight months as a postdoc at Duke University in 1999. His research interests include developing Monte Carlo and quasi-Monte Carlo simulation techniques for the performance evaluation of telecommunication systems, and developing new Internet-pricing schemes and telecommunication-related economic models. He has published more than one hundred papers on those issues. He has also led or participated into several French and European projects, and co-organized several conferences. He is currently Associate Editor for INFORMS Journal on Computing, ACM Transactions on Modeling and Computer Simulation and Mathematical Methods of Operations Research. He has written or co-written two books devoted to simulation: "Rare event simulation using Monte Carlo methods" published by John Wiley & Sons in 2009, and "La simulation de Monte Carlo" (in French), published by Hermes Editions in 2010.