

Application Neutrality and a Paradox of Side Payments*

Eitan Altman
INRIA
2004 Route des Lucioles
Sophia-Antipolis, France
eitan.altman@sophia.inria.fr

Stéphane Caron
Département d'Informatique
École Normale Supérieure
45 rue d'Ulm, Paris, France
stephane.caron@ens.fr

George Kesidis
CS&E and EE Depts
Penn State University
University Park, PA, 16802
kesidis@engr.psu.edu

ABSTRACT

The ongoing debate over “net neutrality” covers a broad set of issues related to the regulation of public networks. This paper contains two separate contributions: (a) an extension of the quadratic-utility framework we proposed in [4] to study the impact of side payments in a system involving pluralities of access and content providers; and (b) a variation of this model to deal with the question of application neutrality. Our analysis of the generalized framework (a) reveals an interesting “paradox” that did not occur with monopolistic players: side payments handicap the providers who perceive them. Application neutrality (b) refers to price discrimination: ISPs charging consumers different fees depending on their use of the network (web surfing, VoIP, file sharing, *etc.*). We analyze the consequences of such discrimination for a simple two-application setting.

Keywords

Net neutrality, side payments, price discrimination, regulation, public networks

1. INTRODUCTION

Different issues have been raised in the context of the net neutrality debate.

- for Tim Berners-Lee¹, it means that “if I pay to connect to the Net with a certain quality of service, and you pay to connect with that or greater quality of service, then we can communicate at that level”;

*A full version of this paper is available at [8].

¹“Net Neutrality: This is serious”, *timbl's blog*, June 2009. <http://dig.csail.mit.edu/breadcrumbs/node/144>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ACM CoNEXT 2009, December 1–4, 2009, Rome, ITALY.
Copyright 2009 ACM X-X-X-X/XX/XX ...\$5.00.

- for Tim Wu², the main idea is that “a maximally useful public information network aspires to treat all content, sites, and platforms equally”; and
- for Hahn *et al.* [1], it “usually means that broadband service providers charge consumers only once for Internet access, do not favor one content provider over another, and do not charge content providers for sending information over broadband lines to end users”.

These definitions raise different questions, including connectivity, non-discrimination of application, type or origin, and network access pricing. It seems therefore more relevant to see net neutrality rather as a topic involving a range of issues regarding the regulation of public networks. With this perspective, [2] discusses five “flavors” of net neutrality: (a) content neutrality, (b) blocking and rerouting, (c) denying IP-network interconnection, (d) network management, and (e) premium service fees. (b) pertains to providers discriminating packets in favor of their own or affiliated content, while (c) is related to agreements between last-mile and backbone providers. (d) has been a central argument for ISPs protesting the enforcement of net neutrality principles: they defend their right to manage their own networks, especially in order to deal with congestion issues (*e.g.*, due to high-volume peer-to-peer (P2P) traffic, see the decision on “Comcast v. the FCC”). They claim that regulations would act as a disincentive for capacity expansion of their networks.

In this paper, we address issues (a): side payments and application neutrality. Both are rooted in historical issues: massive copyright infringements led copyright holders to seek remuneration from ISPs, while congestion due to P2P file sharing led some providers to adopt non application-neutral policies (*e.g.*, Comcast throttling BitTorrent traffic). However, in what follows we abstract from these historical causes, studying side payments in either direction (from ISPs to CPs and con-

²“Network Neutrality FAQ”, http://timwu.org/network_neutrality.html

versely) and considering the impact of non application-neutral pricing independently from congestion.

Our work focuses on *usage-based* pricing: we assume consumers are, to some extent, willing to pay usage-dependent fees, *e.g.*, as overages over fixed monthly fees. Providers are then competing to settle on their usage-based prices, their goal being to maximize revenues coming from these charges. Note that a null price does not mean a provider has no income, but rather that all his monthly revenues are coming from flat-rate pricing. Study of the flat-rate regime is, however, out of the scope of this paper. See, *e.g.*, [3] for a discussion of both regimes.

The rest of the paper is organized as follows. We discuss related work in subsection 1.1 and describe our framework in section 2. In section 3, we study the impact of side payments on the competition between providers. We extend our framework in section 4 to analyze the effect of non application-neutral pricing from the ISPs. We conclude in section 5.

1.1 Related Work

We studied in [4] some net neutrality related issues like side payments and premium service fees (e), limiting our consideration to monopolistic providers. We are now lifting this limitation with an extended model including competition between multiple identical providers (it is actually a generalization of the idea sketched in [4]–VI).

Ma *et al.* [5] advocate the use of Shapley value as a fair way to share profits between providers. This approach yields Pareto optimality for all players, and expects in particular CPs, many of whom receive advertising revenues, to take part in the investments for network capacity. However, this approach is coalitional and there are many obstacles to its real-life implementation.

In [6], the authors address whether local ISPs should be allowed to charge remote CPs for the “right” to reach their end users (again, this is the side payment issue). Through study of a two-sided market, they determine when neutrality regulations are harmful depending on the parameters characterizing advertising rates and consumer price sensitivity.

[7] discusses the net neutrality debate to the light of historical precedents, especially dealing with the question of price discrimination. It also raises an interesting point about the way customers value the network, which is that connectivity is far more important than content. This fact has not been much surveyed in the literature, and unfortunately it is not in the scope of our work either since we focus on a market of identical CPs, not on the network as a whole.

2. PROBLEM SET-UP

Our model is an extension of the one in [4]. It encompasses three groups of players: the Internauts (end users), modeled collectively through their demand response, n_1 last-mile broadband providers (ISPs), and n_2 content providers (CPs). Consumers pay usage-dependent fees for service/content that requires one ISP and one CP. Providers then compete in a game to settle on their usage-based prices, which may turn out to be 0\$/byte, *i.e.*, only flat-rate fees would apply.

2.1 Demand Response

Let us denote by $p_{1i} \geq 0$ (resp. $p_{2i} \geq 0$) the usage-based price of the i^{th} ISP (resp. CP). These prices act as disincentives on consumers’ demand for content/bandwidth. We model this with a simple linear response: the amount users are ready to consume, given that they chose ISP i and CP j , is

$$D(p_{1i}, p_{2j}) = D_{\max} - d_1 p_{1i} - d_2 p_{2j}.$$

We are dealing here with a set of homogeneous users sharing the same response to price variations. The parameter D_{\max} reflects demand under pure flat-rate pricing.

Note that all providers may not measure demand on the same scale: ISPs focus on bandwidth consumption and express demand in bytes, while CPs are concerned with content consumption and/or advertising revenues, thus expressing demand in number of clicks or products sold (books, music albums, *etc.*). However, using a single metric (*e.g.*, bytes) is more convenient, and other metrics can be approximated from this one using an appropriate scaling factor.

In what follows, we furthermore suppose that users are only concerned with the total usage-based price they are charged, *i.e.*, they don’t care whether they are giving money to an ISP or a CP. As a consequence, we have $d_1 = d_2 = d$. Since demand should be non-negative, we can then define a maximum price p_{\max} such that

$$p_{1i} + p_{2j} \leq \frac{D_{\max}}{d} =: p_{\max}.$$

2.2 Customer Stickiness

As we suppose all providers of a given type propose the *same* type/quality of content/service, user decisions are only based on price considerations. For example, if an ISP charges a price significantly lower than the other ISPs, in the long run all customers will choose it and the others will have no choice but to align their prices or opt out of the game. Therefore, our homogeneity hypothesis means all ISPs (and similarly all CPs) have roughly the same prices:

$$\begin{aligned} p_{11} &\approx p_{12} \approx \dots \approx p_{1n_1}, \\ p_{21} &\approx p_{22} \approx \dots \approx p_{2n_2}. \end{aligned}$$

As providers play the usage-based pricing game, first-order differences between these prices may appear (*e.g.*, the

i^{th} ISP reducing his price by δp_{1i} to attract new end users). Consumers are then more likely to go to the cheapest providers of each group, but price differences may be too small to convince all of them to move and some will stay with their current provider. This phenomenon is known as *customer stickiness*, *inertia* or *loyalty*. To model it, we define the fraction σ_{ki} of users committed to the i^{th} provider of the k^{th} group ($k = 1$ for ISPs and 2 for CPs) as a function of $\mathbf{p}_k = (p_{k1}, \dots, p_{kn_k})$, *i.e.*, $\sigma_{ki} := \sigma(i, \mathbf{p}_k)$. Properties expected of the “stickiness function” σ are discussed in [8]. We chose the following model:

$$\sigma(i, \mathbf{p}_k) = \frac{1/p_{ki}}{\sum_{j=1}^{n_k} 1/p_{kj}} =: \sigma_{ki}. \quad (1)$$

The average usage-based price of network/content access for a customer is then $\bar{p}_k := \sum_i \sigma_{ki} p_{ki}$, *i.e.*, the harmonic mean of $\{p_{ki}\}$.

2.3 Non-discriminating setting

In a “neutral” setting with no side payments nor application discrimination, an ISP’s expected usage-based revenue is given by

$$U_{1i} = \sum_{j=1}^{n_2} \sigma_{1i} \sigma_{2j} D(p_{1i}, p_{2j}) p_{1i} = \sigma_{1i} D(p_{1i}, \bar{p}_2) p_{1i},$$

and similarly for the j^{th} CP. Necessary conditions for an interior Nash Equilibrium Point (NEP) are given by $\frac{\partial U_{ki}}{\partial p_{ki}}(\bar{p}_1, \bar{p}_2) = 0$ for $k = 1, 2$. The resulting system is linear and straightforward to solve, and its solution is a local maximum in revenue for all players (*i.e.*, a Nash equilibrium). At the NEP, demand and revenues are given by

$$D^* = \frac{n_1 n_2}{n_1 n_2 + n_1 + n_2} D_{\max},$$

$$U_{ki}^* = \frac{n_{3-k}^2}{(n_1 n_2 + n_1 + n_2)^2} U_{\max},$$

for $k = 1, 2$. As expected, customers benefit from competition among the providers. With 2 ISPs and 2 CPs, demand is only 50% of its potential D_{\max} , while it is about 70% of D_{\max} with 5 ISPs and 5 CPs. This base model also encompasses two expected behaviors: providers from one group benefit from increased competition in the other group, while their revenues are significantly reduced by increased competition in their own group. Note that competition in a provider’s own group has much more impact on his income than competition in the other group.

3. SIDE PAYMENTS

Suppose now that there are side payments between the two groups of providers. We introduce a usage-based fee p_s from the CPs to the ISPs. When $p_s > 0$,

CPs remunerate the ISPs, *e.g.*, to support the bandwidth costs. On the other hand, if $p_s < 0$, ISPs give money to the CPs, *e.g.*, for copyright remuneration. We suppose ISPs or CPs receive side payments collectively and ultimately share the aggregate amount proportionally to their customer shares. Hence, provider revenues become:

$$U_{1i} = \sigma_{1i} D(p_{1i}, \bar{p}_2) (p_{1i} + p_s) \quad \text{for } i \in \llbracket 1, n_1 \rrbracket,$$

$$U_{2j} = \sigma_{2j} D(\bar{p}_1, p_{2j}) (p_{2j} - p_s) \quad \text{for } j \in \llbracket 1, n_2 \rrbracket,$$

where all demand and price factors (resp. D and $p_{ki} \pm p_s$) are deemed non-negative. It is expected that p_s is *not* a decision variable for any player or group of player. Indeed, since utilities are monotonic in p_s , those controlling it would always be incented to increase or decrease it (if they are ISPs or CPs respectively), leading the other players to opt out of the competition. Therefore, p_s would normally be regulated and we will consider it a *fixed* parameter from now on.

Necessary conditions for an interior equilibrium (null first-derivatives) yield:

$$\left[\frac{\bar{p}_1}{p_{\max} - \bar{p}_1 - \bar{p}_2} - \frac{1}{n_1} \right] (\bar{p}_1 + p_s) + p_s = 0, \quad (2)$$

$$\left[\frac{\bar{p}_2}{p_{\max} - \bar{p}_1 - \bar{p}_2} - \frac{1}{n_2} \right] (\bar{p}_2 - p_s) - p_s = 0. \quad (3)$$

With the introduction of non-null p_s , this system is not linear any more. We provide here the complete study of a simplified setting where $n_1 = n_2 = 2$. An extension to an arbitrary number n of ISPs and CPs is given in our working paper [8].

3.1 Interior equilibria

Define $u := (\bar{p}_1 + \bar{p}_2)/p_{\max}$, $v := (\bar{p}_1 - \bar{p}_2)/p_{\max}$ and $s := p_s/p_{\max}$. Equilibrium conditions become:

$$-2sv - 2u^2 - v^2 + u = 0, \quad (4)$$

$$-3uv - 2s + v = 0, \quad (5)$$

where (4) \times (2) + (3) and (5) \times (2) - (3). Equilibrium prices, demand and revenues are now solvable in closed form³. Before we expose the solutions, an important observation we can make at this point is given by the following theorem (we prove it in [8]):

THEOREM 1. *When $n_1 = n_2 = 2$, there is an interior NEP iff*

$$\left| \frac{p_s}{p_{\max}} \right| \leq \max_{x \in [\frac{1}{4}, \frac{1}{2}]} \sqrt{\frac{(1-x)(1-2x)^2(4x-1)}{36x}} \approx 4.64\%$$

In other words, regulated side payments can only occur to a small extent ($|p_s| < 4.64\%$ of p_{\max}), otherwise there will be no interior NEP, which means one of the two

³Yet, these expressions are complicated, long and of no extra-computational interest.




Figure 1: Demand and revenues at NEP_1 .

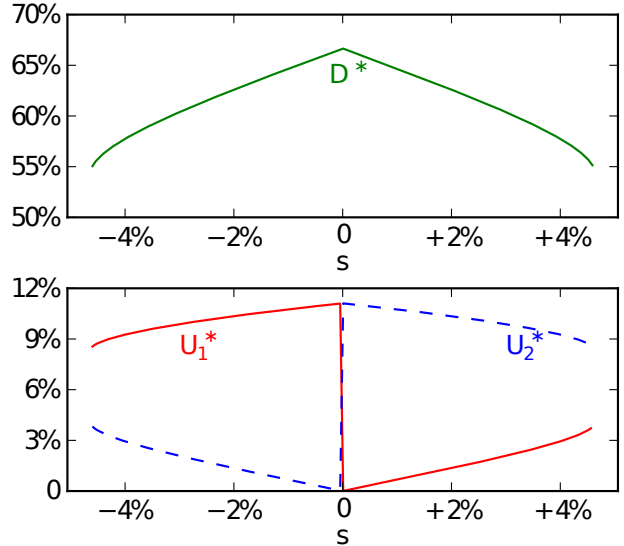


Figure 2: Demand and revenues at NEP_2 .

groups of players will opt out of the usage-based pricing game.

There are two solutions to (2) and (3), and we checked that both of them are Nash equilibria. Demand and revenues at NEP_1 and NEP_2 are shown in Figures 1 and 2. Note that NEP_1 is consistent with the results of the non-discriminating setting (when $s = 0$, $\bar{p}_k^* = p_{\max}/4$, $D^* = D_{\max}/2$ and $U_k^{i*} = U_{\max}/16$ for $k = 1, 2$), while NEP_2 does not exist when $s = 0$ (there is a discontinuity in equilibrium prices at this point). Both equilibrium points share the same “paradox”: providers receiving side payments eventually achieve *less revenue* than the others.

3.2 Convergence to equilibrium

In what follows, we consider $s > 0$. The roles of ISPs and CPs are swapped for $s < 0$. All providers are deemed to act independently under a best-response behavior. Thus, the vector field

$$(\bar{p}_1, \bar{p}_2) \mapsto \left(\frac{\partial U_{1i}}{\partial p_{1i}}(\bar{p}_1, \bar{p}_2), \frac{\partial U_{2j}}{\partial p_{2j}}(\bar{p}_1, \bar{p}_2) \right)$$

is an appropriate indicator of the aggregate “trends” of the system. Computations revealed the following best-response behavior: if $\bar{p}_1 > \bar{p}_1^*(NEP_2)$, the system is attracted by NEP_1 ; otherwise, unless $\bar{p}_1 = \bar{p}_1^*(NEP_2)$, combined actions of ISPs and CPs result in a consensus where all usage-based revenues for ISPs come from side payments. Let us call this consensus NEP_B . Solving (3) with $\bar{p}_1 = 0$ yields

$$\bar{p}_2^*(NEP_B) = \frac{p_{\max}}{6} \left(1 + s + \sqrt{s^2 + 14s + 1} \right),$$

where corresponding expressions for demand and revenues follow directly. At NEP_B , demand is higher than

at NEP_1 or NEP_2 while ISP revenues turn out to be lower (and CPs revenues higher) than at NEP_2 (see Figure 3).

To avoid NEP_2 , ISPs can reduce their prices, relying on increased demand to compensate revenue losses. Yet, CPs can then raise their prices, mitigating the increase in demand and constraining the ISPs to reduce their prices further, down to the point where $\bar{p}_1 = 0$. This is how best-response behavior causes ISPs to eventually prefer NEP_B over NEP_2 , though global optimization shows it is less profitable for them.

3.3 Epilogue: a Poisoned Chalice

For 2 ISPs and 2 CPs, we saw that the introduction of side payments in the model yielded a game with two possible outcomes: if initial prices of side payment receivers are high enough, providers will reach an interior equilibrium where receivers get less revenue than payers (the higher the side payments, the lower the revenue), otherwise, receivers will opt out of the usage-based pricing game, depending only on side payments for their usage-based revenues. These conclusions also hold for an arbitrary number n of ISPs and CPs (see [8]). In both cases, the “paradox” of side payments is that they act as a *handicap* for those who receive them.

4. APPLICATION NEUTRALITY

Now, let us consider a different issue related to the net neutrality debate: to what extent should access providers be allowed to perform packet inspection on the data they transmit? In particular, given that they charge usage-based fees, should they be allowed to perform price discrimination depending on the application at use (*e.g.*, video chat, media streaming, ...)? In this




Figure 3: Demand and revenues at NEP_B.

section, we study the impact of such discrimination in a configuration with two crude example types of applications: web surfing and P2P file sharing.

4.1 Additional problem set-up

We extend our model to a setting with three groups of providers: (1) ISPs, providing last-mile access to the Internauts, (2) Web Content Providers (Web CPs), *e.g.*, search engine portals (recall all providers in a group are deemed identical, so we assume all Web CPs provide the same type of client-server HTTP content as well), and (3) P2P Content Providers (P2P CPs), *e.g.*, private P2P networks operated in cooperation with copyright holders.

Users choose an ISP, a Web CP and a P2P CP. To access web (resp. P2P) content, they pay usage-based fees to both their ISP and their Web CP (resp. P2P CP). We denote by n_1 , n_2 and n_3 the numbers of ISPs, Web CPs and P2P CPs respectively. These groups are not coalitions: in a group, each provider is independent from the others and seeks to maximize their own revenue.

In a neutral setting, the i^{th} ISP charges a single price p_{1i} for all types of traffic, while otherwise he may set up two different prices $p_{12,i}$ and $p_{13,i}$ for HTTP and P2P traffic respectively. Denote by p_{2j} (resp. p_{3j}) the usage-based price of the j^{th} Web CP (resp. P2P CP). We introduce two separate demand-response profiles for both types of content: when ISP i , Web CP j and P2P CP l are chosen, demands for HTTP and P2P content are

$$\begin{aligned} D_2 &= D_{2\max} - d_2(p_{12,i} + p_{2j}), \\ D_3 &= D_{3\max} - d_3(p_{13,i} + p_{3l}), \end{aligned}$$

with $p_{12,i} = p_{13,i} =: p_{1i}$ in the neutral setting. As previously, we write $p_{k\max} := D_{k\max}/d_k$.

The portion of users committed to the i^{th} provider of the k^{th} group is still modeled as

$$\sigma_{ki} := \sigma(\mathbf{p}_k) = \frac{1/p_{ki}}{\sum_j 1/p_{kj}}$$

We will see in 4.3 how to generalize this to ISPs charging two different prices instead of one. Revenues for ISP i , Web CP j and P2P CP l are given by

$$\begin{aligned} U_{1i} &= \sigma_{1i} (D_2 p_{12,i} + D_3 p_{13,i}), \\ U_{2j} &= \sigma_{2j} D_2 p_{2j}, \\ U_{3l} &= \sigma_{3l} D_3 p_{3l}. \end{aligned}$$

Finally, we define the normalized sensitivity to usage-based pricing α and the maximum prices ratio γ ,

$$\alpha := \frac{d_2}{d_2 + d_3} \text{ and } \gamma := \frac{p_{2\max}}{p_{3\max}}, \quad (6)$$

and make the following assumptions:

- $\alpha \geq 1/2 \Leftrightarrow d_2 > d_3$: consumers are more sensitive to usage-based pricing of web content than of file sharing.
- $\gamma < 1 \Leftrightarrow p_{2\max} < p_{3\max}$: customers are ready to pay more for content exchanged on P2P sharing systems (movies, music, *etc.*) than for web pages.

As discussed in section 1, we also suppose that ISPs have enough bandwidth to supply all the demand, *i.e.*, there is no congestion.

4.2 Neutral setting

We now consider the broader setting with non-monopolistic providers (*i.e.*, $n_k > 1$ for $k \in \llbracket 1, 3 \rrbracket$) where application neutrality is enforced. In particular, $U_{1i} = \sigma_{1i}(D_2 + D_3)p_{1i}$. We find necessary conditions for an interior NEP to be:

$$\begin{aligned} (n_1 + 1)\bar{p}_1 + \alpha\bar{p}_2 + (1 - \alpha)\bar{p}_3 &= \alpha p_{2\max} + (1 - \alpha) p_{3\max}, \\ \bar{p}_1 + (n_2 + 1)\bar{p}_2 &= p_{2\max}, \\ \bar{p}_1 + (n_3 + 1)\bar{p}_3 &= p_{3\max}, \end{aligned}$$

which resolution is straightforward. We checked that any solution to this system is also an NEP.

4.3 Non-neutral setting

When application non-neutral pricing is allowed, the ISPs' utilities are $U_{1i} = \sigma_{1i}(D_2 p_{12,i} + D_3 p_{13,i})$, where σ_{1i} refers to the portion of users gathered by ISP i given his prices $p_{12,i}$ and $p_{13,i}$. There are multiple ways to generalize equation (1) to multiple criteria: *e.g.*, one could apply σ to the mean price $(p_{12,i} + p_{13,i})/2$ or model σ_{1i} as a convex combination of $\sigma_{12,i}$ and $\sigma_{13,i}$. We chose

$$\sigma_{1i} := \sigma(i, \tilde{\mathbf{p}}_1) = \frac{1/\tilde{p}_{1i}}{\sum_{j=1}^{n_1} 1/\tilde{p}_{1j}} \quad (7)$$

where $\tilde{p}_{1i} := \sqrt{\alpha\gamma} p_{12,i} + (1 - \sqrt{\alpha\gamma}) p_{13,i}$. In other words, we apply the original stickiness model (1) to a combined price \tilde{p}_{1i} defined as a convex combination of $p_{12,i}$ and $p_{13,i}$. See [8] for a discussion of this choice.

In this model, we found necessary conditions for an interior NEP to be:

$$\begin{aligned} \frac{\alpha}{\sqrt{\alpha\gamma}} (\tilde{D}_2 - \bar{p}_{12}) &= \frac{n_1 - 1}{n_1 \tilde{p}_1} (\alpha \tilde{D}_2 \bar{p}_{12} + (1 - \alpha) \tilde{D}_3 \bar{p}_{13}), \\ \frac{\alpha}{\sqrt{\alpha\gamma}} (\tilde{D}_2 - \bar{p}_{12}) &= \frac{1 - \alpha}{1 - \sqrt{\alpha\gamma}} (\tilde{D}_3 - \bar{p}_{13}), \\ \bar{p}_{12} + (n_2 + 1) \bar{p}_2 &= p_{2\max}, \\ \bar{p}_{13} + (n_3 + 1) \bar{p}_3 &= p_{3\max}, \end{aligned}$$

where $\tilde{p}_1 := \sqrt{\alpha\gamma} \bar{p}_{12} + (1 - \sqrt{\alpha\gamma}) \bar{p}_{13}$. This system can be eventually rewritten as two polynomial equations on \bar{p}_{12} and \bar{p}_{13} . It is solvable in closed form, and any of its solutions is also an NEP.

4.4 Discussion of experimental results

In our numerical experiments, we took $\alpha = 0.8$ and $\gamma = 0.3$. We used Sage for our computations, and all our scripts are available online⁴. Additional plots and results are provided in [8].

The main result we observed is that ISPs and Web CPs prefer the non-neutral setting, while P2P CPs benefit from neutrality regulations. The impact of non-neutral pricing on providers' revenues varies with competition: increased competition brings less benefits for Web CPs and less losses for P2P CPs. Yet, competition has almost no effect on the gains of ISPs (see Figure 4).




Figure 4: Relative variations in revenues between neutral and non-neutral pricing (n is the number of providers in each group, $\alpha = 0.8$ and $\gamma = 0.3$).

⁴<http://www.cse.psu.edu/~kesidis/nn-scripts/index.html>

5. CONCLUSION

We presented a framework to discuss the impact of two net neutrality related issues, side payments and application neutrality, on the three-way interactions between end users, ISPs and CPs. Our model relies on a simple linear demand-response to usage-based prices, and it encompasses customer loyalty.

In section 3, we studied the effect of regulated side payments between the ISPs and CPs. We determined two possible outcomes of the competition, both of them showing the same “paradox”: side payments are actually a handicap for those who receive them, insofar as they reduce their equilibrium revenues.

We addressed the issue of application neutrality in section 4 with a simple setting opposing two types of content: web content and file sharing, the latest showing lower price sensitivity and higher willingness to pay due to the nature of exchanged content. Our analysis suggested that ISPs and Web CPs benefit from application non-neutral practices, while P2P CPs are better off in a neutral setting.

Additional results and detailed computations are available in our working paper [8].

6. REFERENCES

- [1] R. Hahn and S. Wallsten, “The Economics of Net Neutrality”, *Economists’ Voice*, The Berkeley Economic Press, **3**(6), pp. 1-7, 2006.
- [2] R. B. Chong, “The 31 Flavors of Net Neutrality”, *12 Intellectual Property Law Bulletin*, 2008, v. 12.
- [3] G. Kesidis, A. Das, G. de Veciana, “On Flat-Rate and Usage-based Pricing for Tiered Commodity Internet Services”, *Proc. CISS 2008*, Princeton.
- [4] E. Altman, P. Bernhard, S. Caron, G. Kesidis, J. Rojas-Mora and S. Wong, “A Study of Non-Neutral Networks with Usage-based Prices”, June 2010. <http://arxiv.org/abs/1006.3894>
- [5] R. 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”, *Proc. ACM CoNext 2008*.
- [6] 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*, 2009, vol. 8, issue 1.
- [7] A. Odlyzko, “Network Neutrality, Search Neutrality, and the Never-ending Conflict between Efficiency and Fairness in Markets”, *Review of Network Economics*, 2009, vol. 8, issue 1.
- [8] E. Altman, S. Caron and G. Kesidis, “Application Neutrality and a Paradox of Side Payments”, Aug. 2010. Available on <http://arxiv.org>.