Application Neutrality and a Paradox of Side Payments*

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

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

ABSTRACT

The ongoing debate over net neutrality covers a broad set of issues related to the regulation of public networks. In two ways, we extend an idealized usage-priced game-theoretic framework based on a common linear demand-response model [1]. First, we study the impact of "side payments" among a plurality of Internet service (access) providers and content providers. In the non-monopolistic case, our analysis reveals an interesting "paradox" of side payments in that overall revenues are reduced for those that receive them. Second, assuming different application types (*e.g.*, HTTP web traffic, peer-to-peer file sharing, media streaming, interactive VoIP), we extend this model to accommodate differential pricing among them in order to study the issue of application neutrality. Revenues for neutral and non-neutral pricing are compared for the case of two application types.

Keywords

Net neutrality, side payments, regulated public networks

1. INTRODUCTION

Different issues have been raised in the context of the net neutrality debate. For Tim Berners-Lee[1], 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." For Tim Wu[2], the main idea is that "a maximally useful public information network aspires to treat all content, sites, and platforms equally." For Hahn *et al.* [5], 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, based on type or origin, and network access pricing. Net neutrality is a subject involving a range of issues regarding the regulation of public networks [3]: (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 "Comcast v. the FCC" decision [4]). They claim that regulations would act as a disincentive for capacity expansion of their networks.

In this paper, considering usage-based revenues, we address issues from topic (a): side payments among providers and application neutrality. Massive copyright infringements led copyright holders to seek remuneration from ISPs, while congestion due to P2P file sharing led some providers to adopt not application-neutral policies (*e.g.*, Comcast throttling BitTorrent traffic) and to consider usage pricing (as a congestion penalty, for overage of a quota, or for premium service, *e.g.*, [8]). In what follows, we study side payments (from Internet Service (access) Providers (ISPs) to Content Providers (CPs), or in the reverse direction) and consider the impact of not application-neutral pricing independent of

*A full version of this paper is available at [2]. This work was supported in part by the NSF under grant 0916179.

[1] "Net Neutrality: This is serious", *timbl's blog*, June 2009. http://dig.csail.mit.edu/breadcrumbs/node/144

[2] "Network Neutrality FAQ", http://timwu.org/network_neutrality.html

ACM ReArch 2010, November 30, 2010, Philadelphia, USA.

congestion.

That is, we assume consumers are, to some extent, willing to pay usage-dependent fees, *e.g.*, as overages of 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 in the following does not mean a provider has no income, but rather that all their monthly revenues come from flat-rate priced service components. Study of the flat-rate regime is, however, out of the scope of this paper. See, *e.g.*, [7] for a comparison of both regimes for a simple model of congestion management.

The rest of the paper is organized as follows. We discuss related work in subsection 1.1 and describe our problem 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 not application-neutral pricing by the ISPs. We conclude in section 5.

1.1 Related Work

Previously, we considered certain net neutrality related issues like side payments and premium service fees (e), limiting our consideration to monopolistic providers [1]. In the following, we extend this model to include competition between multiple identical providers (actually based on an idea sketched in Section IV of [1]).

Ma *et al.* [9] advocate the use of Shapley values 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 network-capacity investments. However, this approach is coalitional and there are many obstacles to its real-life implementation.

In [10], 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[3].

The net neutrality debate is discussed in [11] in light of historical precedents, especially dealing with the question of price discrimination. A conclusion about the way customers value the network is that connectivity is far more important than content.

2. PROBLEM SET-UP

Our model encompasses three types of players: the Internauts (end users), modeled collectively by their demand response, n_1 last-mile ISPs, and n_2 CPs. Consumers pay providers usage-dependent fees for service and 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 Common Demand Response Model

Let us denote by $p_{1i} \geq 0$ (resp. $p_{2j} \geq 0$) the usage-based price of the i^{th} ISP (resp. j^{th} CP). These prices act as disincentives on consumers' demand for content and 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

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

where d_k is the demand *sensitivity* to price paid to provider of type k (the first subscript $k = 1$ for ISP and $k = 2$ for CP). 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 demand metric is very convenient for our purposes here, 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. Equal demand sensitivies to price ensue, *i.e.*, $d_1 = d_2 = d$. Since $D \geq 0$, define the maximum price

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

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 n_1 ISPs (and similarly all n_2 CPs) have roughly the same prices:

$$\begin{array}{ccccccc} p_{11} & \approx & p_{12} & \approx & \cdots & \approx & p_{1n_1}, \\ p_{21} & \approx & p_{22} & \approx & \cdots & \approx & p_{2n_2}. \end{array}$$

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

[3]As in [6], the outcome essentially depends on end users' price sensitivity, but here it is furthermore related to CP (advertising) revenues.

the cheapest providers of each type, 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} type ($k = 1$ for ISPs and 2 for CPs) as a function of $\mathbf{p}_k = (p_{k1}, \ldots, p_{kn_k})$, *i.e.*, $\sigma_{ki} := \sigma(i, \mathbf{p}_k)$. Properties expected of the "stickiness function" σ include:

(a) $\sigma(i, \mathbf{p}_k) \geq 0$ and $\sum_{j=1}^{n_k} \sigma(j, \mathbf{p}_k) = 1$;

(b) if $\mathbf{p}_k = (p, p, \ldots, p)$, then $\forall i \in \{1, ..., n_k\}$, $\sigma(i, \mathbf{p}_k) = 1/n_k$; and

(c) $p_{ki} < p_{kj} \Rightarrow \sigma(i, \mathbf{p}_k) > \sigma(j, \mathbf{p}_k)$.

In other words, the distribution is uniform if all providers of a given type charge exactly the same price, and ensures otherwise that cheaper providers attract more consumers. We chose the following model which satisfies these properties:

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

The average usage-based price charged by a provider of type k for a customer is then $\overline{p}_k := \sum_i \sigma_{ki} p_{ki}$, *i.e.*, the harmonic mean of $\{p_{ki}\}_i$.

2.3 Non-discriminating setting

In a "neutral" setting with no side payments nor application discrimination, the i^{th} ISP's expected usage-based revenue is given by

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

and similarly U_{2j} for the j^{th} CP. Necessary conditions for an interior Nash Equilibrium Point (NEP) are given by

$$\frac{\partial U_{ki}}{\partial p_{ki}}(\overline{p}_1, \overline{p}_2) \;=\; 0 \;\text{ for } k = 1, 2,$$

i.e., a local maximum in revenue for all players. Solving these equations, we get that total demand and individual player revenues at a NEP are given by

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

$$U^*_{ki} \;=\; \frac{n_{3-k}^2}{(n_1 n_2 + n_1 + n_2)^2} U_{\mathsf{max}} \;\text{ for } k = 1, 2.$$

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 of one type benefit from increased competition among those of the other, while their revenues are significantly reduced by increased competition in their own group. Note that competition in a provider's own group has much greater impact on their income than competition in the other.

3. SIDE PAYMENTS

Suppose now that there are side payments between the two types 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:

$$U_{1i} \;=\; \sigma_{1i} D(p_{1i}, \overline{p}_2)(p_{1i} + p_s), \quad i \in \{1, ..., n_1\},$$
$$U_{2j} \;=\; \sigma_{2j} D(\overline{p}_1, p_{2j})(p_{2j} - p_s), \quad j \in \{1, ..., n_2\},$$

where all demand and price factors are non-negative.

It is expected that p_s is *not* a decision variable for any player or group of players. Indeed, since revenues 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 of revenues) yield:

$$\left[\frac{\overline{p}_1}{p_{\mathsf{max}} - \overline{p}_1 - \overline{p}_2} - \frac{1}{n_1}\right](\overline{p}_1 + p_s) + p_s \;=\; 0, \tag{2}$$
$$\left[\frac{\overline{p}_2}{p_{\mathsf{max}} - \overline{p}_1 - \overline{p}_2} - \frac{1}{n_2}\right](\overline{p}_2 - p_s) - p_s \;=\; 0. \tag{3}$$

With the introduction of non-null p_s, this system is now not linear. We provide here a complete study of the case where $n_1 = n_2 = 2$. See [2] for extensions to arbitrary numbers of ISPs and CPs.

3.1 Interior equilibria

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

$$-2sv - 2u^2 - v^2 + u \;=\; 0, \tag{4}$$
$$-3uv - 2s + v \;=\; 0, \tag{5}$$

where (4) $\propto$ (2) + (3) and (5) $\propto$ (2) $-$ (3). Equilibrium prices, demand and revenues are now solvable in closed form[4]. An important observation we can make at this point is given by the following theorem (proved in [2]):

[4]These expressions are complicated and of no extra-computational interest.

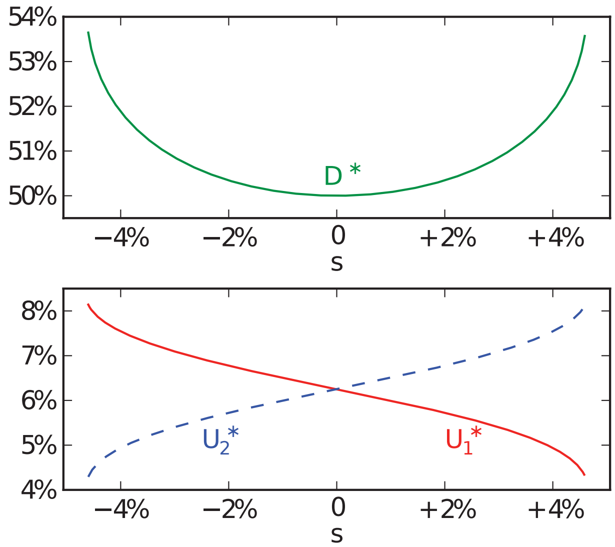

Figure 1: Demand and revenues at NEP_1.

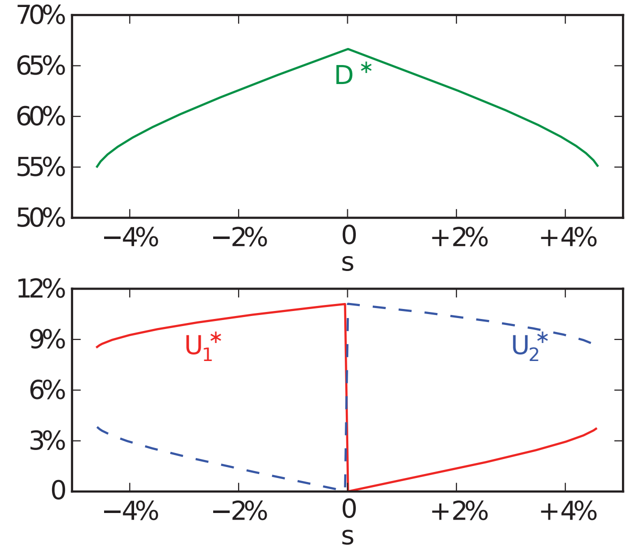

Figure 2: Demand and revenues at NEP_2.

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

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

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

There are two solutions to (2) and (3), both interior Nash equilibria. Generally, additional NEPs may exist on the boundary of the play-action space. 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 $s = 0$, $\overline{p}_k^* = p_{\mathsf{max}}/4$, $D^* = D_{\mathsf{max}}/2$ and $U_k^{i*} = U_{\mathsf{max}}/16$ for $k = 1, 2$), while NEP_2 does not exist when $s = 0$ (there is a discontinuity in equilibrium prices at this point). Both interior NEPs share the same "paradox": providers receiving side payments eventually achieve *less* revenue than the others.

3.2 Convergence to equilibrium

Here we take $s > 0$ (the roles of ISPs and CPs are swapped for $s < 0$). Assume all providers act independently under a best-response behavior. Thus, the vector field

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

is an appropriate indicator of the aggregate "trends" of the system, see Figure 3. So, if $\overline{p}_1 > \overline{p}_1^*(\mathsf{NEP}_2)$, the system is attracted by NEP_1; otherwise, unless $\overline{p}_1$ is precisely equal to $\overline{p}_1^*(\mathsf{NEP}_2)$, the system is attracted to the boundary NEP_B (where usage-based revenues for ISPs come only from side payments), *i.e.*, NEP_2 is an unstable (saddle) point.

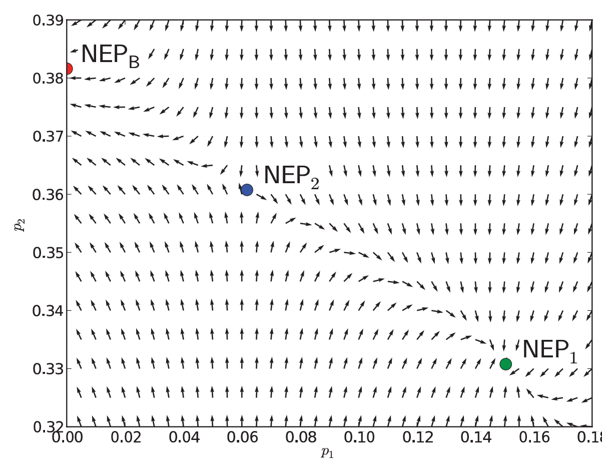

Figure 3: "Macroscopic" trends of the system for $n_1 = n_2 = 2$ and $s = 4\%$.

Solving (3) with $\overline{p}_1 = 0$ yields

$$\overline{p}_2^*(\mathsf{NEP}_\mathsf{B}) = \frac{p_{\mathsf{max}}}{6}\left(1 + s + \sqrt{s^2 + 14s + 1}\right),$$

where corresponding expressions for demand and revenues follow directly. At the boundary NEP_B, demand is higher than at NEP_1 or NEP_2 while ISP revenues turn out to be lower (and CP revenues higher) than at NEP_2 (see Figure 4).

4. APPLICATION NEUTRALITY

Now, let us consider to what extent ISPs should be allowed to perform price discrimination depending on the

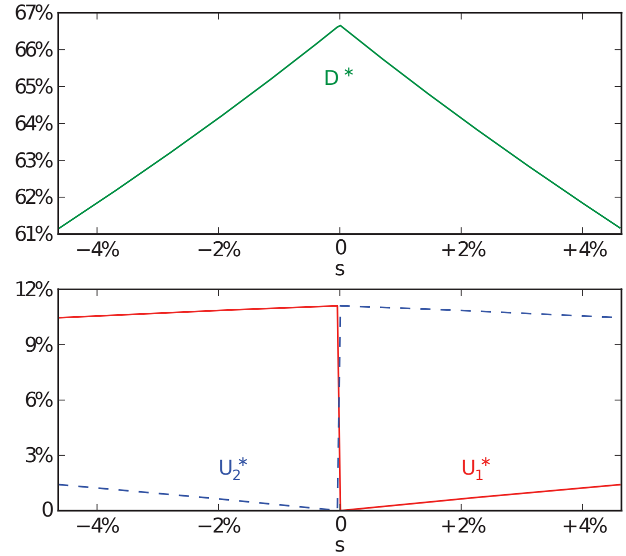

Figure 4: Demand and revenues at $\mathsf{NEP_B}$.

application in use (*e.g.*, video chat, media streaming)? In this 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 types of providers: n_1 ISPs, providing last-mile access to the Internauts; n_2 Web Content Providers (Web CPs), *e.g.*, search engine portals (recall all providers of any given type are deemed identical, so we assume all Web CPs provide the same type of client-server HTTP content as well); and n_3 providers that enable content dissemination by P2P means (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). These groups are not coalitions: in a group, each provider acts independently 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 it 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 charged by the j^{th} Web CP (resp. P2P CP). We introduce two separate demand-response profiles for the two types of content: when ISP i, Web CP j and P2P CP l are chosen, demands for HTTP and P2P content are, respectively,

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

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

The portion of users committed to the i^{th} provider of the k^{th} group is still modeled as (1); 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,

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

and make the following assumptions:

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

4.2 Neutral setting

Here consider the setting of non-monopolistic providers (*i.e.*, $n_k > 1$ for $k \in \{1, 2, 3\}$) where application neutrality is enforced. In particular, $U_{1i} = \sigma_{1i}(D_2 + D_3)p_{1i}$. The necessary conditions for an interior NEP are:

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

whose resolution is straightforward. In [2] we show that any solution to this system is a NEP and also consider the monopolistic case.

4.3 Non-neutral setting

When application non-neutral pricing is allowed, the i^{th} ISPs utility is $U_{1i} = \sigma_{1i}(D_2p_{12,i}+D_3p_{13,i})$, 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 different 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

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

where $\widetilde{p}_{1i} := \sqrt{\alpha\gamma}\,p_{12,i} + (1-\sqrt{\alpha\gamma})\,p_{13,i}$. That is, we apply the original stickiness model (1) to a combined price $\widetilde{p}_{1i}$ defined as a convex combination of $p_{12,i}$ and $p_{13,i}$. This choice, particularly the geometric mean $\sqrt{\alpha\gamma}$ in (7), is motivated by the following considerations: σ_{1i}

still satisfies the properties expected for a stickiness function (see section 2); the weight of $p_{12,i}$ in the combination is increasing in $p_{2\mathsf{max}}$ and d_2, and similarly the weight of $p_{13,i}$ is increasing in $p_{3\mathsf{max}}$ and d_3; and the resulting model is solvable in closed form.

The necessary conditions for an interior NEP are:

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

where $\widetilde{p}_1 := \sqrt{\alpha\gamma}\,\overline{p}_{12} + (1-\sqrt{\alpha\gamma})\,\overline{p}_{13}$ and $\widetilde{D}_k := D_k/d_k = p_{k\mathsf{max}} - \overline{p}_1 - \overline{p}_k$ for $k = 2, 3$. This system can be rewritten as two polynomial equations in $\overline{p}_{12}$ and $\overline{p}_{13}$ which are solvable in closed form. Computations yield a single admissible solution here, which is also a NEP.

4.4 Discussion of experimental results

In our numerical experiments, we compared revenues at this NEP with those of the neutral scenario for $\alpha = 0.8$ and $\gamma = 0.3$. We used Sage for our computations, and all our scripts are available online[5]. Again, additional results are provided in [2].

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 benefit for Web CPs and less loss for P2P CPs. Yet, competition has almost no effect on the gains of ISPs (see Figure 5).

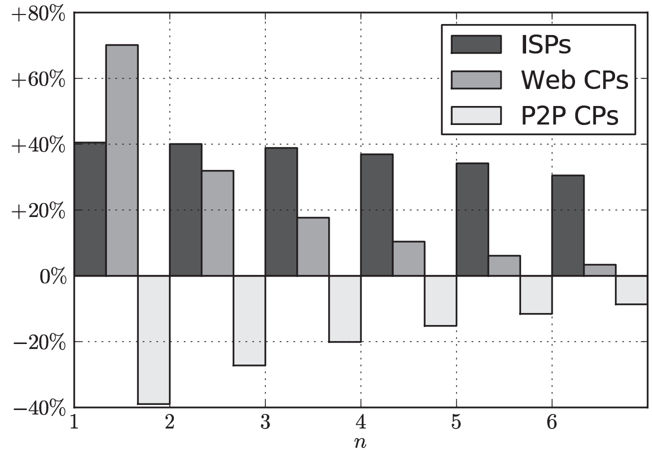

Figure 5: Relative variation in revenue, *i.e.*, the ratio of revenues (non-neutral - neutral)/neutral at the NEP, where n is the number of providers of each type, $\alpha = 0.8$ and $\gamma = 0.3$.

5. CONCLUSIONS

We presented an idealized framework to study the impact of two net-neutrality related issues, side payments and application neutrality, on the interactions among end users, ISPs and CPs. Our revenue model relied on a simple, common linear demand response to usage-based prices, and it accounted for customer loyalty.

We studied the effect of regulated side payments between the ISPs and CPs. The two possible outcomes of the competition both showed the same paradox: side payments are actually a handicap for those who receive them insofar as they reduce Nash equilibrium revenues.

We also studied the issue of application neutrality in a simple setting involving two types of content, web content and file sharing, the latter showing lower price sensitivity and higher willingness to pay under our assumptions of relative demand sensitivity to price. Our analysis suggested that ISPs and Web CPs benefit from application non-neutral practices, while providers that enable content dissemination by P2P means are better off in a neutral setting.

[5] `http://www.cse.psu.edu/~kesidis/neutrality/index.html`

6. REFERENCES

[1] 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. Workshop on Economic Traffic Management (ETM)*, Amsterdam, Sept. 6, 2010. See also: `http://arxiv.org/abs/1006.3894`

[2] E. Altman, S. Caron and G. Kesidis, "Application Neutrality and a Paradox of Side Payments", Technical Report, Aug. 2010. `http://arxiv.org/abs/1008.2267`

[3] R.B. Chong, "The 31 Flavors of Net Neutrality", *12 Intellectual Property Law Bulletin*, v. 12, 2008.

[4] Comcast v. FCC, 600 F.3d 642 (D.C. Cir. 2010).

[5] R. Hahn and S. Wallsten, "The Economics of Net Neutrality", *Economists' Voice*, The Berkeley Economic Press, **3**(6), pp. 1-7, 2006.

[6] P. Hande, M. Chiang, R. Calderbank, and S. Rangan, "Network pricing and rate allocation with content provider participation", *Proc. IEEE INFOCOM*, 2009.

[7] G. Kesidis, A. Das, and G. de Veciana, "On Flat-Rate and Usage-based Pricing for Tiered Commodity Internet Services", *Proc. CISS*, Princeton, 2008.

[8] G. Kesidis, "Congestion control alternatives for residential broadband access by CMTS, *Proc. IEEE/IFIP NOMS*, Osaka, Apr. 2010.

[9] 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.

[10] 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.

[11] A. Odlyzko, "Network Neutrality, Search Neutrality, and the Never-ending Conflict between Efficiency and Fairness in Markets", *Review of Network Economics* **8**(1), 2009.