A model of network neutrality with usage-based prices

**E. Altman · P. Bernhard · S. Caron · G. Kesidis ·
J. Rojas-Mora · S. Wong**

Abstract Hahn and Wallsten (Econ. Voice 3(6):1–7, 2006)
wrote that network neutrality "usually means that broad-
band service providers charge consumers only once for In-
ternet access, do not favor one content provider over an-
other, and do not charge content providers for sending in-
formation over broadband lines to end users." In this paper
we study the implications of non-neutral behaviors under
a simple model of linear demand-response to *usage-based*
prices. We take into account advertising revenues for the
content provider and consider both cooperative and non-
cooperative scenarios. In particular, we model the: impact of
side-payments between service and content providers, con-
sider an access provider that offers multiple service classes,
and model leader-follower (Stackelberg game) dynamics.

E. Altman · P. Bernhard
INRIA, 2004 Route des Lucioles, 06902 Sophia-Antipolis, France

E. Altman
e-mail: eitan.altman@inria.fr

P. Bernhard
e-mail: pierre.bernhard@inria.fr

S. Caron · G. Kesidis (✉)
CS&E and EE Depts, Pennsylvania State Univ., University Park,
PA, 16802, USA
e-mail: kesidis@engr.psu.edu

J. Rojas-Mora
Fac. of Econ. and Bus. Sci., Univ. of Barcelona, 08034 Barcelona,
Spain
e-mail: jrojasmo7@alumnes.ub.edu

S. Wong
Fac. of Law, Univ. of Coruña, 15071 A Coruña, Spain
e-mail: swong@udc.es

We finally study the additional option for one provider to de-
termine the amount of side payment from the other provider.
We show that not only do the content provider and the in-
ternaut suffer, but also the Access Provider's performance
degrades.

Keywords Network neutrality · Pricing

1 Introduction

Network neutrality is an approach to providing network ac-
cess without unfair discrimination among applications, con-
tent or traffic sources. Discrimination occurs when there are
two applications, services or content providers that require
the same network resources, but one is offered better quality
of service (shorter delays, higher transmission capacity, etc.)
than the other. How to define what is "fair" discrimination
is still subject to controversy, and the underlying economic
and policy issues have not been fully debated.[1] A preferen-
tial treatment of traffic is considered fair as long as the pref-

---

[1]The recent decision on Comcast v. the FCC was expected to deal with
the subject of "fair" traffic discrimination, as the FCC ordered Com-
cast to stop interfering with subscribers' traffic generated by peer-to-
peer networking applications. The Court of Appeals for the District of
Columbia Circuit was asked to review this order by Comcast, arguing
not only on the necessity of managing scarce network resources, but
also on the non-existent jurisdiction of the FCC over network manage-
ment practices. The Court decided that the FCC did not have express
statutory authority over the subject, neither demonstrated that its ac-
tion was "reasonably ancillary to the [. . .] effective performance of its
statutorily mandated responsibilities". The FCC was deemed, then, un-
able to sanction discriminatory practices on Internet's traffic carried out
by American ISPs, and the underlying case on the "fairness" of their
discriminatory practices was not even discussed.

erence is left to the user.[2] Internet Service Providers (ISPs) may have interest in traffic discrimination either for technological or economic purposes. Traffic congestion, especially due to high-volume peer-to-peer traffic, has been a central argument for ISPs against the enforcement of net neutrality principles. However, it seems many ISPs have blocked or throttled such traffic even during periods without congestion.

ISPs recently claimed that net neutrality acts as a disincentive for capacity expansion of their networks. In [3], the authors studied the validity of this argument and came to the conclusion that, under net neutrality, ISPs invest to reach a social optimal level, while they tend to under/over-invest when neutrality is dropped. In their setting, ISPs stand as winners while content providers (CPs) are left in a worse position, and users who pay the ISPs for preferential treatment are better off while other consumers have a significantly worse service.

ISPs often justify charging CPs by quantifying the large amount of network resources "big" content providers use. On the other hand, the content and services the CPs offer contribute to the demand for Internet access, and thus benefit the ISPs. Shapley values may be used to obtain fair and Pareto optimal revenue-sharing between different types of players, e.g., [10, 11]. That is, Shapley values can prescribe whether and how (i) CPs should share third-party advertising revenue enabled by subscribers' network access [6], or (ii) ISPs should share subscription revenue enabled by their customers' demand for online content and services.

In this paper, we focus on violations of the neutrality principles defined in [7] where broadband service providers

– charge consumers more than "only once" through usage-based pricing, and
– charge content providers through side-payments.

Within a simple game-theoretic model, we examine how regulated[3] side payments, in either direction, and demand-dependent advertising revenues affect equilibrium usage-based prices. We also address equilibria in Stackelberg leader-follower dynamics.

The rest of the paper is organized as follows. In Sect. 2, we describe a basic model and derive Nash equilibria for competitive scenarios and optimal collaborative scenarios. We consider potentially non-neutral side-payments in Sect. 3 and advertising revenues in Sect. 4, analyzing in each case how they impact equilibrium revenues. We consider an ISP offering multiple service classes in Sect. 5. In Sect. 6, we study leader-follower dynamics. In Sect. 7 we consider the case where one provider can control the amount of side payments from the other one. We show that this results in a dramatic degradation of performance for all actors: at equilibrium, the demand is shown to be zero. We conclude in Sect. 8 and discuss future work.

2 Basic model

Our model encompasses three actors:

– the internauts (users), collectively,
– a network access provider for the internauts, collectively called ISP1, and
– a content provider and its ISP, collectively called CP2.

The two providers play a game to settle on their (usage-based) prices. The internauts are modeled through their demand response.

Consumers are assumed willing to pay a usage-based fee (which can be \$0/byte) for service/content that requires both providers.

Denote by $p_i \geq 0$ the usage-based price leveed by provider i (ISP1 being $i = 1$ and CP2 being $i = 2$). We assume that the demand-response of customers, which corresponds to the amount (in bytes) of content/bandwidth they are ready to consume given prices p_1 and p_2, follows a simple linear model (e.g., [5]):

$$D = D_0 - d(p_1 + p_2). \tag{1}$$

With such a profile, we are dealing with a set of homogeneous users sharing the same response coefficient d to price variations. Infrastructure and operating costs borne by the ISPs and CPs [12] are not considered here. The parameter D_0 corresponds to the demand under pure *flat-rate* pricing ($p_1 = 0 = p_2$).

Demand should be non-negative, i.e.,

$$p_1 + p_2 \leq \frac{D_0}{d} =: p_{\max}.$$

Provider i's usage-based revenue is given by

$$U_i = Dp_i. \tag{2}$$

[2]Nonetheless, users are just one of many actors in the net neutrality debate, which has been enlivened throughout the world by several public consultations for new legislations on the subject. The first one, proposed in the USA, was looking for the best means of preserving a free and open Internet [15]. The second one, carried out in France, asks for different points of view over net neutrality [4]. A third one was presented by the EU during summer 2010, looking for a balance on the parties concerned as users are entitled to access the services they want, while ISPs and CPs should have the right incentives and opportunities to keep investing, competing and innovating [13].

[3]In the European Union, dominating positions in telecommunications markets (such as an ISP imposing side-payments to CPs at a price of his choice) are controlled by the article 14, paragraph 3 of the Directive 2009/140/EC, considering the application of remedies to prevent the leverage of a large market power over a secondary market closely related.

2.1 Non-cooperation

Suppose the providers do not cooperate. An interior Nash Equilibrium Point (NEP) (p_1^*, p_2^*) of this two-player can be obtained from the first order optimality conditions

$$\frac{\partial U_i}{\partial p_i}(p_1^*, p_2^*) = 0, \quad \text{for } i = 1, 2,$$

which lead to $p_1^* = p_2^* = D_0/(3d)$. The demand at equilibrium is thus $D^* = D_0/3$ and the revenue of each provider is

$$U_i^* = \frac{D_0^2}{9d}. \tag{3}$$

2.2 Collaboration

Now suppose there is a coalition between ISP1 and CP2. Their overall utility is then $U_{\text{total}} := U_1 + U_2 = Dp$, and an optimal point (p_1^*, p_2^*) satisfies

$$\frac{\partial U_{\text{total}}}{\partial p_i}(p_1^*, p_2^*) = D^* - d(p_1^* + p_2^*) = 0 \quad \text{for } i = 1, 2,$$

which yields $p^* := p_1^* + p_2^* = D_0/(2d)$. The demand at equilibrium is then $D^* = D_0/2$, greater than in the non-cooperative setting. The overall utility $U_{\text{total}}^* = D_0^2/(4d)$ is also greater than $D_0^2/(4.5d)$ for the competitive case. Assuming both players share this revenue equally (trivially, the Shapley values are $\{1/2, 1/2\}$ in this case), the utility per player becomes

$$U_i^* = \frac{D_0^2}{8d}, \tag{4}$$

which is greater than in the competitive case. So, both players benefit from this coalition.

3 Side-payments under non-cooperation

Let us suppose now that there are *side payments* between ISP1 and CP2 at (usage-based) price p_s. The revenues of the providers become:

$$U_1 = D(p_1 + p_s), \tag{5}$$

$$U_2 = D(p_2 - p_s). \tag{6}$$

Note that p_s can be positive (ISP1 charges CP2 for "transit" costs) or negative (CP2 charges ISP1, e.g., for copyright remuneration).[4] It is expected that p_s is *not* a decision variable

---

[4]In France, a new law has been proposed recently to allow download of unauthorized copyright content, and in return be charged *proportionally* to the volume of the download.

of the players, since their utilities are monotonic in p_s and the player without control would likely set (usage-priced) demand to zero to avoid negative utility. That is, p_s would normally be *regulated* and we will consider it as a fixed parameter in the following (with $|p_s| \leq p_{\max}$).

First, if $|p_s| \leq \frac{1}{3}p_{\max}$, the equilibrium prices are given by

$$p_1^* = \frac{1}{3}p_{\max} - p_s,$$

$$p_2^* = \frac{1}{3}p_{\max} + p_s$$

but demand $D^* = D_0/3$ and utilities

$$U_i^* = \frac{D_0^2}{9d}$$

are exactly the same as (3) in the competitive setting with no side payment. Therefore, though setting $p_s > 0$ at first seems to favor ISP1 over CP2, it turns out to have no effect on equilibrium revenues for both providers.

Alternatively, if $p_s \geq \frac{1}{3}p_{\max}$, a boundary Nash equilibrium is reached when $p_1^* = 0$ and $p_2^* = \frac{1}{2}(p_{\max} + p_s)$, which means ISP1 does not charge usage-based fees to its consumers. Demand becomes $D^* = \frac{1}{2}(D_0 - dp_s)$, and utilities are

$$U_1^* = \frac{(D_0 - dp_s)dp_s}{2d},$$

$$U_2^* = \frac{(D_0 - dp_s)^2}{4d}.$$

Though $p_1^* = 0$, U_1^* is still strictly positive, with revenues for ISP1 coming from side-payments (and possibly from flat-rate monthly fees as well). Furthermore, $p_s \geq \frac{1}{3}p_{\max} \Leftrightarrow dp_s \geq \frac{1}{2}(D_0 - dp_s)$, which means $U_1^* \geq U_2^*$: in this setting, ISP1's best move is to set his usage-based price to zero (to increase demand), while he is sure to achieve better revenue than CP2 through side-payments.

Finally, if $p_s < -\frac{1}{3}p_{\max}$, the situation is similar to the previous case (with $-p_s$ instead of p_s). So, here $p_2^* = 0$ and $p_1^* = \frac{1}{2}(p_{\max} - p_s)$, leading to $U_2^* \geq U_1^*$.

To remind, herein revenues U_i are assumed usage-based, which means there could also be flat-rate charges in play to generate revenue for either party. Studies of flat-rate compare to usage-based pricing schemes can be found in the literature, see, e.g., [8].

4 Advertising revenues

We suppose now that CP2 has an additional source of (usage-based) revenue from advertising that amounts to

Dp_a. Here p_a is not a decision variable but a fixed parameter.[5]

4.1 Non-cooperation

The utilities for ISP1 and CP2 are now:

$$U_1 = [D_0 - d \cdot (p_1 + p_2)](p_1 + p_s), \tag{7}$$

$$U_2 = [D_0 - d \cdot (p_1 + p_2)](p_2 - p_s + p_a). \tag{8}$$

Here, the Nash equilibrium prices are:

$$p_1^* = \frac{1}{3} p_{\max} - p_s + \frac{1}{3} p_a,$$

$$p_2^* = \frac{1}{3} p_{\max} + p_s - \frac{2}{3} p_a.$$

The cost to users is thus $p^* = \frac{2}{3} p_{\max} - \frac{1}{3} p_a$ while demand is $D^* = \frac{1}{3}(D_0 + dp_a)$. Nash equilibrium utilities are given by

$$U_i^* = \frac{(D_0 + dp_a)^2}{9d} \quad \text{for } i = 1, 2, \tag{9}$$

which generalizes (3) and shows how advertising revenue quadratically raises players' utilities.

4.2 Collaboration

The overall income for cooperating providers is

$$U_{\text{total}} = (D_0 - dp)(p + p_a). \tag{10}$$

So, solving the associated first-order optimality equation yields

$$p^* = \frac{p_{\max} - p_a}{2}. \tag{11}$$

The optimal demand is then $D^* = (D_0 + dp_a)/2$, and the maximal total revenue is $U_{\text{total}}^* = (D_0 + dp_a)^2/(4d)$. Assuming this revenue is split equally between the two providers, we get for each provider the equilibrium utility

$$U_i^* = \frac{(D_0 + dp_a)^2}{8d}, \tag{12}$$

which generalizes (4). As before, providers and users are better off when they cooperate.

Thus, we see that $p_a > 0$ leads to lower prices, increased demand and more revenue for *both* providers (i.e., including ISP1).

5 ISP providing multiple service classes

In this section, we suppose ISP1 is offering two types of network access service: a low-quality one l at price p_l, and a high-quality one h at price $p_h \geq p_l$. The role of multiple service classes in a neutral network, i.e., as selected by the users, has previously been explored, e.g., [9].[6] Here, we split the demand D into D_l and D_h: $D = D_l + D_h$ (we will describe later how we implement the dichotomy between D_l and D_h). For now, assume the overall demand still has a linear response profile, i.e.,

$$D = D_0 - d(\underbrace{p_l + p_h}_{\text{formerly } p_1} + p_2). \tag{13}$$

First, we make reasonable assumptions on D_l:

1. *Pricing incentives:* Define $\Delta p := p_h - p_l$. Δp is an incentive for consumers to chose between classes l and h: the higher Δp is, the more likely users are to select l. Thus, if we take $x := 1/\Delta p$ and $y := D_l/D$, we may see y as a function of x and model this pricing response with the following properties:

$$y'(x) \leq 0 \quad (D_l \text{ increases with } \Delta p), \tag{14}$$

$$y(0) = 1 \quad (D_l \uparrow D \text{ as } \Delta p \uparrow \infty), \tag{15}$$

$$y(\infty) = 0 \quad (D_l \downarrow 0 \text{ as } \Delta p \downarrow 0). \tag{16}$$

2. *Congestion incentives:* As D_l approaches D, we assume congestion occurs in the low-quality network, further deterring users to chose it. This motivates the additional assumption that

$$|y'(x)| \downarrow 0 \quad \text{as } x \downarrow 0, \tag{17}$$

that is, D_l *decelerates* as it gets closer to D.

Define

$$\delta := \frac{\Delta p}{\gamma p_{\max}}, \tag{18}$$

where $\gamma > 0$ is an additional users' price-sensitivity parameter. The following demand relation satisfies all conditions (14), (15), (16) and (17):

$$D_l := (1 - e^{-\delta})D. \tag{19}$$

The providers' utilities are then:

$$U_1 = D_l p_l + D_h p_h = D(p_l + \Delta p e^{-\delta}), \tag{20}$$

$$U_2 = Dp_2. \tag{21}$$

---

[5]One may see p_a as the result of an independent game between CP2 and his advertising sources, the details of which are out of the scope of this paper.

[6]Non-neutral class assignment to applications, i.e., not application-neutral networking, is discussed in [14].

Fig. 1 Attraction of the equilibrium line

5.1 Collaboration

If both players cooperate, their overall utility is

$$U_{\text{total}} = D(p_2 + p_l + \Delta p e^{-\delta}).$$

There is no NEP with strictly positive prices $p_i \geq 0$ for $i = 1, 2$. To specify the boundary NEP (where at least one usage-based price is zero), define

$$\phi(x) := (1 - x)e^{-x}$$

and note that ϕ is a bijection of $[0, 1]$.

– If $p_2 = 0$, NEP conditions imply

$$\delta^* = \phi^{-1}(1/2),$$

$$p_l^* = \frac{1}{3}\left(\frac{1}{2} - \gamma \delta e^{-\delta}\right)p_{\max}.$$

Utility at the NEP is therefore

$$U_{\text{total}}^* = \frac{D_0^2}{9d}\left[\frac{1}{2} + 2\gamma \delta e^{-\delta}\right] \cdot [2 + (2e^{-\delta} - 3)\delta\gamma]. \quad (22)$$

In this setting, the value of U_{total} is upper bounded by $\approx 0.162 \frac{D_0^2}{d}$ which is achieved when $\gamma \approx 1.53$ (recall that γ is not a decision variable).

• If $p_l = 0$, then $p_h = 0$ and $p_2 = \frac{1}{2}p_{\max}$, yielding

$$U_{\text{total}} = \frac{D_0^2}{4d}. \quad (23)$$

Hence, irrespective of consumers' sensitivity γ to the price gap Δp, the best solution for the coalition is to set-up usage-based pricing for content only, at price $p_2 = p_{\max}/2$, while network access is subject to flat-rate pricing ($p_l = p_h = 0$).

5.2 Splitting demand-response coefficient

Now consider splitting the demand-response coefficient d into d_l, d_h and d_2, that is:

$$D = D_0 - d_l p_l - d_h p_h - d_2 p_2. \quad (24)$$

If

$$d_2 = d_l + d_h, \quad (25)$$

then the interior equilibrium conditions $\nabla U_{\text{total}} = \underline{0}$ yield:

$$\delta = \phi^{-1}(d_h/d_2),$$

$$p_l + p_2 = \frac{D_0}{2d_2} - \frac{\delta \Delta p_0}{2}\left(\frac{d_h}{d_2} + e^{-\delta}\right).$$

When the demand-response coefficients satisfy (25), we have an *equilibrium line*. Vector field plots of U_{total} suggest it is attractive (see Fig. 1). In this particular setting, providers can thus reach U_{total}^* with non-flat rate pricing.

However, if $d_2 \neq d_l + d_h$, there exists a line of attraction, but with a non-null gradient on it driving players toward border equilibria. Hence, the conclusion of Sect. 5.1 also holds in this more generalized setting.

Fig. 2 Utilities as functions of users' sensitivity to usage-based pricing

5.3 Non-cooperation

When ISP1 and CP2 compete, again there is no interior NEP (with all prices p_i strictly positive). In fact, the condition $\nabla_{p_l, p_h} U_1 = \underline{0}$ implies $p_l = 0 = p_h$ and $D = 0$, so ISP1 has to relax condition $\frac{\partial U_1}{\partial p_l} = 0$ by setting $p_l = 0$ (i.e., *flat-rate* pricing for the best-effort service l). The solution to the two remaining Nash equilibrium conditions is then:

$$p_2 = \frac{1}{4}\left[\sqrt{9\gamma^2 + 2\gamma + 1} - 3\gamma + 1\right] \cdot p_{\max}, \tag{26}$$

$$p_h = \frac{\gamma}{2\sqrt{9\gamma^2 + 2\gamma + 1} - 3\gamma + 2} \cdot p_{\max}. \tag{27}$$

By defining $f_2(\gamma) := p_2/p_{\max}$ and $f_h(\gamma) := p_h/p_{\max}$, we then have

$$U_1^*(\gamma) = f_h(\gamma) \cdot (1 - f_h(\gamma) - f_2(\gamma)) D_0 p_{\max},$$

$$U_2^*(\gamma) = f_2(\gamma) \cdot (1 - f_h(\gamma) - f_2(\gamma)) D_0 p_{\max}.$$

Figure 2 shows utilities at equilibrium (as fractions of $D_0 p_{\max}$). We see that, in any case, CP2 has the advantage in this game: U_2^* is always greater to U_1^*, irrespective of consumers' sensitivity γ to usage-based prices.

Here, $\gamma \to 0$ means users are so sensitive to any usage-based price that they will always choose the best-effort service (which is subject to flat-rate pricing). Users' price sensitivity decreases as γ increases, with the limit $\gamma \to \infty$ corresponding to the setting of Sect. 2 with $\lim_{\gamma \to \infty} U_i^*(\gamma) = \frac{D_0^2}{9d}$.

6 Stackelberg equilibrium

Stackelberg equilibrium corresponds to asymmetric game in which one player is the leader and the other a followers.

That is, actions are no longer taken independently: the leader takes action first, and then the follower reacts.

Though the dynamics of the games are different from the previous study, equations (7) and (8) still hold, with fixed $p_a \geq 0$ and regulated p_s. In the following, we need to assume that

$$p_s \leq \frac{1}{2} p_{\max} + \frac{1}{2} p_a,$$

$$p_a \leq \frac{1}{3} p_{\max} + \frac{1}{4} p_s$$

so that NEPs are reachable with positive prices.

If ISP1 sets p_1, then CP2's optimal move is to set

$$p_2 = \frac{1}{2}(-p_1 + p_{\max} + p_s - p_a).$$

This expression yields $D = \frac{d}{2}(p_{\max} - p_1 - p_s + p_a)$ and $U_1 = \frac{d}{2}(p_{\max} - p_1 - p_s + p_a)(p_1 + p_s)$. Anticipating CP2's reaction in trying to optimize U_1, the best move for ISP1 is thus to set

$$p_1^* = \frac{1}{2} p_{\max} - p_s + \frac{1}{2} p_a,$$

which yields

$$p_2^* = \frac{1}{4} p_{\max} + p_s - \frac{3}{4} p_a.$$

Therefore, when ISP1 is the leader, at the NEP demand is $D^* = \frac{1}{4}(D_0 + dp_a)$ and utilities are:

$$U_1^* = \frac{1}{8d}(D_0 + dp_a)^2, \tag{28}$$

$$U_2^* = \frac{1}{16d}(D_0 + dp_a)^2. \tag{29}$$

Suppose now that CP2 is the leader and sets p_2 first. Similarly, we find:

$$p_2^* = \frac{1}{2} p_{\max} + p_s - \frac{1}{2} p_a,$$

$$p_1^* = \frac{1}{4} p_{\max} - p_s + \frac{1}{4} p_a.$$

These values yield the same cost p^* and demand D^* for the internauts at the NEP, while providers' utilities become:

$$U_1^* = \frac{1}{16d}(D_0 + dp_a)^2, \tag{30}$$

$$U_2^* = \frac{1}{8d}(D_0 + dp_a)^2. \tag{31}$$

Therefore, in either case of leader-follower dynamics, the leader obtains twice the utility of the follower at the NEP (yet, his revenue is not better than in the collaborative case).

7 Further abandoning neutrality

Throughout we assumed that the side payments between service and content providers were regulated. If p_s is allowed to be determined unilaterally by the service provider or by the content provider (as part of the game described in Sects. 3 or 4), then the worst possible performance is obtained at equilibrium. More precisely, the demand at equilibrium is zero, see [1]. The basic reason is that if the demand at equilibrium were not zero then a unilateral deviation of the provider that controls the side payment results in a strict improvement of its utility. (Note that the total demand is unchanged as it does not depend on p_s.)

More generally, assume that an ISP is given the authority to control p_s and that its utility can be expressed as $U = f(D) \times (g(p_s) + h)$ where f is any function of the demand (and possibly also of prices other than p_s), g is a monotone strictly increasing function of p_s, and h does not depend on p_s. Then at equilibrium, necessarily $f(D) = 0$, otherwize U can be further increased by the provider by (unilaterally) increasing p_s.

The same phenomenon holds also in case the CP is given full control over p_s.

8 Conclusions and on-going work

Using a simple model of linearly diminishing consumer demand as a function of usage-based price, we studied a game between a monopolistic ISP and a CP under a variety of scenarios including consideration of: non-neutral two-sided transit pricing (either CP2 participating in network costs or ISP1 paying for copyright remuneration), advertising revenue, cooperation and leadership.

In a basic model without side-payments and advertising revenues, both providers achieve the same utility at equilibrium, and all actors are better off when they cooperate (higher demand and providers' utility).

When regulated, usage-based side-payments p_s come into play, the outcome depends on the value of $|p_s|$ compared to the maximum usage-based price $p_{\max}$ consumers can tolerate:

– when $|p_s| \leq \frac{1}{3} p_{\max}$, providers shift their prices to fall back to the demand of the competitive setting with no side-payments;
– when $|p_s| \geq \frac{1}{3} p_{\max}$, the provider receiving side payments sets its usage-based price to zero to increase demand, while it is sure to be better off than his opponent.

When advertising revenues to the CP come into play, they increase the utilities of *both* providers by reducing the overall usage-based price applied to the users. ISP1 and CP2 still share the same utility at equilibrium, and the increase in revenue due to advertising is quadratic.

We considered in Sect. 5 the implications of service differentiation from the ISP. In our model, when ISP1 and CP2 cooperate, the best solution for them is to set-up usage-based prices for content only and flat-rate pricing for network access. However, when providers do not cooperate, the ISP optimally offers its best-effort service for a flat rate (zero usage-based cost), resulting in more usage-based revenue for the CP.

Under leader-follower dynamics, the leader obtains twice the utility of his follower at equilibrium; yet, he does not achieve a better revenue than in the cooperative scenario.

In subsequent work [2], we explored the effects of content-specific (i.e., not *application* neutral) pricing, including multiple CPs providing different types of content. Also, we considered competition among multiple providers of the same type, including different models consumer stickiness (interia or loyalty). In on-going work, we are also considering providers' infrastructure and operating costs (as in, e.g., [12]), more complex models of end-user demand and their collective social welfare, and the effects of different options for flat-rate pricing (e.g., [8, 16]).

References

1. Altman, E., Bernhard, P., Kesidis, G., Rojas-Mora, J., & Wong, S. (2010). *A study of non-neutral networks*. INRIA Research report Number 00481702, May 2010.
2. Caron, S., Kesidis, G., & Altman, E. (2010). Application neutrality and a paradox of side payments. In *Proc. ACM ReArch*, Philadelphia, Nov. 30, 2010.
3. Cheng, K., Bandyopadhyay, S., & Gon, H. (2008). The debate on net neutrality: A policy perspective. In *Information systems research*, June 2008. Available at SSRN: http://ssrn.com/abstract=959944.
4. Consultation publique sur la neutralité du net. du 9 avril au 17 mai 2010. http://www.telecom.gouv.fr/fonds_documentaire/consultations/10/consneutralitenet.pdf.
5. Economides, N. (2008). Net neutrality: non-discrimination and digital distribution of content through the Internet. *I/S: A Journal of Law and Policy*, 4(2), 209–233.
6. Hande, P., Chiang, M., Calderbank, R., & Rangan, S. (2009). Network pricing and rate allocation with content provider participation. In *Proc. IEEE INFOCOM*.
7. Hahn, R., & Wallsten, S. (2006). The economics of net neutrality. *Economists' Voice*, 3(6), 1–7.
8. Kesidis, G., Das, A., & de Veciana, G. (2008). On flat-rate and usage-based pricing for tiered commodity Internet services. In *Proc. CISS*, Princeton, March 2008.
9. Kesidis, G. (2010). Congestion control alternatives for residential broadband access by CMTS. In *Proc. IEEE/IFIP NOMS*, Ōsaka, Japan, Apr. 2010.
10. Ma, R. T. B., Chiu, D.-M., Lui, J. C. S., Misra, V., & Rubenstein, D. (2008). Interconnecting eyeballs to content: A Shapley value

perspective on ISP peering and settlement. In *Proc. int'l workshop on economics of networked systems (NetEcon)* (pp. 61–66).

11. Ma, R. T. B., Chiu, D.-M., Lui, J. C. S., Misra, V., & Rubenstein, D. (2008). On cooperative settlement between content, transit and eyeball Internet service providers. In *Proc. ACM CoNEXT*.

12. Musacchio, J., Schwartz, G., & Walrand, J. (2009). A two-sided market analysis of provider investment incentives with an application to the net-neutrality issue. *Review of Network Economics*, *8*(1).

13. Neelie Kroes, vice-president of the European Commission, Commissioner for the digital agenda, *Press release RAPID*, April 2010.

14. Network neutrality. Wikipedia. Available at http://en.wikipedia.org/wiki/Network_neutrality.

15. Notice of proposed rulemaking (FCC 09/93). October 2009–April 2010.

16. Odlyzko, A. (2001). Internet pricing and history of communications. *Computer Networks*, *36*(5–6), 493–518.

E. Altman received the B.Sc. degree in electrical engineering (1984), the B.A. degree in physics (1984) and the Ph.D. degree in electrical engineering (1990), all from the Technion-Israel Institute, Haifa. In (1990) he further received his B.Mus. degree in music composition in from Tel-Aviv university. Since 1990, he has been with INRIA (National research institute in informatics and control) in Sophia-Antipolis, France. His current research interests include performance evaluation and control of telecommunication networks and in particular congestion control, wireless communications and networking games. He is on the editorial board of the scientific journals: WINET, JDEDs and JEDC, and served in the journals Stochastic Models, COMNET, SIAM SICON. He has been the general chairman and the (co)chairman of the program committee of several international conferences and workshops (on game theory, networking games and mobile networks). More information can be found at www-sop.inria.fr/members/Eitan.Altman

P. Bernhard was born in 1944. After graduating from Ecole Polytechnique and Ecole des Mines in Paris, he earned a Ph.D. in aeronautical and astronautical engineering at Stanford University, in 1970, under the supervision of John V. Breakwell. He has since been alternately a full time researcher and a university professor, teaching control theory and applied mathematics. He was the director of several French research centers in computer science and Control, including INRIA-Sophia Antipolis of which he was the founder. He is now retired and emeritus senior scientist at INRIA Sophia Antipolis. His main theme of research has been in dynamic games and their applications, such as robust control (he is the co-author, with T. Basar, of a rather well known monograph on H-infinity optimal control) and finance. More recently, he has been interested in evolutionary games and ESS theory. He received the 2008 Isaacs award from the International Society of Dynamic Games.

S. Caron is currently a graduate student in Computer Science at Ecole Normale Superieure (Paris). In 2010, he was a graduate student intern in the CSE Dept at Penn State.

G. Kesidis received his M.S. and Ph.D. in EECS from U.C. Berkeley in 1990 and 1992 respectively. He was a professor in the E&CE Dept of the University of Waterloo, Canada, from 1992 to 2000. Since 2000, he has been a professor of CSE and EE at the Pennsylvania State University. His research, including several areas of computer/communication networking and machine learning, has been primarily supported by NSERC of Canada, NSF and Cisco Systems URP. He served as the TPC co-chair of IEEE INFOCOM 2007 among other networking conferences. He has also served on the editorial boards of the Computer Networks (COMNET) Journal, ACM TOMACS and, currently, IEEE Journal on Communications Surveys and Tutorials. His home page is http://www.cse.psu.edu/~kesidis.

J. Rojas-Mora has a systems engineering degree from the University of Los Andes (Venezuela). His academic career includes an Advanced Studies Diploma (D.E.A.) in Statistics and Operations Research from the University of Santiago de Compostela (Spain) in 2006, and a Ph.D. from the Business Economics and Organization Department of the University of Barcelona (Spain) in 2011. He was a lecturer in networking at the University Fermin Toro (Venezuela), where he also was its networking coordinator. He has been a research engineer at the University of Avignon's LIA and a visiting researcher at INRIA's MAESTRO Group (both in France). Currently, he is a postdoctoral fellow at the UMR CNRS Espace of the University of Avignon. His research includes the clash of information technology's developments with anachronistic law and business models.

S. Wong is a lawyer who graduated "Cum Laude" from the University of Los Andes (Venezuela). Before moving into the academic world, Mrs Wong held various positions in the Prosecutor's Office of the Bolivarian Republic of Venezuela. In 2006, she got her Advanced Studies Diploma (D.E.A.) in Public Freedoms and Fundamental Rights at the University of A Coruña (Spain), where she is pursuing her Ph.D. She was a visiting researcher at the Department of Law Philosophy of the University of Barcelona (Spain), and at the MAESTRO group of INRIA (France). Mrs Wong's research is based on the interaction between fundamental rights and freedoms, intellectual property rights, and information technology (IT).