Cognitive Bias in Network Services

Rade Stanojevic
Telefonica Research
Barcelona, Spain

Vijay Erramilli
Telefonica Research
Barcelona, Spain

Konstantina Papagiannaki
Telefonica Research
Barcelona, Spain

ABSTRACT

The assumption of rationality is fundamental to large part of network economics literature. In this paper, we use a simple definition of rationality based on economic self-interest and test for such behavior using real data on how users purchase and consume mobile network services. If users acted in their best (optimal) interest, then they would opt for the tariff that best suits their demands. However, that need not be the case, as users can fall prey to biases that can lead them to make *seemingly* sub-optimal choices. Such biases are hard to characterize and in this paper we empirically study how end-users purchase *and* use network services.

We find that most customers choose sub-optimal tariffs, and that median and mean overpayment is 26% and 37%, respectively, of the user optimal tariff bill. Additionally, we observe not only that perception of traffic usage biases the tariff choice but also that the choice of tariff biases traffic usage: the traffic demand grows substantially when users switch from pay-as-you-go to a bundle tariff, and that traffic demand on a bundle is not uniformly spread across time.

Categories and Subject Descriptors

K.6.2 [**Management of computing and information systems**]: Pricing and resource allocation

General Terms

Economics

Keywords

Behavioral economics, cognitive bias, rationality, tariff.

1. INTRODUCTION

One of the key assumptions underpinning most of the work done on the intersection of networks and economics is that of *rationality* of various players [1, 7, 12, 13]; individual players have perfect knowledge and always act in their own best interest. However, there is little work done on understanding or observing rationality, or lack thereof, *empirically*.

Our focus in this paper is to understand behavior of users in terms of choosing, and using mobile network service plans. A typical service plan can include voice, SMS and data of specific quantities, at a certain price. Understanding rationality[1] in this context is far from simple. First of all, it is hard to uncover reasons behind choices made by users to offered prices: their response to offered options is often adhoc, based on subjective factors that are hard to model and quantify.

Secondly, most network service providers offer multiple tariffs where the price per unit of service (eg. a minute of voice calls or *Gbyte* of data transfers) decreases if more units are purchased[2]. One would expect that an individual user, with a *perfect* knowledge of her current and future demand, would choose the tariff that would meet her demand. However, humans are rarely capable of errorless prediction of their needs, and hence are prone to make sub-optimal decisions [4, 5, 6, 9].

Lastly, demand is an elastic variable that depends on the choice of tariff and associated prices that complicates the choice of the best tariff for a user. A user that spends 100 minutes of voice calls when the cost per minute is 1 USD, is likely to spend more than 100 minutes if she is on a different tariff with a cost per minute of 0.2 USD. Choosing a different tariff in the future, based on the demand generated while being on a certain tariff, may often lead to a sub-optimal decision.

In order to aid our empirical investigations, we rely on a large dataset that consists of hundreds of thousands of paying customers of a mobile provider in a European country, as well as the entire activity of each customer across different services (voice, SMS, data) over an extended period of time (27 months).

[1]Our definition of rationality is tied to economic self-interest, rather than other factors that can also be construed as self-interest (convenience etc.). We discuss our definition and assumptions in Sec. 3

[2]Typically service units purchased in a discounted bundle have an expiry period of one or several months.

Hotnets '12, October 29–30, 2012, Seattle, WA, USA.

1.1 Main Results

The main findings of the present paper are the following:

- We observe that most users deviate from the strict definition of optimality (rationality) we have chosen; users choose seemingly sub-optimal tariffs when purchasing network services. We quantify the extent of such behavior and several factors that may have impact on the observed sub-optimality of tariff choices.
- We examine how the demand changes when user shifts between PAYG and bundle tariffs and show that in average while user is on a bundled tariff she generates significantly more traffic than while being on PAYG: a factor of 2.7 increase of voice traffic, a factor of 2.5 increase in SMS traffic and a factor of 47 increase in 3G data usage.
- When users purchase service in bundles, the consumption at the beginning of the billing period (one month) is larger than the consumption towards the end of the bundle period, both for the capped service (voice) and for non-capped/unlimited services (SMS and 3G data).

There are a number of reasons that could contribute to these seemingly sub-optimal decisions, including: (1) hard-to-characterize cost of being aware enough to track usage and respond accordingly, (2) the value of convenience and risk-aversion of potential future surprises as well as (3) the uncertainty of future use of the service. We do not attempt to model and argue about these reasons in the present paper, and leave it for future work. We believe that observed findings shed light on how human behavior affects the operations of mobile network service providers and also act as a call-to-arms to study rationality in other network services. At the same time, we would like to point out that our results should *not* be extrapolated to the existence of 'rational' behavior, or lack thereof, in other network services.

1.2 Background

Rational choice theory is the preeminent framework used in microeconomics to study decision making. A central assumption of rational choice theory is that individuals taking the decision are *rational*; they act in their best interests. Many subfields of rational choice theory, including game theory, as well as other fields that use tools from economics, including network economics inherit this assumption [1, 7, 12, 13].

However, this assumption has been questioned, with the belief that humans[3] are often *irrational*, capable of making sub-optimal decisions displaying cognitive bias [6]. In this work we look into such effects with regards to mobile network service plans.

[3]We don't include software agents etc

Tariff	Voice	SMS	Data	Price
PAYG	1★/min	0.6★/txt	$\frac{2★}{20MB}$ per day	0
Bundle 1	60min	unlmtd	$\frac{2★}{20MB}$ per day	50★
Bundle 2	250min	unlmtd	unlmtd	100★
Bundle 3	400min	unlmtd	unlmtd	150★
Bundle 4	800min	unlmtd	unlmtd	200★
Bundle 5	1500min	unlmtd	unlmtd	250★

Table 1: Tariffs. Note that SMS is unlimited across all bundles, while Data is unlimited across most

In commercial network providers, a principal mechanism to regulate demand is *economic* signal: through prices of the service. However, the responses to economic signals depend heavily on the individual requesting the resource and are consequently challenging to study. Most network providers offer a suite of tariffs that often bundle multiple services. Such bundling of service units complicates economic analysis, but is widely adopted by network service providers. Odlyzko [10] argues that, historically, communication services such as ordinary mail, the telegraph, the telephone or the broadband inevitably converge to simpler pricing structures like bundles. Therefore, the tension between the desire for simplicity (fewer bundles) on one hand, and convenience and customization (more bundles) on the other also complicates matters for users and network service providers. More recently Chetty et al. [3] survey the evolution of capped bundles and their effect on broadband usage, that points to the complex interplay between pricing shaping demand and demand shaping pricing.

The phenomena studied here lie in the realm of behavioral economics, the branch of economics that study how social and psychological factors influence the economic decisions such as purchasing or consumption of a good [4, 5]. This work aims at improving our understanding of the behavioral factors that have impact on the operation and revenues of network providers.

2. DATASETS DESCRIPTION

The data we study belongs to a mobile provider operating in a European country. This provider offers standard voice, messaging and data services to its user base, and serves several hundred thousand users. The dataset covers a 27-month period from late 2009 to early 2012. Users can choose from a suite of tariffs described in Table 1. Each tariff is prepaid. In addition to a pay-as-you-go (PAYG) option in which each service unit is charged separately, the user has a choice of purchasing one of five service bundles, in which a bundle with a certain quantity of service units is purchased and the units expire a month from the day of purchase. A bundle can be purchased at any time, but not before the expiration of the previously purchased bundle. Thus, a

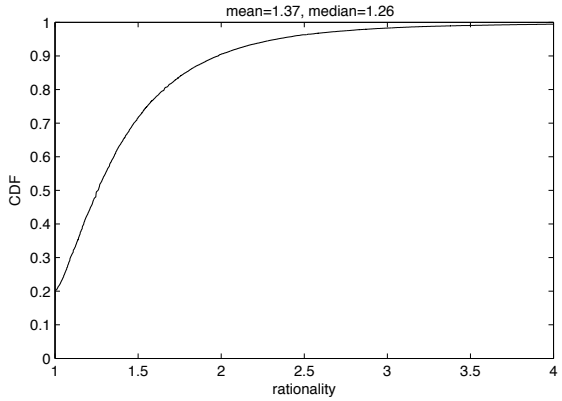

Figure 1: Distribution of per-user rationality. Only users with $optimal\_cost > 500$★ considered.

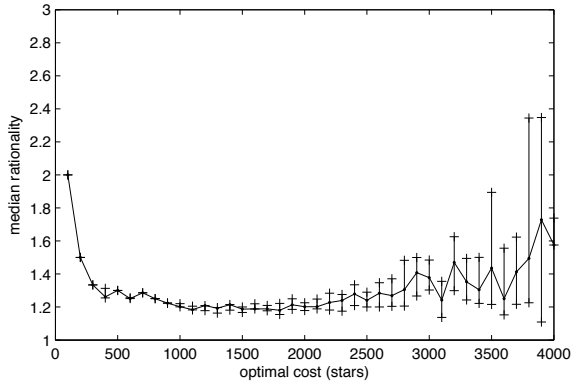

Figure 2: Median rationality as a function of $optimal\_cost$.

user that consumes the service units from a bundle must pay the PAYG prices for the service until the bundle expires, after which she can choose whether she wants to continue with the PAYG tariff or buy a new bundle. There are *no* switching costs imposed in changing between bundles and the procedure to change bundles is exactly the same for every customer, irrespective of the bundle they subscribe to. Everything is carried out online, and hence the effort required to stay with the same bundle and renew is the same as the effort (or convenience/inconvenience) to switch to a different bundle. We therefore assume the switching costs to be zero.

For each user we have the following information:

- Calls: every call made and received by the user with the time-stamp and the duration ($\sim$ 1B call records)
- SMS/MMS messages: every SMS/MMS message sent and received by the user with the time-stamp ($\sim$ 1B records)
- Data usage: every data session generated by the user with the time-stamp, data-volume and the duration of the session ($\sim$ 1B records)
- Payments: every payment made by the customer stating whether the payment is purchase of a bundle or PAYG top-up credit.

We point out again that the provider we analyze does not offer any long-term contracts that many other mobile providers use to lock customers within a (potentially sub-optimal) tariff for a long period of time. And we also stress that the bundle offerings and the associated prices have not changed during the duration of the collected dataset.

3. RATIONALITY OF TARIFF CHOICE

In this section, we present evidence on rationality of users, defined in terms of the choices they make with regard to different service plans or bundles. Intuitively, a user who acts in her own best interest and is 'rational' would choose a plan that is commensurate with her demand and pay accordingly. Our definition of 'rational' behavior is based on economic self-interest and is also myopic. We do not consider other notions of self-interest such as 'convenience' or agileness in our definition. Convenience/agileness indeed carry certain value/cost, which is indeed hard to translate into a tangible economic quantity. Not taking into account such notions of self-interest may limit the generality of our results, but as we will argue later, some properties of our dataset helps us exclude simpler explanations of the observed behavior.

There are numerous reports that point to the underusage of the service bundles sold by network operators [11, 2]. In addition, bundles are structured in a way that it is often financially advantageous for customers to purchase a bundle that offers more units than what they consume. For example, let us consider Alice, Bob and Carl who are customers of the same network provider, with the pricing plans described in Table 1. All three of them use only voice services but consume different amounts: Alice uses 50 minutes per month, Bob uses 150 minutes and Carl uses 420 minutes per month. It is not hard to figure out that Alice would be best off purchasing the PAYG service, Bob would be best off by purchasing the 250-min bundle while Carl would be best off with the 400-min bundle and spending the extra 20 minutes on the PAYG tariff.

Choosing the 'right' tariff, however, is not as simple as solving a simple numerical problem. Individual customers find it hard to accurately predict their monthly consumption and their needs. Choosing the optimal tariff is further exacerbated when we consider that multiple services (voice, txt and data) are billed to the same account making the problem of predicting consumption harder. It is no surprise, then that most customers prefer to be risk averse, and consistently purchase a larger-than-optimal bundle to *insure* themselves

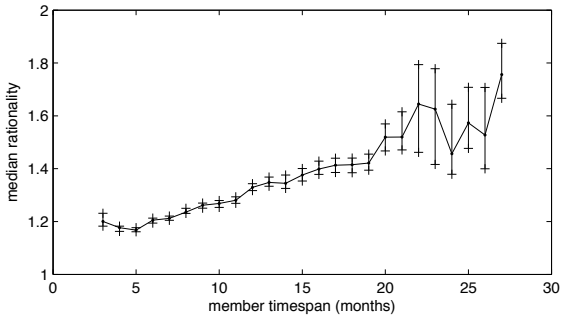

Figure 3: Median rationality as a function of user subscription duration. Only users with $optimal\_cost > 500$★ are considered.

against a higher bill due to overconsumption beyond the limits of a smaller bundle.

In order to quantify how far tariff choices are from the optimal in the mobile provider data described in Sec. 2, we first calculate the $optimal\_cost(u)$ for every individual customer u as the minimal value needed to service all traffic generated by the customer from the day she joined the network. We define the *rationality* of a user u as the ratio between actual payments ($actual\_cost(u)$) she made since joining the network and the $optimal\_cost(u)$, amount needed for servicing her demand:

$$rationality(u) = \frac{actual\_cost(u)}{optimal\_cost(u)}. \quad (1)$$

In Fig. 1 we report the CDF of the set of users with optimal cost[4] of at least 500 ★ (which is equivalent to 500 voice PAYG minutes). We *filter out the light users*, as they are either relatively new customers or those that use the service only occasionally and very irregularly (eg. as a second phone)[5]. This enables us to focus on rationality as a long-term behavioral trait. From the figure we can observe that median and mean rationality is 1.26 and 1.37 respectively. Thus, half of all the customers pay 26% more than necessary to service the demand they generated. While most customers are relatively rational (rationality index close to 1) there is a nontrivial fraction of users ($\sim 10\%$) that pay a factor of 2 or more of optimal amount needed to service their demand.

We see that sub-optimality of tariff choice is an inherent property that most customers experience to some extent. A natural question to ask then: does rationality vary across different groups of customers? In Fig. 2 and Fig. 3 we depict the median rationality index (along with the 95-th percentile confidence bars) as a function of two parameters: the *optimal_cost* (an indicator of the demand of the user; heavy or light) and user duration (the duration the user has been with the provider), respectively. We observe that, in general, heavier users have lower rationality index, a property that can be explained by the non-linear bundle pricing, and the fact that relative difference in the bundle prices are smaller for the larger bundles. Once the *optimal_cost* passes 500★, the median rationality index remains flat indicating relative insensitivity of rationality to the volume of the user. This is somewhat counter-intuitive as heavy users who normally opt for expensive plans can be taken to be price-insensitive compared to users who opt for cheap plans. This doesn't appear to be the case. Note that while relative overpayment exhibits a decreasing trend along the *optimal_cost*, converging to around 20%, in terms of absolute overpayment the heavier users in average overpay more than the lighter.

From Fig. 3, we can observe that new users (those that have joined in the last several months), appear to be more aware of the pricing options and in general have lower rationality index. The users that are longer with the provider, appear to be less optimal, which is a phenomenon worth deeper investigation.

Discussion: The results we have shown so far portray a complex picture of rational behavior. Users are often not rational when it comes to purchasing bundles, and users belonging to different groups have different behavior. In relative terms, heavy users tend to be more rational, as do the ones who are relatively light. In absolute terms, however, heavy users do overpay more than the light users, on average. Understanding and attributing the reasons behind such behavior, whether it be the proclivity of users to be risk-averse, or other reasons is left for future work. However, given the properties of the datasets; no switching costs, freedom to switch between bundles at any time etc., we posit that simple explanations will not suffice.

4. IMPACT OF TARIFF CHOICE ON TRAFFIC

In the previous section, we looked into how users choose different plans based on their demands. In this section, we look into the flipside of this equation; we look at the extent to which the choice of a particular bundle with an associated tariff drives demand.

We begin our analysis examining the usage characteristics of users that switch between PAYG tariff and the bundles. Bundles 2-5 offer significantly cheaper voice, SMS and data services, and hence one could expect that a user consumes more of a particular service while using a bundle than while being on the PAYG tariff. To examine the difference between the consumption at PAYG tariff and a bundle we consider the set $\mathcal{L}$ of all subscribers that have spent at least 3 months on PAYG and at least 3 months on one of the bundles 2-5[6]. By comparing the demand levels of users that use both the

[4]Total optimal cost, not monthly.

[5]Such light users correspond to only a small fraction of the total traffic and revenues.

[6]We exclude Bundle 1, as it offers negligible discounts for

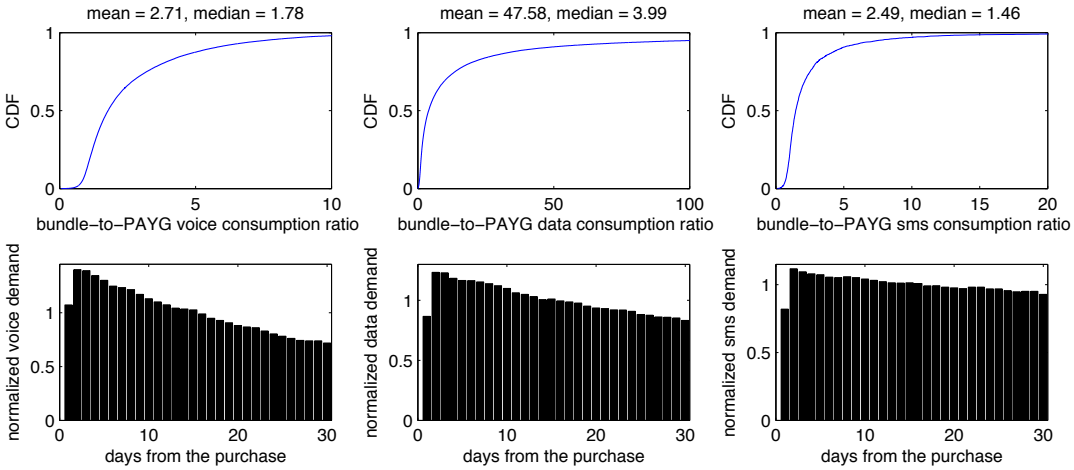

Figure 4: Top: Distribution of per-user ratios of bundle and PAYG demand. Bottom: Dynamics of usage during the month of the bundle. Left: voice. Middle: 3G data. Right: SMS.

PAYG and bundle tariffs we can measure elasticity of the demand for a particular service.

For every user u in $\mathcal{L}$, let us denote with $D_{PAYG}(u)$ and $D_{bundle}(u)$ the average daily demand while the user u is on PAYG tariff and a bundle, respectively. The bundle-to-PAYG demand ratio for user u is

$$B2P(u) = \frac{D_{bundle}(u)}{D_{PAYG}(u)}.$$

Thus for a user v with same average daily demand on both PAYG and bundle tariffs, $B2D(v)$ is equal to 1. For a user that consumes more minutes on PAYG, $B2D(v)$ is less than 1, while for the others $B2D(v) \geq 1$. We note here that these findings provide evidence for the price elasticity of demand, but not fully prove it.

In Fig. 4, we depict the distribution of $B2P(u)$ for users in $u \in \mathcal{L}$, for all the services: voice, 3G data and SMS. In all cases the majority of users ($\sim 90\%$) have, as expected, $B2P(u) > 1$. The median (mean) $B2P$ is 1.78 (2.71), 3.99 (47.58) and 1.46 (2.49) for voice, 3G data and SMS, respectively. However, a nontrivial fraction ($\sim 10\%$) of users consume less when the service is cheaper. Understanding the root-causes for this counterintuitive behavior is out of scope of present work and an interesting open research problem.

The second property of service usage under bundles is the time-variable nature of the usage demand during the bundle period (one month). In Fig. 4, we also plot the aggregate daily demand for all the users and all purchased bundles (excluding PAYG) for each day of the bundle duration[7]. It turns out that the average demand in the days immediately after the bundle is purchased is highest and it decays as the month progresses. The reason for the decrease in voice usage can perhaps be explained by the risk-averse nature of most users as they approach their last days of their respective cycle; having lower units to use can lower usage. However, in case of 3G data and SMS services, all considered bundles have unlimited usage caps. Hence, it is hard to reason on possible causes of variability in daily usage during the bundle period – why should the usage decrease towards the end of the period? There could be multiple reasons for this – using the device more for calls encourages data usage as well, and hence if the device usage for calls drop, data consumption drops as well. Or perhaps users do not realise they are on unlimited data plans. Evaluating these conjectures is left for future work.

The observed dependence between the time-since-bundle-purchase and demand may have critical impact on the way the provider bills for service. For instance, many providers prefer billing the monthly usage on a fixed day every month for accounting reasons. This synchronizes the billing period for all users using the bundle. Such synchronization of the billing periods may undesirably lead to a traffic peak during the first few days of the bundle and can be non-negligibly larger than the average traffic rate, than otherwise if the billing period of different customers were out of sync, and uniformly distributed over a month.

5. OPEN PROBLEMS

In the context of price-demand relationship of network services, there are a number of problems that remain open.

Sub-optimal tariff choices: root causes. In Sec. 3, we observed that many customers choose a tariff that is sub-optimal (in an economic sense) with regards to their consumption footprint. What are the factors that cause such decisions? Is the multiple-service nature of the bundle responsible for the degree of sub-optimality?

voice and data. Additionally, only a very small fraction of users choose this tariff.

[7]To protect sensitive information, we normalized the demand values to have mean equal to 1.

How does behavior change over time? Are users risk averse with regards to the unpredictable cost of PAYG; do they insure themselves against the unpredictable costs? Is user rationality correlated with volatility of daily usage time series of the user? How is sub-optimality influenced by user's over/under-estimation of her consumption?

What happens with other providers/domains? The observed behavioral phenomena in tariff choices and network usage indeed depends on the studied network and its user base. The provider we analyze is representative in terms of user base and the tariff structure, and thus we are reasonably confident that observed qualitative results are likely to hold in other mobile networks. In contrast, the behavioral patterns in other types of network services such as residential broadband networks or wholesale IP transit remain unknown. A detailed empirical study of behavioral patterns in different domains and providers would help understanding the generic and particular behavioral factors that affect network service usage and revenues.

Price elasticity of demand (PED) of bundled network services. PED is a metric economists use to quantify the flexibility of demand of a product, subject to changes in prices. In Sec. 4, we demonstrate the elasticity of basic mobile network services. However, when multiple services are sold in a bundle it is not obvious how to quantify the effect that each of the services have on the purchasing decision, and hence evaluation of PED (including a proper methodology for studying such phenomenon) in the context of bundled services remains an open research problem.

Efficiency of bundles. In a recent paper [13], authors argue that in the context of IP transit market, a low number of bundles (or so called tiers), say three or four, yields near-optimal profit for the ISP selling the IP transit service, under the assumption that all the players are rational. In the mobile network services, the price elasticity of demand follows different laws, multiple services are bundled together and users choose sub-optimal bundles. Hence, the question of how many bundles (and how to price them) are necessary to achieve near-optimal profit for the mobile service provider stands wide open.

6. SUMMARY

In this paper, we looked into rationality of users when it comes to choosing network service plans. Service pricing is a primary mechanism for demand and revenue control of commercial network providers. However, the way end-users respond to different prices and services is largely unknown. In this paper, we investigated several phenomena related to how end users purchase and use network services. Faced with an array of tariffs to choose from, users must decide what tariff is the most appropriate for them. Looking at the data from a mobile provider offering voice, SMS and data services, and serving several hundred of customers, we observe that a significant fraction of end-users make suboptimal decisions: the mean and median overpayment are in the 37% and 26% of the optimal charge, respectively. On the other hand, the decision on the tariff affects the user demand not only through the growth of demand when the service is cheaper, but also through intriguing temporal variability of demand through the billing period.

Overall, we believe the results in this paper present a more complex view of users, that they often act in ways that do not appear to be rational. In addition to trying to understand the reasons behind such behavior, as well as the implications on the network, we also hope that the network research community takes into account this complex, yet more realistic, view of human behavior in their efforts.

7. ACKNOWLEDGEMENTS

We would like to acknowledge Richard Gass and our partners at the operator that provided us with data. We also thank the anonymous reviewers for the constructive comments, and an extra special thanks to our shepherd, Emin Gün Sirer who went above and beyond to help us polish the paper.

8. REFERENCES

[1] G. J. Armitage. "Revisiting IP QoS: why do we care, what have we learned?". ACM SIGCOMM CCR, vol 33(5),2003.
[2] Cell Phone Shopper's Guide. Available online: www.citizensutilityboard.org/CellphoneShoppersGuide.html
[3] M. Chetty, R. Banks, A.J. Bernheim Brush, J. Donner, R. E. Grinte. "Youre Capped! Understanding the Effects of Bandwidth Caps on Broadband Use in the Home". In Proc. of ACM CHI 2012.
[4] S. DellaVigna. "Psychology and economics: Evidence from the field". Journal of Economic Literature, 47(2), 2009.
[5] S. DellaVigna, U. Malmendier. "Paying not to go to the gym". American Economic Review vol. 96, 2006.
[6] T. Gilovich, D. W. Griffin, D. Kahneman. "Heuristics and Biases: The Psychology of Intuitive Judgement". Cambridge University Press, 2002.
[7] F.P. Kelly, A.K. Maulloo, D.K.H. Tan. "Rate control for communication networks: shadow prices, proportional fairness and stability". Journal of the Operational Research society vol. 49(3), 1998.
[8] A. Lambrecht et al. "Price Discrimination in Service Industries". Mark Lett vol. 23, 2012.
[9] N. Lathia, L. Capra. "Mining mobility data to minimise travellers' spending on public transport". In Proc of ACM KDD 2011.
[10] A. Odlyzko. "Internet pricing and the history of communications". 2001.
[11] OMIO Report. Available online: www.omio.com/media_centre/docs/WastedMinutes.pdf
[12] J. Shu, P. Varaiya. "Pricing Network Services". In Proceedings of IEEE INFOCOM 2003.
[13] V. Valancius, C. Lumezanu, N. Feamster, R. Johari, V. Vazirani. "How Many Tiers? Pricing in the Internet Transit Market". In Proc. of ACM SIGCOMM, 2011.