

Collaborative Consumption for Mobile Broadband: A Quantitative Study

Juan Camilo Cardona^{†,‡}, Rade Stanojevic[‡], Nikolaos Laoutaris[‡]
Telefonica Research, Barcelona, Spain[‡]
IMDEA Networks Institute, Madrid, Spain[†]

ABSTRACT

Mobile broadband is predominantly priced following tiered plans that involve a certain prepaid *commit volume* and additional metered volume priced at a higher *penalty rate*. An individual's demand, however, may vary wildly from month to month and thus users inevitably purchase packages that are either too small or too large for their needs. By collaborating in predefined closed (*e.g.*, family) or open groups (*e.g.*, through tethering) users can reduce both the amount of paid-but-left-unused capacity or the high penalty rates. In this paper we present a quantitative study of collaborative consumption using data from 40K mobile subscribers and tariffs from ten operators around the world. We show that small 2-person family plans offer modest expected savings in the range of 3% to 14%, whereas getting more substantial savings requires rather impractically large groups of approximately 10 people or more. Going over to open groups, where users can freely trade their data capacity, we characterize the impact of the secondary market price on the user costs and the operator revenues, and show that Telcos might be better off to embrace secondary markets (*e.g.*, let them integrate with billing), and thus have some control on, as opposed to letting them operate unsupervised through tethering. The latter may severely harm the revenues of a Telco, in an uncontrolled manner, especially in densely populated areas.

Categories and Subject Descriptors

K.6.2 [Management of computing and information systems]: Pricing and resource allocation; J.4 [Social and behavioral sciences]: Economics

General Terms

Economics

Keywords

Cellular Networks; Collaborative Consumption; Data Plans; Mobile Broadband.

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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

CoNEXT'14, December 2–5, 2014, Sydney, Australia.

Copyright is held by the owner/author(s). Publication rights licensed to ACM. 978-1-4503-3279-8/14/12 ...\$15.00.

<http://dx.doi.org/10.1145/2674005.2674997>.

1. INTRODUCTION

Over the last two decades, cellular networks have been instrumental to connecting billions of people. Mobile broadband currently corresponds to a significant fraction of the cellular operators' revenues (currently \$1000 billion, or 1-2% of world-wide GDP) and is projected to be the dominant source of revenues in the near future [37]¹. Mobile broadband retails through tiered plans that involve a predefined volume at a certain *commit rate*, and additional (variable) metered volume at a higher *penalty rate* [35].

As we will see later, the user demand for mobile broadband services is extremely volatile, making the decision process of choosing the 'right' data package non-trivial. Consequently, many users either underuse their quota, or occasionally use more than their committed quota, paying large overage fees. With users on both sides of the fence, those that have unused capacity and those that may need some extra capacity, mobile broadband can be a fertile land for a *sharing economy*. In the recent past, sharing economies have emerged in a number of domains sharing various items or services from cars [13], to houses [1], to books [4], and not surprisingly such trend have been picked up in the realms of mobile broadband by researchers [7, 15, 24], app developers [9, 19, 14] as well as the cellular operators themselves [8, 23, 40]. We can classify mobile broadband sharing systems in two groups:

Telco-driven: A number of operators have devised family plans [8, 34]. Such plans are *closed*, meaning that the group of customers that share their caps is closed/predetermined. In some cases, Telcos have even permitted users to freely sell their leftover capacity in an *open* market to any other user willing to buy it [23, 40]. In both cases, the collaborative consumption is implemented through the billing system of the Telco. By allowing plan sharing, Telcos improve customer retention/acquisition and help customers feel they truly own the packages they purchase.

User-driven: Collaborative consumption can be implemented even without the involvement (or blessing) of the Telco through tethering. This is a typical example of the *sharing economy* driven by the users themselves. Several apps have been developed that allow users to share mobile broadband [14, 19]. The large-scale success of these apps, however, depends on the density of users (per km^2) using them; users in low density areas would have infre-

¹We have witnessed a similar trend in the landline telecom market in the previous decade, where the revenues from the voice communications have been almost completely replaced by the broadband services.

quent opportunities to buy/sell unused mobile broadband, while those in denser areas would gracefully have more opportunities for efficient sharing.

Although it is intuitive that collaborative consumption is beneficial for users, it remains largely unknown by exactly how much, and at what impact on Telco revenues. Our goal in this paper is to quantify the potential of collaborative consumption of mobile broadband and for this we use real consumption data from 40K mobile broadband customers as well as pricing data from 10 mobile broadband operators geographically spread across developed and developing markets. We begin by asking “*Is there a case for a sharing economy for mobile broadband?*”. We subsequently ask “*What are the benefits that can be expected from family (closed group) plans?*”. We then move to open group plans and examine both Telco- and user-driven approaches. We ask “*How the price of leftover bandwidth on the secondary market impacts user bills and operator revenues?*”. We examine the previous question under different levels of “liquidity” in terms of the ability of users to exchange bandwidth (either through operator supported billing yielding perfect liquidity or through user driven tethering yielding partial liquidity depending on the density of users). We finally ask “*What is the impact of market liquidity on open collaborative schemes?*”.

The main contributions of our work are the following:

- We demonstrated that the volatility of user demand makes it hard for the users to pick the ‘right’ package, effectively leaving significant fraction of purchased packages unused, with occasional spikes in demand that go beyond purchased commit volumes. We show that this holds even if users are able to adjust their contract on a monthly basis using various package selection strategies.
- Looking at the smallest closed (“family”) plans involving $k = 2$ users, we see that expected savings are 3-14%. Increasing the group size yields greater savings but the extent always depends on the tariff structure. For $k = 10$ the savings can range from 11% to 45%. Such groups are rather large, however, for most practical uses. Generally the savings achieved through collaboration are higher for tariffs that involve a greater difference between commit and penalty rate.
- We turn to open sharing model and start with full liquidity as realized for example by Telcos that help their users sell their leftover capacity using their billing system. We characterize the relationship between the steady-state supply, demand and the price of the mobile data in such secondary market. Overall users could expect significant savings in the range of 19% to 71%. The exact number depends on the relationship between the commit, penalty, and secondary market rates. Knowledge of this relationship is important for a Telco in order to decide whether it wants to permit a secondary market, whether it wants to control it by setting the secondary price, and also to help the operator to design the tariffs in the presence of such secondary market.
- In the case of user-driven collaborative consumption through tethering, liquidity plays a paramount role. Unlike in the Telco-supported case, it is no longer guaranteed that the leftover capacity of a user can be brought to the market. We employ census density

maps to estimate the liquidity in one representative European country and characterize the impact of the density of sharing users on the sharing economy. Our results suggest that even in low-density areas of 3 neighbors per user the potential user savings can be significant, in the range of 8-23%, while in the dense, urban, areas the opportunity for sharing via tethering approximates the case of full liquidity.

For both open and closed groups, we report a number of factors that influence the efficiency of sharing. Explicit recommendations on how to choose tariffs, or how to choose partners in closed groups or buyers/sellers in the open groups is out of scope of the present paper.

2. DATA

Here we briefly introduce the datasets used in the rest of the paper.

2.1 Per-user demand

The mobile broadband usage data analyzed in this paper are from a small national European cellular operator. The dataset contains a record for each mobile broadband session with the anonymized user id, timestamp and the volume of the data session in bytes. To facilitate the empirical study in the following sections, we extract the set of all those users who have used mobile broadband at least once every month over a 7-month period spanning the end of 2011 and beginning of 2012. There are around 40 thousand such users and we have a detailed history of their mobile broadband usage over the 7 month period. An important property of the dataset is that mobile broadband in this operator was not metered during the period over which the data was collected, *i.e.*, there are *no* volume caps. This means that the demand generated by users is their natural one, without any bias from pricing, *e.g.*, self-regulation due to volume caps and penalty rates.

We observe some seasonal variability in the data and a relatively small growth of the overall mobile broadband demand. Fig. 1 (left) shows the per-user median and mean mobile broadband demand over $m = 7$ months covered by our dataset, from which we can observe that in the final months of the dataset the overall demand is some 20% larger than in the first months of the considered period. Also, we note that mean is greater than the median by a factor of 5, indicating a skewed distribution in per-user demand.

Denoting the mobile broadband volume of user u in month t by $d_u(t)$, we capture her temporal demand volatility through the ratio between the minimum and maximum of the monthly demand series:²

$$variability_index(u) = \frac{\min_{t \leq m}(d_u(t))}{\max_{t \leq m}(d_u(t))}. \quad (1)$$

The empirical CDF of *variability_index* across all 40K studied users is depicted in Fig. 1 (right). It reveals a large amount of volatility in monthly demand; for around 60% of all the considered users the maximal monthly usage is $\times 5$ or more of their minimal monthly usage (*variability_index* < 0.2).

Such temporal volatility far exceeds the volatility observed in shared backbone links where statistical multiplexing of

²Note that our relatively small sample (7 months) precludes using measures like standard deviation to quantify volatility.

Operator	Tariff 1	Tariff 2	Tariff 3	Tariff 4	Tariff 5	Tariff 6	Tariff 7	Tariff 8	Tariff 9
Vodafone IND	300/100/2	1000/250/2	1500/375/2	2000/450/2	3000/650/2	4000/750/2	5000/850/2	8000/1250/2	12000/1600/2
Telstra AUS	250/20/0.2	700/30/0.2	1000/40/0.2	3000/50/0.2	4000/80/0.2	6000/100/0.2	12000/180/0.2		
Vivo BRA	150/30/0.4	250/35/0.1	500/60/0.1	2000/90/0.1	5000/120/0.1	10000/200/0.1			
Claro ARG	125/49/0.1	512/69/0.1	1000/99/0.1	3000/129/0.1	5000/159/0.1	10000/199/0.1			
Olleh KOR	500/34/.04	1000/42/.04	1500/52/.04	3000/62/.04	5000/72/.04	7000/85/.04	10000/100/.04		
Unicom CHI	150/46/0.3	300/66/0.3	400/126/0.3	500/156/0.3	750/226/0.3	950/286/0.3	1300/386/0.3	2000/586/0.3	3000/886/0.3
E-plus GER	250/10/.05	750/15/.01	2000/20/.01	5000/25/.01					
Vodafone UK	250/3/.1	500/7.5/.015	2000/15/.015						
Etisalat UAE	100/29/1	1000/99/.5	5000/249/.5	10000/449/.5					
ATT US	300/20/.067	3000/30/.01	5000/50/.01						

Table 1: The tariffs for prepaid mobile broadband from 10 representative mobile operators. The compact representation (x, y, z) refers to the cap of x Mbyte, price y in local currency and z overage rate in local currency per Mbyte.




Figure 1: Left: mean/median monthly usage (scale not shown to preserve confidential information). Right: empiric CDF of the ratio between the min and max of the per-user monthly demand.

individual flows makes daily or monthly aggregate volumes almost periodic (off by few percentage points [17]). Interestingly, even individual consumption for utilities like water and electricity [3] appears to be more predictable than mobile broadband consumption.

2.2 Pricing

In addition to mobile broadband usage data, we also collected data on the pricing packages of 10 different mobile operators in 10 different countries across 5 continents (listed in Table 1) [11]. This dataset depicts a wide range of relationships between commit and penalty rates, including offers with low caps and relative high penalty rates, which are usually referred to as pay-per-byte packages. Plan diversity, as will be shown later, plays an important role in package selection and are critical in understanding the potential of collaborative consumption in mobile broadband.

Since we started working on this paper in the autumn of 2013 there were virtually no changes in the tariffs among those 10 studied operators.

3. USER PACKAGE SELECTION

Before we look at the potential of collaborative consumption in the mobile broadband, we will first in this section demonstrate that the high temporal volatility of demand makes the selection of an appropriate plan hard, thus leaving many users effectively (heavily) under-using their package with occasional spikes in demand which go beyond the cap of the purchased cap.

We use $d_u(t)$ to denote the demand of user u in month $t = 1, 2, \dots, m$. At the beginning of every month,³ users

³We assume that the billing happens in the monthly cycle. Both shorter (e.g., daily or weekly), or longer (yearly) plans are being offered by some providers, though the monthly billing cycle is the most widely used around the world [35], and hence we focus on them exclusively in this paper.

choose a plan among the set of available plans \mathcal{P} . The plan $i \in \mathcal{P}$ is determined by the triplet $(cap_i, price_i, \mu_i)$: the volume with a cap (in Mbyte), the committed price of the plan (in \$), and overage price – what user has to pay after the volume cap is reached (in \$⁴ per Mbyte), respectively. If the plan $i \in \mathcal{P}$ is chosen the charge at the end of the month during which the user consumed volume d (in Mbyte) is

$$C(d, i) = price_i + \max(0, d - cap_i) \cdot \mu_i.$$

At the beginning of every month t , the user u has to choose the plan i that matches her demand. If she knew her demand up-front, she could choose the plan that minimizes the charge for the month t , and we use $o_u(t)$ to denote the minimal possible charge and refer to this as the **a-posteriori optimum**

$$o_u(t) = \min_{i \in \mathcal{P}} C(d_u(t), i).$$

However, the demand $d_u(t)$ is rarely known upfront, and users can use various heuristics that depend on history of consumption in the previous months, to choose the plan during month t based on the estimated demand in that month t . For that purpose, we use the following possible heuristics:

Max rule (MR): in month t the user u chooses the plan i_{MR} that minimizes the charge in the month of the maximal demand prior to t .

$$i_{MR} = \arg \min_{i \in \mathcal{P}} \left(C \left(\max_{1 \leq \tau \leq t-1} (d_u(\tau)), i \right) \right). \quad (2)$$

For user u that chooses the plan i_{MR} at the month t the charge is

$$h_u^{MR}(t) = C(d_u(t), i_{MR}) \geq o_u(t)$$

Optimal-expectation rule (OER): in month t the user u chooses the plan i_{OER} that minimizes the expected charge assuming that the demand in month t is equal to the demand in one of the previous months with uniform probability $(1/(t-1))$.

$$i_{OER} = \arg \min_{i \in \mathcal{P}} \left(\frac{1}{t-1} \sum_{\tau=1}^{t-1} C(d_u(\tau), i) \right). \quad (3)$$

For user u that chooses the plan i_{OER} at the month t the charge is

$$h_u^{OER}(t) = C(d_u(t), i_{OER}) \geq o_u(t).$$

We also experiment with other heuristics that output the plan that optimizes the cost, by using the estimate for the next month demand equal to the average and median of

⁴By \$ we mean local currency.

the previous month’s demand, and observe the average costs comparable to MR and OER. Thus we omit them for brevity. Sophisticated time-series analysis is unlikely to provide significant improvements over these simple estimators, since our dataset contains relatively small number of data points (only 7 samples per user). However, the impact of more advanced estimators, remains to be studied in the future.

Remark 1. We would like to note that modeling the user decision process regarding the tariff choice is highly challenging. We settle here for the two intuitive strategies, OER and MR, as two possible ways to model such decision process. Whether such models reflect the tariff selection process (employed by end-users) remains to be validated in practice.

Remark 2. Similarly, modeling the impact that volume caps may have on user demand is a hard problem [22][5][28]. In this paper, we assume that users’ demand is invariant and influences the decision on the tariff as well as the overall cost. Alternatively, one could analyze another model in which users adjust their demand under caps paying an implicit (non-monetary) cost in terms of reduced traffic consumption. Models with invariant user demand can lead to higher cost per user with respect to the self-regulation models. However, our base assumption allows us to compare different types of collaborative systems directly, without the need of quantifying the indirect cost experienced by users when they cannot satisfy their original traffic demand.

3.1 User demand vs. capacity

Intuitively, the high volatility of demand combined with large penalty fees, implies that many (rational) users would often, but not always, choose a tariff that leaves them some spare unused capacity as a cushion against the risk of going over the cap and paying large overage fees. We evaluate mismatch between the user demand and the capacity she purchases (under OER or MR), and in Figure 2 we plot the data volume (right) and fraction of users under-the-cap and over-the-cap (left). We observe that in most⁵ of the cases the amount of traffic (and fraction of the users) under-the-cap is significantly larger than the traffic (and fraction of the users) over-the-cap, indicating an opportunity to ‘offload’ the over-the-cap traffic through those users that under-utilize their cap and avoid paying the large penalty fees. In later sections, we will study how sharing of underutilized packages may affect the bill of the mobile broadband users. While over-the-cap volume may appear to be small, the amount of revenues operator could collect from those overage charges is nontrivial, and in the case of OER it ranges from 11% for ATT tariffs to 42% in Vodafone India tariffs. The revenue coming from overage fees is indeed smaller in MR than for OER, as by default customers purchase larger packages and have smaller over-the-cap volumes.

3.2 Heuristics vs. a-posteriori optimum tariff

This paragraph is not directly related to sharing economy, but it does represent an interesting result on its own, and partly answers the question raised by [28] on why users choose ‘inappropriate’ data plans. Namely, we quantify how much do heuristics (OER and MR), based on the historic demand pattern, differ from the a-posteriori optimum tariff, for the real users with volatile demand pattern.

⁵The only exception is China Unicom, which has extremely low penalty rate, which basically discourages users to purchase large packages.




Figure 2: The amount of total unused and over-the-cap traffic (left) and fraction of users under- and over-the-cap (right) for the 10 operators under OER and MR tariff selection. Significant amount of unused packages, could potentially serve all of the over-the-cap traffic.

To compare the cost max rule h_u^{MR} , and optimal-expectation rule h_u^{OER} , against the a-posteriori optimum o_u we use diverse pricing policies, described in Sec. 2 to study the effect of different available tiers and over-the-cap rates. In Fig. 3 we plot the ratio between the average⁶ charge based on our heuristic rules (MR and OER) and the average o_u , in the final month of our dataset $t = m = 7$.

From the Fig. 3 we can learn several lessons. First, the relative difference between the history-based heuristics and the a-posteriori optimum largely depends on the structure of the available plans, and specially the relative price per *Mbyte* of the committed cap and the overage/penalty rate. For operators with very large penalty price compared to the price of *Mbyte* in a committed cap, such as Vodafone India or Etisalat UAE, the heuristics perform much worse than the a-posteriori optimum with customers using the OER heuristic paying 70% and 45% more than a-posteriori optimum. For operators in which the penalty rate does not differ much from the committed rate, the cost of the heuristic rules does not significantly differ from the a-posteriori optima, as is the case with China Unicom and ATT for which the difference between the OER bill and the a-posteriorly optima is under 5%. We also observe that the optimum-expectation rule and maximum rule show similar average costs, with OER being marginally better in all of the ten studied policies.

4. CLOSED GROUP SHARING

In the previous section, we showed that demand volatility and diverse commit-to-penalty ratios make the task of selecting an appropriate plan for an individual user hard, leaving many users selecting plans much larger than their needs. One way to address such difficulties is through *collaborative purchase plans* shared among multiple users. The key advantage of shared plans is that the variability of the aggregate of independent demands is smaller than the sum of individual variabilities of its constituents. This effectively

⁶Average over the whole user-base.




Figure 3: Ratio between the average cost of Max rule (optimal-expectation rule) and average a-posteriori optimum.

allows a group to pick collectively a better suited plan than what members can do on their own. Indeed, there exist several proposals for protocols that facilitate ad-hoc sharing of resources like voice call minutes [29] and mobile broadband [19] between users. Likewise, several operators offer shared plans for families and friends [8]. Our analysis is focused on characterizing the gains from collaboration – to the best of our knowledge, this is the first such empirical study.

The scenario we study is the following. A group G is composed of k users that purchase individual mobile broadband plans. When a user exhausts her volume, she may consume from the volume of any of the other users in the group. In order for such a scheme to work in practice, users must be in a shared data plan offered by some of the operators, or be in close proximity for a significant part of the billing period and thus be able to implement ad-hoc sharing over WiFi or bluetooth [19, 29, 14]. If k users in the group G have chosen plans $I_G = (i_1, \dots, i_k)$, and generated the demand $D_G = (d_1, \dots, d_k)$, respectively, then the cost for delivering that traffic is:

$$C(D_G, I_G) = \sum_{s \in G} price_{i_s} + \max(0, \sum_{s \in G} d_s - \sum_{s \in G} cap_{i_s}) \cdot \min_{s \in G} (\mu_{i_s}).$$

Users can purchase plans (for each month t) either in a:

- (1) non-coordinated manner using one of the heuristics described in the previous section, say OER, or in a
- (2) coordinated-OER manner by choosing the tariffs $I_G = (i_1, \dots, i_k)$ in month t to minimize the expected cost aggregated among all partners in the group.

$$I_G^{OER} = \arg \min_{I_G \in \mathcal{P}^k} \frac{1}{t-1} \sum_{\tau=1}^{t-1} C(D_G(\tau), I_G). \quad (4)$$

4.1 Setting up groups for volume sharing

There are many ways in which the users can engage in collaborative cap sharing. We briefly describe three dimensions of the problem of group selection that most collaborative arrangements face at some stage.

First, we need to answer how to select the group for collaborative sharing, which can be done in many ways. For example, the group selection can be done based on (1) *location* where users residing at the same location can use short-distance communication to share via tethering; or (2) *social relationships* where users with close social ties collaborate; or (3) *random* where random users collaborate opportunistically.

To evaluate the gains we use the random group matching. In Appendix A we demonstrate that groups formed using




Figure 4: The average cost in the groups with coordinated tariff selection, relative to the cost of individual a-posteriori optima.

the social relationship or home location result in statistically identical benefits as those with random grouping.

The next question would be to decide how large the group should be. Smaller groups are easier to set-up and coordinate, while larger ones are more cumbersome but may provide greater benefits from statistical multiplexing. Finally, once a group is formed, one needs to decide on how to select individual plans – in a non-coordinated manner (simple but sub-optimal) or in a coordinated one (more complicated but efficient).

In the remainder of this section, we study several facets of this design space. We study how the group size affects the cost of the group $C(D_G, I_G)$ compared to the cost of delivering the same demand without sharing the caps $\sum_{u \in G} C(d_u, i_u)$, when the groups are selected randomly. We also evaluate how much the coordination between the users in choosing the next-month-tariff can help in reducing the bill, when the groups are selected randomly. We conclude the section with the analysis of how the demand volume similarity affects the expected statistical multiplexing and costs.

4.2 Effect of group size

Intuitively, larger groups yield more predictable aggregate demand by statistically multiplexing more independent sources. In order to test this hypothesis and evaluate the impact of the group size on the overall cost incurred by the group $C(D_G, I_G)$, we vary the group size k , and for each k generate $N = 10000$ random groups of k members from our dataset, and evaluate the ratio between the average cost per group where the packages are chosen in a coordinated-OER manner (optimizing Eq (4)) in the final month of our dataset ($t = m = 7$), and the average cost of sum of the individual packages when no sharing is in place. We plot our findings in Fig. 4. With $N = 10000$ sampled groups, the standard error in all cases is less than 1% and, thus, the confidence bars are omitted for clarity.

Expectedly, the groups with 2 or more members can reduce the group cost. Having a single partner in the group is likely to bring the average cost of the 2-partner group down for 3%-14% compared to the cost of purchasing the plans individually using the OER and not sharing them $\sum_{s \in G} h_s^{OER}(m)$. For groups of size $k = 2$, the largest reduction in cost (around 14%) occurs with plans that impose high penalty rates (e.g., Vodafone India/UK, Etisalat




Figure 5: The empirical CDF of (coordinate) group savings in Vodafone India. Group size of 2 and 10 shown. The other group sizes and operators follow similar trend.

UAE and Telstra AUS tariffs). With larger groups, *e.g.*, $k = 10$, the benefits grow more and can range between 11% (for Olleh South Korea tariffs) and 45% (for Etisalat UAE tariffs). Another way to appreciate the gains is to look at the corresponding aggregate cost reduction for the entire group. Figure 5 depicts the CDF of the group savings for Vodafone India tariffs, and group sizes of 2 (pairs) and 10. For the group size of $k = 2$ (pairs), 60% of all pairs would see no benefit in collaborating, while around 20% of pairs could expect the cost reduction of 30% or more by sharing their caps. For larger groups of $k = 10$ partners, the distribution of expected savings is smoother, with the mean and median at around 35%.

While groups of 10 or more partners may bring considerable benefits for the involved partners, they are rather difficult to setup and maintain. With the open sharing analyzed in Section 5 such concerns of group creation and maintenance disappear.

4.3 Intra-group coordination




Figure 6: The ratio of the average group cost between the coordinated and the non-coordinated tariff choice. For small group sizes, coordination gives small benefit, which grows for larger groups.

In the previous section, we assumed that once several users engage in collaborative sharing of their caps, they choose their packages in a coordinated manner (to minimize Eq. (4)) based on their historic use pattern. However, if




Figure 7: The cost (relative to the a-posteriori optima $\sum_{u \in G} o_u$) of the pairs sharing the cap for different similarity indices. Users with similar demand (low similarity index) complement each other's cap better than the users with large difference in demand.

users are not willing to share their history with the others in the group, but rather purchase their packages individually, in a un-coordinated manner (using for instance, the optimal-expectation rule Eq. 3) they can still benefit from sharing, although the benefit of such group sharing is likely to be lower than in the coordinated case. In Fig. 6 we depict the ratio between the average group cost between the coordinated and the non-coordinated plan choice obtained by randomly selecting $N = 10000$ groups with k users from our dataset, and evaluating the average cost of the group with and without coordinated plan selection. We can conclude that coordination brings relatively small benefit (5% and less) for very small groups of 2 partners and brings higher benefit for the larger groups. Indeed, the expected benefit of coordination is the pricing-policy dependent.

4.4 Volume-based group selection

A natural question in the context of collaborative volume sharing is related to whom should one partner with. Next, we demonstrate that similarity of demand, in terms of total volume per month, should be taken into account when constructing groups. In particular, a group should combine users that consume similar volume across the month because in that case the statistical multiplexing is likely to provide the largest benefit to the cooperating partners.

The average demand per user is very skewed, covering multiple orders of magnitude. For a group G of users, we define the *similarity* metric as the ratio between the maximum and minimum average demand among the members in the group.

$$similarity = \frac{\max_{u \in G} \sum_{\tau < m} d_u(\tau)}{\min_{u \in G} \sum_{\tau < m} d_u(\tau)}.$$

To evaluate the relationship between the demand *similarity* and the impact of sharing the cap on the cost, we select $N = 10000$ random pairs of users from our datasets, and split them in 7 different sets depending on which of the following 7 segments the pair demand similarity falls into: $[1, 2)$, $[2, 4)$, $[4, 8)$, \dots , $[32, 64)$, $[64, \infty)$. For each of the 7 segments, we evaluate the average cost of the pair sharing the cap in the last month of the dataset, divided by the a-posteriori optima ($\sum_{u \in G} o_u$). We report our findings in Fig. 7 based on the plans of Vodafone India. The other plans and group sizes follow similar pattern and are omitted. From this figure, we can conclude that it is most beneficial

to collaborate with the users with similar demand volumes. Note that users with similar demand would produce a low *similarity* index.

We conclude this section reiterating that closed (family) sharing can reduce the bill of the involved parties, with small savings for small 2-partner groups that grow when group has more partners, with similar volumes and with a coordinated plan selection.

5. OPEN SHARING

The close collaborative groups that we study in the previous section, although beneficial for users, could be too restrictive, and may not allow efficient sharing of unused caps. The gain from the closed-group sharing becomes substantial only when the many users partner-up, they generate similar traffic volumes and coordinate the plan selection. Moreover, coordination and the issues related to how the cost/savings should be shared in large groups may be cumbersome. For that reason, we look at a more flexible model, which we refer to as open sharing model (or simply open model), in which anyone with a mobile broadband subscription may buy or sell mobile data according to their demand and capacity.

Open sharing systems are supported by some major mobile broadband operators, like China Mobile Hong Kong [40] and SK Telecom Korea [23], which allow their subscribers to freely resell unused mobile data. Similarly to (closed) family sharing plans, these offers help operators attract/retain customers by allowing them better control and ownership of their packages.

For the customers of the operator that do not enable sharing/reselling to others, there is an option of sharing their unused caps via tethering to users in their close proximity. Several applications, such as Airmobs [19] or Hotspotio [14], allow users to share their data for capacity credit or other type of rewards.

In the rest of this section we first study the open sharing where users are allowed to trade their unused data freely via an operator-controlled market (in Section 5.1) and then explore the potential of the open sharing under the proximity constraint in Section 5.2.

5.1 Telco-assisted open sharing

5.1.1 System model

As we mentioned above, our goal here is to understand the potential of the mobile broadband sharing in which the users are free to resell their unused cap for a fee. We denote with \hat{p} the price per *Mbyte* of mobile broadband on such secondary market. In the case of telco-controlled market, \hat{p} can be either controlled by the operator or by the free market based on the supply and the demand. As we will see, the value of \hat{p} , has a critical impact on the user plan selection, and hence on the demand D and supply S of the mobile broadband on the secondary market.

For a user selecting plan i and with data consumption of d (in *Mbyte*), the total cost of their data plan at the end of the month is:

$$C_{TS} = \begin{cases} \begin{aligned} &price_i \\ &+ \max(0, d - cap_i)\mu_i(1 - b) \\ &+ \max(0, d - cap_i)\hat{p} \cdot b \\ &- \max(0, cap_i - d)\hat{p} \cdot s \end{aligned} & \text{if } \hat{p} < \mu_i \\ \begin{aligned} &price_i + \max(0, d - cap_i)\mu_i \\ &- \max(0, cap_i - d)\hat{p} \cdot s \end{aligned} & \text{if } \hat{p} \geq \mu_i \end{cases} \quad (5)$$

with s we denote the fraction of the unused capacity that the user could sell to others in the trading system. Likewise, b denotes the fraction of the capacity over the cap that the user could acquire from others. These fractions depend on the demand and the supply on the secondary market. If the supply is greater than the demand, $s = 1$, otherwise $b = 1$.

Given the price \hat{p} on the secondary market, and supply/demand parameters, s and b , users can select the plan according to the rules described in Section 3 to minimize their monthly cost. In the rest of the section we use the optimal-expectation rule: users select the plan i_{OER}^{TS} as following

$$i_{OER}^{TS} = \arg \min_{i \in \mathcal{P}} \left(\frac{1}{t-1} \sum_{\tau=1}^{t-1} C(d_u(\tau), i, \hat{p}, s, b) \right). \quad (6)$$

We define of the supply S in the open-sharing market to be the sum of unused caps across all the customers participating in the market:

$$S(\hat{p}, s, b) = \sum_u (\max(0, cap_{c_u} - d_u))$$

Similarly, the demand D of in the open-sharing market is given by the sum of demand on top of the purchased cap across all users:

$$D(\hat{p}, s, b) = \sum_u (\max(0, d_u - cap_{c_u})) \text{ where } \mu_{c_u} > \hat{p}$$

Note that a user under a plan i with a μ_i smaller than \hat{p} does not contribute to D , as these users would obtain their over-the-cap capacity directly from the operator. The overall amount of traffic users can share in this market is simply $\min(S, D)$. We distribute equally the amount of capacity exchanged among all users; therefore the fraction of unused/extra traffic users can sell/buy on the sharing market will be respectively:

$$s = \frac{\min(S(\hat{p}, s, b), D(\hat{p}, s, b))}{S(\hat{p}, s, b)} \quad (7)$$

$$b = \frac{\min(S(\hat{p}, s, b), D(\hat{p}, s, b))}{D(\hat{p}, s, b)} \quad (8)$$

Users participating in the market make individual decisions on which package to purchase based on their consumption, the value of \hat{p} , and the parameters s and b (optimizing Eq. (6)). On the other hand, such decisions determine supply and demand. Hence for a given \hat{p} , the supply and the demand would self-regulate to stable values $s(\hat{p})$ and $b(\hat{p})$, which solve the system of equations (7)-(8). In general, the smaller \hat{p} result in larger demand (as user purchase smaller packages), while for larger \hat{p} supply on the secondary market dominates the demand; see Fig. 8 (top).

The final charge per user is then

$$h_u^{OER, TS}(t) = C(d_u(t), i_{OER}^{TS}, \hat{p}, b(\hat{p}), s(\hat{p}))$$

5.1.2 Benefits of trading systems

We use the equilibria supply/demand parameters ($s = s(\hat{p})$ and $b = b(\hat{p})$) to estimate the impact of \hat{p} on the cost of users and the revenues of the operator. By choosing the equilibria parameters, we are analyzing steady state scenarios in which users actually obtain from the system what they expect. Unstable scenarios can occur at the starting phases of the system and continue until users adapt their s and b




Figure 8: Supply and demand (top), average user savings (middle), and operator revenues ORR (bottom) for varying \hat{p} in Vodafone India tariffs.

to the steady ones⁷. Figure 8 shows the steady state supply ($S(\hat{p})$) and demand ($D(\hat{p})$) for a range of \hat{p} for Vodafone India tariffs in the last month of our dataset. Also, we report the average savings per user and the revenues of the operator. The reported revenues of the operator are relative to the earnings without trading system:

$$ORR(\hat{p}) = \frac{\sum_u h_u^{OER,TS}}{\sum_u h_u^{OER}}$$

When \hat{p} is equal to the penalty rate (in case of Vodafone India⁸ it is $\mu_0 = 2$), each user makes the decision to purchase a cap, which is the same as in the case where there is no secondary market. Hence the $\sim 19\%$ savings that happen when $\hat{p} = \mu_0$ come not from users purchasing smaller packages but rather exclusively from sharing the packages they would purchase if no sharing was enabled (by optimizing eq. (3)). We report the average per-user savings and the overall operator revenue ORR , for $\hat{p} = \mu_0$ for all of the 10 studied operators in Table 2. We observe that the ORR is greater than 80% in all but one operator (Vodafone UK). This means that if the Telco allows sharing through their billing system, but controls the price to be the penalty rate, the impact on its revenues can be somewhat controlled and kept relatively low.

In Vodafone India, a \hat{p} slightly over 0.1282 drives the system to the largest average user savings, which is around 47%. We denote that optimal point with \hat{p}^* . This \hat{p}^* is also the one in which $S(\hat{p}^*) = D(\hat{p}^*)$ (i.e. where the supply and demand are identical) and where the relative revenues of the operator are the lowest (near 36%). We report the average per-user-

⁷An example of an unstable scenario is when most users select the plan with the lowest cap (usually the cheapest), probably expecting to acquire cheap capacity in the secondary market for most of their demand. This, however, creates a buyer's market, in which capacity is not available in the sharing system; thus, compelling users to obtain capacity under the penalty rate of the mobile operator.

⁸We note that each time we refer to a particular telco, e.g. Vodafone India, we actually refer to their tariffs not the telco itself.

savings and relative operator revenue for other 9 operators in Table 2 for $\hat{p} = \hat{p}^*$. From the table, we observe how users under tariffs with high penalty rates, such as Vodafone India and Etisalat UAE, receive large benefits from the trading system while the telco revenues are strongly impacted by it. On the other hand, for ATT and Olleh South Korea plans, which possess plans with relative low penalty rates, operators' suffer a low impact on their revenues even under the \hat{p}^* . Additionally, the users in such low-penalty tariffs experience the least amount of savings from the trading system. China Unicom is again an exception for our analysis. The tariff system offered by China Unicom is somewhat atypical with commit rate being cheapest in the second smallest package⁹, which creates some unusual dynamics.

Regarding the relative savings per user, the ECDF of per user savings with Vodafone India tariffs for the case $\hat{p} = \hat{p}^*$ is depicted in Figure 9. From the figure we observe that per-user savings are widely distributed with a peak around 38%, which corresponds to a large fraction of users which move from one package to another. In Figure 10, we plot the average relative savings for several groups of customers based on their average consumption for $\hat{p} = \hat{p}^*$. The ways savings are distributed among the customers is very tariff-dependent, and indeed differ from one operator to the other. In general, the heavier the customer is, the larger is her expected savings, even though in some tariffs this is not the case (Vodafone UK and China Unicom). As a large percentage of users have a relative low consumption (less than 200MB), the savings from these users drive, in a large percentage, the global benefits provided by the collaborative system.




Figure 9: ECDF of relative savings per user with Vodafone India tariffs ($\hat{p} = \hat{p}^* = 0.128$).

5.2 Open sharing via tethering

If the network operator does not enable reselling of the unused data caps, the users can still share via short-range tethering with others in their close proximity. Such proximity constraint may significantly limit the potential of sharing in the rural areas, where the density of mobile users may not be enough to allow meaningful cap sharing. However, in the dense urban environments, the opportunities of finding a buyer/seller match are much more likely. In this section we aim to quantify the potential of the open sharing model under the proximity constraint.

⁹In all other telcos, the lowest commit rate is in the largest package.




Figure 10: Average savings for users under different consumption intervals for $\hat{p} = \hat{p}^*$.

Operator	ORR $\hat{p} = \hat{p}^*$	Avg. user savings $\hat{p} = \hat{p}^*$	ORR $\hat{p} = \mu_0$	Avg. user savings $\hat{p} = \mu_0$
Vodafone India	36%	47%	81%	21%
Telstra Australia	51%	33%	84%	14%
Vivo Brazil	58%	28%	89%	3.9%
Claro Argentina	68%	22%	95%	2.1%
Olleh South Korea	80%	14%	94%	4.5%
China Unicom	81%	54%	81%	54%
E-Plus Germany	72%	22%	83%	14%
Vodafone UK	51%	47%	68%	32%
Etisalat UAE	29%	53%	93%	5.9%
ATT US	79%	17%	84%	13%

Table 2: Lowest average user cost and lowest relative revenue found in stable systems ($\delta = \hat{\delta} = \delta_e$) for all operators.

The first question we ask is how many cellular users can one expect to communicate with over a short-range channel. To answer that question, we use the census data from the country of the operator we study. For each ward, the census data reports the area (in km^2) and number of people per ward from which we can calculate the density of people per km^2 . To estimate the number of people a cellular user can communicate with over a short-range channel, we calculate the expected number of people in the circle with $20m$ radius for each ward. The $20m$ radius was taken as a standard WiFi indoor range. Different technologies would indeed have different ranges, but for the purpose of quick first order estimate, we use the $20m$ range. The CDF of the number of neighbors within $20m$ range is depicted in Figure 11. In this particular country, a cellular user has a median of 3.18, and an average of 4.11 other cellular users in their $20m$ radius. Note that these numbers are rather conservative in that they assume that humans are spread uniformly in the area covered by the ward. In practice both in urban, and especially in rural, areas a large fraction of space is non-occupied (e.g. parks, highways or agriculture land), and even in the populated land, humans tend to be clustered; hence the denominator area in the density calculation is likely much smaller, and the number of proximal neighbors is likely to be larger than what the CDF in Fig. 11 would suggest.

Next, we study the relationship between the density of an area (represented by the expected number of users a user can communicate with) and the potential savings users may expect in the open sharing among their neighbors. For this analysis, we let users select their cap based on their historical consumption, optimizing their bill as if full market liquidity was in place with various \hat{p} , but allow them to trade capacity with other users in a near range of them. The number of users within the tethering range is defined by a parameter η we call *density*, which we vary between 0 and 20 following the official census data report. To understand the impact of *density* on the sharing, we employ the following procedure. We take $N = 40000$ users for which we have a detailed history of data sessions described in Section 2. Time is slotted in seconds, and each user tracks the total consumption and spare capacity in the billing circle. At the time of each data session with volume V , user u checks its total consumption. If the total consumption is smaller than the capacity of her package, she uses it; otherwise she picks randomly $\eta = \text{density}$ other users as neighbors. In case any of them has a spare capacity greater than V , the one with the largest spare capacity responds by sharing it with the user u and adjust its own consumption by $+V$. In case none of the neighbors has enough spare capacity, user u purchases it on the primary market from the operator paying the penalty rate. Users decide on which package to purchase optimizing Eq. (6) for various \hat{p} . Varying \hat{p} , allows to control the supply and the demand of the secondary market.

For each of the ten operators and each density value, we experiment with several \hat{p} values and report the maximal *total relative savings* in Figure 12. For low densities, in the optimal state the supply is much larger than the demand of the secondary market, allowing users which have an excessive demand to find a supplier with high likelihood. In the case of high densities, the system converges to the case with full liquidity studied in Section 5.1 where in the optimal state the supply of the secondary market is equal to the demand.

From Figure 12, we can observe that the savings of users in environments with average density population (4 people) are in the range of 9-25%. For the highly populated areas, with 20 people per tethering circle, the relative savings increase to 18-48%. The results are encouraging and complementary, from an economical perspective, previous endeavors of measuring technical benefits of collaborative systems for the mobile users [18].

We conclude this section with the remark that our analysis suggests that in some cases the operator may be better off (have higher *ORR*) by embracing the sharing economy and assisting it though its own billing system while controlling the price on the secondary market than allowing the users to create their own secondary market via tethering. For example, the operator relative revenues *ORR* of Claro Argentina is 0.95 when \hat{p} equals the penalty rate μ_0 . However, when users can choose their own \hat{p} , even with the constraint of tethering they can save more than 5% of their bills as long as they have sufficient density of sharing neighbors; in dense areas with 20 neighbors the revenues of the operator could be reduced by as much as 19%. Our results indicate the existence of a fine tradeoff between the tariff structure, the population density and the potential of the sharing economy with and without Telco assistance.




Figure 11: Expected number of cellular users in 20m radius.




Figure 12: Total user savings (1-ORR) for varying density values.

6. RELATED

Our work is related to pricing and economics of mobile broadband. Recent work has looked into the merits and demerits of usage based pricing (UBP) in cellular networks [2, 20, 26]. Likewise, authors of TUBE [12] suggest time dependent pricing can help in reducing peak-hour congestion. Authors of [30, 31] investigate the existing 3G/4G billing systems and demonstrate a number of loopholes in them. Our contribution is to investigate how volatility (hence, lack of predictability) in demand can lead to sub-optimal decisions in terms of choosing plans. A closely related recent work investigated the notion of ‘irrationality’ of users in choosing mobile broadband plans – users often pay more than they consume [28] for different services. Our work builds up on the findings to investigate variability in demand and suggest how one can pick close to optimal plans.

A number of authors have analyzed the sharing in cellular networks. In [29] authors suggest sharing unused the minutes and sms and designed a prototype of the system that enables such sharing. In [15, 21] authors apply game theoretic tools to study crowdsourced architectures (including collaborative consumption) for mobile data access. Authors of [24] propose *kibbutz*, a system that leverages mobile link sharing to improve the energy consumption and connection performance of users. We complement these works by providing quantitative insights on the viability and the impact of collaborative consumption in the mobile data access.

There has been recent surge of services that enable collaborative consumption of various resources including apartments [1], cars [13], books [4], etc. Results of our work here suggest that collaboration between end-users can be beneficial for users of mobile broadband, providing economic incentives for technical solutions like Airmobs [19] or Hotspots [14].

7. DISCUSSION

In this section, we briefly discuss some of the issues that may influence our results.

A major assumption that we rely on is that the user demand is independent of the pricing signal and market structure. This is a fairly strong assumption, since it is well known that pricing signal may have a rather important impact on how much demand one generates. However, we believe that the demand data we use here offers a useful input for a first-order approximation of the demand mobile broadband users put on the network. The effects of the pricing signal on the mobile broadband remain to be studied in the future.

As stated earlier, our analysis is performed assuming pre-paid plans, which are the norm in large parts of the world [11]. Having said that, long-term contracts are popular as well and we intend to extend our analysis to such plans. Likewise, mobile broadband prices could vary depending on the technology mobile technology employed (4G/3G) or if they are bundled with other mobile services (e.g. voice or sms). These options can play a role in the choice of a mobile broadband plan, but we do not incorporate such externalities in our analysis. However we expect that mobile communications will follow the trend seen in fixed-line communications where residential broadband has replaced a large part of traditional fixed-line business (voice telephony).

Our analysis is purely economic. In the context of opportunistic traffic sharing (via tethering), such sharing may impact the end-to-end performance (e.g. increased bandwidth from multiple radios [36]), quality of experience (e.g. TCP issues with multihoming [33]), network coverage (e.g. users in close proximity may have different cellular signal strengths) and energy consumption [24]. These considerations are out of scope of this paper.

We believe our analysis can form the basis for designing incentives for risk management with and without user collaboration. Various solution (cost/revenue-sharing) concepts¹⁰ from cooperative game theory as well as non-cooperative game theory could be of great use in designing such systems.

In the context of open sharing, an important element for enabling such system are micropayments. Our largest motivation for this type of systems comes from innovative Asian operators that facilitate the exchange of data using in-house applications [23, 40]. These solutions should include the necessary features to facilitate safe transactions among users. On the other hand, applications could also leverage tethering to provide a similar service [19, 29]. These applications should be enhanced to support micropayment without compromising the security and privacy of users. Solutions along the lines of [39] could be used to eliminate the effect of possible fraud in such a p2p micropayment system.

In this paper, we assume that operators do not react to the appearance of collaborative systems. Although we used a wide range of existing pricing plans to explore different scenarios, finding the effects of long-term strategies from mobile providers under this environment is still an open problem. Operators could, for instance, try to tune their penalty rates or cap structure to discourage the use of collaborative systems, while also minimizing revenue impact. Also, operators with control of the collaborative system could try to find a secondary market price that maximizes its revenues. Since

¹⁰Eg. Shapley value, Nash bargaining, core, nucleus, etc.

the domain space of possible strategies and pricing plans is enormous, studying these cases in a general way is very challenging (either analytically or empirically) and forms an interesting line for future research.

Recently, it has been recognized that time and the location are two important dimensions, which may influence the value of mobile data [12]. For example a crowded cell in the peak hour may price mobile data higher than a sparse cell over off-peak hours. Such time/location based pricing may improve the efficiency of cellular networks, albeit most existing cellular operators do not distinguish the price of the mobile data based on time nor location. Incorporating time/location in the models we study would be an interesting direction for future research.

8. CONCLUSIONS

Collaborative consumption is an exciting new trend of sharing various goods or services for economic and social benefits. In the context of mobile broadband, collaborative consumption can either be user-driven (e.g. via tethering) or Telco-driven (integrated through the Telco's billing system). We quantified the economic impact collaborative consumption could have both on the Telco and the customers. The methodology we develop can assist the operators to design the tariffs and/or control the secondary market to optimize its revenues in the presence of mobile broadband sharing.

9. ACKNOWLEDGEMENTS

This work has benefitted from a number of discussions we had with Vijay Erramilli, whose input significantly influenced the early versions of this paper.

10. REFERENCES

- [1] Airbnb. <http://airbnb.com/>.
- [2] J. Blackburn, R. Stanojevic, V. Erramilli, A. Iamnitchi and K. Papagiannaki. "Last call for the buffet: economics of cellular networks". In *Proc. of MobiCom 2013*.
- [3] L. Bravo. "Aqualogy". Personal communication. 2013.
- [4] Bookmooch. <http://bookmooch.com/>.
- [5] M. Chetty, R. Banks, A.J. Bernheim Brush, J. Donner, R. E. Grinte. "You're Capped! Understanding the Effects of Bandwidth Caps on Broadband Use in the Home". In *Proc. of ACM CHI 2012*.
- [6] S. DellaVigna. "Psychology and economics: Evidence from the field". *Journal of Economic Literature*, 47(2), 2009.
- [7] N. Do et al. "CrowdMac: A Crowdsourcing System for Mobile Access". In *Proc. of the 13th International Middleware Conference 2012*.
- [8] K. Fitchard. Att and Verizon shared data plan. <http://gigaom.com/2012/07/18/att-new-shared-smartphone-data-plans-undercut-verizon/>. GigaOM, 2012.
- [9] FON. <http://corp.fon.com/en>.
- [10] T. Gilovich, D. W. Griffin, D. Kahneman. "Heuristics and Biases: The Psychology of Intuitive Judgement". Cambridge university press 2002.
- [11] Google. International Broadband Pricing Study: Dataset for public use. <http://policybythenumbers.blogspot.com/2012/08/international-broadband-pricing-study.html>
- [12] S. Ha et al. "TUBE: time-dependent pricing for mobile data". In *Proceedings of ACM SIGCOMM 2012*
- [13] M. Hart, M. J. Roberts, J. D. Stevens. "Zipcar: refining the business model". Harvard Business School Pub., 2005.
- [14] Hotspotio. <http://hotspotio.com/>
- [15] G. Iosifidis et al. "Enabling Crowd-Sourced Mobile Internet Access". *Proc. of Infocom 2014*.
- [16] A. Lakhina, K. Papagiannaki, M. Crovella, C. Diot, E. D. Kolaczyk, and N. Taft. "Structural analysis of network traffic flows" In *Proc. of ACM SIGMETRICS'04*.
- [17] N. Laoutaris et al. "Delay tolerant bulk data transfers on the Internet". In *Proc. of ACM SIGMETRICS'09*.
- [18] S. Liu and A. Striegel. "Exploring the potential in practice for opportunistic networks amongst smart mobile devices". In *Proc. of MobiCom 2013*.
- [19] A. Lippman, H. Holtzman and E. Toledano. "Air Mobs". MIT Media Lab, 2013
- [20] D. Lyons. "The Impact of Data Caps and Other Forms of Usage-Based Pricing for Broadband Access". Preprint 2012.
- [21] M. H. Manshaei et al. "On Wireless Social Community Networks". In *Proc. of INFOCOM 2008*.
- [22] M. Muraven and R. F. Baumeister. "Self-regulation and depletion of limited resources: Does self-control resemble a muscle?." *Psychological bulletin* 126.2. 2000.
- [23] Money Today. "Smart phones data business, 1GB price". <http://news.mt.co.kr/mtview.php?no=2013020713333217104>
- [24] C. Nicutar, D. Niculescu, C. Raiciu. "Using Cooperation for Low Power Low Latency Cellular Connectivity" In *Proc. of CoNEXT'14*.
- [25] A. Odlyzko. "Pricing and architecture of the internet: Historical perspectives from telecommunications and transportation". In *TPRC 2004*.
- [26] A. Odlyzko et al. "Know Your Limits: Considering the Role of Data Caps and Usage Based Billing in Internet Access Service". Public Knowledge 2012.
- [27] R. Stanojevic, N. Laoutaris, and P. Rodriguez. "On economic heavy hitters: shapley value analysis of 95th-percentile pricing". In *Proc. of ACM IMC '10*.
- [28] R. Stanojevic, V. Erramilli, K. Papagiannaki. "Cognitive bias in network services". In *Proc. of ACM HotNets 2012*.
- [29] H. Pan et al. Sharing airtime with Shair avoids wasting time and money. In *Proc. of ACM HotMobile 2009*
- [30] C. Peng et al. "Can we pay for what we get in 3G data access?". In *Proc. of ACM MOBICOM, 2012*.
- [31] C. Peng et al. "Mobile data charging: new attacks and countermeasures". In *Proc. of ACM CCS 2012*.
- [32] M. Roughan et al. "Experience in measuring backbone traffic variability: Models, metrics, measurements and meaning." *Proc. of the 2nd ACM SIGCOMM Workshop on Internet measurement*. ACM, 2002.
- [33] C. Rossi et al. "3GOL: Power-boosting ADSL using 3G OnLoading" In *Proc. of CoNEXT, 2013*.
- [34] G. Shaw. "Telus announces new two-year talk and text, data plans." <http://blogs.edmontonjournal>.

com/2013/07/15/telus-announces-new-two-year-talk-and-text-data-plans/

- [35] S. Soumya et al. “Pricing Data: A Look at Past Proposals, Current Plans, and Future Trends”. Preprint 2012.
- [36] C. Tsao, R. Sivakumar “On effectively exploiting multiple wireless interfaces in mobile hosts”. CoNEXT, 2009.
- [37] Telco 2.0 Initiative. “European Mobile: The Future’s not Bright, it’s Brutal”. http://www.telco2research.com/articles/EB_European-mobile-future-brutal-telco_Summary, 2012.
- [38] 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.
- [39] B. Yang, H. Garcia-Molina. “PPay: micropayments for peer-to-peer systems”. In Proc. of ACM CCS 2003.
- [40] ZDNet. “China Mobile Hong Kong subscribers can sell unused data capacity”. <http://www.zdnet.com/china-mobile-hong-kong-subscribers-can-sell-unused-data-capacity-7000024541/>
- [41] Vodafone India, tariffs. https://www.vodafone.in/3gworld/pages/tariff_plans.aspx?cid=del
- [42] Vodafone UK, tariffs. <http://www.vodafone.co.uk/shop/internet/dongle/index.htm>
- [43] E-Plus Germany, tariffs. <http://www.eplus.de/Tarife/E-Plus-internet-flats/E-Plus-internet-flats.asp?WT.ac=pk/Tarife/Internet-Flats/T2/130401>
- [44] Telstra, Australia, tariffs. <http://www.telstra.com.au/internet/mobile-broadband/prepaid/rates/>
- [45] Claro, Argentina, tariffs. <http://claro.com.ar/wps/portal/car/pc/personas/movil/planes-claro/planes-claro-control>
- [46] Olleh, South Korea, tariffs. <http://shop.olleh.com/weblogic/EventServlet?pname=EventLTE>
- [47] Etisalat, UAE, tariffs. <http://www.etisalat.ae/eportal/en/personal/mobile-mydata-packages.html>
- [48] ATT, US, tariffs. <http://www.att.com/shop/wireless/data-plans.html#fbid=aE3VuzRgpe1>

APPENDIX

A. GROUPING BY LOCATION OR SOCIAL RELATIONSHIP

As described in Section 4.1, users can form collaboration groups based on their location or social relationships. Our dataset includes additional data from all users, such as billing information or call registries, which we can leverage to build these types of groups. In this section, we show that the benefits of these groups are similar to the ones from random groups, which we analyzed previously.

To form groups based on location, we use the residence information of our dataset. Users living in the same postal code¹¹ are placed in the same location group. On the other hand, we use the call registry from the users to build social relationships groups. The call registry allows us to find pairs

¹¹In the country we analyzed, postal codes have very high granularity, with around 35 people in average per postal-code.




Figure 13: Number of groups of different sizes for social and location groups.




Figure 14: Average cost of random grouping vs. social and location grouping. The value of coordination (left). The impact of group size (right).

of users who have called each other at least once. We refer to these pairs of users as peers. We create the social group of every user by combining all his/her peers. Figure 13 shows a histogram of group size for social and location groups we obtained from the 40K users in our dataset. Note that the y -axis is in the log-scale: the number of groups per group size decreases rapidly for each group selection type.

Figure 14 (left) depicts the cost of the coordinated-OER rule for localization and social groups relative to the cost obtained for random groups. We observe how the benefits obtained by group size are similar to the ones described in Section 4.2. We show the results until groups of size five, as we could not obtain a statistically significant number of groups with larger sizes. In Figure 14 (right) we compare the costs for users when coordinated and non-coordinated tariffs are chosen relative to the ratios for random groups. Similarly, the results agree with the analysis of random groups that we describe in detail in Section 4.3.