

# Internet Demand under Different Pricing Schemes\*

Jörn Altmann      Björn Rupp      Pravin Varaiya

Department of Electrical Engineering & Computer Sciences  
University of California at Berkeley

## Abstract

The proliferation of electronic commerce has emphasized the need for multiple qualities of service (QoS) that meet user requirements for, e.g., best-effort services, transactional guarantees, prioritized services, and bandwidth reservations. When Internet users will be able to choose from a large selection of network services, an important question is whether this flexibility will be appreciated by the user and how much value is placed on network resources. This paper focuses on user acceptance of different usage-based pricing models for network services as well as on which users benefit from dynamic pricing. The analytic results presented here are based on pricing experiments conducted within the INternet Demand EXperiment (INDEX). The pricing experiments comprise per-minute and byte-volume pricing. After describing the experimental design and characterizing our subject pool using demographic data, we analyze user demand for Internet access by comparing the three experiments in terms of traffic generation, QoS selection, and expenditure. Finally, we show which user groups will benefit from such services.

## 1 Objectives

Apart from offering Virtual Private Networks to corporations, Internet Service Providers (ISPs) presently do not segment their market in regard to network services. All users experience the same service quality regardless of their usage profile. However, there are numerous alternatives that are both technologically and economically viable. Different pricing structures [3] can meet requirements of different user groups, resulting in a higher overall value of the network. In this paper, we present an analysis of aggregate data from the INDEX Project that support these claims and show how different pricing schemes affect user behavior.

INDEX, the INternet Demand EXperiment, is a market trial for quality-differentiated Internet services. INDEX provides Internet access over ISDN lines to a group of users from the Berkeley campus community. Users select network services from a menu of QoS-price offerings and pay for their

usage. An overview of the technology, experimental setup and design of INDEX can be found in [2]. INDEX' long-term objectives are to identify the key parameters for user perception of service quality, quantify the correlation between application type and service demand, and measure the economic value that individual users place on distinct activities and network resources. This information will help to define market segments for quality-differentiated network services that accommodate future demand and applications with different requirements.

We analyze three experiments. The first experiment, *Variable Symmetric Bandwidth*, gives users the choice between six different bandwidths at different prices. The data is used to estimate the price elasticity of demand for connection speed to the Internet. In the *Variable Asymmetric Bandwidth* experiment, in-bound and out-bound bandwidth can be chosen separately. Data from this experiment help to determine whether users value bandwidth for incoming traffic more than for outgoing traffic, and also reveal differences in effective line utilization. In the *Byte Volume* experiment, users face per-byte charges. Analysis of aggregate data from these three experiments yields insights into how different pricing schemes influence user behavior and utilization of network resources.

The remainder of the paper is organized as follows: The second section characterizes the subject population and illustrates design issues of the experiments conducted. The third section deals with the analysis of the aggregated data. After presenting some measurement data, we compare the experiments in terms of the average number of bytes transmitted and the amount of time our subjects were connected to the Internet. This comparison allows us to infer the overall acceptance of usage-based pricing models for Internet services. Finally, we show the extent of user segmentation and document the benefits to be gained from flexible pricing of network services.

## 2 Experiment Descriptions

### 2.1 Population

So far, INDEX has recruited 80 subjects affiliated with the University of California at Berkeley (students, faculty, staff). As of March 1999, there is data on the first 65 subjects recruited for the project, all from the UC Berkeley campus community. The data covers the period from April 1998 to December 1998. All subjects go through the same sequence of experiments. Because not all of the subjects joined INDEX at the same time, some of them had not yet completed

\*This research was supported by grants from the National Science Foundation, Cisco Systems, SBC Communications, the California State MICRO Program and Hewlett-Packard.

all experiments when the analyses for this paper were conducted. The number of users that we report below for the earlier experiments is therefore moderately higher than for the later experiments.

INDEX recruitment procedures aim to obtain a suitable variation in field of study, expected Internet usage, computer sophistication, travel distance to campus, and general demographic characteristics. In order to compensate for sampling bias, we undertook a detailed demographic survey of our subject pool. Many of its questions are taken from a representative, population-projectable study conducted by Nielsen Media Research [1]. Once the entire subject pool has been recruited, an extrapolation model will be calibrated in order to generalize findings from the INDEX Project’s demographic base to the general U.S. population.

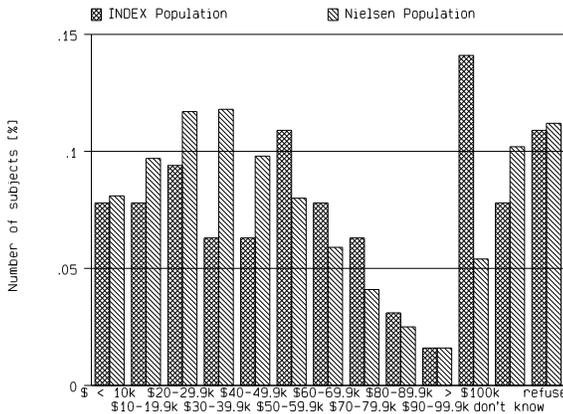


Figure 1: Demographics: Income Distribution

The demographic data permits a general characterization of the current subject pool. As expected, the subjects have a high level of education: 40 % have a post-graduate degree, 16 % some post-graduate school experience and 20 % a bachelor’s degree. 19 % are currently in college or have some college experience. High School represents the highest level of school for only 5 % of the subjects. Most INDEX subjects are also experienced Internet/computer users: 90 % first used the Internet three or more years ago (compared to only 9 % in the Nielsen survey), and 86 % (Nielsen 49 %) have been computer users for at least five years. These numbers, combined with other data on computer and Internet usage, suggest that the current INDEX subject population is more representative of tomorrow’s user generation than of the present general U.S. population. However, the general demographic characteristics of our subject pool span a wide range of attribute values. The income distribution is remarkably widespread, as can be seen in Figure 1. Significant differences between INDEX data and Nielsen data can only be found in the income ranges from \$30k to \$39.9k, \$40k to \$49.9k and > \$100k. The subjects’ age distribution ranges from 20 to 72 years with an average age of 35 and a median of 29 years.

## 2.2 Design Details

Whereas non-experimental studies are forced to rely on cross-sectional variation in price and demand to infer the price elasticity, price is varied during INDEX experiments to measure the demand response for each individual. All

experiments are preceded by a free trial week that provides reference data and allows the subjects to become familiar with the pricing structure.

### 2.2.1 Variable Symmetric Bandwidth Pricing

The first experiment faced by all subjects, *Variable Symmetric Bandwidth*, allows users to choose between six different access link bandwidths (8 kbps, 16 kbps, 32 kbps, 64 kbps, 96 kbps, 128 kbps). The selection can be changed instantaneously, even during an active session. Subjects are charged a per-minute rate that depends on the selected connection speed. Prices increase with bandwidth. The first five weeks of this six-week long experiment feature weekly changes of the price schedule. During the last week of this experiment, prices change daily.

The 8 kbps service is priced at zero cents per minute. The price for each bandwidth above 8 kbps is the sum of the price for the next lower bandwidth and an independent random price increment. The exact algorithm for generating price increments was designed to produce a vector of strictly increasing prices:  $F(p) = p_{min} + \gamma \cdot \alpha p^\theta$  with  $p_{min} = 0.1$ ,  $\alpha = 2.8$ ,  $\theta = \frac{10}{3}$  where the factor  $\gamma$  equals 0.6 for the increment from 8 kbps to 16 kbps, 0.7 for the increment from 16 to 32 kbps and 2.0 for all other increments.  $p$  is a long-period uniform random deviate between 0 and 1 exclusive of the endpoint values. Setting  $\theta$  at  $\frac{10}{3}$  produces a distribution with most mass farther away from  $p_{min}$  and a shorter right hand tail. This formula yields a distribution of prices from a minimum of 0.1 c/min for 16 kbps service to a theoretical maximum of 20.94 c/min for 128 kbps service with smoothness in the price intervals. It provides low-priced alternatives for low-speed choices while still allowing exploration of the high-price range.

### 2.2.2 Variable Asymmetric Bandwidth Pricing

The *Variable Asymmetric Bandwidth* experiment differs from the first experiment in only one aspect: Subjects can choose different bandwidths for traffic from the Internet and to the Internet separately. This experiment is motivated by access technologies with different data rates for down-stream and up-stream traffic (e.g. ADSL and CATV). In addition, the experiment seeks to determine whether individuals value bandwidth for in-bound traffic differently than bandwidth for out-bound traffic.

The pricing structure is based on the algorithm of the *Variable Symmetric Bandwidth* experiment. To ensure that a subject would have the same bill for the same exact service in the absence of behavioral changes, and so make the results from the two experiments directly comparable,  $\gamma$  is conceptually cut in half. The resulting prices are then applied separately to bandwidth for in-bound and out-bound traffic, i.e. the price for a given service class (8 kbps, 16 kbps, 32 kbps, 64 kbps, 96 kbps, 128 kbps) amounts to half of the price in the first experiment.

### 2.2.3 Byte Volume

In the *Byte Volume* experiment, subjects face per-byte charges that change weekly during the six weeks of the experiment. Some pricing proposals justify this pricing structure on the basis that service degradation occurs due to network traffic. This pricing structure also constitutes a departure from traditional telephone-oriented, per-minute

pricing to which individuals are accustomed. The experiment seeks to explore whether users understand the basis for such prices, and the price sensitivity of the amount of traffic generated by users.

The pricing structure is straightforward. Users have two options: Free 8 kbps service or 128 kbps service billed on a per-byte basis. Prices are varied during the experiment from 0.1 to 20 c/mbyte. This range is divided into three segments (low/medium/high). In order to make sure that a subject faces prices out of all three segments, two draws from a uniform distribution within each segment are performed. The resulting six prices are then permuted and presented to users in a random order.

### 3 Analysis

To investigate the acceptance of usage-based pricing, we analyze the change in user demand for Internet services under different pricing structures. In this context, some interesting insights can be drawn from our subjects' responses to some related questions from our demographics survey. On a scale of 1 to 5, with 1 meaning "strongly disagree" and 5 meaning "strongly agree", subjects could indicate how much they agreed or disagreed with the following statements: (1) "There are sufficient network resources *right now* so that Internet transmission quality is very good"; (2) "There will be sufficient network resources *in the future* so Internet transmission quality will be very good"; (3) "Available network resources should be distributed in such a way that the users *who value them most* should get the best quality, while users who value them less should get worse quality"; and (4) "Available network resources should be distributed in such a way that the users *who are willing to pay the most* should get the best quality, while users willing to pay less should get worse quality". The last two questions have been the subject of intense debate. Economists generally assume the terms "valuation" and "willingness to pay" to be identical; we split the issue into two separate questions to test whether this assumption would hold in this specific environment.

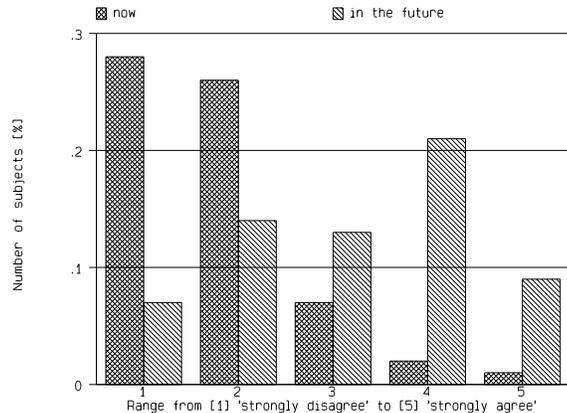


Figure 2: Demographics Survey: Sufficient network resources {right now/in the future} so that Internet transmission quality {is/will be} very good

As Figure 2 shows, almost all of our subjects agree that today's Internet provides insufficient network performance, but in their opinion the situation will take a turn for the

better in the future. As illustrated in Figure 3, the answers to questions (3) and (4) are less definite. However, the majority of the subjects did express their willingness to accept changes in the way Internet resources are distributed. Looking at deviations among the two questions, we did indeed find differences. There seems to be a slightly more visible reluctance towards "valuation"-based network resource allocation when it is taken to its final consequence of being made financially accountable. Because a survey can only collect data on stated preferences, these results can however be affected by strategic misrepresentation or justification bias. Therefore, the differences in the distribution of acceptance patterns do not seem significant enough to permit a more general conclusion. Although more subjects seem to believe that user valuation can allocate resources better than a performance-dependent service charge, 41 % still were more inclined towards approval of price-based resource allocation (vs. 39 % that to some extent disagreed).

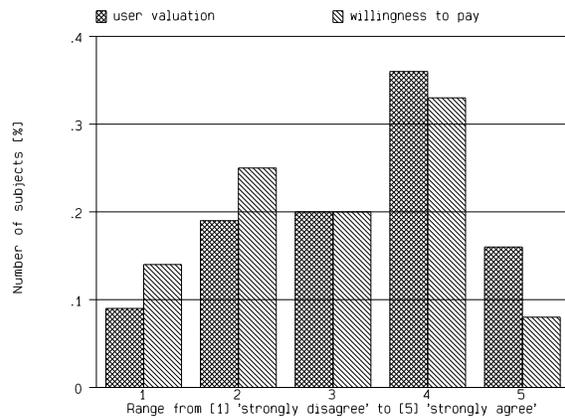


Figure 3: Demographics Survey: Approval of allocation mechanisms for distribution of network resources

Thus, although our subjects have not yet experienced service-dependent prices when taking the survey, they are already cautiously inclined towards such a potential solution for today's Internet congestion problems. The revealed preference data presented in the next section will show the subjects' acceptance of usage-based pricing in greater detail.

#### 3.1 Acceptance of Experiments

The subjects' response to the offered service prices is reflected in the average cost per megabyte and the average cost per minute they are willing to pay. The average cost per megabyte is defined as the ratio of total expenditure to number of bytes transmitted at priced services. The average cost per minute is the ratio of total expenditure to the number of minutes connected to the Internet at priced services. It should be noted that the aggregate data presented here hides the large variation in prices faced by individual users and, consequently, the large variation in their demand.

| experiment | cost per mb | expenditure | bytes sent |
|------------|-------------|-------------|------------|
| sym. bw.   | 17.9 c      | 160151 c    | 8927.9 mb  |
| asym. bw.  | 17.1 c      | 111202 c    | 6482.2 mb  |
| byte vol.  | 8.6 c       | 69344 c     | 7994.2 mb  |

Table 1: Cost per megabyte

| experiment | cost per minute | expenditure | connect time |
|------------|-----------------|-------------|--------------|
| sym. bw.   | 1.1 c           | 160151 c    | 8430143 sec  |
| asym. bw.  | 1.2 c           | 111202 c    | 5531777 sec  |
| byte vol.  | 0.2 c           | 69344 c     | 19936876 sec |

Table 2: Cost per minute

As seen in Table 1 and Table 2, the average cost per minute of being online and the resulting derived average cost per megabyte are nearly equal for the *Variable Symmetric Bandwidth* and the *Variable Asymmetric Bandwidth* experiments. There is a significant difference between the per-minute experiments and the per-byte experiment, because the average offered price in the *Byte Volume* experiment is lower.

The expenditures and connect time for different experiments cannot be compared directly. Therefore, we normalize the averages by dividing the measured values by the number of *user-days*. A *user-day* is defined as a day during which the user was actually using the service. Table 3 shows the number of users and *user-days* in each experiment.

| experiment     | free t.p. | sym.bw. | asym.bw. | byte v. |
|----------------|-----------|---------|----------|---------|
| # of user-days | 843       | 2082    | 1451     | 1201    |
| # of users     | 63        | 64      | 55       | 51      |

Table 3: Number of user-days

### 3.2 Demand Change under Usage-Based Pricing

The change in demand under usage-based pricing is an indication of preferred pricing schemes. The first change in user demand is visible in connection utilization, the percentage of purchased connection capacity that is actually used. Connection capacity is defined as the amount of bytes which could have been sent by a user fully utilizing all purchased bandwidths. Table 4 displays the connection capacity, the number of bytes actually sent, as well as the connection utilization for the free trial periods and the three experiments.

| experiment | utilization | capacity | in-b byte | out-b byte |
|------------|-------------|----------|-----------|------------|
| free t. p. | 2.0 %       | 5.3e+11  | 9.7e+09   | 1.0e+09    |
| sym. bw.   | 7.5 %       | 1.1e+11  | 7.4e+09   | 1.5e+09    |
| asym. bw.  | 10.4 %      | 6.2e+10  | 6.0e+09   | 4.4e+08    |
| byte vol.  | 1.2 %       | 6.3e+11  | 7.0e+09   | 9.5e+08    |

Table 4: Connection utilization

As depicted in Table 4, the connection utilization differs considerably between the per-minute pricing experiments, the free trial periods, and the per-byte pricing experiment. Interestingly, the connection utilization in the free trial periods is higher than in the per-byte pricing experiments. The utilization for the per-minute pricing experiments is 3.5 to 5 times higher than for the free trial periods. Evidently, when subjects face per-minute prices, they economize by switching to the disconnect state or the free 8 kbps service after using the Internet. In the per-byte pricing experiment and the free trial weeks, there is no financial incentive for the subjects to switch back, so utilization is naturally lower. Of particular interest is the difference between connection utilization in the *Variable Symmetric Bandwidth* experiment and the *Variable Asymmetric Bandwidth* experiment. Subjects took advantage of the offered flexibility to increase utilization in the *Variable Asymmetric Bandwidth* experiment. These results are encouraging – they show that our subjects understand and exploit the flexibility of advanced pricing schemes if these schemes are accompanied by the proper economic incentives. Considering the 7:1 ratio between total in-bound

and out-bound traffic, the potential savings to be gained from lowering the out-bound bandwidth are obvious.

The change in demand can also be seen by comparing the connect times for the free 8 kbps service and the services at higher bandwidths (16 kbps to 128 kbps). Figure 4 illustrates the users' average daily connect time (see section 3.1 for remarks on the number of observations).

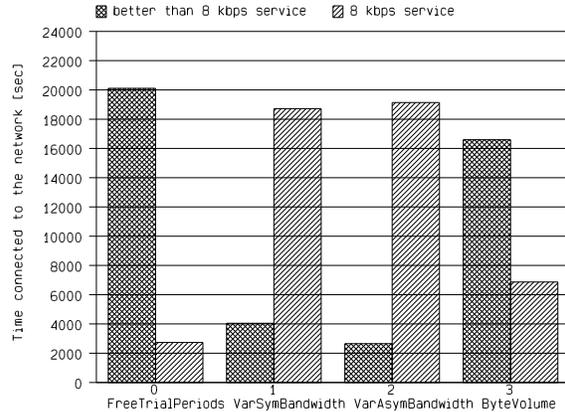


Figure 4: Average daily connect time

As expected, in the free service experiment, subjects stay connected longer at higher bandwidths than in all other experiments. When subjects face per-byte pricing, the change in behavior is only minor: There is a scant decrease in connect time at bandwidths above 8 kbps and a slight increase of connect time at 8 kbps. However, the behavior under per-minute pricing is dramatically different. For both time-based experiments, the time spent online at 8 kbps approximately equals the connect time for services at higher bandwidths in the free trial period. In other words, the ratio of paid vs. unpaid connect time has effectively been reversed. This also implies that the total connect time per *user-day* does not vary across the different experiments.

Although the overall time that users set aside for Internet usage on a given day is more or less constant, the variations in connect time over different access speeds and the change in connection utilization demonstrate that subjects adapt their usage pattern to the different pricing structures. In the per-minute pricing experiments, they increase utilization and reduce priced connect time by switching back and forth from the free 8 kbps service to the priced services.

The impact of usage-based pricing on user behavior is also observed in the average number of bytes transmitted per *user-day* (see subsection 3.1). Figure 5 shows the average number of bytes sent per experiment and per service class (8 kbps vs. all other bandwidths).

In the free service experiment, on average 12.9 mbyte are transmitted at high bandwidth services, versus 6.7 mbyte in the per-byte pricing experiment. Also, under per-byte pricing, the number of bytes sent at 8 kbps increased. Comparing the *Variable Symmetric Bandwidth* experiment and the *Variable Asymmetric Bandwidth* experiment, we can point out a notable property of user behavior. We see that the subjects consumed more bytes per day in the *Variable Asymmetric Bandwidth* experiment. Yet the average expenditure per day in the two experiments is almost identical. This apparent paradox can be easily resolved when further analyzing the expenditure pattern for both experiments. As

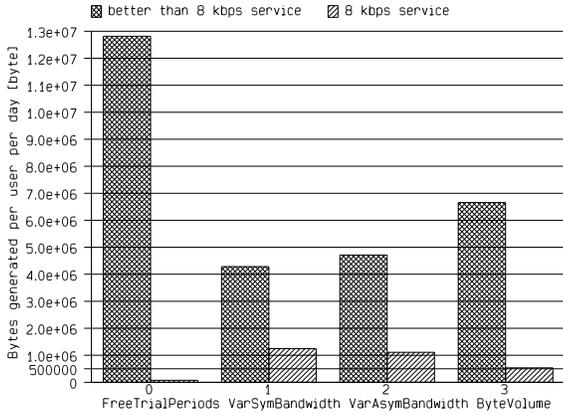


Figure 5: Average number of bytes transmitted

depicted in Figure 6, the distribution of expenditure for in-bound and out-bound traffic changes. Subjects spent 50 % less money per *user-day* on out-bound traffic than in the *Variable Symmetric Bandwidth* experiment. Moreover, as we can also see in Figure 6, they also spent the money saved this way on in-bound traffic.

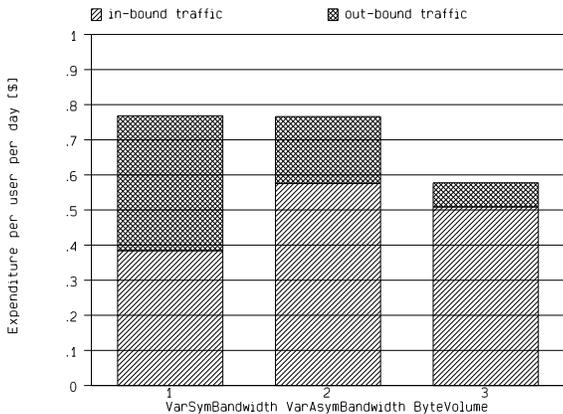


Figure 6: Average expenditure

### 3.3 User Segmentation

Whenever the marginal cost of network resource utilization is zero (like under a flat-rate pricing scheme), light users effectively subsidize heavy users that are not made financially accountable for their consumption of resources. The extent of this cross-subsidy depends on the variation in demand among users. In this subsection, we show that this variation in demand is indeed large and that consequently, the majority of Internet users should benefit from usage-based pricing. To demonstrate that a fair pricing mechanism needs to have a usage-based component, we examine user expenditure under flat-rate and usage-based pricing schemes. A subsequent analysis studies the distribution of traffic among users.

To show the variation in resource consumption, we first calculate the cumulative expenditure for the *Variable Symmetric Bandwidth* experiment. After aggregating and normalizing the expenditure data on a per-user basis, we rank

users by their traffic volume, processing heaviest users first. The resulting upper curve in Figure 7 plots cumulative expenditure versus cumulative traffic, starting from the heaviest users close to the graph's origin and proceeding to the lightest users on the far right. Each dot represents one subject. We now compare this actual distribution of expenditures with the distribution of a derived flat-rate tariff. We impute a flat-rate expenditure to each user by dividing the total actual revenues by the number of users participating in this experiment. The lower curve in Figure 7 represents the resulting flat-rate tariff.

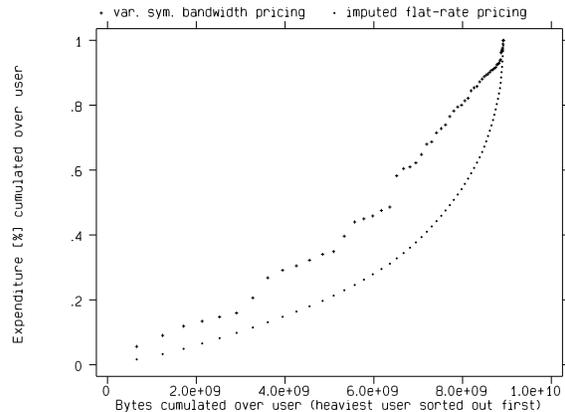


Figure 7: Comparison between flat-rate and *Variable Symmetric Bandwidth* pricing

A comparison of those two curves shows that 13 of the 15 heaviest users (23.4 % of the subject population) had higher actual expenditures than they would have had under flat-rate pricing. As illustrated in Figure 7, these 13 users accounted for about 5 gbytes of actual traffic and about 40 % of total expenditures. But there are also 8 of the remaining 49 subjects that did not generate that much traffic but paid more than the imputed flat-rate tariff. However, 41 of the 49 subjects with lighter traffic volume had to pay less. In total, 67 % of all subjects are better off in the *Variable Symmetric Bandwidth* pricing experiment.

We also made an equivalent investigation for the *Byte Volume* experiment. We used the same procedure as in the *Variable Symmetric Bandwidth* experiment and obtained the curves shown in Figure 8. The results of this investigation are even more convincing. The 12 heaviest users out of all 51 subjects (23.5 % of the subject population) that had completed the experiment had to pay a higher charge compared to the imputed flat-rate pricing. All remaining 39 subjects (76.5 %) were better off under *Byte Volume* pricing. The distribution is even more skewed than in the *Variable Symmetric Bandwidth* experiment – the 5 heaviest users in the *Byte Volume* experiment account for about 4 gbytes of traffic and half of total expenditures.

Another advantage of usage-based pricing can be illustrated by looking at the cumulative distribution of traffic and the effect of usage-based pricing on traffic generation. Figure 9 shows the cumulative distribution for the free trial periods, the *Variable Symmetric Bandwidth* experiment and the *Byte Volume* experiment. The x-axis represents the number of all subjects (in percent) that participated in a particular experiment. We again ranked users by their traffic volume, sorting the heaviest Internet users out first. The

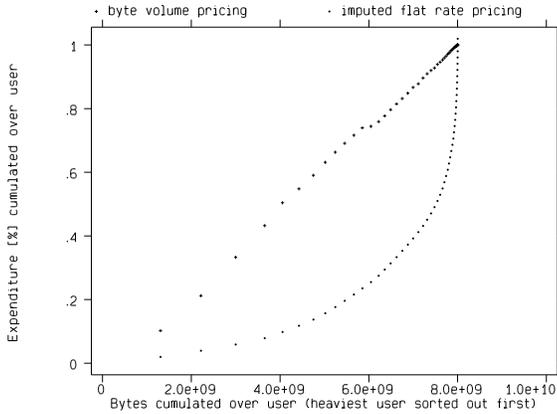


Figure 8: Comparison between flat-rate and *Byte Volume* pricing

y-axis represents the cumulative percentage of bytes transmitted during an experiment.

All three curves illustrate that the amount of bytes generated by all users is not equally distributed, even under usage-based pricing: In all experiments, priced or unpriced, 40 % of the subjects are responsible for at least 80 % of all traffic. If these subjects faced a flat-rate pricing structure, the low-usage subjects (60 % of all subjects) would have to pay 200 % more than they should have paid considering the lower level of resources they used.

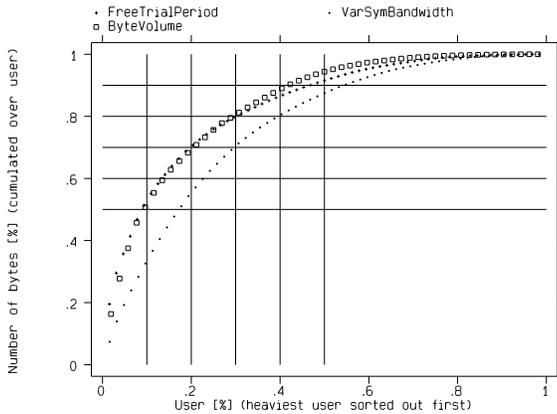


Figure 9: Cumulative Distribution of Traffic

Furthermore, Figure 9 shows that the introduction of time-based pricing reduces the number of bytes generated by the high-volume users significantly. On the other hand, while the *Variable Symmetric Bandwidth* curve is located well below the curves of both the free trial periods and the *Byte Volume* experiment, the two latter curves are almost identical. That means that all our subjects reduced their consumption in the *Byte Volume* experiment proportional to the total amount consumed during the free trial weeks.

#### 4 Conclusion

This paper presented results from the INternet Demand EXperiment (INDEX). We illustrated the change in individual

demand for Internet access under different usage-based pricing structures. In particular, we compared connection utilization, bytes transmitted, and connect time for three experiments: The *Variable Symmetric Bandwidth* experiment, the *Variable Asymmetric Bandwidth* experiment, and the *Byte Volume* experiment.

We have shown that demand is very sensitive towards different pricing structures. While the overall time that our subjects set aside for Internet usage was almost constant, we were able to demonstrate that the subjects understood and exploited the flexibility of advanced pricing schemes for their own advantage. Our subjects were able to increase the connection utilization in the *Variable Asymmetric Bandwidth* experiment although it required some knowledge about in-bound and out-bound traffic. In the per-minute pricing experiments, they significantly reduced the time they were connected to the Internet at priced bandwidths compared to the *Byte Volume* experiment.

Our findings also confirm the existence of large variations in user demand. The disparity between heavy and light Internet users is likely to widen as Internet penetration increases. Under such circumstances, equitable pricing mechanisms are crucial for avoiding large inefficiencies and cross-subsidies. We demonstrated that usage-based, quality-differentiated pricing fulfills these requirements. It is a fair way to charge people and leads to efficient use of network resources. Under flat-rate pricing, about 70 % of our subjects would have to subsidize the traffic generated by the remaining 30 %. Those 70 % are better off under usage-based pricing. We conclude that quality-differentiated services, available to both light and heavy users on an on-demand basis, increase the overall value of the network.

Future experiments will offer our subjects a choice of different pricing schemes with the objective of estimating user preferences for a variety of pricing structures. Other analyses of INDEX data seek to estimate individual user preferences. This real-world data can aid future network provisioning decisions on the way towards an Integrated Services Internet that will provide a universal platform for ever more users and services.

#### References

- [1] CommerceNet/Nielsen Media Research [1997]: *The Spring 1997 Internet Demographic Survey, Full Report*. New York, 1997.
- [2] Rupp, Björn, Richard Edell, Harish Chand, Pravin Varaiya [1998]: INDEX: A Platform for Determining how People Value the Quality of their Internet Access. In: *Proceedings of the 6th IEEE/IFIP International Workshop on Quality of Service*, pp. 85-90, Piscataway, NJ: IEEE Press, 1998.
- [3] Shenker, Scott, David Clark, Deborah Estrin, Shai Herzog [1996]: Pricing in Computer Networks: Reshaping the Research Agenda. In: *Telecommunications Policy*, vol.20(3), pp.183-202.