Providing Internet Access: What We Learn From INDEX*

Richard Edell and Pravin Varaiya
edell@eecs.berkeley.edu, varaiya@eecs.berkeley.edu
Department of Electrical Engineering & Computer Science
University of California, Berkeley, CA 94720

Abstract

The Internet Demand Experiment or INDEX is a market and technology trial. Its objective is to determine how much users value different qualities of service for Interent access. Findings from the trial imply that today's system of flat-rate pricing by ISPs is very inefficient. Flat-rate pricing wastes resources, requires light users to subsidize heavy users, and hinders deployment of broadband access. INDEX is a prototype of an alternative ISP model that offers differentiated-quality service on demand, with prices that reflect resource cost. In this alternative ISP consumers pay less, suppliers increase profits, and the deployment of broadband access is facilitated.

1 The unavailability of differentiated quality service

A major contribution of network engineering is the development of techniques (signalling protocols and associated algorithms) for using the same set of network resources (links and switches or routers) to simultaneously provide different services to end users. This is sometimes called an integrated services model to emphasize network support of different applications such as real-time voice and video and non-real time data transfer. One may also say that these techniques allow provisioning of differentiated quality service. The latter characterization emphasizes the flexible transport capabilities of the network and frees users to select the best way to match service quality to the demands of their application, time, and budget.

Although it is possible to provide it, differentiated quality service is not sold in the marketplace. There may be for two reasons for this, one economic, the other technical. To offer differentiated quality requires the design and testing of economically viable alternatives to the flat-rate pricing model adopted by virtually all Internet Service Providers (ISPs). It also requires a technology to package such a service in forms that users can purchase, provide the means for users to express their demand, signal the network to provision the requested quality, and generate accounting and billing records. The INDEX trial tests alternative pricing models and a technology to implement those models.

*Research supported by National Science Foundation, State of California MICRO program, CISCO, and SBC. We thank Jörn Altmann, Karyen Chu and Hal Varian for use of their unpublished results. A version of this paper was presented as the keynote talk at Infocom '99.

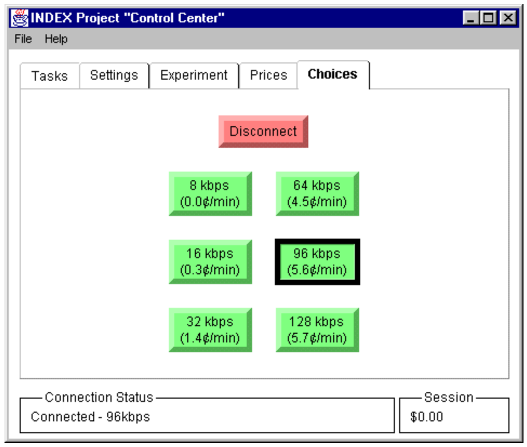

Figure 1: In the first experiment a user instantaneously selects her access speed and pays per minute of connect time. The control panel also shows the accumulated charges. Prices are selected randomly.

2 The INDEX trial

INDEX offers its subjects (customers) differentiated quality Internet access at home. The ongoing trial started in April 1998, and this paper reports findings based on data from the trial. The 70+ customers are students, faculty and staff of the University of California at Berkeley. There has been some turnover among customers due to changes in residence.

Every customer participates in a sequence of *service plans* or *experiments*, each lasting six to ten weeks. A service plan is characterized by a menu of service quality/price combinations. So each service plan implements a particular pricing model, parametrized by quality choices and prices. Customers pay by credit card. The value to a customer of a particular service quality is measured by how much money and time she spends consuming it.

INDEX collects detailed statistics on customer behavior. The basic data include byte counts in each direction, quality choice, accounting, and connection information. The time granularity is the minimum of one minute and change in quality choice. These data, with adequate protection of customer privacy, are available for study by researchers. The website: www.INDEX.Berkeley.EDU gives the details.

Figure 1 is a snapshot of the 'Choices' panel of the INDEX 'Control Center' or CC for the first experiment. The CC is the interface running on the desktop through which a customer controls access to the Internet. In this particular service plan a customer can select one of six different speeds (8, 16, 32, 64, 96, 128 Kbps) at a cost ranging from 0 cents for 8 Kbps to 5.7 cents per minute for 128 Kbps. The selected service quality (speed in this case) is provisioned virtually instantaneously.[1] The lower left corner of the CC indicates the current status of the connection—"connected at 96 Kbps" in this snapshot.

[1]The highest quality service that INDEX provides is a permanent 128 Kbps ISDN channel to the UC Berkeley campus network. There are plans to provide broadband service over ADSL.

The lower right corner of the CC shows a ‘spending meter’ that can be toggled to reveal the cost incurred until now during the current session, day, or month. (The figure shows the cost of the session to be $0.00, indicating the session has just started.) The spending meter is updated each minute, so customers can, if they wish, be aware of the cost of the resources they are consuming. Flat-rate charges, by contrast, deny users all information about their resource consumption. As a result, ISPs have to place inefficient restrictions on subscribers to limit resource consumption.

Another distinguishing feature is that an INDEX customer can *instantaneously* shift between different service qualities with no effort (beyond a mouse click). The quality may be changed during a session, if the user wishes. By contrast, when an ISP does offer different access speeds, the service is segmented into tiers: users must pick a single speed, and it is not possible to change that speed.

How much an INDEX customer values different speeds is measured directly by her purchase of the different speed options. The prices that are seen in Figure 1 are randomly selected and varied each week. Different customers face different prices. The price variation is large so that the demand estimates are robust.

Service quality has many dimensions besides speed, and INDEX service plans are designed to explore these dimensions. Furthermore, how a service is packaged into a commodity and priced makes a difference in demand, and some plans are designed to explore those differences.

By March 1999 the earliest subscribers participated in six experiments: (1) Symmetric bandwidth, (2) Asymmetric bandwidth, (3) Volume pricing, (4) Volume plus capacity charge, (5) Self-selecting tariff, (6) Enhanced flat-rate tariff. Figure 1 illustrates the choices in experiment (1), but it must be remembered that the prices are randomly selected for each customer and for each week. The aim is to determine sensitivity of demand for a certain speed to the prices of that speed and its substitutes. From INDEX data we learn that the price sensitivity is very high.

Experiment (2) is similar except that customers separately choose and pay for bandwidth in the upstream and downstream directions. The experiment is motivated by cable TV and DSL access, both of which have asymmetric speeds. The objective is to find out if users are aware of the asymmetry in their traffic pattern and make use of this asymmetry to reduce their bill. The answer is most are aware, and they do reduce their bill.

In experiment (3) users can select 128 Kbps and pay so many cents per megabyte of upstream and downstream traffic, or select free 8 Kbps service. We can compare user behavior when the commodity is megabytes of data transfer versus minutes of connect time. It turns out that the connect time goes up dramatically under volume pricing compared with connect time charges.

In experiment (4) a user incurs a volume or per megabyte charge plus a per minute connect time charge. This form of charge can reflect the bandwidth and buffer resources needed to support the user’s traffic. (See [1, Chapter 8].)

In experiment (1) the plan offers a per minute connect time price, and in experiment (3) the plan offers a per MB price (at 128 Kbps). In experiment (5) users are allowed to pick a convex combination of an offered per minute connect time price for 128 Kbps and an offered per MB price. The user must choose the combination at the beginning of each week and the combination is then fixed for the rest of the week. If the traffic to be generated during the week can be predicted, the least cost option for the user is either the pure volume or the pure connect time charge. The aim of the

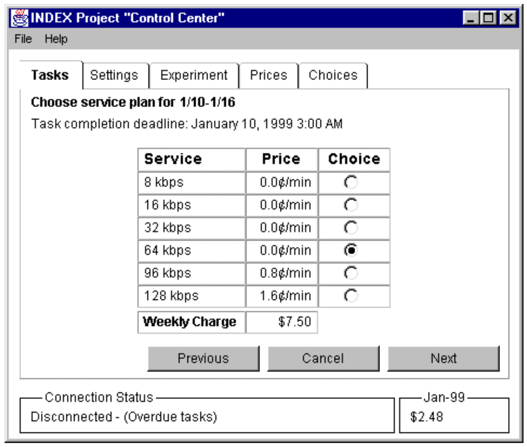

Figure 2: Customer choice panel for Flat Rate experiment. This customer is choosing to pay $7.50 for one week of unlimited 64 Kbps access. This choice is made once per week. During the week, the customer may select higher speeds at an additional per minute charge. The experiment combines flat rate and usage-based charges.

experiment is to find out how well users can predict their traffic pattern.

Experiment (6) is designed to test how much users value flat-rate pricing. At the beginning of each week consumers purchase unlimited usage for one week at a particular speed. They may connect at a higher speed for an extra cost per minute. Figure 2 shows the CC for a user who was offered one week of unlimited usage at a charge ranging from $0 for 8 Kbps to $15.00 for 128 Kbps and chose to pay $7.50 for 64 Kbps, while retaining the option of selecting 96 Kbps at 0.8 cents and 128 Kbps at 1.6 cents per minute.

Future experiments are aimed at estimating how much users value reduced blocking in modem pools, reduced congestion, as well as how much users shift their demand over time in response to time-of-use pricing.

The different INDEX service plans are alternatives to the prevailing dominant ISP model of flat-rate pricing.

3 Today's ISP

Figure 3 depicts an ISP network. Subscriber traffic, generated over dial-up 28 Kbps or higher speed DSL and cable TV modems, is aggregated at the ISP's point of presence or PoP. The ISP directs user datagrams over the Internet backbone through a Network Access Provider or NAP. Subscribers may connect to ISP servers that provide e-mail, news, web caching.

Users are charged a monthly flat rate of $20 for 28 Kbps dial-up access, $50 - $200 for access over dedicated 128 Kbps–1.5 Mbps DSL lines, and $40 for shared access over cable TV. More than 90 percent of residential subscribers use 28 Kbps dial-up modems, and in addition pay $15 per month for a telephone line. In case of cable TV, there also is a $25 monthly TV subscription charge.

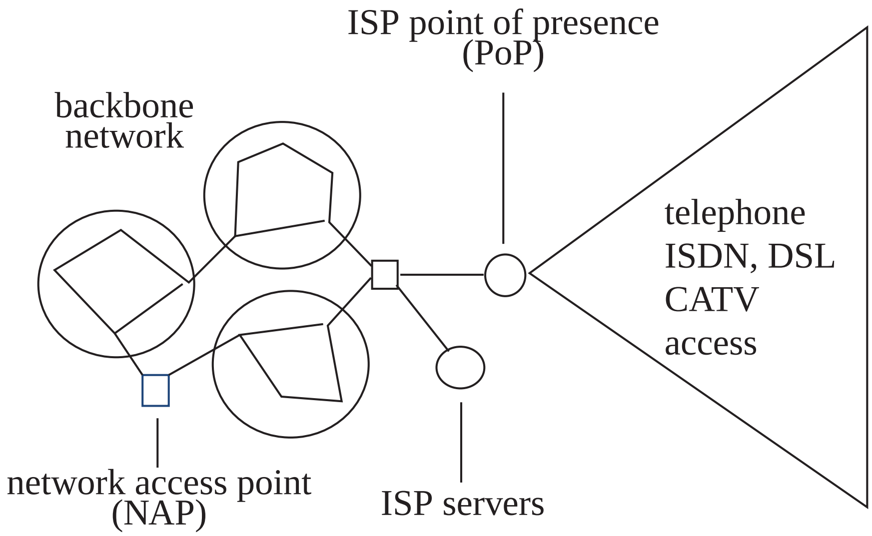

Figure 3: An ISP aggregates subscriber traffic and directs it to a Network Access Provider or NAP.

The ISP must purchase a modem pool, pay the telephone company for leased trunks to the modem pool, and the NAP for access to the Internet backbone. Assuming a monthly cost of $25 for one modem and one trunk, and a concentration of 10 subscribers per modem, this amounts to $2.5 per month per subscriber. (An ISP that charges $40 per month for access over cable in turn pays $30 to the cable operator.) Typical NAP charge is under 1 cent per MB of backbone traffic.

So the subscriber's cost of data transfer is high. The average amount of data transfer via 28 Kbps modems is 60 MB per month, so at a monthly cost of $20, the subscriber is paying 33 cents/MB. If you include the monthly $15 for a telephone line, the subscriber is paying 58 cents/MB. The light user who transfers a fifth of the average traffic pays an exorbitant $1.65–$2.91 per MB, and the heavy user who transfers fives times the average, pays 6.6–11.6 cents/MB. So light users subsidize heavy users. Data show that the 10 percent of heaviest users generate 30 times as much traffic as the 10 percent of the lightest users, so the subsidy is large.[2]

Table 1 compares the distribution by protocol of subscriber traffic in one ISP and in INDEX. The first column is the rank of the protocol used by the ISP's subscribers. The third column is the percentage of the traffic by protocol for the ISP, and the fourth column is for INDEX. The most obvious difference is that the top four applications (HTTP, NNTP, e-mail or POP3+SMTP, FTP) account for 96 percent of all ISP traffic, whereas INDEX users use a much wider set of protocols. The difference is explained by the fact that INDEX customers are more familiar with the Internet and they are more experienced computer users than the average Internet user interviewed in the national Nielsen-CommerceNet telephone survey. In a few years more Internet users will resemble today's INDEX subjects.

We now present the case against flat-rate pricing, beginning with the theoretical argument.

[2]The average connect time of AOL's subscribers is 21 hours per month, or $1 per hour. The median subscriber, whose connect time is much lower, is paying correspondingly more.

Rank	Protocol	ISP %	INDEX %
1	HTTP	76	43
2	NNTP	12	2
3	POP3	3	3
4	FTP	3	11
5	SMTP	2	0.5
6	HTTPS	2	2
18	Telnet/SSH	0	4
N/A	X-Windows	N/A	2

Table 1: Protocols used in ISP vs INDEX

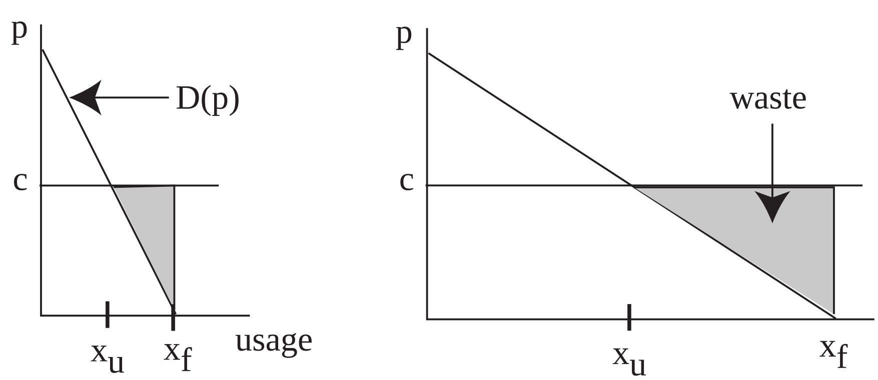

Figure 4: At a unit price of p, a customer will consume $D(p) = x_u$ units; under a flat-rate charge she will consume x_f. The shaded area is the waste.

4 What's wrong with flat-rate pricing - theory

We analyze a model of consumer demand and show that flat-rate pricing (1) encourages waste and increases cost, (2) forces light users to subsidize heavy users, and (3) can introduce differentiated service quality only by inefficient segmentation in quality tiers. The next section presents empirical evidence supporting the model.

Waste

Suppose Internet access is sold at a price of p per unit of usage, measured in MB of data transferred or minutes of connect time. Then a user's demand is modeled as a function $D(p)$. This is the decreasing curve in Figure 4. The meaning of the demand curve is that each of the $D(p)$ units of access is worth at least p to the user, and every additional unit is worth less than p.

If access is sold at the ISP's incremental cost of c per unit, the user will purchase $x_u = D(c)$ units. The value to her of consuming $D(c)$ units minus her cost, $cD(c)$, equals the area of the triangle above the horizontal line, $p = c$. This area is called the consumer surplus.

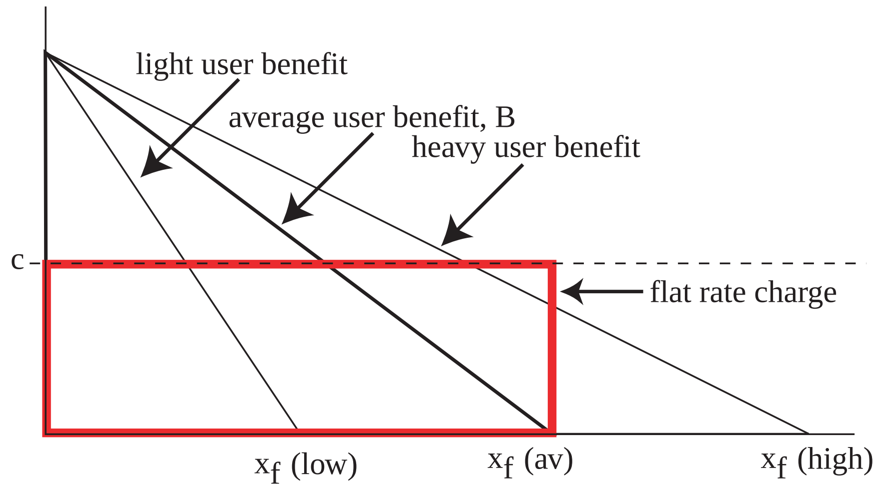

Figure 5: At a unit cost of c, the flat-rate charge is the rectangle, the small triangle is the value to the light user and the large triangle is the value to the heavy user.

But if the ISP charges a flat rate for unlimited use, the user's cost of additional consumption is zero, and she will consume $x_f = D(0)$. The shaded area is waste: it is the amount by which the cost of providing $x_f - x_u$ units, namely $c(x_f - x_u)$, exceeds their value to the user, $\int_0^c [D(p) - D(c)]dp$. The consumer surplus is now lower: it is the large triangle (the value of consuming x_f units) minus the cost cx_f. The surplus may even become negative.

Different users have different demand functions. Let $D(p)$ be the average demand function, and suppose a particular user's demand function is $\alpha D(p)$. So $\alpha < 1$ for a light user, and $\alpha > 1$ for a heavy user. The waste is proportional to α: heavy users cause more waste.

Thirdly, the waste will be larger if demand is more sensitive to price. If the price elasticity, $p/D(p) \times |\partial D/\partial p|$, is high, the demand curve will have a long "tail" and the shaded area will be larger. INDEX data indicate that the price elasticity for connect time is very high.

User cross-subsidy

If the flat-rate charge equals the cost of serving the average user this charge must be $c \times x_f(av) = c \times D(0)$. This is the rectangular area in Figure 5. Every subscriber pays this amount.

The benefit or value to the average user of consuming $D(0)$ is the area under her demand curve. This is the middle triangluar area. Suppose this area equals B. If the demand curve is $\alpha D(p)$, the user's benefit is αB. So light users ($\alpha < 1$) subsidize heavy users ($\alpha > 1$). The larger is the spread of α, the greater is the subsidy. Data indicate that this spread between demands of heavy and light users is on the order of 30.

The benefit to very light users, αB, may be smaller than the flat rate charge $c \times D(0)$, and so they will not subscribe. If the ISP wishes to retain these light subscribers it must either set the charge below average cost (and incur an operating loss) or it must restrict usage. ISPs engage in both

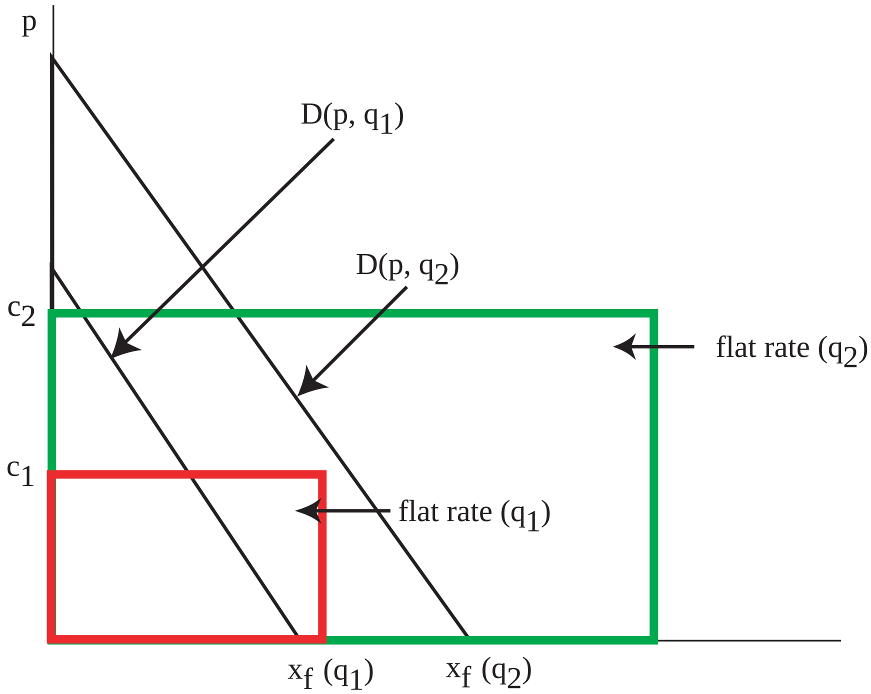

Figure 6: Demand increases with quality. Light users will not subscribe to higher-quality service tiers.

practices.

Suppose the ISP charges a usage-based price c instead of a flat rate. Then usage will decrease, all except the heaviest users will pay less and enjoy greater net benefits, even light users will subscribe, and the ISP's net revenue will increase.

Tiered quality

Suppose an ISP offers service of different quality, q, with larger q indicating better quality. (To fix ideas, think of quality as access speed.) The demand function $D(p, q)$ will depend on q and it will increase with q. This is shown in Figure 6 where the higher quality demand $D(p, q_2)$ is to the right of $D(p, q_1)$.

Suppose the ISP offers both service qualities at a flat rate. Then the service must be segmented into tiers: users can only subscribe to a single quality. The flat rates for the two qualities are the two rectangular areas. If a subscriber chooses the higher quality tier, she will consume $x_f(q_2)$ units, her benefit will be the area of the large triangle and her cost will equal the large rectangle. If she chooses the lower quality tier, her benefit will be the small triangle and her cost will be the small rectangle. If these areas area as shown, this subscriber will choose the lower quality tier.

However, if the ISP offered both qualities at unit prices c_1, c_2, as in the INDEX trial, the subscriber would consume *both* service qualities.

Thus the introduction of quality tiers leads to two kinds of inefficiency. First, the high quality tier subscribers will tend to be the heavier users and the waste will be higher. The larger waste increases the high quality tier charge (the large rectangle). So light users will tend to exclude themselves from the high quality tier. The excluded users lose benefits and the ISP loses revenues. One consequence is that the market for higher speed access will be unnecessarily limited. Second, both types of customers lose because their options are limited to one or the other tier.

INDEX data show that *all* users purchase high quality service sometimes if the cost is proportional to usage. ISP flat-rate charges for broadband access (384 Kpbs-1.5 Mbps) are about 5 to 10 times higher than for 28 Kbps. Such a large price differential is likely to discourage most users from subscribing to the high quality tier.

5 What's wrong with flat-rate pricing - evidence

We now present findings from INDEX that support the theoretical arguments above.

Responsiveness to prices

In the first experiment (see Figure 1) users select their instantaneous speed and pay per minute of connect time. The following demand model is estimated:

$$\log X_i = b_i + \sum_j a_i^j \log P_j, \quad i,j \in \{16, 32, 64, 96, 128\}, \tag{1}$$

where X_i is the number of minutes of connect time of speed i Kbps purchased by a consumer during a week when facing a price of P_j cents per minute for speed j. With this 'log-log' model, the coefficient a_i^i is the own-price elasticity, i.e. a_i^i is the percent change in demand X_i for speed i due to a one percent change in its price P_i; and a_i^j is the cross-price elasticity, i.e. the percentage change in the demand X_i due to a one percent change in the price P_j of speed j. The prior expectation is that $a_i^i < 0$, demand for speed i will drop if its price increases, and $a_i^j > 0$, demand for i will increase if the price of a substitute speed increases.

The estimated demand equation is

$$\begin{array}{rcrrrrrr} x_{128} & = & \mathbf{4.1} & \mathbf{-1.65} p_{128} & +0.44 p_{96} & \mathbf{+0.55} p_{64} & -0.12 p_{32} & +0.00 p_{16} \\ x_{96} & = & \mathbf{2.6} & \mathbf{+1.23} p_{128} & \mathbf{-3.34} p_{96} & +1.17 p_{64} & +0.23 p_{32} & +0.00 p_{16} \\ x_{64} & = & \mathbf{2.7} & +0.08 p_{128} & \mathbf{+0.84} p_{96} & \mathbf{-1.71} p_{64} & +0.47 p_{32} & +0.55 p_{16} \\ x_{32} & = & \mathbf{2.33} & +0.48 p_{128} & -0.58 p_{96} & \mathbf{+0.88} p_{64} & \mathbf{-1.10} p_{32} & +0.08 p_{16} \\ x_{16} & = & 0.52 & +0.42 p_{128} & -0.26 p_{96} & +0.18 p_{64} & \mathbf{+0.97} p_{32} & \mathbf{-1.29} p_{16} \end{array} \tag{2}$$

where $x_i = \log X_i$, $p_i = \log P_i$. The estimated coefficients in (2) whose t-static (i.e. estimate divided by its standard error) is larger than 3 are written in bold.[3] Observe that the demand for a particular speed is very sensitive to its own price and to the price of the next higher speed. The own-price elasticity is between -1 and -3 and the cross-price elasticity is around 1. The high own-price elasticity implies that when usage is connect-time, waste induced by flat-rate pricing is large.

[3] The demand estimates are due to Karyen Chu and Hal Varian.

Second, the large cross-price elasticity indicates that if users are offered differentiated quality of service, they would indeed demand more than one service quality. Direct evidence of this in INDEX data also comes from the fact that on average, each customer selects 3.5 out of the 5 available priced speeds (in addition to the free 8 Kbps speed) each week. Since consumers do value multiple service qualities, there is a loss to providers and consumers when ISPs only offer tiered quality service as they do today.

In experiment (3) users have two service choices: a free 8 Kbps and a *traffic volume*-priced 128 Kbps service. Figure 7 shows the average weekly usage (in megabytes) of each service as function of the 128 Kbps price. This figure also shows sensitivity to price, although the elasticity is lower than the estimates (2) for experiment 1. The likely explanation for the lower price sensitivity is two-fold.

In experiment 1, users are offered quality options that are close substitutes. In experiment 3 the only available substitute for 128 Kbps is 8 Kbps, and the very small increase in use of 8 Kbps despite a 30-fold increase in price of 128 Kbps, indicates that users don't consider 8 Kbps an effective substitute.

Second, under volume pricing, a user is better able to control her expenditure which depends only on the web or ftp transfer she executes. Under connect-time charges, her expenditure depends also on how much time she is connected to the network while she is thinking but not transferring data, and also on the actual rate of data transfer which depends on network congestion as well as on the (peak) rate she has selected. These less controllable aspects of connect-time charges are likely to lead to greater price sensitivity.

Despite the smaller price elasticity, note that the average volume transferred is 30-40 percent less even under the smallest usage-sensitive charge compared with when the transfer is free (as in the flat-rate option).

Quality-sensitivity of demand

INDEX data show that the user's demand increases strongly with quality as Figure 6 suggests. To explain the data, rewrite the demand equation (1) as

$$X_i = e^{b_i} \times \prod_j P_j^{a_i^j}, \tag{3}$$

so that higher values of the constants b_i corresponds to a multiplicative outward shift in the demand curve.

From (2) the estimated constants are,

$$e^{b_{128}} = 60.3,\ e^{b_{96}} = 13.5,\ e^{b_{64}} = 14.9,\ e^{b_{32}} = 10.3,\ e^{b_{0.52}} = 1.7.$$

Thus increasing the speed from 32 Kbps to 128 Kbps leads to a six-fold increase in average demand for connect time. The same order of magnitude is evident in Figure 7 if we compare the volume transferred at 8 Kbps versus 128 Kbps during the week when the latter service is free.[4]

[4]Demand models have been proposed in which a user selects one of several different qualities to transfer an exoge-

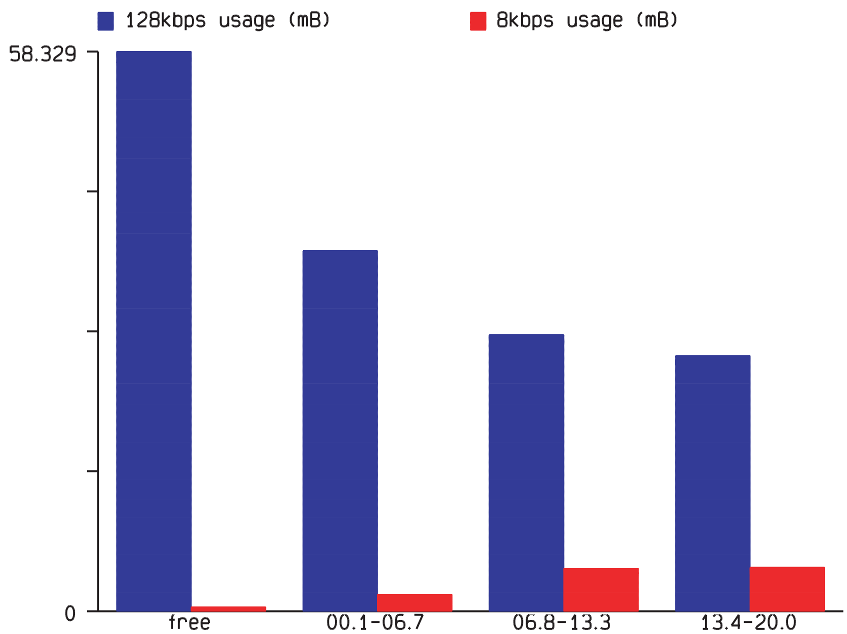

Figure 7: Variation in average weekly usage in MB of data transferred at 128 Kbps and at free 8 Kbps as a function of price per MB at 128 Kbps.

There seem to be two fundamental reasons for the high quality-sensitivity. One reason is that certain applications evoke a better subjective experience at higher speed than at lower speed, and so users demand for those applications will increase with access speed. Browsing of some web objects and interactive sessions (telnet, X-windows) are examples. A more subtle reason has to do with the fact that the full cost to a user is the sum of the money cost of access plus the value of her own time spent. So the total cost at the lower speed may well be *greater* than at the higher speed.[5]

The large quality-sensitivity of demand implies that consumers and suppliers lose on two counts from the current ISP practice of tiered quality service. Since tiered service is flat-rated, and demand is quality-sensitive, the waste at higher speeds is greater, and the flat rate charge is correspondingly higher. Also, many light users will not subscribe to flat-rated higher quality service, even though they would occasionally subscribe if the charge were usage-based. So those users are denied the benefits of higher quality service, and producers are denied the revenues from those subscribers. In turn, the reduced market for broadband access will lower the pace of equipment cost reduction and the diffusion of quality-sensitive applications.

Variation among users

Figure 5 indicates that inter-user subsidy and waste under a flat rate charge increase with the variation in demand among users. INDEX data show that the variation is indeed very large. To measure

nously specified amount of data. The results here indicate that the amount of data transferred depends on the available quality, i.e the amount is endogenously determined.

[5] Karyen Chu and Hal Varian have estimated INDEX users valuation of time.

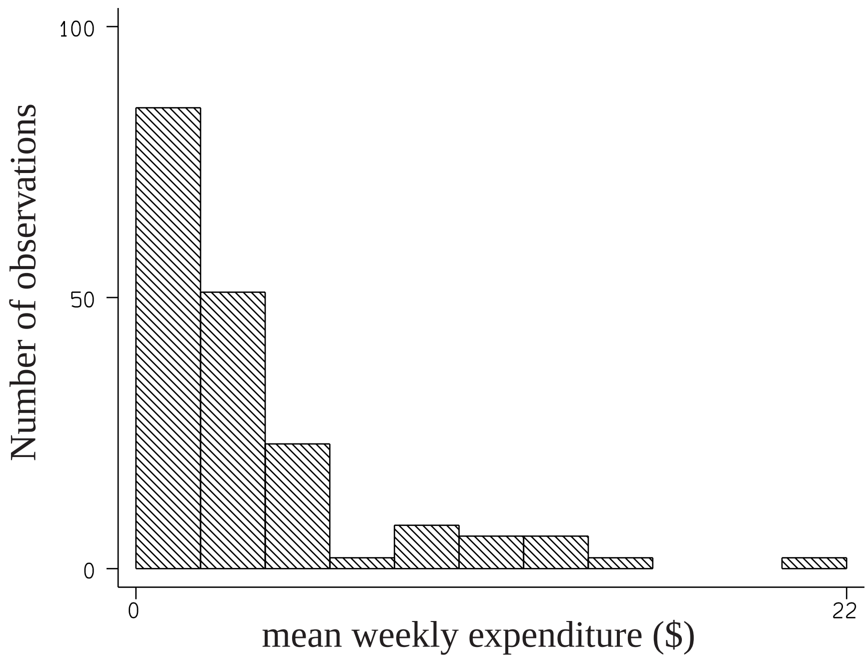

Figure 8: Variation in average weekly expenditure among users.

this variation, we re-estimate the connect-time demand equation (2) with a user-dependent constant term,

$$\log X_i^k = b_i^k + \sum_j a_i^j \log P_j, \quad i, j \in \{16, 32, 64, 96, 128\}, \tag{4}$$

where X_i^k is the consumption of user k, and b_i^k is the constant that depends on k. So user demands can differ by a multiplicative constant.

The R^2 of the estimated equation (2) ranges between 0.15 and 0.30, and the R^2 of the estimates of (4) ranges between 0.97 and 0.98. Moreover the estimates of the coefficients a_i^j do not change significantly.[6] The results imply a very strong user-specific multiplicative component in the demand.

The estimated coefficients b_i^k differ among the heaviest and lightest users by 3 or 4, so that their demands differ by a factor of $e^3 = 20$ to $e^4 = 54$.

The difference in demand is also indicated in Figure 8 which shows that the heaviest users spend 15 to 20 times more than the lightest users.

We can estimate the extent of inter-user subsidy.[7] We evaluate the distribution of actual expenditures among users over the six weeks of the volume-priced experiment 3. We compare it with what that distribution would have been for a flat rate tariff, assuming that user traffic is unchanged and the tariff is set to recover the same revenues. Figure 9 is obtained as follows. We rank users by their traffic volume, highest users first. For each n we plot the cumulative expenditures of the first n users versus their cumulative traffic volume. This gives rise to the upper 'dots' in the figure. The lower

[6]This is expected since prices are varied randomly.

[7]These results are due to Jörn Altmann.

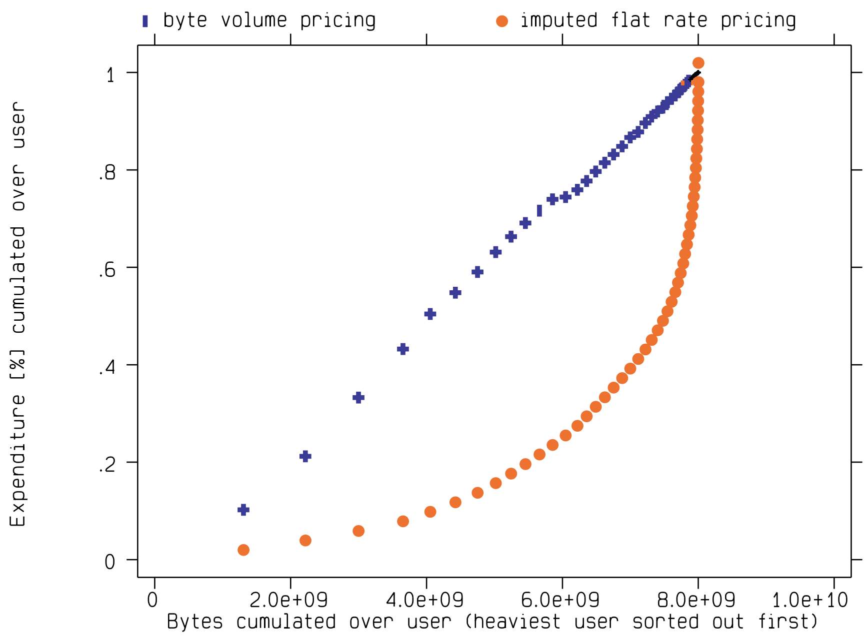

Figure 9: Expenditures versus usage under usage-based and flat rate pricing.

dots give the plot for an ‘equivalent’ flat rate tariff.

Thus, for example, the three heaviest users account for about 3 GB of data and 35 percent of total expenditures in the actual experiment. Under a flat rate tarrif, however, these three users would only account for 3/60 or 5 percent of the total expenditures, represented by the lower dot.

The volume-based pricing scheme is clearly more fair since users pay almost in proportion to their usage.[8] Second, the heaviest 30 percent of users would pay less under the flat rate scheme, whereas the remaining 70 percent would pay more. In fact, this comparison understates the subsidy involved under a flat rate because it assumes that users will generate the same amount of traffic under a flat rate system. But we know that in this experiment they generate twice as much traffic (see Figure 7), and the additional traffic is of less value. So the percentage of users who will be worse off under a flat rate tariff is likely to be 80 percent or more.

6 Why ISPs shun usage-based pricing

In December 1996, AOL changed from usage-based pricing ($9.95 per month, including 5 hours of connect time, plus $2.95 for each additional hour) to a $19.95 per month flat rate (later increased to $21.95). As a result, average monthly connect time per subscriber jumped from 6.4 hours to 22.1 hours in 1998, revenue per hour of connect time declined, as did the company’s operating margin ([2]).

[8] The upper dots do not lie strictly along a straight line because different users face different, randomly chosen prices and because the use of the free 8 Kbps service varies across users.

The shift to flat-rate pricing implies two fundamental changes in an ISP's business. First, it creates an incentive for the ISP to passively or actively degrade service quality, since per subscriber usage and cost decrease with worse quality but revenue remains the same. Passive degradation is acheived, for example, by limiting the size (increasing blocking probability) and speed of the modem pool. Active degradation is acheived (especially at higher speeds) by placing contractual and administrative limits on the ways customers' can use network access. For example, ISPs that provide access over cable modems limit streaming and access speeds. Other common practices are to prohibit users from installing a web server or a LAN in their homes, and to deny permanent IP addresses. (Providers also limit speed in order to segment the market into tiers.)

The only incentive to limit service degradation is the threat of loss of subscribers to other ISPs. ISPs reduce this threat by increasing the cost of switching to other providers. For example, in order to switch, a subscriber would have to reconfigure her computer which she may find difficult to do, and her e-mail would not be forwarded. ISPs incur significant costs to recruit subscribers. Typically, these costs are promotional free trials, and free services in the form of news clips, bulletin boards.

Thus competitive, flat-rate pricing increases per subscriber recruitment and retention cost and lowers quality. Providing network access thereby becomes an unprofitable business and so ISPs seek other revenue sources.[9] The largest ISPs now are 'portals', seeking to collect website rents in the form of advertising revenues and sales commissions. The idea is borrowed from TV and real estate. The more viewers a program has, the larger is the advertisement rent the TV station can charge. And the larger the number of shoppers in a city district, the greater is the sales per unit area, and so the higher is the rent that landlords in that district may charge.

The goal of becoming a profitable portal diverges from that of providing network access. Portal rents increase the longer subscribers stay at the ISP's own website. So the ISP has the incentive to discourage its subscribers from exploring the Internet. The most common approach is to provide better access to the ISP's own websites (in order to raise advertisement revenue) and the websites of merchants from whom the ISP can collect a sales commission.

Whether the change to portal will be profitable, it is too soon to tell. But the change has led to the disappearance of ISPs whose principal product is to provide Internet access to diverse users. One model for such an ISP is presented next.

7 An alternative ISP

The INDEX trial suggests an alternative ISP model that provides differentiated-quality network access to consumers priced according to usage, and reflecting resource cost. Cost of access would be significantly lower and quality higher than in the flat-rated regime. Customers of this alternative ISP would receive feedback about the cost of their resource consumption. They would have better control over their expenditures. We have already seen in Figure 1 that INDEX customers can find out how much they are spending by the minute, day, or month to date.

INDEX customers have two additional control mechanisms. First, the Control Center (CC) has a

[9] Under a recent agreement, AOL's Internet service provision is now handled by UUNet, so AOL may properly be said not to be an ISP any more.

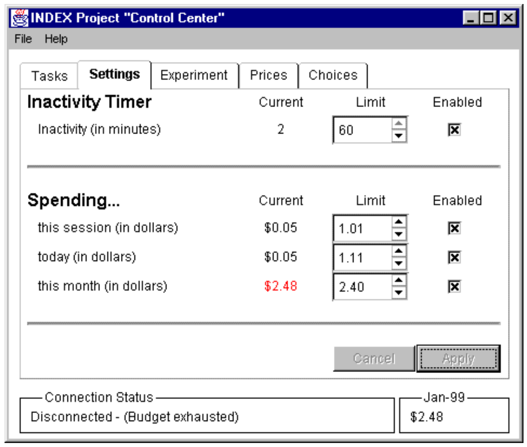

Figure 10: The settings panel allows the consumer to set expenditure limits by session, day or month.

settings panel, Figure 10, in which a customer can set separate limits on her expenditure during a session, day, or month. If any limit is exceeded, the CC places the connection in "disconnected" mode, the settings panel pops up, and the violated spending limit is highlighted in red. The customer may then stop accessing the network or raise her spending limit. (In the figure, the consumer's monthly limit was exceeded.)

The second mechanism is the detailed record of her transactions available on-line. A portion of this monthly record, which resembles a long-distance telephone bill, is shown in Figure 11. This particular user, during December 1998, particpated in Experiments 5 and 6 described earlier. The top part of the record is for Experiment 6. It shows that the consumer on 12/18/98 chose the 128 Kbps flat rate of $6 for the week of 12/20/98 and on 12/25/98 chose the 128 Kbps flat rate of $2 for the week of 12/27/98 (recall that prices vary randomly each week). The lower part of the record is for Experiment 5. The consumer selected to pay a pure volume charge. Two transactions are shown: the first on 12/01/98 cost 58 cents for transferring 14.5 MB (at 4 cents/MB) and the second transaction on 12/02/98 cost $2.11 for transferring 52.8 MB. Consumers that wish to understand their usage pattern can do so using these records.

Because the cost of using better quality provided by the alternative ISP would be small if their usage is small, most users would try better quality, and the market for broadband access would grow faster than is likely at present under tiered access. The INDEX trial shows that the technology for creating an alternative ISP exists, as does a market that would make it economically viable.[10] The alternative ISP would have the incentive to encourage Internet access (rather than to confine users to their own rent-producing websites), and this may restore some of the experimentation that characterized the earlier days of the Internet.

[10] INDEX technology is not discussed here. For a full discussion, see [3].

Usage 12/01/98 - 12/31/98

Transaction Date	Description	Total ($)
12/18/1998	128kbps service plan for 12/20-12/26	6.0000
12/25/1998	128kbps service plan for 12/27-1/2	2.0000

Transaction Date	Service/Time Quantity	Units	Price/unit	Price	Total ($)
12/01/1998	128kbps service 03:00:49-12/02 01:00:59				
	1320.17	minute	0.000	0.0000	
	14.548	megabyte	0.040	0.5819	0.5819
12/02/1998	128kbps service 08:36:53-21:47:24				
	790.53	minute	0.000	0.0000	
	52.817	megabyte	0.040	2.1127	2.1127

Figure 11: A portion of a user's detailed billing records for December 1998. The upper part of the bill is for Experiment 6 in which this user chose the flat rate option for each of two weeks of 128 Kbps service. The lower part of the bill is for two transactions during Experiment 5 in which this user selected to pay a pure volume charge.

8 Conclusion

INDEX reveals the structure of consumer demand for Internet access. The most important findings to date are: (1) the demand is very sensitive to price and to quality; (2) differences in demand among users are persistent and large; (3) the commodity form of the service (e.g. transport of traffic by volume or by time) has a large impact on demand.

The INDEX findings show that ISPs offering flat-rate service create large social inefficiencies in the form of waste (usage whose value to consumers is below cost), inter-subscriber subsidies, and tiered service. In the longer term, tiered, flat-rate service will limit deployment of broadband service. The change to flat-rate service has transformed the ISPs business into one that seeks site rents with incentives to degrade service quality and inhibit network access beyond revenue-raising sites.

Finally, INDEX is a market and technology trial of an alternative ISP that offers differentiated service quality with usage-based prices that reflect resource cost. Such service will promote the adoption of broadband access. The trial shows that such an alternative is technically feasible and economically viable, provides gains to providers and most consumers, and promotes the diffusion of broadband access.

References

[1] J. Walrand and P. Varaiya. *High-Performance Communication Networks.* San Francisco: Morgan-Kaufmann. 1996.

[2] America Online 1998 Annual Report. Available at www.aol.com/corp/inv/reports/

[3] R. Edell. *The Internet Demand Experiment*, Ph.D. thesis, University of California, Berkeley, June 1999, forthcoming.