The Demand for Bandwidth: Evidence from the INDEX Project

Hal R. Varian*
University of California, Berkeley

September 2001
(revised: February 11, 2002)

Abstract

The Internet Demand Experiment (INDEX) was a project to measure how much people are willing to pay for different kinds of Internet quality of service (QoS), such as bandwidth or volume of bits transferred. In the bandwidth experiments, which took place during 1998-99, subjects were offered high-speed residential Internet connections. Each week the price of bandwidth was changed, and experimenters could observe the users' responses to these prices. The resulting data allowed us to measure how consumers valued bandwidth in that environment. We found that most users placed a surprisingly low value on bandwidth, though certain kinds of users were willing to pay substantially more than others.

The INDEX project was a set of experiment designed to estimate how much people are willing to pay for various kinds of Internet Quality of Service (QoS), most prominently bandwidth. The INDEX designers architected the system to provide different QoS's on demand and to

*This work was supported by a National Science Foundation grant titled "Demand for quality-differentiated network services." I would like to thank Karyen Chu and Jorne Altmann for help with data extraction and manipulation and the entire INDEX team for helpful discussion and critiques of this work. My special thanks go to Pravin Varaiya and Richard Edell for designing and implementing this experiment.

record the usage of each different QoS by each user. Users were able to change their requested QoS instantaneously as often as they wished and were billed monthly for their usage.

From April 1998 to December 1999 we provided approximately 70 users at UC Berkeley with residential ISDN service through the INDEX Project. This paper summarizes some of the results of our study, and relates it to the current discussion about broadband deployment in the United States.

Here is a one-paragraph summary of what we found. (1) People aren't willing to pay very much for bandwidth, at least with the currently available set of applications. (2) The group that is willing to pay the most for bandwidth consists of technical and professional workers who work at home at least some of the time. (3) People are willing to pay a substantial premium for flat-rate, as opposed to metered, service. The implications for these findings for broadband deployment are discussed in Section 12.

1 Experimental design

The technical design is illustrated in Figure 1. Users were given always-on, 128 Kbs ISDN lines provided by Pacific Bell. These gave them direct access to a high-speed switched ethernet in a UC Berkeley laboratory, which in turn was connected to the campus fiber backbone and the public Internet.

The user's packets were routed through a "billing gateway" (BGW) which provided two functions. First, it measured various aspects of user behavior which were recorded in a database (DBMS). Second, it could offer various qualities of service by selectively degrading how the packets were handled.

For example, even though the actual ISDN bandwidth was always 128 Kbs, the billing gateway could buffer packets so that actual throughput was any desired speed less than 128-Kbs. In our experiments, we limited ourselves to 6 speeds: (8 Kbs, 16 Kbs, 32 Kbs, 64 Kbs, 96 Kbs, 128 Kbs).

Once a month each user's credit card was charged based on various aspects of his or her usage, as recorded in the database. Users could access their billing records at any time, and could view the charges as they accumulated in real time.

We wrote a small software application that by default ran in minia-

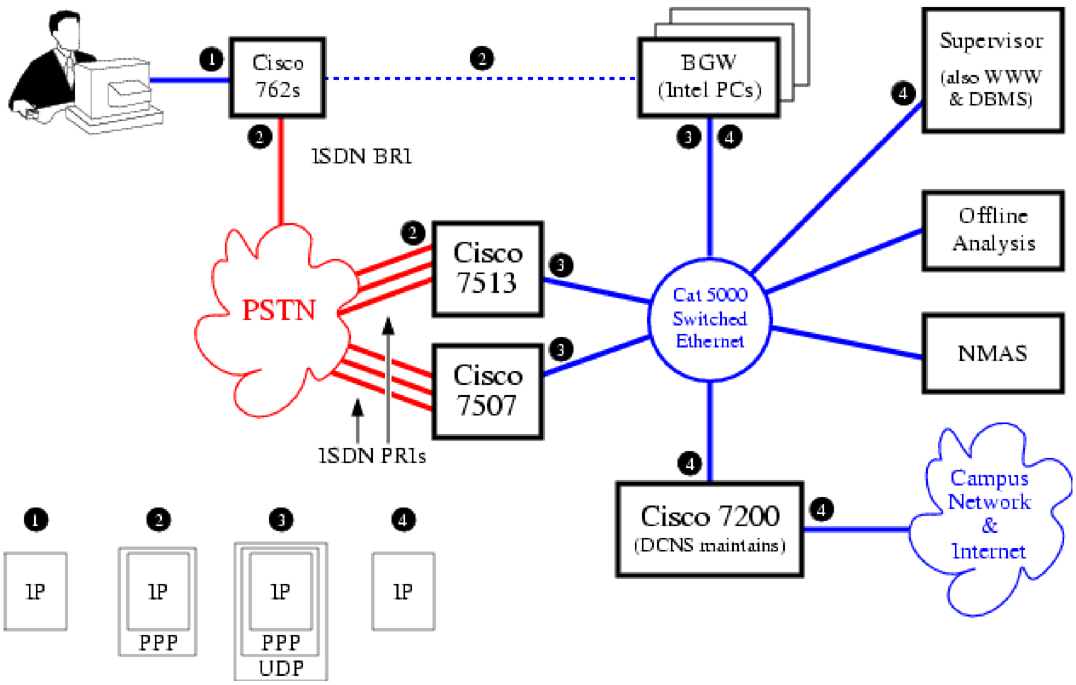

Figure 1: Design of INDEX experiment.

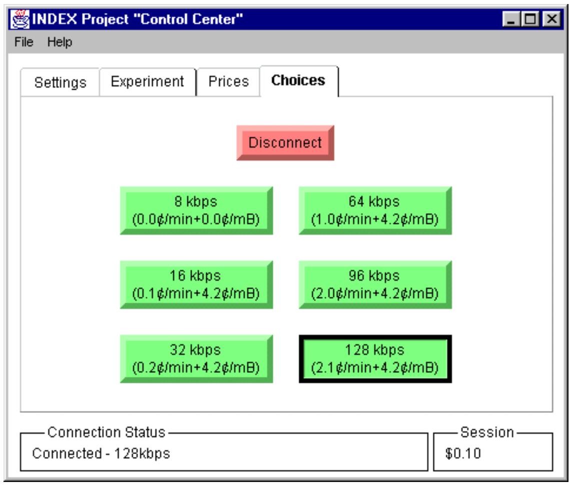

Figure 2: A typical menu of choices.

turized mode, but could be opened any time by a mouse click. This application showed the choices facing the user at that time, which choice was operative, and what charges were being incurred. The actual layout of the choices varied with the experiment, but Figure 2 shows a typical set.

In the example shown, a user can choose one of 8 different bandwidths, and incur different charges by doing so. The bottom of the screen shows the connection status and the charges being incurred in this session. Clicking on the session window switches it to charges per month or per day.

Note that the users could change their bandwidth instantaneously, simply by clicking one of the buttons. Hence the INDEX project allowed for much more flexibility in bandwidth choice than one would normally receive by signing up for a high-speed Internet provider and paying a flat monthly fee.

See Edell and Variaya [1999] for motivation on why and how the INDEX experiment was created, and see Edell [2001] for a detailed description of the technology that went into it.

2 Description of experiments

We ran a number of different experiments over the 20 months that we were in operation. A typical experiment would be implemented on a Sunday evening. Users were generally given a week to experiment with the different service qualities, without incurring any charges. Then we would run 6 weeks of priced service with that particular set of choices, varying the prices for the choices each week.

We ran the following experiments:

Symmetric bandwidth. Users paid the same prices for bandwidth for upload and download.

Asymmetric bandwidth. Users paid different prices for bandwidth used for upload and download.

Buyout pricing. Users were presented with a set of prices for bandwidth, and a rate at which they could buy out from the metered pricing. At the start of the week they could choose whether they wanted to face the metered schedule or purchase unmetered schedule.

Volume pricing. Users were charged for the cumulative bytes transferred over the course of the week.

Fixed mixtures. Users were charged based on a weighted average of bytes transferred and bandwidth chosen, with the weights chosen by the experimenters.

Variable mixtures. Users were charged based on a weighted average of bytes transferred and bandwidth chosen, with the weights chosen by the users at the start of each week.

One experiment that we had planned but were unable to implement offered the users different delays to connect. This was intended to measure the willingness to pay for "always on" service.

Users ran though experiments asynchronously, so that each user started with the first experiment, regardless of when they were recruited. Users were recruited by ads posted around campus and in campus newsletters. We attempted to get a reasonably balanced mix of faculty, staff and students, but all of our subjects had to be part of the Berkeley community in order to use the campus Internet service.

In general, our subjects were experienced and heavy users of the Internet, as compared to the population as a whole. This is indicated by the following statistics:

- 91% had used the Internet for more than 3 years [in 1998]
- 86% had used computers for more than 5 years
- 58% characterized their Internet use as "above average"
- 56% considered themselves "computer profesionals"

This sample of users is clearly not representative of the population as a whole, but may be representative of "early adopters" of new technology.

Our subjects, of course, were volunteers. They found participation in the experiment attractive since we subsidized a portion of the costs. In particular, they did not have to pay for the ISDN modems, they did not have to pay a setup charge to Pacific Bell, and they were, on the average, charged below market rates for the service. The equipment and setup charges amounted to several hundred dollars, and the rates were chosen to be, on the average, around 75% of the market price. The actual rates the users faced were typically chosen randomly, subject to a monotonicity requirement that higher quality had to cost as least as much as lower quality.

Each user was required to fill out a survey before their service was started. Several of the questions were identical to the CommerceNet-Nielsen survey which claims to be a representative sample of the U.S. population, so we could determine just how different our population was from the national averages. In particular, users reported their occupation, Internet experience, income levels, and who was paying for the service (e.g. whether they were spending their own money, or someone else paid for their access.)

3 Demand for bandwidth

Here I report on one set of INDEX experiments designed to measure the willingness to pay for bandwidth. In these experiments users were offered the choice of 6 different bandwidths, ranging from 8 Kbs to 128 Kbs. Users could choose 8 Kbs service for free at any time. Each Sunday a new set of prices were chosen for the other bandwidths, ranging from 0.1 cents to 12 cents per minute of use. The INDEX system measured how much bandwidth subjects consumed at each different price, allowing experimenters to estimate demand for different bandwidths as a function of the price vector. See Varian [2001] for a more

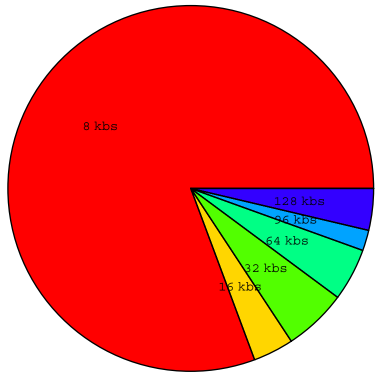

Figure 3: Bandwidth usage.

detailed econometric analysis of this dataset, and see Beckert [1999] for a more detailed structural approach to estimation.

4 Reduced form demand estimates

Figure 3 depicts a pie chart of total usage. About 3/4 of the usage was 8 Kbs service. Since 8 Kbs was free, users tended to keep it on all the time. Usage was roughly equally divided among the other five priced bandwidths.

Table 1 depicts the output of regressing the log of total minutes used at each bandwidth on the log of the five different prices. Observations with zero usage were omitted. No attempt was made to restrict coefficients across equations since this regression is intended only to provide a succinct description of the observed patterns of demand.

The coefficients in these log-log regressions can be interpreted as price elasticities of demand. Coefficients printed in bold are statistically significant at the 95% level.

Note that the diagonal terms (the own-price effects) are all negative and statistically significant. The subdiagonal terms are the cross-price effects for lower bandwidths. The positive numbers indicate that one-step lower bandwidths are perceived as substitutes for the chosen bandwidth.

This sign pattern is quite plausible. It is also worth noting that the implied elasticities are rather large. The regression for 96 Kbs

Bandwidth	p128	p96	p64	p32	p16
128	$-\mathbf{2.0}$	+.80	+.25	−.02	−.16
96	$\mathbf{+1.7}$	$-\mathbf{3.1}$	$\mathbf{+.43}$	+.19	+.18
64	+.77	$\mathbf{+1.8}$	$-\mathbf{2.9}$	+.59	+.21
32	+.81	−1.0	$\mathbf{+1.0}$	$-\mathbf{1.4}$	+.15
16	+0.2	−.29	+.04	$\mathbf{+1.2}$	$-\mathbf{1.3}$

Table 1: Reduced form estimates. All own price effects are significantly negative; the cross-price effects for one-step lower bandwidths are positive.

Bandwidth	With ISE	No ISE
128	.95	.11
96	.93	.25
64	.92	.18
32	.95	.14
16	.90	.17

Table 2: Regression R^2. The R^2s with individual specific effects are large.

service implies that a 1% increase in the price of 96 Kbs leads to a 3.1% drop in demand, and a 1% increase in 128 Kbs service leads to a 1.7% increase in the demand for 96 Kbs service.

We ran these regressions with and without dummy variables for the individual users, with little change in the estimated coefficients. Table 2 depicts the R^2s for these regressions.

Roughly speaking about 20 percent of the variance in demand is explained by price variation, about 75 percent of the variance in demand is explained by individual specific effects, and about 5 percent is unexplained. These fits are remarkably good, giving us some confidence that the subjects are behaving in accord with the traditional economic model of consumer behavior.

5 Structural demand estimates

The reduced form estimates given above suggest that the users are behaving in an economically sensible way. Hence it makes sense to try to model their choice behavior in more detail so we can extrapolate to other environments.

I adopt a very simple behavioral model, and assume that users get utility from the bits transferred ($u(x)$) and the time (t) it takes to transfer them. The cost of transfer time has two components: the subjective cost of time (c), which varies according to users and circumstances, and the dollar cost, which depends on the price of the chosen bandwidth ($p(b^*)$). If b^* is the chosen bandwidth, optimization implies that

$$u(x) - [c + p(b^*)]t \geq u(x) - [c + p(b)]t, \tag{1}$$

for all bandwidths b.

Since bandwidth is by definition bits per unit time, we have $t = x/b$. Making this substitution and canceling the xs, we have

$$[c + p(b^*)]\frac{1}{b^*} \leq [c + p(b)]\frac{1}{b}, \tag{2}$$

for all bandwidths b.[1]

It follows from simple algebra that

$$\min_{b^*<b} \frac{p(b^*)b - p(b)b^*}{b^* - b} \geq c \geq \max_{b^*>b} \frac{p(b^*)b - p(b)b^*}{b^* - b}. \tag{3}$$

This gives us observable upper and lower bounds for c, the user's subjective cost of time.

Figure 4 depicts these bounds graphically. Define the "total cost of time" by

$$K(c) = [c + p(b)]\frac{1}{b}, \tag{4}$$

and plot these affine functions for each bandwidth b. A user with subjective time cost c will choose the bandwidth b with the lowest total cost. Conversely, an observed choice of b implies that the time cost must be bounded above and below as indicated in Figure 4. Note that a choice of the lowest available bandwidth only yields an upper bound on time cost, and a choice of the highest available bandwidth only yields a lower bound on time cost.

[1] If users waste some of their bandwidth we could write $t = ax/b$, where $a > 1$. As long as the fraction wasted is constant across bandwidths, the cancellation of ax can still be performed.

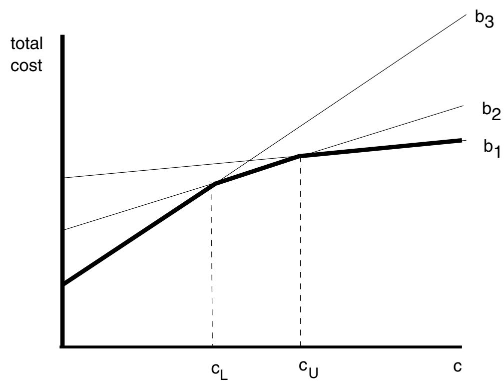

Figure 4: The straight lines are the "total cost of time" at different bandwidths. If we observe a particular bandwidth being chosen, we can calculate upper and lower bounds on the subjective time cost c.

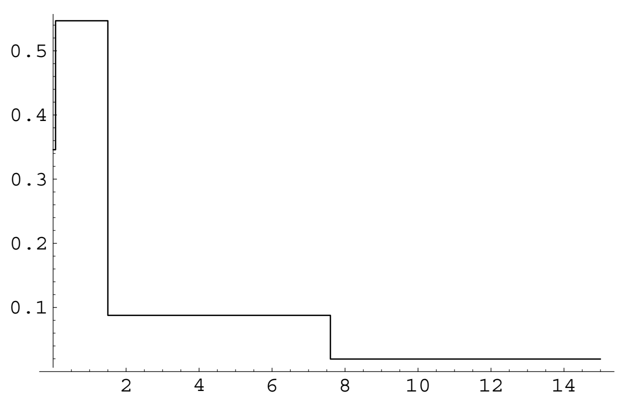

Figure 5: Histogram illustrating the fraction of the time that a particular user's time cost falls in the indicated region in a particular week.

Range	0	1	2	3	4	5	6	7	8	9	10	11
Upper bound	39	8	3	4	1	2	2	1	2	0	3	0
Lower bound	63	3	0	0	1	0	0	0	0	0	0	0
Average	47	7	2	3	3	3	1	3	1	1	0	0

Table 3: Frequency with which time cost in given range is observed.

6 Nonparametric estimates of the value of time

We assume that the user's time cost is a random parameter, drawn from a distribution $p(c)$. Sometimes the user is in a hurry, which means he or she has a high cost of time. Sometimes they are patient, which means the user has a low cost of time. This distribution of time cost is summarized by the probability distribution $p(c)$ and our objective is to estimate this distribution.

Each weekly menu of prices and bandwidths gives us a set of upper and lower bounds. Since we observe the frequency with which the user chooses each bandwidth b during a week, we can construct a histogram for each user for each week illustrating the implied time costs. An example for a particular user in a particular week is given in Figure 5.

7 Distribution of the time costs

Table 3 shows the frequency with which the upper and lower bounds fall in a give range. For example, 39 of the users, or about 60%, have an average upper bound on the time cost of less than 1 cent a minute, 8 of the users, or about 12%, have an average upper bound greater than 1 cent a minute, but less than 2 cents a minute and so on. The last line in this table is the distribution of a simple average of the upper and lower bounds, which is a rough-and-ready nonparametric estimate of the distribution of time cost across the population. Figure 6 depicts the same information in a bar chart.

The remarkable thing about Table 3 and Figure 6 is the low values that users place on their time. Most of the users have a time cost of less than 1 cent a minute, with only a few users having higher time costs.

The obvious question is whether we can predict which users have

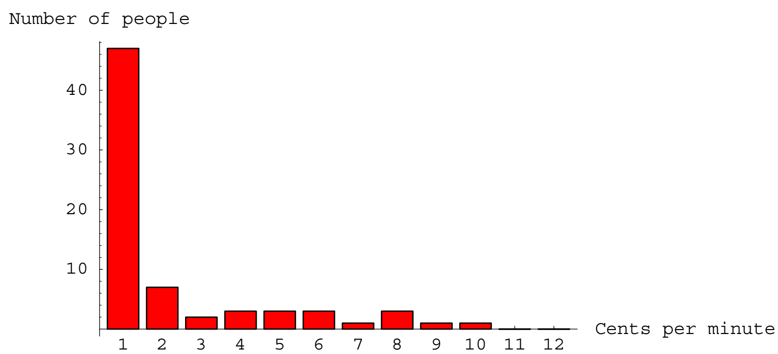

Figure 6: Histogram of number of people with different willingnesses to pay.

higher time value. Relevant variables available are occupation type, income, and whether the employer or the user pays for the service. We found that occupational dummies do a pretty good job of explaining the time costs using the following regression:

$$\text{c} = .86 \text{ professional} + 2.4 \text{ technical} + 7.02 \text{ admin} + .91 \text{ student}.$$

All coefficients are statistically significant and the R^2 for the regression is .646. Adding in both "income" and "who pays" yields an R^2 of .652, a negligible increase, suggesting that the best predictors of willingness to pay are the occupation dummies.

Figure 7 shows the distribution of time values by occupational classification, which tells essentially the same story as the regression.

To sum up: users placed a remarkably low value on their time, on the order of half-a-cent per minute. However, users who classified themselves as being in technical and administrative jobs tended to place a significantly higher value on their time. It appears that many of the technical and administrative users were telecommuting, suggesting that this group of users have particularly low demand elasticities, or, equivalently, particularly high time costs and therefore particularly high willingness to pay for broadband.

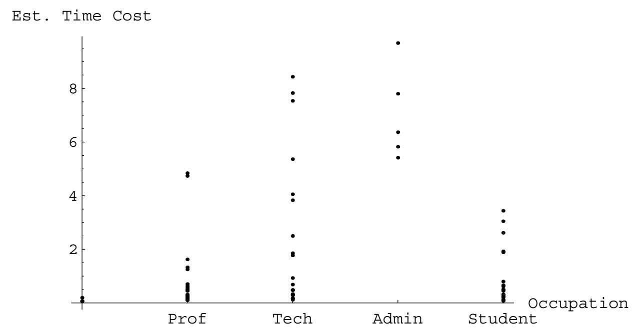

Figure 7: Time cost versus occupational category.

8 Why is the time cost so low?

The estimated time costs seem remarkably low, and it is worth considering this is is so. Several hypotheses suggest themselves.

Users are non-representative. This is likely part of the story. Our users are volunteers, many are students, and it is apparent from Figure 7 that certain occupations have much higher time valuations that our representative user. However, the fact that our panel volunteered for broadband and that they are primarily drawn from a university population suggests that, if anything, our users should have a *higher* value for broadband than the population as a whole.

Other uses of time. Not all of the time that a user is "waiting for a download" is wasted since it is common to engage in alternative activities. Indeed, we have already mentioned that users tend to leave 8 Kbs service on all the time so that email could be downloaded in the background. Because of this multitasking capability, the value of "time saved" could easily be lower than one might think. A closely related point is that certain activities, such as Web surfing, tend to involve bursts of activity, followed by a period of time spent in absorbing the acquired material. In this situation, the bandwidth per se is not necessarily the constraining factor in acquiring and absorbing information.

Service quality on rest of Internet. We can only control the quality of service on the link from the user's residence to the ISP. We have no control over bandwidth elsewhere on the Internet. If the user is accessing a site that is highly congested, an increase in speed on the residential-ISP link could have no value since it would not increase overall throughput.

Truncation at the high end. We could not charge more than commercially available ISDN services, and our methodology (or any methodology) cannot determine how much *more* people would be willing to pay than the highest price that they face. Hence we may be underestimating the willingness to pay for high-end users. However, note from Table 3 that there are no observations at the very high end of the distribution, so this explanation does not appear to be a strong one.

Can only measure value of existing applications. We can only measure how the user values time given the existing mix of applications. If the user had access to high bandwidth at low cost, there could easily be applications that are infeasible at current bandwidths that the user could find valuable. We cannot measure the value of such hypothetical applications using the methodology at our disposal. Note, in particular, that our experiments were run before Napster became popular in college communities. It may easily happen that if broadband were widely deployed new and compelling applications could emerge.

9 Buyout pricing

Another interesting set of experiments examined how much people were willing to pay to avoid being metered. Earlier evidence, as summarized in Odlyzko [2001], showed that a substantial fraction of local telephone users preferred flat rate service, even when they would have been better off under metered service.

It has often been argued that the same thing is true of Internet users. There is a psychological cost to being metered, and it is of interest to know how much people are willing to pay to avoid facing this cost.

The experiment we designed to examine and quantify this issue was 10 weeks long. Users faced the price schedule shown in Table 4. Each week prices were drawn in a range from $1 to $20, which measured the cost to the users of purchasing unlimited 128 Kbs usage. The column

Kilobits per second	Cents per minute	Buyout fraction
16	.4	0.125
32	.8	0.250
64	1.6	0.50
96	2.4	0.75
128	3.2	1.00

Table 4: Buyout from metering experimental design.

labeled "buyout fraction" in Table 4 shows how much it cost to buy unlimited service up to and including other speeds. For example, if it cost $10 to buy unlimited 128 Kbs speed, it would cost $5 to buy unlimited 64 Kbs speed (since the "buyout fraction" for 64 Kbs is 0.5). If the user wanted to consume higher bandwidth than he or she had bought out, they could do so by incurring the indicated per minute charges.

9.1 Analysis of buyout experiment

Chu [1999] has conducted a preliminary analysis of the data generated by 40 subjects with 337 person-weeks, and is in the process of conducting a detailed analysis based on all the data.

She found that users chose the buyout option during roughly 80% of the 337 weeks. Of the weeks that were bought about, 26 percent of the subjects would have been better off facing the metered pricing, given their *ex post* behavior.

Define the *buyout premium* by the total amount paid for Internet access during a given week (usage+buyout charge) minus the minimum cost of purchasing that same pattern of use. Chu [1999] found that the buyout premium ranged from $9 to 65 cents, with a median of about $2. The median expenditure per week was about $4, so $2 is quite a large premium relative to expenditure.

The buyout premium can be viewed either as an optimization error, or the additional benefit users felt from not having to be metered. Since the demand studies presented earlier suggest that users seemed to have done a good job at optimization, we lean toward the latter interpretation.

Behavior was quite different between the metered and unmetered weeks. The volume of data transferred during the 265 bought out weeks was about 11 megabytes per day, while the volume of data transferred during the 72 non-bought-out weeks was roughly 1.25 megabytes per day. Revenue collected was higher during bought out weeks, but the revenue per megabyte transferred was much lower: roughly 8 cents per megabyte during bought-out weeks and 30 cents per megabyte during non-bought out weeks.

Whether flat or metered pricing is more profitable for the provider depends critically on the assumed cost structure. Total bytes transferred was roughly 9 times larger during the bought-out/flat priced weeks, but people were willing to pay, on average, a 50 percent premium to face a flat price.

10 Summary of INDEX findings

The implications of the two INDEX experiments discussed here are the following.

- Users are not willing to pay very much for higher bandwidth when using applications available in 1998-99.
- Administrative and technical users, who appear to be telecommuting, are willing to pay a significantly higher price that the average user.
- People were willing to pay a substantial premium to face unmetered pricing, but they also placed much larger demands on the system than when they were metered.

11 Comparison with recent studies

According to General Accounting Office [2001], roughly 50% of the population of Internet users has access to broadband as of Fall 2000, but only 12% of those have purchased the service. This fact, along with the INDEX findings, suggests that the problem with broadband is not *access* but *applications*. Ordinary users need a good reason to pay a premium of roughly $25 a month to get broadband access.

Much has been made of the fact that South Korea has achieved a fifty percent penetration of households using broadband. What has achieved less publicity is the fact that the incremental cost of broadband in South Korea is less than $23 per month—about the cost of dialup service in the U.S.[2]

Additional evidence for significant price sensitivity comes from Charter Communications, a U.S. cable TV company owned by Microsoft billionaire Paul Allen. Charter offers cable modem service at 256 Kbps and 512 Kbps at price points of $23 and $39.95, respectively. According to Blumenthal [2002] the take-up is split 60%-40% in favor of the $23 service.

Finally, Kridel et al. [2002] estimate elasticities for U.S. cable modem purchase and find the elasticity at $29.95 per month to be 1.075 while the elasticity at $49.95 per month is 1.793. These estimates demonstrate considerable price sensitivity, particularly at the high end. Unfortunatley, this is the direction in which prices appear to be moving, as we describe below.

Given the set of applications available today, it appears that users aren't willing to pay a big premium for broadband. What are the prospects for new applications increasing demand?

Available studies are somewhat inconclusive. According to Berchtold et al. [2001], broadband users spend 27 percent more time on-line and average 37 percent more sessions a month. This study found that broadband users spent proportionally *less* time on the Web than the same users did when they only had narrowband access. The users spent more time on email, chat, and downloading music. However, the music downloads dropped dramatically after Napster suspended operation. This report also found that one category of Web sites that saw a significant increase in use were game sites.

However, Rappoport et al. [2002] time spent on the Internet by broadband and narrowband users was not that different. It may be that the difference in the usage patterns detected by these two studies is due use of Naspster and gaming applications.

These two studies, taken together, suggest that users do not change their behavior dramatically when switching to broadband, particularly now that Napster is gone.

The fact that usage patterns are relatively stable suggests that the results from INDEX remain valid: users are not willing to pay much of a premium for broadband, given the set of applications currently

[2]See Paltridge [2001] for a discussion of the early days of the Korean experience, along with experiences of other OECD countries.

available. Of course, it may well be that some must-have killer app will arise when broadband is sufficiently widespread.

There is a serious chicken-and-egg problem with broadband deployment: telecommunications providers don't want to pay for network upgrades if the demand isn't there, and if the infrastructure isn't there, software developers will have little reason to deploy broadband applications and content.

In an environment where such indirect network externalities are prominent, it makes sense to offer low "penetration prices" to get the positive feedback going.[3]

Unfortunately, prices in the U.S. appear to be heading higher. According to Ames [2002], cable broadband access rose 12% in 2001, from $39.40 per month at the beginning of the year to $44.22 per month by the end of the year. DSL prices rose 10 percent over the same period, from $47.18 to $51.67.

These trends lead me to conclude that unless new compelling applications are forthcoming or the price of broadband connectivity falls significantly, we should not expect to see a large surge in demand for broadband in the U.S. in the near future.

12 Implications for ISPs

What can be done about this relatively pessimistic forecast? Let us use the INDEX estimates to make some rough calculations about demand. As we have seen dialup access costs about $25 a month, and broadband access currently costs about $50 a month, which indicates that the marginal broadband user is apparently willing to pay at least $25 for the additional utility provided by broadband.

The GAO figures cited above indicate that about 12% of Internet users have signed up for broadband. If this group values their time at about $1.20 an hour (roughly 4 times the average in the INDEX population) and broadband saves them roughly 20 hours a month of waiting per month (1 hour per business day), the $25 premium for broadband makes economic sense.

However, the average user in our population valued time at only 36 cents per hour. If they saved 20 hours of waiting time per month, this

[3]See Shapiro and Varian [1998] and Rohlfs [2001] for discussions of business strategy in the presence of network effects.

would imply a willingness to pay of only \$7.20 premium for broadband, which is much lower than current rates.

The numbers we used above are for metered service. We have seen that the premium people are willing to pay for flat-rate service may be as high as 50 percent, which would make the these figures significantly larger.

The question facing those who wish to sell broadband and complementary services is: what can be done to make broadband more economically attractive to users and providers?

First, since technical and administrative users appear to have the highest willingness to pay for broadband, ISPs would be well advised to think about ways to do some market segmentation. One approach would be to devise a range of premium services that would appeal primarily to this market. Examples would be things like Virtual Private Networks (VPNs) that provide secure connection to the corporate intranet, online backup services, home office support, and so on. This may be the best route to profitable deployment of broadband.

Secondarily, those who wish to see more rapid deployment of broadband should think about how they might nurture the development of new applications. The INDEX study was conducted before Napster became popular. There is ample evidence to suggest that this single application dramatically increased the demand for bandwidth on college campuses and among cable and DSL users.

Of course, both the bandwidth and the content was free to the end users, so we can't infer much about actual willingness to pay for online music. Providers are talking about \$9.95 as a monthly subscription for the content alone, so one could easily imagine music fans willing to pay an extra several dollars for broadband download of music.

If compelling content and applications that have mass market appeal become available, and price for the service declined, the prospects for broadband could change quite dramatically.

References

Sam Ames. Study: Broadband fees climbed in 2001. *CNET News.com*, 2002. `http://news.com.com/2100-1033-818013.html`.

Walter Beckert. Estimation of stochastic preferences: An empirical analysis of demand for Internet services. Technical report, Department of Economics, UC Berkeley, 1999.

Joseph Berchtold, Veit V. Dengler, Bonnie M. Johnson, and Srividya Prakash. What do broadband consumers want? *McKinsey Quarterly*, 4, 2001. `http://www.mckinseyquarterly.com`.

George Blumenthal. High speed internet access is all about price, not bandwidth. Technical report, NTL Powerpoint Presentation, 2002.

Karyen Chu. User reactions to flat rate options under time charges with differentiated quality of access: Preliminary results from INDEX. Technical report, Department of Economics, UC Berkeley, 1999. `http://www.marengoresearch.com/isqe/agenda_m.htm`.

Richard Edell. *Internet Demand Experiment: Technology and Market Trial*. PhD thesis, UC Berkeley, 2001.

Richard Edell and Pravin Variaya. Providing Internet access: What we learn from the INDEX trial. *IEEE Network*, 13 (5), 1999. `http://www.path.Berkeley.EDU/~varaiya/papers_ps.dir/networkpaper.pdf`.

General Accounting Office. Telecommunications: Characteristics and choices of Internet. *GAO Reports*, GAO-01-345, 2001. `http://www.gao.gov/docdblite/form.php?entry=1`.

Don Kridel, Paul Rappoport, and Lester Taylor. The demand for high-speed access to the Internet. In David G. Loomis and Lester D. Taylor, editors, *Forecasting the Internet: Understanding the Explosive Growth of Data Communications*, Topics in Regulatory Economics and Policy. Kluwer Academic Publishers, Boston, 2002.

Andrew Odlyzko. Internet pricing and the history of communications. In Lee W. McKnight and John Wroclawski, editors, *Internet Services: The Economics of Quality of Service in Netowkred Markets*. MIT Press, Cambridge, MA, 2001. `http://www.research.att.com/~amo/doc/networks.html`.

Samuel Paltridge. The development of broadband access in OECD countries. Technical report, Directorate for Science, Technology, and Industry, OECD, 2001. http://www.oecd.org/pdf/M00020000/M00020255.pdf.

Paul N. Rappoport, Donald J. Kridel, and Lester D. Taylor. Alternative approaches to analysis and modeling of residential broadband demand. In Robert Crandall, editor, *Broadband Communication: Overcoming the Barriers*. Brookings Institution, Washington, DC, 2002.

Jeffrey Rohlfs. *Bandwagon Effects in High-Technology Industries*. MIT Press, Cambridge, MA, 2001.

Carl Shapiro and Hal R. Varian. *Information Rules*. Harvard Business School Press, 1998. http://www.inforules.com.

Hal R. Varian. Estimating the demand for bandwidth. In Lee W. McKnight and John Wroclawski, editors, *Internet Services: The Economics of Quality of Service in Networked Markets*. MIT Press, Cambridge, MA, 2001. http://www.sims.berkeley.edu/~hal/people/hal/papers.html.