La Trobe University
Online Media Program

Research Report No. 3

Australia's Peak Demand for Internet Bandwidth: A modelling and forecasting methodology

Colin Richardson

December 1996

La Trobe University Online Media Program

The primary aim of the La Trobe University Online Media Program is to undertake social and policy research into the development of online media services. The Research Program has three research streams. They are:

Research Stream 1: New Media and Communication Environments
Organisational, regulatory and technological change is occurring at such a rapid pace that it is difficult decide on the outcome of change. Service providers are as uncertain as potential users about the viability or usefulness of new services. This stream of research examines the role of the new media in people's lives.

Research Stream 2: Emerging industry structures
The shape of the media, communication, publishing and computing industries is rapidly changing. Most of those changes are being brought about through strategic alliances between discrete elements of multi-functional and often competing organisations, and linkages with smaller start-up companies which have developed innovative products and services. For the first time "carriers" are having to pay attention to "content", software-based services and consumer electronics. These industry changes are occurring as the various industries "internationalise" and as governments remove layers of regulatory control. As a consequence the restructuring of media, communication, publishing and computing industries raise important questions for the industry participants and government policy makers.

Research Stream 3: Organisational innovations in research
New media and communications services have led to the emergence of organisations where management responsibilities are distributed globally. This stream of the program explores the use of online services in a national and international research context.

For additional details contact Associate Professor Peter B. White, Director, La Trobe University Online Media Program, Department of Media Studies, School of Arts and Media, La Trobe University, Bundoora, Victoria 3083, Australia. Phone +61 3 9479 2785 Fax +61 3 9479 3638 Email:pbwhite@latrobe.edu.au http://teloz.latrobe.edu.au/teloz/

The La Trobe University Online Media Program gratefully acknowledges the generous financial support of the Telstra Fund for Social and Policy Research.

Additional copies of this and other Research Reports published by the La Trobe University Online Media program can be obtained at a cost of $60 per copy from the Online Media Program or free of charge from the WWW at http://teloz.latrobe.edu.au/teloz/

La Trobe University Online Media Program Research Reports
Series Editor: Peter B. White

ISBN 1 86446 204 3 hard copy version
ISBN 1 86446 205 1 online version - http://teloz.latrobe.edu.au/teloz/

Revision: August 22, 1998

La Trobe University
Online Media Program

Research Report No. 3

Australia's Peak Demand for Internet Bandwidth: A modelling and forecasting methodology

Colin Richardson

December 1996

About the Author

Colin Richardson is a Research Associate in the La Trobe University Online Media Program. He is the principal of Internet Economic Research Associates, consulting economists specialising in the economics of the Internet and in economic strategies for cities, regions and organisations. His education in economics, statistics and accounting was gained at Monash University, Melbourne and Queens' College, University of Cambridge.

Over the past 30 years he has worked for the Malaita Shipping Company in Honiara (General Manager), the Committee for Melbourne (Economic Strategist), Her Majesty's Treasury in Whitehall (Economic Adviser on the Transport and Energy Industries), the Department of Management and Budget in Victoria (Chief Economist), the Parliament of Victoria (Economic Adviser to the Leader of the Opposition), the Gas and Fuel Corporation in Melbourne (Market Researcher), and the Royal Australian Navy in destroyers (War Service in Malay Peninsula, Borneo and Vietnam).

During 1990, Colin was President of the Economic Society (Victoria Branch) and he is a founding member of the Cambridge Society of Australia. His apers, articles and reviews have been published in *The Melburnian* and *The Australian Rationalist* magazines; the *Economic Papers, Regional Economics, Australian Economic Review*, and *Australian Universities' Review* journals; the *Solomons Voice* newspaper; the *Bisnis Nuis* newsletter; and the papers and proceedings of various Australian economic conferences.

Colin can be reached *via* electronic mail at colin@teloz.latrobe.edu.au, *via* telephone at (03) 6391-2623 and *via* postal service at "The Toll House", 72 Wellington Street, Longford, Tasmania 7301, Australia.

Contents

Executive Summary

This paper focuses on theoretical and methodological issues associated with forecasting import demand for bandwidth on Australia's principal trans-oceanic data communications link with the world Internet. "TelstraLink" now connects Paddington in Sydney with Hay in San Francisco. In 1988 it was pioneered by the University of Melbourne and NASA as a satellite circuit with a bandwidth of 64 kilobits per second. Between 1990 and 1995, it was progressively upgraded to a 4.5 megabits per second (Mbps) fibre optic circuit by the Australian Vice-Chancellors' Committee and CSIRO, which also funded the Australian Academic and Research Network (AARNet) domestic backbone.

Now jointly owned by the Telstra (Australia) and MCI (United States) Corporations, both the AARNet backbone and TelstraLink itself have been operated by Telstra Internet since July 1995. Despite massive investment to increase its bandwidth capacity more than nine-fold to 44 Mbps in just sixteen months, TelstraLink nowadays is facing excess demand pressures and Telstra Internet is receiving customer complaints during peak periods[1]. The situation would be far worse except for a number of other Australian network providers who also have been commissioning trans-oceanic Internet links. Though individually smaller, these private links in aggregate provide Australia with at least a further 46 Mbps of international capacity.

Demand forecasting for bandwidth capacity planning at present is done as a "bottom-up" and "domestic-to-international" exercise by Australian network providers such as Telstra Internet. Typically, they add capacity at points-of-presence on their domestic Internet backbones as customer requirements become known, then sum the bandwidths and expand the overseas links *pari passu*. This paper lays theoretical foundations and develops methodology for a new approach, *ie* a "top-down" and "international-to-domestic" forecasting tool is presented to complement and enhance the present system of Internet capacity planning in Australia.

Theory drawn from the disciplines of economics and marketing suggest that Australia's consumer and business requirements for TelstraLink bandwidth at the peak hour of each month can be modelled as a single demand function. This econometric equation may have up to seven "driver variables" for which long monthly time-series of data can be readily obtained. More problematic is the dependent variable for unconstrained peak bandwidth demand, which in principle can be constructed from information on peak monthly capacity, demand, quality-of-service, and load.

The research proposes that Australia's peak demands for TelstraLink bandwidth are driven by up to seven influences, *viz* the Telstra Internet tariff schedule, gross domestic product (GDP), the prices of computer/modem equipment needed to access the Internet, the increasing "content-richness" of World Wide Web pages, the overall size of the world Internet (presently 13 million computers are reachable), the level of employment within Australia, and average weekly earnings. This information was used to program a forecasting tool and make (demonstration only) projections of peak monthly unconstrained import demand for TelstraLink bandwidth, with the forecasting horizon being November 1997.

The Base Case forecasting run was supplemented by a set of Perturbed Case forecasts (effectively, sensitivity analyses) based on changing the monthly growth rate test data characterising all seven driver variables in the Base Case. When more information becomes available, good estimates of the parameters of our econometric demand equation will replace the dummy own-price, income, cross-price, and other elasticities used for this set of mere demonstration forecasts.

[1] Link capacities are being constantly upgraded. Link capacities quoted in this document refer to December 1996, unless otherwise stated.

It is suggested that a Stage 2 project be initiated. With access to appropriate demand information, this methodological exercise could form the basis of an authoritative empirical study. The benefits of deriving accurate estimates for parameters such as price, income and cross elasticities of demand include making forecasts of Australia's peak demand for TelstraLink bandwidth based on several plausible scenarios. This can only improve the practice of capacity planning and lead to more efficient outcomes.

Australia's Peak Demand for Internet Bandwidth: A modelling and forecasting methodology

Chapter 1

The Internet

Introduction

This chapter provides essential background information for understanding the theoretical and methodological issues associated with forecasting peak import demand for bandwidth on Australia's principal trans-oceanic data communications link with the world Internet ("TelstraLink"), which now connects Paddington in Sydney with Hay in San Francisco. TelstraLink is jointly owned by the Telstra (Australia) and MCI (United States) Corporations and is operated by Telstra Internet, which also controls Australia's principal Internet backbone network.

Eleven other Internet network providers (INPs) compete with Telstra Internet. The Appendix shows that, individually, these INPs operate smaller domestic backbones and overseas links, but in aggregate the international bandwidth capacity they provide is more than that of TelstraLink.

Demand forecasting for bandwidth capacity planning at present is done as a "bottom-up" and "domestic-to-international" exercise by Australian INPs such as Telstra Internet. Typically, they add capacity at points-of-presence on their domestic Internet backbones as customer requirements become known, then sum the bandwidths and expand the overseas links *pari passu*. This paper lays theoretical foundations and develops methodology for a new approach, *ie* a "top-down" and "international-to-domestic" forecasting tool is presented to complement and enhance the present system of Internet capacity planning in Australia.

LANs and the PSTN

The simplest local area network (LAN) uses a hardware cable to connect two computers which communicate *via* software protocols. These determine how the two machines structure, interpret and deal with the streams of binary data passing between them. Today, the most popular LAN system is Ethernet, typically connecting as many as 100 machines in a bus topology using half-inch thick coaxial cables up to 500 metres in length. By bridging several such clusters together, an extended LAN can be engineered to service all floors in a skyscraper or all buildings on a campus. Networking beyond such sites, however, requires either dedicated telecommunications links or a dial-in connection to the telephone system.

Such "internetworking" relies heavily on the world-wide public switched telephone network (PSTN), owned and operated by telecommunications companies (telcos) like Telstra, MCI, British Telecom, Cable and Wireless, Sprint, and AT&T. It is these telcos which lease to customers the long-distance data circuits needed for permanent internetworking and which provide the means for "data over voice" sessional dial-in connections.

All the world's telephone devices (handsets, fax machines, modems) are interconnected *via* the global PSTN. Its copper wire (twisted-pair, coaxial cable) and fibre optic cables run underground, overhead and subsea. Its microwave and radio links carry voice and data conversations between antennae on towers and transponders on geostationary satellites. By dialling its unique "address" (country code + area code + telephone number), any one of more than a billion telephone devices can be reached from any other.

It is "circuit switching" which establishes a physical communications path between two devices for the duration of a call session. Circuit switching typically takes place within telephone exchanges situated *en route* from the calling to the called device. Once a circuit is established, the voice, fax or data conversations are carried as a mix of analog and digital signals which modulate an electronic carrier wave - or else stimulate photonic pulses along fibre optic cables.

While both ends of the circuit (the "subscriber loops" between the devices and the first and final telephone exchanges) are dedicated to the two participants in the call, it is a different story between the *en route* telephone exchanges. By using statistical multiplexing (MUX) plus analog-to-digital (ADC) and digital-to-analog (DAC) conversion, numerous voice and data conversations can be made to share each physical circuit. These intermediate links between telephone devices become *virtual circuits*.

Technical progress - including virtual circuits, digitisation of telephone exchanges and replacement of copper wire with fibre optic cable and wireless links - has brought about a long-term decline in the real unit costs of operating the global PSTN.

WANs

Many corporations, governments, universities, and research institutions have LAN clusters in different suburbs, cities, regions, and even countries. By interconnecting these LANs using leased lines (between heavy-traffic pairs) and dial-in connections (for light-traffic pairs), they can create a wide area network (WAN), with significant commercial benefits.

In the first place, the decline in real PSTN costs has been reflected in "tariff rebalancing" (rather than overall tariff reduction) by the telcos, *ie* they have reduced the price of local calls while keeping long-distance and international prices artificially high[1]. Large geographically dispersed organisations can avoid these costs by leasing lines to handle most of their intra-company voice, fax and data traffic. Secondly, by adopting the networking model of "client-server computing" over a wide area, they can "warehouse" their data and have all departments in all locations feeding off company-wide databases held on huge fileserver machines, thereby saving the costs of duplicating common marketing, customer, sales, purchases, inventory, accounting, and technical information.

Finally, if there is spare capacity on their WAN, the organisation also can become a network provider and sell data communication services to the public. Of course, once WAN technology became available at reasonable cost, some companies were set up as pure network providers from the outset. Originally, "online service providers" like America Online, CompuServe, OzEmail, and the Microsoft Network offered proprietary services developed in-house, but nowadays their extensive networks have become integrated with the Internet.

The Internet

The Internet is the world's largest and fastest-growing WAN. It is defined by its constituent networks adhering to one particular suite of 100-plus open-standard protocols. This is known as Transmission Control Protocol/Internet Protocol (TCP/IP) and was named after the two key component "layers" of the suite.

TCP/IP had been under development since 1974 by Robert Kahn and Vinton Cerf, two scientists working for the Advanced Research Projects Agency (ARPA) of the United States Department of Defense. It was officially adopted as that agency's standard networking protocol suite in January 1983. The Internet (*nee* ARPANet) had been born in September 1969, out of Cold War fears that vital circuit-switched communication links could be disrupted by nuclear weapons. This original ARPANet "internetwork" was expected to interconnect no more than 128 separate computer networks.

Kahn and Cerf came up with a solution that was both simple and elegant. They decided to overlay the circuit-switched PSTN with a "packet-switched" network, to handle messages that have been chopped up into tiny, individually-addressed "packets" of binary data. In place of telephone exchanges, packet-switched networks like the Internet have "routers" - special computers loaded with routing table software, which shunt trains of these data packets in the general direction of the router serving their ultimate destination. This often involves numerous "hops" *en route* along a chain of intermediate routers. If a damaged physical circuit is encountered, the packets simply flow around the obstruction *via* alternate paths, selected from the routing table held in the latest machine encountered. Nowadays there are "learning routers" which can sense when a new WAN, LAN or computer has been added nearby, update their own routing tables, then propagate the information to other routers on a "need-to-know" basis.

To be part of the Internet, a computer must have a TCP/IP software "stack" to implement the layered architecture of this protocol suite. At the topmost applications layer, a user's software program (say, her E-mailer) sends its data down through the local computer's stack, which encapsulates them in packets of up to 1,500 bytes (12,000 binary digits or "bits", including source and destination addresses, sequence number, *etc*) then transmits these *via* the lowest physical layer. After several router hops, these packets arrive at the remote computer's physical layer and travel up through its stack, eventually reaching the topmost applications layer and appearing as an E-mail message. The intermediate IP layers in the stacks at each end simply address and send data packets; it is the TCP layers which assemble them into correct sequence and ensure that all lost or corrupted packets are retransmitted.

To be part of the Internet, a computer also must have an "IP address" and a "domain name". At present, all Internet-enabled machines have four-byte (*ie* 32-bit) IP addresses. These are the binary equivalents of four-part decimal number codes such as 131.172.26.102 and are used to identify the source and destination of each data packet passing across the Internet. The more user-friendly domain names are of the form computer.owner.subdomain.domain (*eg* teloz.latrobe.edu.au). These are resolved into their corresponding IP addresses by a local Domain Name Server (DNS) machine whenever commands are issued to named computers from a user's keyboard or mouse.

The Internet today is a sprawling global internetwork of more than 134,000 addressable packet-switched computer networks and every 30 minutes another LAN or WAN is connected to it[2]. These networks all interconnect, both within and between 175 of the world's 237 nation-states[3]. It is the largest WAN on the planet, comprising almost 13 million addressable machines (called "hosts") organised into some 488,000 addressable operating clusters (called "domains")[4] to which around 36 million users should have access by December 1996.[5]

The number of hosts has been doubling every 12 months since 1988,[6] following almost 20 years of slow but steady growth from ARPANet's original 1969 base of four university mainframe computers communicating *via* Network Control Protocol (NCP) software and Information Message Processor (IMP) hardware. NCP was the forerunner of TCP/IP and the IMPs evolved into today's Internet routers.

The Matrix and Cyberspace

While the ARPA scientists were working on TCP/IP, other networking protocols were under development by computer enthusiasts in the university, corporate and even household sectors. This resulted in the international Unix-to-Unix Copy Program (UUCP), BITNet and FidoNet networks - plus several others. All these packet-switching networks have E-mail "gateways" to and from the Internet but are not really part of it. FidoNet is perhaps the most interesting development. Individuals interested primarily in messaging and computer gaming began writing software to create electronic "bulletin boards", to which participants could dial-in for uploading and downloading games programs. Eventually, many Bulletin Board Systems (BBS) throughout the world were linked together after the FidoNet protocol suite was developed in 1983 - the same year that TCP/IP replaced NCP as the universal Internet protocol suite.

The global network of all computer networks which can swap E-mail messages (and share some other store-and-forward services, like FTPmail) is known as “the Matrix”, a term borrowed from the science fiction novels of William Gibson. The Matrix has around 39 million users[7] (including the 36 million Internet users) in 186 nation-states[8]. “Cyberspace” is simply the universe of all networked computers. This includes extensive WANs like Minitel, which is independent of the Internet and the Matrix. Minitel grew out of France Telecom’s decision to replace all its printed telephone directories with a simple keypad computer terminal. According to Francois Olibet, Director General of France Telecom Intelmatique, “Right now we have 6.5 million Minitels in France. This means we have about 20 million users, assuming that there are 3 people per Minitel. So we have a penetration which is almost 30% of the total population of France.” Minitel is used by more than 10,000 companies to supply about 25,000 different services[9].

The World Wide Web

Following the birth of the Internet in September 1969, several client-server protocols and applications came into general use. There were *store-and-forward* services like electronic mail (E-mail), mailing lists (Listserv, Listproc, Majordomo, etc) and newsgroups (Usenet - originally developed on the UUCP network). There were *real-time* services like remote login (Telnet and Rlogin), Internet relay chat (IRC and Talk) and multi-user domains/object-oriented (MUDs and MOOs) for playing interactive computer games such as Dungeons and Dragons. And there were *software distribution* services like archive searching (Archie) and anonymous file transfer protocol (FTP).

But it was *information discovery and retrieval* services that spawned numerous client-server protocols for such tasks as remote library catalogue searching (Hytelnet), document databases (Gopher), wide area information servers (WAIS), and document database searching (Veronica). In March 1989, Tim Berners-Lee of CERN (the European Laboratory for Particle Physics in Geneva) cut this Gordian knot by proposing a World Wide Web (WWW) of cross-linked or “hypertext” documents, successfully demonstrating his concept in December 1990.

His hypertext transfer protocol (HTTP) was developed to serve up documents that have been modified using the hypertext mark-up language (HTML). World Wide Web documents contain a unique feature called “hyperlinks” - specially marked-up portions of text with “pointers” to related documents. It was rather like footnoting in scientific papers, except that one could retrieve the full text of the reference, no matter which of the world’s Internet servers physically held the hyperlinked document. Each HTML document was assigned a Uniform Resource Locator (URL) which resembled a DOS or UNIX filepath (*eg* http://www.w3.org/default.html) and one could retrieve it simply by clicking on this URL - or some highlighted/underlined text representing it.

Soon, the Common Gateway Interface (CGI) concept made its appearance. CGI scripts allowed one to create URLs for virtually all existing Internet protocols (*eg* ftp://, news://, telnet://, gopher://, wais://) so that most services could now be accessed from a common “Web browser”, instead of from highly-specialised software clients. The original CERN browsers were text-only, but in June 1993 Marc Andreeson of NCSA (the National Centre for Supercomputing Applications in Illinois) demonstrated the first graphical user interface (GUI) Web browser, named “Mosaic”. It was released as “freeware” in September 1993 and thousands of copies were downloaded by users, marking the start of the World Wide Web’s rise to dominance as the key component of today’s Internet.

With the development of “socket” communications software, computers with non-UNIX operating systems like DOS, Windows, MacOS, and OS/2 could be Internet-enabled, *ie* they too were fitted out with TCP/IP stacks. Today, the owner of virtually any modem-equipped PC can open an account with an Internet service provider (ISP), install the necessary socket/browser software and join the Internet as a dial-up user. Descendants of the original Mosaic browser (Netscape Navigator, Microsoft Explorer, etc) nowadays have E-mail and Usenet newsreader clients built in. Individuals also can use the HTML language to establish a personal "Website" (one "home page" plus numerous other pages of text/graphics/audio/video information) on a Web server residing on their own or their ISP's computer.

"Webspace" is the universe of all hyperlinked Web pages held on these servers throughout the Internet. Early attempts to navigate Webspace were made with the assistance of directories that had been classified by human intervention, the best-known of these being "Yahoo", which was developed at Stanford University. Nowadays, the indexes are built automatically by swarms of tiny "robot" software programs (often called Webcrawlers or Spiders), which examine every Web page and Usenet posting, then file the information resources they have discovered back at their home base. This will consist of a "disk farm" (an array of hard disk drives) to store the index and a computer with a server called a "search engine" to interrogate the index each time a remote client submits a search request structured according to Boolean logic, *eg* find all Web pages containing the words "Melbourne" and ("capital" or "city") near "Victoria" but not ("Lord" or "Viscount" or "Florida"). One of the largest search engines (called Alta Vista) handles 21 million queries each weekday from people searching its index of 15 billion words on 30 million Web pages and four million postings to14 thousand Usenet newsgroups[10].

While text and graphics capabilities are standard on all Web servers and their browser clients, new multimedia Internet protocols and services are being announced on almost a monthly basis. Examples include streaming audio and video (RealAudio and CUSeeMe), application programs that execute in the browser rather than the server (Java applets), virtual reality "worlds" (Live3D), content rating schemes (PICS), voice telephony (Webphone), and electronic commerce tools for authentication (Digital Signatures), security (PGP Encryption) and payments (Cybercash). The associated client software can be seamlessly integrated with one's current browser as either a "helper application" or a "plug-in". When the next version of the browser is released, the most popular helper apps typically become plug-ins and the "killer application" plug-ins are upgraded to "built-in" status.

Most browsers and other clients are simply given away as freeware for unrestricted downloading from Internet software archives by individuals and not-for-profit organisations. The software developers gain most of their revenue by (a) licensing corporations and governments to use these same Web clients and (b) selling them the associated Web server software. It is an unusual and attractive business model. By offering free copies of the client, the software development company strives to establish its new application as a *de facto* World Wide Web standard, thereby stimulating the (a) and (b) markets for the new product.

Webspace presently comprises some 90 million Web pages[11] held on more than 825,000 Web servers[12]. The number of Web sites is doubling every six months[13]. Compared with the 13 million Internet hosts, this does not seem very many. But after allowing for numerous routers, gateways and other unattended hosts in the "support Internet", most of the remaining machines are part of the "consumer Internet". Users of this consumer Internet certainly interact *via* E-mail, IRC and many other protocols, but mostly they download information from hosts comprising the smaller "core Internet" - the realm of ISPs, electronic commerce, search engines, software archives, and information-providers - whose fastest-growing segment is the World Wide Web.

Once every minute, a new Web page is added to this globally distributed hypermedia information system. The present generation of search engines and Webcrawlers seems incapable of clearing the backlog, let alone indexing every new Web page as it appears. Alta Vista covers about one-third of the Web and newer search engines like HotBot[14] and Infoseek Ultra have indexed more than 50 million of the 90 million pages that comprise Webspace.

Competing and Complementary Services

Internet applications are beginning to compete with a range of traditional (usually networked and geographically dispersed) services, along a spectrum of product substitution possibilities ranging from very strong to fairly weak.

There is competition between E-mail and traditional postal delivery services. Streaming audio applications are being used to take AM (and even FM) radio station feeds onto Web sites and streaming video is in the early stages of doing the same with television programmes. Already there have been many specialised Webcasts of news events, music concerts and sporting matches, while IRC competes with talkback radio as a highly interactive electronic medium. There are now Websites entirely devoted to gambling on the world's classic sporting events. With hundreds of daily newspapers and monthly magazines now on the Web, publishers are widening their audience reach. During 1997, Web advertising in the United States is bidding fair to absorb more dollars than radio advertising. The PSTN infrastructure is being exploited daily by the developing technology of Internet telephony, with early adopters making international calls at a fraction of the price charged by telcos.

In the scientific and cultural spheres, E-journals have in some cases completely replaced the less timely and more expensive paper versions, while full-text versions of classic books in the public domain are available on numerous Websites. The world's cultural archive of paintings, cinema films, sculptures, and *objets d'art* is gradually being digitised and placed on Web servers. In virtual reality, one can take tours of famous museums, monuments and art galleries. Governments are digitising their records and making them available on the Internet, along with the capability of submitting electronic forms to make registrations or access services.

Electronic commerce is the major area of complementarity with traditional modes of doing business. Electronic Data Interchange (EDI) between companies is now commonplace. Once the problems of authentication, security and payments have been solved, small wholesale and retail companies with fine products but limited capital will be able to sell directly into the global marketplace. Already, such goods as computers, software, compact disks, books, and wines can be ordered *via* the Internet for delivery anywhere in the world. The financial sector - banking in particular - will be a major participant in (and beneficiary of) E-commerce when it takes off.

Convergence of Computing, Communications and Content

The common thread linking those traditional services which are in the strongest competition with those available through the Internet is that they are two-tiered. Take publishing, for example. The literary and music publisher historically has acted as a "gatekeeper", exercising a form of quality control over the actual content-providers (authors, journalists, musicians, *etc*) - often to the chagrin of this lower tier of producers. The Internet, however, shows signs of reducing the significance of the upper tier. Today, any content-provider with a TCP/IP connection can set up a personal Web site and practice "vanity publishing" at a fraction of the previous cost of this activity.

New protocols like the Platform for Internet Content Selection (PICS), originally designed to help parents and teachers filter out material unsuitable for young children, are now being used to rate content in many other dimensions - quality, acceptability to certain religious groups, "political correctness", and so on. Meanwhile, independent writers, artists, musicians, scholars, cartoonists, radio announcers, film directors, software developers, and other creatives are hailing "triple convergence" for the new freedoms it offers.

Trend Towards Increasing Filesize and Bandwidth

Users typically operate the Internet by downloading files, which then are stored in a "cache" on their own hard disk drives or viewed on their screens - usually both. Filesizes range from perhaps one kilobyte (kB) for a short E-mail message to more than 10 megabytes (MB) for a one minute full-motion video clip with photo-realistic colour and CD-quality surround sound. The maximum speed at which files can be downloaded depends partly on whether queues have developed at (or *en route* from) the remote server. But it mainly depends on the "bandwidth" of the slowest communication device or link between server and client.

If a 1500 byte packet (which contains 12 kilobits of data) passes a point on the circuit in one second, the bandwidth of that device or link is said to be 12 kilobits per second (12 kbps). Many fibre optic links have a bandwidth of 155 megabits per second (Mbps) and "gigabit networking" will soon be quite common on the Internet. But if the packet-switched circuits multiplexed into these links start transmitting ever-larger files, then overall quality-of-service must fall: a rising percentage of the packets are dropped then re-transmitted by the TCP layer in the server's stack. This markedly increases network "latency", the time it takes for any one packet to traverse the chain of devices and links connecting client to server.

Typically there will be a modem in the subscriber loop at one or both ends of the circuit. The most common telephone modems have speeds of 14.4 or 28.8 kbps and quite often these are the bandwidth-limiting devices. This bottleneck can be eased by subscribing to the Integrated Services Digital Network (ISDN) service which most telcos now offer as a PSTN overlay, because ISDN terminal adaptors have bandwidths of 64 kbps - and multiples thereof. By convention, anything below 64 kbps is "narrowband" and anything above it is "broadband". With the roll-out of broadband hybrid fibre coaxial (HFC) networks in Australian cities by Telstra and Optus comes an opportunity for connecting to the Internet *via* cable modems with bandwidths of (initially) 10 Mbps and (eventually) 30 Mbps.

There is a definite trend towards increasing filesize and bandwidth. Virtually every newly-developed Internet application requires larger software clients and handles files containing far more kilobytes of information than any earlier client, protocol or service. The trend towards increasing bandwidth is even more marked. Despite a significant degree of data compression (176-to-one is claimed for streaming audio), these ever-growing files will not download at a rate acceptable to users unless ever-higher bandwidth capabilities are provided to the telephone/cable modems and ISDN terminal adaptors which feed the information onto the user's screen and hard disk drive.

It seems that strong consumer preference for richer content is fuelling the industry's relentless drive for faster connections on the one hand and bandwidth conservation on the other. Efforts to husband scarce bandwidth have progressed way beyond the early "netiquette" practices of using abbreviations in E-mail messages and downloading software from the geographically closest FTP server. Innovations since then have included hardware and software file compression/decompression (codec), technologies like Asymmetric Digital Subscriber Line (ADSL) to convert narrowband copper wire into a broadband medium at speed ranging from 1.5 to 10 Mbps, file caches within both clients and servers, domestic mirror sites for the most popular overseas FTP, Gopher and Web servers, Java applets which execute on the client (not the server) host, and the TeleVisual Modelling Language (TVML). Just as a short set of MIDI instructions can recreate locally a music performance that otherwise would have to be downloaded as a bloated WAVE soundfile, so a TVML server avoids sending a completely rendered 3-D virtual reality "world" to the client by simply transmitting a short file of instructions to instantly recreate it on the user's screen.

Future of the Web, Internet and Matrix

The top five Internet services in order of traffic carried on TCP/IP backbone networks are HTTP, FTP, Usenet, E-mail, and Gopher[20]. Significantly, these and virtually all other applications can be launched from a single client: the multimedia Web browser. Electronic commerce will become ubiquitous once the problems of authentication, security and payments have been solved and, significantly, all the relevant E-commerce protocols are being developed as Web browser built-ins, plug-ins and helper apps. Thus we may expect that, in time, the World Wide Web will completely colonise the Internet.

TCP/IP enjoys a “network externality” - a prospective new networker can reach more hosts using this than proprietary protocol suites like Novell’s Internetwork Packet Exchange (IPX). Also, in the 27 years since their invention, the heterogeneity of packet-switched network systems has increased markedly. Ethernet, HFC, ISDN, and ADSL have been mentioned already, but others include Token Ring, Fibre Distributed Data Interface (FDDI), Frame Relay, Switched Multimegabit Data Service (SMDS), Asynchronous Transfer Mode (ATM), and Synchronous Digital Heirarchy (SDH).

The best proven approach to internetworking this diverse range of LAN and WAN technologies on the backbones, tailends and links of complex, multi-topology networks is the open-standard TCP/IP suite of protocols. These considerations suggest that the Internet, in time, should completely colonise the Matrix as well. However, apart from an eventual gateway to France Telecom’s unique Minitel online network, there seems no compelling reason for the Web/Internet/Matrix to extend its reach into the remainder of Cyberspace, the domain of experimental research networks and “extranets”, *ie* highly-secure TCP/IP networks having no connection with the Internet itself.

The present wave of telco commercialisation, privatisation, internationalisation, and withdrawal of state monopoly privileges should see the benefits of technology-driven unit cost reductions translated into lower PSTN tariffs for trunk and international calls. As the Full Service Digital Network (FSDN) based on fibre optic cables and photonic switching devices washes out the last lingering traces of electronic analog telephony, today’s voice/fax/data distinctions will become meaningless, with “data dialtone” being continuously available at the wall socket. Bandwidth should become a commodity, whose price will track its ever-reducing cost and the International Telecommunication Union cartel may collapse. Continued high growth of the World Wide Web is likely in the developed world, levelling off from around 2005, leaving only Africa and Latin America (plus the Pacific and Caribbean islands) with low levels of narrowband Internet access.

Chapter 2

Australia's Tran-Oceanic Internet Links

The AARNet Link

Australia first joined the Internet in 1988, when the University of Melbourne leased 64 kbps of bandwidth on a packet-switching data communications satellite link to the NASA Ames Research Centre in California. Ames is a laboratory complex of the National Aeronautics and Space Administration, which funded the link jointly with the Australian Government. This became Australia's first gateway to NSFNet, the United States Internet backbone network originally developed by the National Science Foundation to link five supercomputing centres. The government-subsidised NSFNet had replaced the original military ARPANet backbone and, from May 1995, was itself shut down in favour of several large national and regional WANs that had been developed by private enterprise.

In 1990, the University of Melbourne link was acquired and enlarged by the Australian Vice-Chancellors' Committee and the Commonwealth Scientific and Industrial Research Organisation. This pioneering international link soon was supplemented with an extensive domestic backbone network called AARNet, which interconnected all Universities, the CSIRO and certain other research institutions. Other organisations in the government, corporate and education sectors which agreed to abide by an "Acceptable Usage Policy" were permitted to connect their own LANs with AARNet as well.

TelstraLink

This developing commercial side of Australia's Internet backbone was taken over by Telstra Internet from July 1995. At this time the international link was experiencing massive excess demand, despite its bandwidth having been expanded by AARNet to 4.5 Mbps. By November 1996, Telstra Internet had built up TelstraLink's capacity to 42 Mbps, a more than nine-fold expansion over just 16 months. Telstra Internet also has commissioned a further 2 Mbps of capacity to New Zealand.

Private Links

Side-by-side with these developments came a rash of new entrants into the business of providing Internet backbone networks for Australia's ISPs and other large customers from the government, corporate and education sectors. Some of these Internet network providers (INPs) decided to supplement their access to TelstraLink by leasing trans-oceanic capacity from Telstra, Optus and overseas telcos like MCI, British Telecom, Sprint, and AT&T. The Appendix suggests that the capacity of all twelve trans-oceanic Internet links between Australia and the rest of the world lay somewhere between 90 and 128 Mbps.

All of Australia's smaller domestic Internet WANs are interconnected with the Telstra Internet backbone network and utilise TelstraLink, despite several INPs operating separate trans-oceanic links of their own (see Appendix). Thanks to methods which include automatic routing arbiter protocols, any excess of demand over installed capacity tends to be shunted in the direction of Telstra Internet, which has become the *de facto* "network provider of last resort" to Australia's smaller Internet network providers.

Size and Growth of the Australian Internet

Australia's domestic Internet today comprises almost 400,000 hosts, organised into some 10,000 domains[15] and operated by up to one million users[16]. The size of the Australian (.au) "top-level domain" has been more than doubling every 12 months, *ie* it is growing faster than the global Internet. In terms of number of hosts, Australia is now the sixth-largest country of the 175 nation-states on the Internet[17], but in population per host we rank equal second with Sweden[18] and in hosts per million US dollars of gross domestic product we rank second after Iceland[19].

Of all infrastructure elements put in place to guarantee effective and efficient operation of the Australian Internet domain, it is TelstraLink which has generated the most public controversy. The problem is excess demand. Despite the more than nine-fold capacity expansion from 4.5 to 44 Mbps by Telstra Internet since it took over from AARNet, many commentators expect TelstraLink to experience further congestion with consequent higher packet losses, increases in latency and reductions in quality-of-service.

Satellites and Submarine Cables

Australia's international telecommunications are carried by geostationary satellites and submarine cables. The satellites typically are parked over the Indian and Pacific Oceans and are owned by INTELSAT, INMARSAT or by private enterprise - often in association with global media interests. However, the principal earth stations or "teleports" in Perth and Sydney (for up/downlinking the satellite transponders) are controlled by Telstra. The submarine cables connect Australia with vast trans-oceanic networks built by international consortia of telcos (including Telstra in some cases) and operating at speeds of 2.5 gigabits per second (Gbps) or more using fibre optic technology.

Telstra Corporation is a founding member and major investor in the International Telecommunications Satellite Organisation, which operates a global network with 24 INTELSAT satellites in orbit. Also, a new generation of INMARSAT satellites are now being launched, with the final three due by October 1997. Private satellites accessed by Australia (principally for broadcasting) include PanAmSat-2 and AsiaSat-2.

The future of satellite communications seems to lie with "necklaces" of Low Earth Orbit (LEO) satellites encircling the planet. Pocket-sized cellular telephones with built-in satellite antennas will allow voice and data communications between even the most remote areas. Plans for LEO satellite networks are being developed by ICO Global Communications, Odyssey, Globalstar, Motorola ("Iridium"), and Teledesic - the latter having participation by Microsoft Corporation and calling for a flight of 840 satellites to be launched into low earth orbit.

Telstra also has substantial investments in submarine cable systems. These include the 16,500 kilometre South Pacific Cable Network linking Australia with major destinations in Asia, North America and Europe. In a joint venture with Optus, Telstra soon will commission the 2,800 km JASURAUS cable, linking Port Hedland with Indonesia at 5.33 Gbps and thence *via* the Asia-Pacific Cable Network to Hong Kong, Japan, South Korea, Malaysia, Philippines, Singapore, Taiwan, and Thailand.

Two fibre optic submarine cables are of key importance for the Internet in Australia. PacRim-West runs from Sydney to Los Angeles *via* Guam and Tokyo. PacRim-East runs *via* New Zealand and Hawaii. TelstraLink normally uses PacRim-West, with satellite backup should (say) the Sydney-Guam sector go down due to a cable break. The Tasman 2 cable connects Australia with New Zealand. These digital submarine cables also are backed up by the older Tasman, AIS and A-PNG analogue cables.

While most of these overseas links are owned and operated by Australia's licensed telcos (in consortium with overseas partners), significant capacity is leased by them to those interests who need secure international telecommunication circuits, *eg* the military, diplomatic corps, news services, entertainment companies, and banks. In some cases, the lessees act as

Internet network providers, offering connectivity to Australian ISPs. These INPs include the IBM Global Network, CompuServe Pacific, Connect.com.au, and AccessOne.

Fibre Optics

Whereas copper wire and radio/microwave links use electronics for signalling, fibre optics employs photonics. A laser beam sends pulses of coherent light (representing binary zeroes and ones) along hair-thin glass fibres encased in a protective sheath. Signal attenuation can be as little as one-millionth that of an equal length of copper wire, which requires a chain of repeater stations to maintain the original signal strength. All that prevents pulses travelling at the speed of light along the full length of a fibre optic cable are carrier conversions within the *en route* switching devices, which are mainly electronic. Photonic switches, however, should replace these in a matter of years.

Modern submarine cables like FLAG employ this technology. This Fibreoptic Link Around the Globe project was launched in 1996 by a consortium of AT&T Submarine Systems and Japan's NDD Submarine Cable Systems. The $US1.5 billion project will undertake to lay undersea fibre optic cables from England to Japan, with landing points in Europe, the Middle East, Africa, and Asia. Capacity will be 120,000 circuits, each of 64 kbps. Some 50 companies from around the world have agreed to use the FLAG infrastructure.

Fibre optic technology also is being used to replace the narrowband copper wire links between domestic telephone exchanges. On the Sydney/Melbourne sector, there already are 2 x 20-fibre 560 Mbps fibre optic cables and 4 x 36-fibre 2.4 Gbps cables. In addition, Telstra recently built a 5,000 km fibre optic cable linking Brisbane, Sydney, Canberra, Melbourne, Adelaide, and Perth - the first installment of a 2.5 Gbps ring around the continent, with an additional north-south link and a connection to JASURAUS at Port Hedland. Scheduled for completion by the end of 1998, it will be capable of carrying 32,000 simultaneous telephone calls or transmitting the text of 150,000 A4 pages every second.

At the present time, both Optus Vision and Telstra Multimedia are rolling out new hybrid fibre coaxial (HFC) cable networks, principally to carry pay TV signals downstream to subscriber premises. From the television studio to the local nodes to the hundreds of homes served by each, the sum total of downstream signals is a voltage changing at frequencies up to 750 megahertz (MHz). This voltage is converted into a similarly changing brightness of coherent light for optical fibre, then back into a voltage for coax to arrive in the home, where the TV signals are tuned and demodulated by a set-top box. The Optus cable also will carry telephone conversations and both telcos are going to offer broadband data communications services for the bandwidth-hungry Internet.

The cable modems presently available in Australia are rated at a speed of 512 kbps, but a device capable of 10 Mbps downstream and 19.2 kbps upstream is to be released later this year. Later there will be faster *asymmetric* modems rated at 30 Mbps downstream and 8 Mbps upstream. The Internet community in Australia is urging both Telstra and Optus to switch to high-speed *symmetric* cable modems in the interests of interactivity. If all users were to have server as well as client software on their machines, the distinction between the Core and Consumer Internet would break down.

Telcos in Australia and Overseas

Until the market is opened to competition in July 1997, only three telcos have been granted licences to operate within Australia: Telstra, Optus and Vodafone. Telstra owns and operates Australia's "legacy" PSTN, but newer entrants Optus and Vodafone access it under "interconnect agreements". All three companies are upgrading and/or expanding the nation's telecommunications infrastructure, *eg* the broadband HFC networks mentioned above. All three telcos have built urban cellular telephone networks as well. Telstra is expected to fully digitise its circuit-switched terrestrial network and introduce high technology "fixes" like Asymmetric Digital Subscriber Line (ADSL) to squeeze more bandwidth out of it.

The foreign telcos presently of most significance for the Internet in Australia are partners in the nation's trans-oceanic submarine cable and satellite links. These include MCI, AT&T, British Telecom, Sprint, UUNet, and New Zealand Telecom. Optus is part-owned by BellSouth (United States) and Cable & Wireless (United Kingdom). From July 1997, all these players should expand their operations and many new entrants are expected.

Chapter 3

Australia's Peak Demand for Internet Bandwidth

What is the Product?

Like Internet users in the rest of the world, Australians operate the Internet by commanding local clients to download datafiles or datastreams from remote servers. These users desire *information* which is highly specific to their own needs and it may take the form of text, graphics, sound, video, or multimedia. Ultimately, where the data resides on computers outside Australia, from all their individual keystroke and mouse-click commands is derived an aggregate demand for *bandwidth* on Australia's trans-oceanic Internet links. This is the product which is the focus of our study, so let us define it rigorously.

Suppose the time is 11.00 pm on a typical Sunday night. Australian users, as a group, are downloading simultaneously so many and such large files from servers located overseas that their demand for bandwidth exceeds the capacity of TelstraLink. Data packets start queueing up at the San Francisco router. Those not sent towards Sydney within a certain time frame (the "packet loss percentage") are simply dropped. This won't slow down TelstraLink itself, but it does mean that all users must wait longer than is technically necessary for their *particular* downloads to begin. Once commenced, such downloads also will take longer than is technically necessary as the TCP layers at each end of TelstraLink communicate to organise retransmission of all those lost packets.

Some users simply will hit the stop button on their Web browsers, out of sheer frustration and/or because the opportunity cost of their time is high. Much of this *constrained* peak demand will be shifted to other (*ie* offpeak) periods.

Demand which has been so artificially limited cannot be used to model consumer behaviour. We need some measure of *unconstrained* demand for bandwidth. Such demand is growing month-by-month but fluctuates markedly day-by-day and even hour-by-hour. Data pipes like TelstraLink are most economically sized to meet some large fraction of *peak* (not average or offpeak) unconstrained demands. Finally, we must choose a standard time period during which these demand peaks can be measured.

In accordance with the aims of this study, we formally define the product as: Australia's monthly peak unconstrained demand for bandwidth on TelstraLink, measured in megabits per second (Mbps) of dataflow. Demand is considered to be unconstrained whenever packet losses stand at one percent or less. A method will be proposed for using the observed packet loss percentages to convert the (sometimes capacity-limited) measured *loads* into the unconstrained *demands* that Australian Internet consumers actually exert on TelstraLink, in response to a range of influences that "drive" their consumer behaviour.

Who are the Consumers?

Since TCP/IP packet-switched networks typically ride upon the PSTN infrastructure, one might expect to find a simple Internet "value chain" running straight from the telcos through the Internet network providers (INPs), who wholesale bandwidth to Australia's 250-plus Internet service providers (ISPs), who in turn retail Internet connectivity to their customers. But things are never so simple in practice. In Australia, Telstra and Optus lease domestic - and sometimes international - telecommunications capacity to INPs operating high-speed backbone WANs. Small proprietary WANS - plus computers and LANs owned by the ISPs - connect to these backbones, with connection, access and usage fees accruing to the Internet network providers.

Australia's premier INP is Telstra Internet, a subsidiary of our largest telco, Telstra Corporation. Another subsidiary (Telstra On Australia) is one of our biggest ISPs. This gives Telstra a vertically-integrated presence in Australia's Internet industry.

Most other INPs also are large ISPs (OzEmail, CompuServe Pacific, Microsoft Network, IBM Global Network, AccessOne, Connect.com.au, *etc*) who retail to small-medium enterprises and households. Australia's biggest organisations (governments, statutory authorities, corporations, universities, schools, *etc*) tend to buy their bandwidth wholesale from the INPs but make it available *gratis* to certain departments and individuals. These large organisations are best viewed as ISPs who "retail" Internet connectivity to their subsidiaries, staff, students, associates, and favoured clients at zero cost.

Below this tier are the majority of ISPs, ranging down in size to small local operators with a 28.8 kbps router port into the Telstra Internet backbone, one computer and a small rack of modems, to which some tens of customers dial in for access to the Internet. Furthermore, one or two of these customers may *themselves* be tiny local ISPs, who have co-located a dedicated modem in their colleague's rack for a financial consideration. But, from the largest to the smallest, all ISPs typically charge a once-only connection fee and offer four methods of paying for access: per hour, per block of hours, per month (with a monthly hours limit), and per month (unlimited). There are discounts for paying monthly or annually in advance. Hourly charging may be favoured by ISP customers who need Internet connectivity for short periods, *eg* for checking their E-mail daily. Most consumers, however, use their ISP accounts for time-intensive research, entertainment and general "Web browsing", which leads them to choose one of the block or monthly charging schemes.[21]

On the face of it, the answer to the question "Who are the consumers?" might seem to be "The retail customers of the ISPs" - all those households and small businesses plus people accessing the Internet *via* their organisations' LANs. Yet, when considering the influences that drive consumer behaviour generally, one of the most significant is always found to be the *marginal cost* of consumption - the price that must be paid to purchase and consume one extra unit of any commodity. However, in Australia virtually all these retail consumers face zero marginal cost! Probably one-half are using an organisation as their ISP and the other half have chosen block or monthly charging schemes offered by commercial ISPs.[22]

So, it is the ISPs themselves who are the real decision-making Internet consumers - those who must take steps to adjust their bandwidth requirements whenever significant changes occur in whatever influences are found to drive consumer behaviour in this industry. They do this by imposing rationing on their retail customer base whenever bandwidth needs to be conserved and by relaxing restrictions whenever the driver variables indicate that consumption may be increased. What kind of "bandwidth management" practices do these ISPs actually use?

Rationing of customers by ISPs is often direct. Shutting down overloaded Internet servers (those handling IRC or MUDs/MOOs are soft targets), taking fewer newsgroups in the Usenet feed, delaying connections for newly-signed customers, not renewing the subscriptions of particularly bandwidth-hungry customers, and letting the service generally degrade are all options. There is a lot of "churn" in this industry as customers restlessly search for ISPs offering the best quality-of-service - fewer engaged tones on dial-in, faster response to browser commands and less time spent waiting for downloads. Those network managers who act as internal ISPs to Australia's largest organisations have the easiest rationing task. They can hold back on extending LANs or simply rewrite the rulebook concerning who gets access to which services for how long.

All ISPs have the option of further managing their customers' bandwidth consumption by investing in low-cost technological fixes. Storage is cheaper than bandwidth. ISPs can install hard disk "caches" to capture the most recently and frequently downloaded WWW pages, Usenet postings, FTP files, *etc*. A "firewall" could be thrown up to prevent customers using public WWW, Usenet, FTP, and other servers in preference to the ISP's bandwidth-conserving caches. Codec software or hardware for performing file compression-decompression might be installed as well. These and numerous other technological fixes are available to relieve pressure on ISPs to purchase extra bandwidth from INPs like Telstra Internet and AccessOne, while simultaneously improving their customers' perceptions of quality-of-service. "Bandwidth management" has been developed into a fine art by Australia's ISPs. It is they who are the *de facto* consumers and this study is concerned with modelling and forecasting their demand behaviour in response to a set of independent or exogenous driver variables.

The Marketing Approach

Although the Internet has been in constant use by the military and scientific communities since 1969, it was only following development of the World Wide Web concept (1989) and, most significantly, the Mosaic browser (1994) that global computer networking began to capture the imagination of a critical mass of users. Effectively, therefore, the Internet is a "new product" and the mainstream economics approach - most applicable to mature commodities traded in established markets - may have to be supplemented by insights from the more youthful discipline of marketing.

Empirical marketing data typically trace out sigmoid "diffusion curves" of cumulative sales volume against time, as sales of each new product progressively (1) build up from slow initial acceptance, (2) pass through a phase of rapid growth, then (3) decelerate as the market approaches saturation. Presently the Internet is in the second of these three stages of growth. Frank Bass [1969] likened the diffusion process to an "epidemic" sparked off by a "contagion" of consumers with the new product. Earlier attempts, such as that of Everett Rogers [1962], had taken a purely sociological perspective. Those who came later, surveyed by Zettelmeyer & Stoneman [1993], tried to incorporate insights from other disciplines, including economics.

There are two constants in the Bass model. These parameters are called the coefficient of imitation (p) and the coefficient of innovation (q). Ignoring replacement demand, cumulative sales volume at time t during the diffusion process is V, which in future should rise to N once the market finally becomes saturated with the new product.

During a certain time period (say, one month) we can model sales volume (dV/dt) as the sum of two terms:

$$dV/dt = p.(N-V).(V/N) + q.(N-V)$$

Assuming each consumer generates one sale, the first term may be interpreted as some proportion (p) of those who have not yet purchased the new product (N-V), multiplied by their probability of meeting someone who already has (V/N). This could be called the "imitation effect". The second term can be viewed as some proportion (q) of those who have not yet purchased the new product (N-V). These consumers do not rely on word-of-mouth but get their new product information from external sources, such as the mass media. This may be termed the "innovation effect".

Mainstream economics cannot handle such "learning" by consumers, neoclassical partial and general equilibrium models being shackled by an axiom of perfect information. Under the highly "visible hand" of advertising, promotion and personal recommendation, the marketing approach shows how consumers may be led to redraw their psychological "preference maps". As a result, consumers' previous equilibrium positions are disturbed and this causes changes in the established pattern of purchasing as the new product diffuses throughout a market economy.

The Economics Approach

The economics approach to consumer demand - as pioneered by Alfred Marshall [1890] then developed by John Hicks [1946] and many others - relies upon the method of comparative statics. The analysis begins at a point of "consumer equilibrium" - a stable situation from which no one feels any particular need to change their pattern of purchasing - but then some fundamental disturbance, perhaps a change in relative commodity prices, is injected. Consumers begin adjusting to the new realities in the marketplace and, eventually, a new point of equilibrium is established.

Empirical economic data typically trace out downward-sloping "demand curves" of own-price against quantity demanded in commodity markets - provided all other influences driving demand (money income, relative prices, buyer expectations, income distribution, and preference maps) are held constant and supply conditions are not permitted to distort the statistical picture. The Law of Demand states that lowering the price of any commodity (its "own-price") will result in an increase in quantity demanded of that product, *ceteris paribus*. This phenomenon is a true "iron law of economics", to which there are very few minor exceptions.

The response of quantity demanded to a change in own-price is known as the "price elasticity of demand". If an x per cent rise (fall) in own-price leads to an x per cent fall (rise) in quantity demanded, then that segment of the market demand curve is said to have "unit elasticity", *ie* an elasticity of $e = 1$ in absolute terms. Segments having $e < 1$ are said to be "price-inelastic" while those with $e > 1$ are said to be "price-elastic". If all segments of the demand curve have the same elasticity, then the curve must be a rectangular hyperbola. "Constant elasticities" is a common assumption adopted by econometricians when estimating the parameters of demand curves and functions. Business firms are vitally concerned to know whether $e = 1$ or not, because if demand for their products is price-elastic, a price cut will increase the firm's sales revenues. Conversely, sales revenues will decrease when prices are cut, should demand happen to be price-inelastic.

Relaxing our *ceteris paribus* condition converts the demand curve into the following "demand function":

$$D = D(P, Y, R, T, E, B)$$

This relates our dependent or endogenous variable (quantity demanded) to six potential driver variables rather than simply one, *ie* own-price (P), but also income (Y), prices of closely-related goods (R), preferences (T), buyer expectations (E), and income distribution (B). Conventionally, elasticities also are computed for income and for two classes of closely-related goods: the "substitutes" for and "complements" of the product in question. In principle, elasticities also could be calculated for preferences, buyer expectations and income distribution.

The response of demand to a change in income is known as the "income elasticity of demand". Most commodities are "normal goods", for which a rise in income leads to an increase in demand. If the percentage rise in demand exceeds the percentage increase in income, then this elasticity must be greater than 1 and the product is said to be "income elastic" and a "luxury". It is called an "income inelastic" product and a "necessity" if this elasticity is less than unity. For "inferior goods", the income elasticity is negative, so that a rise in income actually leads to a decrease in demand - as when poverty-stricken consumers get more income and suddenly start eating less offal but more chops.

Steak and chops are substitutes, while steak and eggs are complements. If the price of steak goes up significantly, demand for chops will rise as consumers substitute these now-relatively-cheaper cuts of meat in their diets. But they also will demand fewer eggs, because this complementary product typically is consumed in tandem with steak only. The response of demand for the product in question to a change in the price of some closely-related commodity is called the "cross elasticity of demand".

Consumer Demand for Internet Bandwidth

By combining insights from the marketing and economics disciplines, while taking account of the special features of Internet bandwidth *qua* commodity, we can specify a demand function for this new product with a view to using econometric techniques to estimate its elasticity coefficients. Studies of telecommunications demand by econometricians such as Lester Taylor [1994] suggest that "externalities" are quite important when the product is delivered via a network.

In economics, externalities are unrequited benefits conferred (or costs imposed) on some by the actions of others. An apiarist setting beehives near a grove of fruit trees confers an external benefit on the orchardist, yet receives no recompense. An upstream factory releasing waterborne pollutants imposes an external cost on downstream producers and consumers, yet the sufferers are not compensated. Taylor has identified two external benefits associated with the PSTN: the "call externality" and the "network externality".

The call externality refers to a presumed increase in satisfaction when the called party is telephoned by the calling party, who pays for the entire call. Until Internet telephony is further developed, about the only one-to-one call externality seems to reside in E-mail messaging. Most other Internet services resemble either one-to-many broadcasting (*eg* FTP and Web browsing) or many-to-many conferencing (*eg* IRC and Usenet), both modes having a highly diffuse pattern of external costs and benefits. But Taylor's second external benefit, by contrast, is highly significant to the Internet - under both the economics and marketing approaches.

The network externality is best appreciated by considering the effect on the welfare of existing users each time the number of Internet hosts doubles - as it has been doing every 12 - 15 months since 1988. Suddenly, there is a far wider menu of services and information to choose from. Many more people with diverse interests and viewpoints have joined the unending worldwide electronic "conversation" that is the Internet. One's chances of finding valuable information, more software, better entertainment, and linking up with people sharing one's own special interests have improved markedly. Thus, increases in the size, scope and reach of the Internet must confer a positive external benefit on most users.

The pervasive influence of the network externality can be recognised by adding one more term to our Internet demand function:

$$D = D(P, Y, R, T, E, B, V)$$

Thus, demand for bandwidth will rise whenever the network externality (V) increases in value. A good proxy for this would be the number of hosts connected to the Internet.

This also happens to be a proxy for the cumulative sales volume (V) variable in our discussion of the marketing approach. Although, theoretically, the ultimate cumulative sales volume (N) also should figure in this Internet demand function, it has been left out. No one really knows what the ultimate extent of diffusion of this attractive new product will be. There are 13 million computers on the Internet now, compared with more than one billion telephone handsets in the world. If, as is likely, the number of Internet hosts exceeds the number of telephones early in the third millenium, our (V/N) term becomes too small and the (N-V) term unworkably large for an econometric estimating equation.

Business Demand for Internet Bandwidth

Thus far, our discussion seems to have been concentrated on general consumer demand, apparently neglecting special features of those businesses operating in the commercial, education, government, and not-for-profit sectors. In principle, one could specify a function showing how "derived demands" for certain inputs - including Internet bandwidth - stem from the production of outputs by businesses in response to perceived "final demands" for a range of goods and services in the economy.

Demand for Internet bandwidth, however, is but a minor subset of all business telecommunications demand - and even this has never been modelled effectively. For instance, Taylor [1994, p 65] states that

> "Telecommunications serve a variety of needs in a business, some complementary to other inputs, others substitutes, some related to production, others related to marketing, some internal to the business, others external. To treat telecommunications as a single input is to gloss over this variety of uses and probably to fail to capture the real structure of business telecommunications demand."

Just as a nation needs a stock of "social overhead capital" like roads and bridges, so also, it seems, access to telecommunications networks (both public and private) is really part of the infrastructure necessary to underpin business operations.

Therefore, it is our contention that business demand for Internet bandwidth is driven by the same seven independent variables in the demand function developed above. However, Taylor [1994, p 83] notes that the primary determinant of business telecommunications usage (for a given technology and service configuration) is the number of people employed. He also states [p 68] that "the use of telecommunications to make labor and capital more efficient has been, and will continue to be, the major force driving business telecommunications demand." Taylor notes that, as real wages increase, the time to be saved by more intensive use of telecommunications also increases in value.

Our consideration of business demand for Internet bandwidth, therefore, suggests that two further exogenous variables should be added to the demand function as developed so far, *viz* the number of full-time employed persons (L) and their average real wage rate (W).

An Internet Bandwidth Demand Function for Australia

At this point in our theoretical discussion, the Internet demand function stands as

$$D = D\ (P, Y, R, T, E, B, V, L, W)$$

which has no fewer than nine driver variables, for each of which an elasticity parameter, in principle, could be econometrically estimated. However, our aim is to explain much by little, not little by much. Clearly, some variables on the right hand side will have to go. Also, those that remain must be made operational and quite specific to the Australian situation.

The buyer expectations (E) variable refers to expectations that prices may change in the near future. One can see this influence in operation annually as the Budget Speech approaches and worried consumers start buying up stocks of cigarettes, wine and beer in advance of increases in indirect taxation of these commodities that they have been led to expect. However, the history of infrequent changes in the AARNet/Telstra Internet tariff schedule suggests that we can drop buyer expectations as an explanatory variable for Internet bandwidth demand in Australia. Surveys consistently throw up a profile of the typical Internet user as being male, aged 20 to 44 years, with tertiary qualifications and a higher-than-average income. But the lamentable lack of any annual (let alone quarterly) time series on how Australia's income distribution (B) has shifted along the rich-poor spectrum means we also must omit this possibly important variable from the specification of our econometric model.

Thus, our final *theoretical* demand function now stands as

$$D = D\ (P, Y, R, T, V, L, W)$$

and features seven explanators of Australia's monthly peak unconstrained demand for bandwidth on TelstraLink, measured in megabits per second (Mbps) of dataflow. Doubtless our econometric tests (see Chapter 4) will establish that the final *empirical* demand function will have fewer critical driver variables; the fewer the better, in fact, because this maximises the statistical degrees of freedom available to our parameter estimation procedure, given the number of monthly observations available.

In the meantime, we must precisely specify operational definitions of demand and its putative seven determinants. The longer the monthly time series of the values of these variables, the better. Six years from January 1991 to December 1996 (*ie* 72 monthly observations on each of the variables discussed below) would yield more than enough degrees of freedom to give us confidence in the results of our econometric modelling.

The Demand Variable

We have formally defined the product as Australia's monthly peak unconstrained demand for bandwidth on TelstraLink, measured in megabits per second (Mbps) of dataflow, with demand being considered as being unconstrained whenever packet losses stand at one percent or less. Yet only the largest component of total bandwidth demand (import demand by Australians for downloading files from foreign servers) is being modelled and only our largest network provider's trans-oceanic Internet link is mentioned. There are two points to be made here. The first is obvious: to use the results of our econometric model and forecasting program for capacity planning on TelstraLink, an independent forecast of export demand (by foreigners for downloading files from Australian servers) will have to be obtained. The second point is much more subtle.

The demand equation developed above "features seven explanators of Australia's monthly peak unconstrained demand for bandwidth on TelstraLink", with no recognition of competition from trans-oceanic Internet links operated by several other network providers in Australia. Some might argue that our dependent variable (D) really represents import demand on all trans-oceanic Internet links, not just the largest one. Alternatively, because all links are such close substitutes, there should be an eighth explanatory variable (the price of non-Telstra Internet bandwidth) in our demand equation to recognise a strong cross-elasticity effect on its dependent variable (D).

However, it is our contention that Telstra Internet would be looked upon as a *price leader* by these competing network providers, even if they all had completely separate networks and links. The fact that their own Internet backbones are so intimately interconnected with the premier Telstra Internet network and link only strengthens this view. Recall that Telstra Internet is Australia's *de facto* "network provider of last resort". With all Australian users of the Internet *effectively* facing the Telstra Internet tariff for bandwidth, it seems quite acceptable for demand on TelstraLink (rather than demand on all trans-oceanic Internet links) to remain on the left-hand-side of our equation.

We now propose a method for using observed packet loss percentages to convert the (sometimes capacity-limited) peak monthly import *loads* into the unconstrained *demands* that Australian Internet consumers actually exert on TelstraLink. Such a time series could be created for the behavioural variable D from parts of the following set of nine observed monthly datasets needed for this project.

For each of the 72 months (January 1991 to December 1996) we would need observations on:

L Total Load (in terabytes);
E Export Load (in terabytes);
I Import Load (in terabytes), defined as L - E;
T Date/Hour when Total Bandwidth peaked;
K Bandwidth Capacity at time T (in Mbps);
B Peak Total Bandwidth at time T (in Mbps), where B cannot exceed K;
X Peak Export Bandwidth at time T (in Mbps);

M	Peak Import Bandwidth at time T (in Mbps), defined as B - X; and
p	Hourly Packet Loss at time T (in percent).

Hourly Packet Loss (p) is the average of all packet loss measurements made during the peak hour of the month in question. Telstra Internet monitors packet losses on TelstraLink 12 times in every hour of every day by sending a train of 50 ICMP Echo 450-byte test packets to MCI's San Francisco router at Hay Exchange (IP address: 204.70.204.5), which feeds Internet downloads requested by Australians into the TelstraLink datapipe for onward transmission to the Paddington Exchange router (IP address: 139.130.249.228) in Sydney.

To get from observed import load (M) to unconstrained import demand (D) for bandwidth, one could perform the following simple calculation:

$$D = M / [(100 - p - 1) / 100]$$

Defined in this fashion, D will not always measure the full extent of unconstrained import demand. For instance, the variable M mostly may represent complete files imported to the satisfaction of consumers prepared to wait for them to download. But when p is very large (normally at those peak times when B = K) there also, no doubt, will be many incomplete and largely unusable files - hence some very dissatisfied users.

This phenomenon is due to frustrated Internet users in Australia hitting the stop button on their Web browsers to cancel those downloads they perceive as being way too slow. To model this effect, we would need Echo data on TelstraLink latencies, psychological data on time delays that are acceptable to users and economic data on the opportunity costs of users' time. Neither of these latter datasets are likely to be readily available. Capacity planners should view D (as calculated above) as a lower limit of the likely true level of unconstrained import demand for TelstraLink bandwidth.

Quite apart from its relevance for capacity planning, there are sound econometric reasons for using unconstrained (D) rather than observed (M) data when estimating the parameters of our single-equation demand function. In all markets, the quantities actually sold (M) may be viewed as part of the solution of a multi-equation simultaneous system, including at least one demand function (D with its own list of driver variables) and one supply function (with another list of variables, this time driving K). Were we to go ahead and estimate the single equation

$$M = M (P, Y, R, T, V, L, W)$$

we could never be sure exactly which *behavioural* function (demand or supply) had been identified, if either! By using unconstrained demand to estimate our econometric equation, we can *finesse* this troublesome "identification problem", thereby preventing our parameter estimates from becoming contaminated with "simultaneity bias".

The Seven Driver Variables

Own-Price Variable For each of our 72 months, this will be derived as a single price (in dollars/gigabyte) from the Tariff Schedules, Service Agreements and Customer Handbooks that have been published from time to time by AARNet and Telstra Internet. The time series will be lagged by one month to reflect the normal commercial billing period. One way of computing $/GB may be to express the Import component of Total Load (L) as a percentage utilisation of TelstraLink, then apply the most popular connectivity option of the multi-part tariff. The time series will be deflated by the Consumer Price Index (CPI) to express this own-price variable in real terms.

Income Variable Income will be specified as gross domestic product (GDP), since this measure of value-added includes income accruing to households as net wages and dividends, to businesses as retained profits and to governments as taxes less subsidies. The GDP deflator will be used to express our income variable in real terms. These measures are available as quarterly time series from the Australian Bureau of Statistics (ABS), as is the

CPI. Intermediate months will have to be interpolated to convert these to monthly time series, as outlined in Chapter 4.

Complement-Price Variable Information on the most popular makes, models and specifications in each of the past 72 months should be available from PC importers and from local modem manufacturers like NetComm, who also will have data on list prices of such equipment and the degree of discounting practised by their distributors. Inferred “street prices” will be checked against contemporary advertisements placed in metropolitan newspapers by discount retail chain stores. Street prices of the preferred PC+Modem combinations will be used to develop a dataset showing the declining cost per kbps of Internet connection speed. Street prices are favoured because the list prices published by manufacturers and retailers are almost meaningless in this industry. As "big-ticket items", computers and modems are heavily discounted in the corporate & government (CoGo), small office/home office (SoHo) and household markets. The time series will be deflated by the CPI to express our complement-price variable in real terms.

Preferences Variable Consumers value richness-of-content, download speed and user-friendliness, *ie* what's out there (files), how fast can they get it (modems) and how easily can they access it (clients)? Text supplemented with graphics, sound, animation, video, virtual reality, *etc* is as attractive to businesses and governments as it has been to households. The trend towards ever-increasing filesizes and bandwidths noted in Appendix A will be used to create our preferences variable.

Network Externality Variable The number of hosts connected to the Internet - as disclosed by the Domain Name Surveys carried out at least semi-annually by Mark Lottor [1996] of Network Wizards - will be used as a proxy for the network externality. Intermediate months will have to be interpolated to convert these to monthly time series, as outlined in Chapter 4.

Employment Variable The ABS monthly survey figure of full-time employed persons will be used as our employment variable.

Wage Rate Variable The ABS monthly survey figure for average weekly earnings of full-time employees will be used. It will be deflated by the CPI to express our wage rate variable in real terms.

Next Steps

Now that we have a handle on the theoretical determinants of Australia’s peak demand for TelstraLink bandwidth, the next step is to specify a formal model of precisely how movements in the driver variables affect the dependent variable. What is the functional form of this dependency relationship? Once this question is settled, the tried and tested techniques of econometric analysis can be applied to the data in order to estimate key parameters such as the price, income and cross elasticities of demand. Chapter 4 recommends a methodology for achieving these *modelling* aims, while Chapter 5 demonstrates how such model results could be used to reach our *forecasting* goals.

Chapter 4

Modelling Australia's Peak Demand for Internet Bandwidth

Theoretical Demand Function

In Chapter 3 we developed the following theoretical demand function for modelling Australia's peak monthly unconstrained import demand for bandwidth on TelstraLink:

D = D (P, Y, R, T, V, L, W)

where the dependent variable is

D	=	Peak Demand for Import Bandwidth	(in Mbps)

and the independent variables are

P	=	Real Own-Price	(in 1989-90 dollars/unit)
Y	=	Real Income	(in 1989-90 dollars)
R	=	Real Complement-Price	(in 1989-90 dollars/unit)
T	=	Preferences	(in MB filesize)
V	=	Network Externality	(in '000 Internet hosts)
L	=	Employment	(in '000 persons)
W	=	Real Wage Rate	(in 1989-90 dollars/week)

The (unconstrained) peak import bandwidth demand was inferred from the (constrained) peak import bandwidth load by making use of the hourly packet loss percentage:

D = M / [(100 - p - 1) / 100]

Our task is now fourfold: (1) to construct monthly time series for those independent variables which have been measured at longer intervals, (2) to specify a functional form for our econometric equation, (3) to estimate its parameters so as to derive price, income, cross, and other elasticities of demand for use in a forecasting program, and (4) to perform economic, econometric, statistical, and forecasting tests on the results of our estimation procedures.

Degrees of Freedom

In most of our statistical and econometric tests, the higher the number of *degrees of freedom* we have in estimating the demand equation's parameters, the more confident we can be that the true underlying price, income, cross, and other elasticities have been identified. To maximise degrees of freedom, we need a small number of independent variables and a large number of monthly observations on demand and its determinants. With 72 monthly observations and only eight parameters to estimate, we have an enviable 64 degrees of freedom available.

Monthly Time Series

Our econometric model will be estimated from time series of monthly observations on one dependent and seven independent variables. However, measurements of Real Income and the Consumer Price Index are made quarterly and the Network Externality is, at best, measured semi-annually. Clearly, we will need to use standard interpolation methods for the two missing monthly observations in each quarterly time series and for the five missing monthly observations in each semi-annual time series.

The Econometric Estimating Equation

Most empirical demand studies have adopted the *log-linear* or *constant elasticities* form because it is simple to estimate, its parameters can be interpreted directly as elasticities and these elasticities are constant for all values of their associated driver variables. Our chosen functional form, therefore, is

$$D_t = \alpha P_t^{\beta 1} Y_t^{\beta 2} R_t^{\beta 3} T_t^{\beta 4} V_t^{\beta 5} L_t^{\beta 6} W_t^{\beta 7} e^{ut}$$

where	α	=	Scaling Constant
	β_1	=	Price Elasticity (*aka* Own-Price Elasticity)
	β_2	=	Income Elasticity
	β_3	=	Cross Elasticity (*aka* Complement-Price Elasticity)
	β_4	=	Elasticity with respect to Preferences
	β_5	=	Elasticity wrt Network Externality
	β_6	=	Elasticity wrt Employment
	β_7	=	Elasticity wrt Wage Rate
	e	=	Base of Natural Logarithms
	u_t	=	Error Term
	t	=	Time, ie month number 1, 2, ..., 72

We can transform this multiplicative formulation into a linear estimating equation by taking the natural logarithms (*ln*) of both sides, *viz*

$$\ln D_t = \beta_0 + \beta_1 \ln P_t + \beta_2 \ln Y_t + \beta_3 \ln R_t + \beta_4 \ln T_t + \beta_5 \ln V_t + \beta_6 \ln L_t + \beta_7 \ln W_t + u_t$$

where β_0 = Intercept Parameter ($\ln\alpha$)

The error term (u_t) captures the effect on demand of three things: missing independent variables, errors in measuring the seven independent variables and numerous random influences causing consumers to deviate (erratically and temporarily) from their planned usage levels. Missing independent variables may be both known, such as Australia's changing income distribution, and unknown, due to the *corpus* of economic and marketing theory being incomplete.

Parameter Estimates and Elasticities

We plan to estimate our peak unconstrained import demand for bandwidth equation using the Ordinary Least Squares (OLS) method. Our estimates of the intercept parameter (β_0) and of all seven elasticities (β_1, β_2, ..., β_7) - together with their t-statistics - will be discussed in full and confidence intervals computed. The coefficient of determination (R^2) for the equation as a whole will be calculated, to show what percentage of variation in the demand variable is being explained by variations in its seven driver variables. When adjusted for degrees of freedom, R^2 will fall slightly and be used to assess the equation's goodness-of-fit to the data.

Tests of Economic Plausibility

The economic plausibility of our model is determined by its consistency with the principles and postulates of economic and marketing theory and by whether the theoretically expected signs and sizes of the estimated elasticities are obtained.

Identification Problem If the particular datasets used will not suffice to identify a demand function, then estimation of our equation's parameters is meaningless. Our demand function *is* considered to be identified for two reasons. First, the AARNet/Telstra Internet monthly own-prices have *never* been set intentionally to achieve a classic demand-supply equilibrium in a competitive market for Internet bandwidth. Secondly, the dependent variable (D) is a behavioural measure of pure unconstrained demand - rather than a technical measure of the observed (and sometimes constrained) load - that has been constructed using an indicator of excess demand on TelstraLink, *viz* the packet loss percentage above one percent.

Simultaneity Bias The dependent variable (D) is functionally dependent on our seven independent variables, but if the level of demand feeds back onto the determination of even one of these seven explanators, then simultaneity bias is present and our estimates will be inconsistent. Demand for TelstraLink bandwidth is too small to affect any of these five driver variables: income, preferences, network externality, employment, or wage rate. In principle, it *might* affect own-price, but that is ruled out because P has never been set to achieve a competitive market equilibrium. That leaves only complement-price. If rapid growth of the Internet stimulated demand for entry-level personal computers and modems to such an extent that it raised R we would have a feedback situation. However, these prices are affected far more by the persistent trend of falling real hardware costs than by demand pressures due to the Internet's popularity, so simultaneity bias is not likely to present any problem for our analysis.

Elasticity Estimates From economic theory, we would expect the signs and sizes of the parameters to be:

Parameter	**Variable**	**Sign & Size**		
β_0	Intercept	+ve		
β_1	Own-Price	-ve	and	< |1.0|
β_2	Income	+ve	and	> |1.0|
β_3	Complement-Price	-ve		
β_4	Preferences	+ve		
β_5	Network Externality	+ve		
β_6	Employment	+ve		
β_7	Wage Rate	+ve		

There is no *a priori* expectation about the magnitude of any parameter except for the income elasticity (which should exceed unity since Internet access and usage are luxuries, not necessities) and the own-price elasticity (which should be less than unity in absolute terms, from the Law of Demand).

The own-price elasticity will tend to be inelastic because of:

(1) the availability of few close substitutes: while there are many distant substitutes for most Internet services, there are few close substitutes and the hypermedia aspect of the World Wide Web (retrieving relevant information simply by clicking on a hyperlink) is unique;

(2) the importance of being unimportant: Internet access and usage represents only a tiny proportion of the budget of most consumers;

(3) its large number of uses: the Internet's text, graphic, audio, animation, video, and virtual reality capabilities - plus its flexibility as a one-to-one (unicast), one-to-many (broadcast) and many-to-many (multicast) medium - means that there are thousands of actual and potential uses; and

(4) its addictive qualities: these are attested to by the comments and activities of new and many established Net surfers, with the phenomenon now under active study by psychologists.[23]

Homogeneity of Degree Zero From economic theory, we would expect no change in quantity demanded if all prices and total expenditure (*cf* income) change proportionately. So, the sum of our first three elasticity parameters should be zero. We will use Fisher's F-test for a linear restriction to check whether $\beta_1 + \beta_2 + \beta_3 = 0$.

Tests of Econometric Assumptions

The ordinary least squares (OLS) method is known to be the best linear unbiased (BLU) and minimum mean-square-error (MSE) estimator. Our chosen OLS estimator potentially possesses all the desirable

* small sample estimator properties of (1) unbiasedness, (2) least variance, (3) efficiency, (4) least mean-square-error, and (5) sufficiency; and

* large sample estimator properties of (1) asymptotic unbiasedness, (2) consistency and (3) asymptotic efficiency,

as noted by Koutsoyiannis [1988], Chapter 6.

But our OLS parameter estimation procedure will exhibit these properties if and only if all the following seven key assumptions hold true:

Randomness The error term is a random variable. Its value in any particular month may be positive, negative or zero, but this must be the outcome of pure chance. There is no way of testing this assumption; it must be established on *a priori* grounds. We consider that random behaviour of our error term is to be expected for three reasons:

(a) consumer behaviour is naturally erratic in that there can be numerous reasons for temporary deviations from planned Internet usage levels;

(b) our seven independent variables have been measured by different agencies, so the inevitable errors should be unrelated; and

(c) our known missing independent variables (Australia's changing income distribution, the prices of numerous related goods and consumer expectations of own-price behaviour) should not introduce any systematic element.

Zero Mean The error term has a mean value of zero. The positive and negative errors should cluster around a zero average. This is another assumption which must be established *a priori*, for it cannot be tested. We expect that the error term will have a zero mean for two reasons:

(a) microeconomic and marketing theory assures us that all important explanatory variables (except income distribution) have been included in the equation; and

(b) errors of measurement are likely to be offsetting, with no bias towards either the positive or negative side of zero.

Homoscedasticity In all months, the dispersion of positive and negative errors around their zero means must be the same, irrespective of whether our seven independent variables take on high or low values. There are *a priori* grounds for believing that this assumption of *homoscedasticity* is violated in many econometric models, *eg* errors of measurement are more likely when the independent variables have larger values. So, given that most explanatory variables get larger (in absolute terms) the usual case of heteroscedasticity is that the error term's variance increases with $t = 1, 2, ..., 72$. We plan to check for the presence of heteroscedasticity using Spearman's rank-correlation test, the Goldfeld & Quandt test and the Glejser test.

Normal Distribution In all months, the errors should exhibit a bell-shaped frequency distribution about their zero means. Small absolute errors should have a higher probability of being observed than large ones; extreme values of the error term should become more unlikely the more extreme they get. This normality assumption cannot be tested, so econometricians usually appeal to the Central Limit Theorem, which states that even if the population from which the sample was drawn is *not* normally distributed, the sampling distribution of the mean tends to the normal distribution as the sample size increases.

No Autocorrelation The error terms in any two months are independent of each other. The size and sign of the error in any particular month is not affected by the behaviour of errors belonging to earlier months. If successive values in the 72-months-long time series of errors are found to be systematically correlated, then *autocorrelation* is a problem. The standard errors of our parameter estimates are then not a reliable criterion for evaluating the statistical significance of the elasticities. We plan to use the Durbin-Watson statistic (d) to test for the presence of autocorrelation in the time series of errors. Within its possible range of (0, 4) a value of d = 2 would establish the absence of any autocorrelation.

Independence The error term is independent of all seven explanatory variables. Errors must not vary in tandem with one or more of the variables on the right hand side of the demand equation. If this assumption is violated, the OLS parameter estimates will be both biased and inconsistent. This is not considered a problem in our particular equation because (a) there are no lagged dependent variables on the right hand side, (b) the equation is not part of any system of simultaneous equations - apart from that describing the total Australian economy - and (c) our seven explanatory variables are not stochastic (unlike the error term) - nor are they riddled with systematic measurement errors, as far as we can tell.

No Multicollinearity None of the seven driver variables are strongly linearly correlated with each other. Otherwise our datasets are afflicted with *multicollinearity*. If two or more explanators are intercorrelated (*eg* employment with income or wage rate with employment), OLS will not be able to separate out their separate influences and assign these to the elasticity estimates. Ideally, each of our seven exogenous variables should be *orthogonal* to all the others, *ie* their covariances should be zero. We plan to test for the presence of multicollinearity by applying Frisch's Confluence Analysis and the Farrar-Glauber test.

If any of these seven assumptions underpinning the OLS procedure are wrong then our elasticity estimates are not BLU/MSE, or the statistical tests become meaningless, or both. We plan to apply tests to check out three of these key assumptions; the others have to be taken on faith. Whenever a model fails its econometric tests, it is customary (despite the suspicion of "data mining") for variable(s) to be transformed, substituted, added, or omitted until the model does pass muster. We fully expect that fewer than the seven explanatory variables suggested by theory will be found to have practical significance for the determination of peak import demand for bandwidth on TelstraLink. As a minimum, we believe that own-price, complement-price and preferences will be confirmed as three of the most important drivers of demand.

Tests of Statistical Reliability

Our model may turn out to be economically plausible and econometrically sound, yet be statistically unreliable. The demand equation may not fit the demand data very well or, if it does, this finding may prove to be statistically insignificant. The true elasticities may lie within a wide band around the estimated elasticities. Even if the band is found to be narrow, this also may prove to be statistically unreliable because of large standard errors of the estimates.

Goodness-of-Fit of the Equation: Coefficients of Determination, *ie* R^2 and R^2 adjusted for degrees of freedom.

Significance of the Equation: Fisher's F-test of R^2 at the 1% and 5% levels.

Significance of the Estimates: Student's t-test of every parameter at the 1% and 5% levels.

Confidence Intervals for the Estimates: 99% and 95% confidence intervals for every parameter.

Having established that our model will have 64 degrees of freedom (*ie* 72 observations reduced by the 8 parameters), we plan to apply these four statistical tests and provide a full discussion of the results.

Tests of Forecasting Ability

Time-Invariance of Parameters: An econometric model may pass all the above tests, yet still be unsuitable for prediction because its parameters are changing over time, not constant. We plan to test for the invariance over time of our elasticity estimates using several subsets of our 72 monthly observations. The model will be re-estimated for three separate 24-month time series then for 24 and 48 months of data. We plan to use the Chow test for equality between the elasticities obtained from these different samples.

Extra-Sample Performance: Finally, we plan to make a "live forecast" of bandwidth demand for three extra-sample months, *ie* for $t = 73, 74, 75$. Data for these months will have been reserved and kept separate from those observations used in the econometric work. We plan to statistically test the differences found between actual and forecast values before offering any opinion on the ability of our model to make defensible forecasts of Australia's monthly peak unconstrained import demand for bandwidth on TelstraLink.

Chapter 5

Forecasting Australia's Peak Demand for Internet Bandwidth

The TelstraLink Forecasting Tool

Our forecasting tool is a simple computer program. Users input a known peak unconstrained import demand figure for the base month (Q) and a desired forecasting horizon T months ahead.

The program operates by plugging Q and T into the following monthly growth equation:

$$D = Q*[(1 + G) \wedge T]$$

where * is the multiplication operator;

^ is the exponentiation operator; and

$$G = B0 + B1*GP + B2*GY + B3*GR + B4*GT + B5*GV + B6*GL + B7*GW$$

The program obtains eight dummy parameters (B0, B1, ... , B7) from an elasticities file and test data on the independent variables from a drivers file. The latter file contains projected monthly growth rates of our seven driver variables (GP, GY, ... , GW). These files may be edited, for the purpose of (a) inserting more credible data before a Base Case is run or (b) performing sensitivity analyses by perturbing values of the driver variables.

If we were making production (rather than mere demonstration) forecasts, these data files would contain proper econometric estimates of the elasticity parameters of our peak import demand equation and projections of growth rates for its seven exogenous variables developed by recognised authorities who deal with these data on a day-to-day basis. For instance, the Commonwealth Treasury's budget forecasts might be used for the likely future growth of Income, Employment and the Wage Rate. Likewise, the Internet Global Penetration Study by Killen & Associates might be used to project the Network Externality variable.

Dummy Parameter Values

In the absence of proper econometric estimates, we have placed the following dummy parameter values in the elasticities file:

Code	**Elasticity**	**Sign & Size**
B0	Intercept	+0.01
B1	Own-Price	-0.4
B2	Income	+1.2
B3	Complement-Price	-1.3
B4	Preferences	+0.8
B5	Network Externality	+0.7
B6	Employment	+0.6
B7	Wage Rate	+0.5

NOTE: No reliance should be placed upon these dummy elasticities in any production forecasting exercise.

Base Month and Forecasting Horizon

We have chosen November 1996 as our forecasting base, this being the month when Telstra Internet last commissioned extra capacity. This brought the bandwidth of TelstraLink up to its present level of 44 Mbps. Suppose one-quarter of this bandwidth (11 Mbps) was being absorbed during peak hour by foreign demand. That would have left 33 Mbps available to meet Australia's peak demand for Internet bandwidth last August.

Suppose, further, that peak hour in November 1996 exhibited a packet loss percentage such that excess demand of 2 Mbps was indicated. So, the figure we shall assume for Australia's peak monthly unconstrained demand for TelstraLink bandwidth in our base month is Q = 35 Mbps. This is, of course, a mere test datum and no reliance should be placed on this figure in any production forecasting exercise.

We have chosen one year as our forecasting horizon, so that T = 12 months.

Test Data for Forecasts of Growth in Driver Variables

Until authoritative projections can be obtained, we have placed the following test data values in the drivers file to characterise the Base Case in our demonstration forecasts:

Code	Driver Variable	Growth
GP	Real Own-Price	-0.2 %
GY	Real Income	+0.3 %
GR	Real Complement-Price	-0.3 %
GT	Preferences	+3.0 %
GV	Network Externality	+7.0 %
GL	Employment	+0.1 %
GW	Real Wage Rate	+0.4 %

NOTE: No reliance should be placed upon these test data in any production forecasting exercise.

Demonstration Forecasts of Australia's Peak Demand for Internet Bandwidth

Base Case After running the forecasting program with Q = 35 Mbps and T = 12 months, it produced a projection of peak monthly unconstrained import demand for TelstraLink bandwidth in November 1997 of 102.8 Mbps, an increase of 194 per cent in total or more than nine per cent per month. This Base Case assumes an inflation rate of 2.4 per cent for the year ended 30 November 1997. So, nominal prices currently quoted in the Telstra Internet tariff schedule (*eg* $190/GB for all data received on a 2 - 10 Mbps connection) are assumed not to increase, even though their real value will have been reduced by 2.4 per cent over the 12-month forecasting horizon. In the Perturbed Cases which follow, all sensitivity analyses are made relative to the Base Case. For each run, one monthly growth rate was changed temporarily, with all other test data in the drivers file remaining as shown above.

Obviously, our Base Case results cannot be used in any production forecasting exercise.

Tariff Increase We ran the forecasting program again after changing the Base Case projection of a 0.2 per cent monthly decrease in the real own-price of bandwidth to a 0.2 per cent real monthly increase, implying that nominal Telstra Internet tariffs would rise from $190 to $199 per GB over the next 12 months. The effect was to reduce the Base Case forecast of peak monthly unconstrained import demand for TelstraLink bandwidth from 102.8 to 101.0 Mbps.

Tariff Decrease We ran the forecasting program again after changing the Base Case projection of a 0.2 per cent monthly decrease in the real own-price of bandwidth to a 0.6 per cent real monthly decrease, implying that nominal Telstra Internet tariffs would fall from $190 to $181 per GB over the next 12 months. The effect was to increase the Base Case forecast of peak monthly unconstrained import demand for TelstraLink bandwidth from 102.8 to 104.6 Mbps.

Income Boost We ran the forecasting program again after doubling the Base Case projection of a 0.3 per cent monthly increase in real income to a 0.6 per cent increase. The effect was to increase the Base Case forecast of peak monthly unconstrained import demand for TelstraLink bandwidth from 102.8 to 106.9 Mbps.

Cheaper Equipment We ran the forecasting program again after changing the Base Case projection of a 0.3 per cent monthly decrease in the real price of a PC+Modem combination income to a 0.5 per cent real monthly decrease. The effect was to increase the Base Case forecast of peak monthly unconstrained import demand for TelstraLink bandwidth from 102.8 to 105.7 Mbps.

Larger Filesizes We ran the forecasting program again after raising the monthly rate of growth of (uncompressed) Internet files - a measure of the ever-richening content preferred by consumers - from 3 to 5 per cent. The effect was to increase the Base Case forecast of peak monthly unconstrained import demand for TelstraLink bandwidth from 102.8 to 122.3 Mbps.

Bigger Internet We ran the forecasting program again after projecting the number of Internet hosts to increase by 9 instead of 7 per cent per month, *ie* a more powerful network externality. The effect was to increase the Base Case forecast of peak monthly unconstrained import demand for TelstraLink bandwidth from 102.8 to 119.7 Mbps.

More Employment We ran the forecasting program again after doubling the monthly rate of employment growth from 0.1 to 0.2 per cent. The effect was to increase the Base Case forecast of peak monthly unconstrained import demand for TelstraLink bandwidth from 102.8 to 103.4 Mbps.

Higher Wages We ran the forecasting program again after raising the rate of real wage growth from 0.4 to 0.6 per cent per month. The effect was to increase the Base Case forecast of peak monthly unconstrained import demand for TelstraLink bandwidth from 102.8 to 103.9 Mbps.

Obviously, these sensitivity analyses of our Perturbed Cases cannot be used in any production forecasting exercise.

Future Directions

We demonstrated above how simple it is to make Base Case and Perturbed Case forecasts of demand once we have a set of price, income, cross, and other elasticities to work with. Armed with some base month observation of Australia's peak monthly unconstrained import demand for bandwidth on TelstraLink, we can apply these elasticities to independently-projected growth rates for each of our seven driver variables to make predictions of future levels and growth rates of demand.

At present, however, we are using dummy values for these key elasticity parameters, so can make only demonstration forecasts. To prepare defensible production forecasts, we need to have accurate econometric estimates of these parameters. This would require access to the appropriate datasets, as discussed in a section of Chapter 3 titled "The Demand Variable". If this was to occur our illustrative-only demonstration forecasts could be replaced with usable production forecasts. This would mean that an alternative "top-down" and "international-to-domestic" forecasting tool would become available to complement and enhance the present “bottom-up” and “domestic-to-international” methods of projecting peak bandwidth demands and planning Australia’s trans-oceanic Internet link capacities.

The benefits of deriving accurate estimates for parameters such as price, income and cross elasticities of demand include making forecasts of Australia's peak demand for TelstraLink bandwidth based on several plausible scenarios. This can only improve the practice of capacity planning and lead to more efficient outcomes.

Appendix

Capacity of Australia's Trans-Oceanic Internet Links at November 1996

Capacity[1]	Links terminate in ...	TI Connect	PoPs	Notes
Telstra Internet (TI)				
44 Mbps	California, New Zealand	na	8	Australia's principal link
AccessOne				
8.128 Mbps	California, Hong Kong	2.128 Mbps	41	
Connect.com.au				
7.256 Mbps	California, Singapore	2 Mbps	20	
Telstra On Australia				
6 Mbps	California	1 Mbps?	16	
IBM				
6 Mbps	California, Japan	1 Mbps	11	800+ PoPs in 45 countries
AUSNet Services				
5 Mbps	Oregon	1 Mbps?	7	
Geko				
4 Mbps	California, Hong Kong, Japan	1 Mbps?	1	
TPG Internet				
4 Mbps	California	2 Mbps	42	
OzEmail				
2 Mbps	California	11 Mbps	22	
Magna Data				
2 Mbps	New Jersey	1 Mbps	5	
CompuServe Pacific				
1 Mbps?	California	1 Mbps?	10	???+ PoPs in ?? countries
Hutchison Internet Services				
1 Mbps?	California	1 Mbps?	5	
90.384 Mbps[2]		24.128 Mbps		Min. estimate only[3]

Notes:

1. All other INPs and ISPs offering a "private international link" (*ie* as well as or instead of TelstraLink) are believed to be sub-leasing capacity from one or more of those tabulated.

2. However, according to Thomas P. Koltai [tomk@koltai.com], "When we add Telstra's 32 Mbps to MCI's 18, UUNet's 29, AT&T's 6, Sprint's 6, BT's 12, IBM's 6, DEC's 6, and NZT's 1.5 we find a total of 116.5 Mbps that actually travels across the Australian backbone (either coming or going) before it moves overseas." As Telstra has 44 rather than 32 Mbps, an updated Koltai estimate (summation of telco links) would be 128.5 Mbps.
3. The Richardson estimate (summation of INP link) is 90.4 Mbps but this is clearly an underestimate because not all INPs (*eg* Compuserve Pacific) will disclose their link capacity.

Notes

1. "At present, international calls are often a monopoly or at least a cartel. Usually only one operator - the national giant - has the right to hand calls across the international gateway that separates one country's telephone network from another's. At the gate, predictably enough, it pockets a charge. This lucrative custom is backed up by the system of international accounting rates, which allocates the cost of a call between countries. The two operators agree on a price for handling the call and split it, usually down the middle. If one operator puts through more calls to another than it receives, it hands over a settlement payment to even things up." (Christopher Anderson, "A Survey of the Internet" in The Economist (1 July 1995) at http://www.economist.com/). Also, "... trans-oceanic calls ... are priced far higher than cost. (ITU figures show actual costs of [3 US cents] per hour for voice circuits!)" Brian Kahin, "The Internet Business and Policy Landscape" in Twenty-FourthTelecommunications Policy Research Conference Papers, Section III, October 1996,

2. "Every 30 minutes, another network is connected."(Internet Society at gopher://gopher.isoc.org:70/)

3. Internet Domain Survey (July 1996) by Mark Lottor of Network Wizards at http://www.nw.com/

4. Mark Lottor, *op. cit.*

5. Mark Lottor considers "... the numbers presented in the domain survey to be fairly good estimates of the *minimum* size of the Internet." Lottor found 12,881,000 hosts split roughly 60:40 (and falling) between the USA and the rest of the world in July 1996. If there are, conservatively, 1.5 users per host in the USA and 2.5 users per host in the rest of the world - and the number of hosts is still doubling annually - then by December 1996 there should be at least 36 million Internet users worldwide. (John Quarterman is reported [in HotWired at http://www.hotwired.com/Lib/Extras] as claiming that the real factor for users per Internet host is about 3.5, rather than 7.5 or even 10 users/host as is assumed by other researchers.) Support for an estimate exceeding 30 million users comes from Nua Ltd at http://www.nua.ie/ who commented upon an influential US survey of Internet users: "Conducted in August 1995, the Nielsen/CommerceNet survey found that there were 23 million US and Canadian users [and] to get worldwide usage of the Internet you could increase the above figures by approximately one half." Support also comes from Killen & Associates at http://www.killen.com/ who used 13 different sources to estimate that there were 30 million Internet users worldwide in January 1996. However, the author is aware of other estimates ranging from fewer than 20 million to more than 80 million Internet users worldwide.

6. "The Internet has been growing at about 100% per year since 1988. The Internet doubles in size each year. That means that each year there are as many new people on the Internet as all the people who were on it before." (John Quarterman, "Sizes of the Internet in October 1995 from the Third MIDS Internet Demographic Survey" at http://www.mids.org/ids3/)

7. In October 1995 there were about 3.6 million users of BITNet, UUCP and FidoNet, all of which can exchange E-mail with the Internet, but this figure continues to fall as the Internet slowly absorbs the Matrix. (John Quarterman, *op. cit.)*

8. Lawrence H. Landweber and Internet Society (15 June 1996) at ftp://ftp.cs.wisc.edu/

9. Killen & Associates at http://www.killen.com/

10. See http://www.altavista.digital.com/

11. The Infoseek Ultra search engine at http://ultra.infoseek.com/ has found more than 80 million Web pages.

12. The AltaVista search engine claims to have indexed 30 million Web pages found on Web sites, so there are about 109 Web pages per Web site. If there are 90 million pages in Webspace, this implies more than 825,000 Web sites.

13. "For the second half of 1993, the Web had a doubling period of under 3 months, and even today the doubling period is still under 6 months." (Matthew Gray, "Measuring the Growth of the Web: June 1993 to June 1995" at http://www.mit.edu/people/mkgray/)

14. See http://www.hotbot.com/

15. Mark Lottor, *op. cit.*

16. "Telstra unveiled the results of a privately conducted survey in June that suggested that more than one million Australians regularly use the Internet ... Other figures released by Telstra included the suggestion that 7% of households have access to the Internet, and 11% of all people over the age of 15 had accessed the Internet for one hour in the week prior to the survey. A total of 16% have accessed the Internet at some point ..." (Internet Australasia (August 1996) at http://www.interaus.net/1996/8/cover.html). Also, the July 1996 AGB McNair survey of Australians aged 18+ years found that 1,178,000 can access the Internet from home and that 2,707,000 (one in every five) had used the Internet at some point. (AGB McNair at http://www.agb.com.au/online/)

17. The top five Internet host countries are USA, UK, Germany, Japan, and Canada. (Mark Lottor, *op. cit.)*

18. Finland (24), Australia and Sweden (59) with the world average being 850 persons/host. (Killen & Associates at http://www.killen.com/ipf.htm)

19. Iceland (1.67), Australia (2.28), Norway (2.53), Finland (2.62), and USA (2.79) with the world median being 6.18. (Internet Society at gopher://gopher.isoc.org/)

20. In April 1995, HTTP traffic over the former US Internet backbone (NSFNet) exceeded FTP traffic for the first time. (See ftp://nic.merit.edu/nsfnet/statistics/) Since then, HTTP traffic has grown to account for more than one-third of all Internet traffic. (Christopher Anderson, *op. cit.)*

21. In the United States, "... the market seems to be tilting strongly towards flat-rate pricing" (p. 63) and "... the retail market for access to the global Internet is moving toward a zero marginal cost model, paid for by a fixed monthly fee" (p. 64) (Brian Kahin, *op. cit.)*

22. The author has observed the same pricing trend among Australian ISPs.

23. "The Center for On-line Addiction has ... conducted extensive empirical research ... on the subject of Internet Addiction. Internet Addiction has now gained credibility among mental health professionals as a clinically significant disorder which negatively impacts social, occupational, family, and financial functioning. In addition, the center has identified problem areas that mental health professionals, educators, and human resource managers are experiencing due to individuals' becoming addicted to the Internet in the same manner one would become addicted to drugs, alcohol, or gambling." (Dr Kimberly S. Young at http://www.pitt.edu/)

Bibliography

Bass, F.M. "A New Product Growth Model for Consumer Durables" in *Management Science*, 1969, vol 15, 215-227

Hicks, J R "Value and Capital" (Oxford: Oxford University Press, 1946)

Koutsoyiannis, A "Theory of Econometrics: An Introductory Exposition of Econometric Methods" (London: Macmillan, 1988)

Marshall, Alfred "Principles of Economics" (London: Macmillan, 1890)

Rogers, E M "Diffusion of Innovations" (London: Free Press, 1962)

Taylor, Lester D "Telecommunications Demand in Theory and Practice" (Dordrecht: Kluwer Academic Publishers, 1994)

Zettelmeyer, F & Stoneman, P "Testing Alternative Models of New Product Diffusion" in *Economics of Innovation and New Technology*, 1993, vol 2, 283-308