

**Networks, Noise and Web Navigation:
Sustaining Metcalfe's Law through
Technological Innovation**

by

Paul Windrum *

MERIT
University of Maastricht

and

G. M. Peter Swann

Manchester Business School
University of Manchester

January 1999

* Please address all correspondence to Dr. Paul Windrum, MERIT, University of Maastricht, Tongersestraat 49, PO BOX 616, NL-6211 LM Maastricht, The Netherlands.

E-mail: P.Windrum@merit.unimaas.nl

We gratefully acknowledge funding from European Community Contract No. SOE1-CT97-1062 (DG12-SOLS), for the *Colline* Project, part of the TSER Programme.

Abstract

The paper examines the relationship between network externalities and the value of the World-Wide Web. It is shown that value depends on two offsetting effects. First, as the Web grows in size, so existing users gain from the additional content provided by new users. This is the standard concept of a network externality. However, at the same time, a large and expanding network can suffer from congestion. Congestion can manifest itself in the time taken to download information from the Web. But an arguably more important problem is navigating the very large number of webpages on offer and to identify the particular page containing information that is of interest. These navigation issues have serious implications for the continued growth of the Web, as described by Metcalfe's Law.

The paper reviews the current state-of-the-art in navigation technology, identifying the weaknesses that exist and where innovative effort is likely to be focused in the near future. The review forms the basis of a formal model that analyses how the overall value of the Web changes as the size of the network increases and how improvements in navigation can increase the value obtained from the Web. Finally, these are contrasted with other options that are available to policy-makers.

Keywords: World-Wide Web, Metcalfe's Law, network externalities, congestion costs, search engines.

1. Metcalfe's Law and the Growth of the Web.

The spectacular growth of the World-Wide Web (WWW), and indeed the Internet as a whole, has forced politicians, industrialists and academics alike not only to rethink telecommunications but the future nature of communications, economics and society itself. Estimates of the size and growth of the Internet vary.¹ Examining various US data sources (which currently accounts for some 60 to 70% of all users and host computers), Coffman and Odlyzko (1998) estimate Internet growth to be around 100% per year in the US compared to a growth rate of 8% for voice traffic over traditional telecoms networks. Should the Internet continue to grow at the current rate, Coffman and Odlyzko observe, data traffic will overtake voice traffic in the US by the year 2003 and in other countries sometime thereafter.

While the Web and the Internet are not identical, the growth of the Internet is clearly related to the increasing popularity of the Web. Indeed the rapid spurt in Internet traffic observed in 1995 and 1996 coincided with two web-related developments. The first was Netscape's introduction of a graphical browser interface which, for the first time, facilitated the use of illustrations within websites. The second was the development of user-friendly search engines which "made the Internet more valuable than it would otherwise have been" (Coffman and Odlyzko, 1998).

Measuring the growth of the Web is as difficult as measuring the growth of the Internet, not least because of the linked 'hypertext' format: where does one site end and another begin? (Windrum and Flanagan, 1997). According to Gray (1997) the number of websites (defined as those sites with a distinctive host name) grew from 130 in mid-1993 to 650,000 at the beginning of 1997. Between 1993 and 1995, the proportion of traffic routed through the US NSFNET backbone accounted for by web-related transmissions rose from 0.5% to 23.9%². By mid-1997 it was estimated that content of the Web could be as high as 150 million pages (Pike, 1997). The proportion of websites classified as commercial has risen dramatically, with the percentage of sites with Uniform Resource Locators (URLs) ending in '.com'³ increased from 1.5% in June 1993 to 62.6% in January 1997 (Windrum and Flanagan, 1997).

¹ The variation in estimates partly reflects the nature of the Internet as a loose collection of private and public networks in which the vast majority of traffic goes through more than one peering point. However the chief reason is that private sector carriers do not release detailed information about their networks. Since no institution is responsible for collecting data on the Internet, analysts must make estimates on the basis of piecemeal information that is largely provided by public sector organisations. Measurement problems are further compounded by a lack of consensus over what constitutes an appropriate set of statistical indicators. Coffman and Odlyzko, for example, include transmission capacity as well as traffic (measured in bytes) as indicators.

² Data after April 1995 is unavailable, a consequence of NSFNET handing over responsibility for the provision of the US Internet backbone to private service providers.

³ This estimate excludes many commercial sites located outside the US.

The Domain Survey (July 1998) conducted by Network Wizards <<http://www.nw.com/>> indicates the growth of the Web continues unabated. It estimates some 36,739,000 computers connected in 156 countries, up from 26,053,000 in July 1997 (an increase of 41%). The Domain Survey attempts to discover every host on the Internet by doing a complete search of the Domain Name System. It should be noted that some of these are clusters of machines. On the other hand, many systems will be uncounted because they are hidden behind firewalls. Figure 1 presents the Network Wizards' half-yearly estimates for the period January 1991 to July 1998 and indicates the quite astounding growth of the Internet over the past seven years.

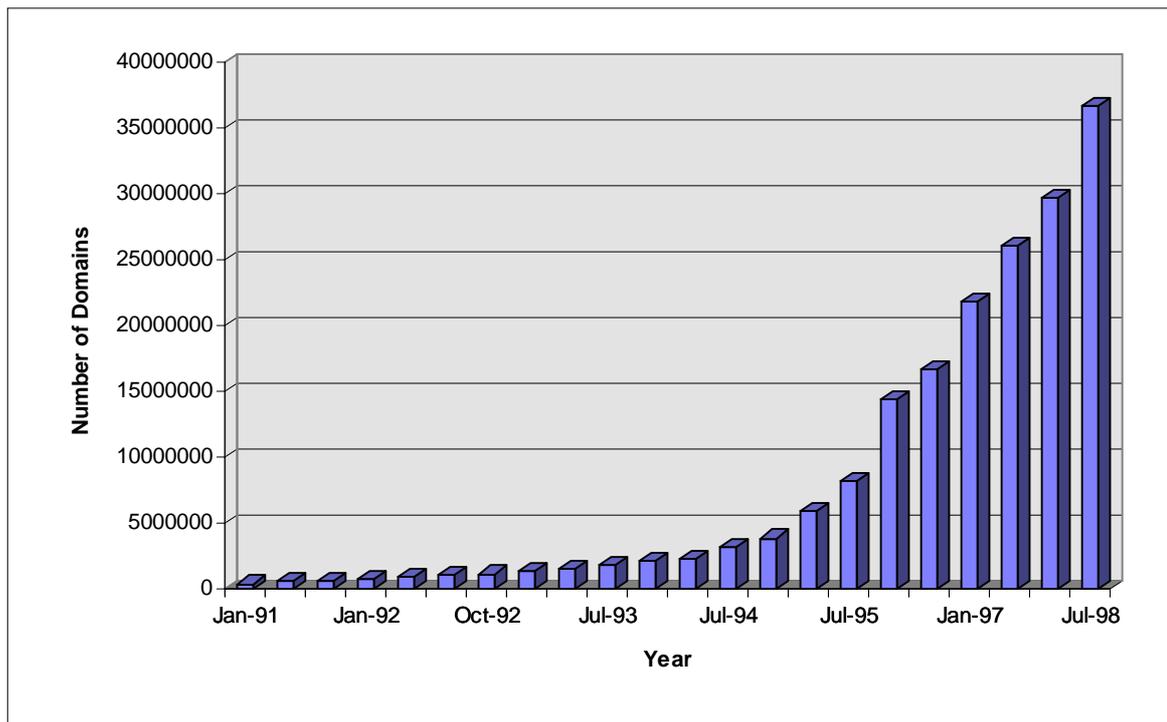


Figure 1 Growth of domain names (January 1991 to July 1998)
 Source: Network Wizards <<http://www.nw.com/>>

This annual doubling in size is commonly ascribed to *Metcalfe's Law*⁴ by the computer community. The 'law' states that the value of a network is proportional to the square of the number of its users (see Gilder, 1993). Hence it observes a positive feedback between the value of a network and the number of people using and contributing to the network. As such it is an example of the network externalities concept developed by Farrell and Saloner (1985), Katz and Shapiro (1985), David (1985) and Arthur (1989) in the economics literature on standards⁵.

⁴ The law is named after Robert M. Metcalfe, inventor of the Ethernet and founder of 3Com.

⁵ See David and Greenstein (1990) and Economides and Encaoua (1996) for a discussion of subsequent developments in this literature.

There are two problematic issues associated with this rapid expansion, one of which has received attention by economists and others, the other which has not. The first is due, ironically, to the popularity of websites containing both images and text. The large bandwidth required to deliver images imposes a significant strain on the Internet. The debate on how to deal with limited bandwidth usually revolves around establishing a pricing mechanism that will efficiently distribute scarce bandwidth resources amongst different user groups (e.g. MacKie and Varian, 1995; Odlyzko, 1997; Kelly, 1997).

The second issue concerns the extent to which additional users add value to a network. Two key underpinning the Metcalfe Law;

- (i) that each additional user adds value (or is at least is of equal marginal value) to the network;
- (ii) the ability to communicate within the network is not adversely affected by increasing numbers of users.

While these assumptions have not been seriously investigated in the literature, their validity is equally important for the continued growth and development of the Web. A web authoring company contacted by Windrum (1998) suggested that - contrary to the prediction of Metcalfe's Law - there exists a threshold beyond which increasing numbers of users leads to a *decline* in the value of the network. Indeed this dissenting voice claimed that the value of the Web is

“already beginning to *decline*... [it] is becoming awash with rubbish. This is not information but noise. So while there is value in a network, it doesn't mean that simply increasing numbers will add more value... it is the *application* of the network which really matters and not the technology in itself” (quoted in Windrum, 1998).

The distinction between successful communication, which requires meaning and usefulness of a message on the part of the receiver, and mere information processing is clearly expounded by Luhmann (e.g. Luhmann, 1997, pp.1088-1095). Not only are the acts of sending and receiving different but the perceived value of a communication typically differs between the sender and the receiver. While the individual creating a new webpage may believe it to contain interesting data, it may have no appeal to anyone else, and so is treated as an unwanted message by other Web users. This generation of unwanted messages is what the web author describes as “noise”. By contrast, a message that *is* perceived to be of interest by others is described as containing “information”.

Returning to the key assumptions underpinning Metcalfe's Law, the web author is claiming that the assumption (i) no longer holds because the contribution of new users is negative. Although the validity of this claim is clearly a matter for empirical testing, it does appear rather strong. We note that the research conducted by Coffman

and Odlyzko (1998) suggests that growth over the last three years is not simply due to increasing numbers of new users but also to increasing volumes of transmissions by established users.

A weaker claim would be that the degree of common interest between users starts to decline when a network becomes very large. It is likely that the information provided by newer users will have a more limited audience and so, as the network continues to grow, the probability of each user wishing to communicate with every other user will at some point begin to decline. In the limit there may be an upper limit to the size of a sustainable network. Interestingly, Metcalfe himself acknowledges the possibility of an upper bound for networks containing a large number of users (N), observing that Metcalfe's Law

“might overestimate the value of a network for very large N. A user equipped to communicate with 50 million other users might not have all that much to talk about with each of them. So maybe the growth of systemic network value rolls off after some N” (Metcalfe, 1995).

While Metcalfe concedes the existence of an upper bound, he argues that this remains a long way off for the Web. Still, Metcalfe's Law is not quite clear of troubled waters. Although the upper bound may be some way off, more immediate restrictions to growth may be due to problems associated with assumption (ii). Indeed it is possible that a whole series of factors exist which could impose their own limits to growth - limits that have thus far been offset through technological innovation.

Many of the respondents interviewed by Windrum (1998) observed that rapid growth in the number of webpages is making Web navigation increasingly difficult. Searching for one page amongst 150 million pages is very much like looking for a needle in a haystack. If continued expansion of the Web adversely affects the capacity of each user to identify potentially useful data then the second key assumption of Metcalfe's Law's is undermined.

There are a number of possible measures that could offset this problem. Some users try to reduce search costs by localising their search activities to a limited number of 'favourite sites'. While the individual user will certainly reduce search costs in this manner, the efficiency of this particular strategy is questionable. Far more effective would be a restructuring of the Web to facilitate new forms of social communications. For example, a numbers of distinct areas or 'cells', each containing a set of users with specialist interests, could be developed. Common membership of a cell would then denote a degree of similarity between users. At the same time, breaking down the network into a series of smaller cells effectively reduces the search task since the membership of any one cell (n) is smaller than the membership of the entire Web (N). *Ceteris paribus* this effectively reduces the number of sites that need to be considered, increasing the efficiency of information search.

Portal Sites are one means of reconstructing communications currently being explored by Netscape and Microsoft. Each company has been trying to turn its site into a major information hub, allegedly offering visitors a quicker and more direct access to the content of the Web. Both companies argue that Portal Sites will go beyond the Yahoo index, which simply adds articles to an index, because they will assess the quality of content and provide 'top 10' type lists. What is more, through the use of a cookie - a user identifier which runs in the background - the Portal providers will be able to update information on what type of pages a user has accessed in the past and can display new pages which may be of interest next time the user logs on. In return for these services, Netscape and Microsoft stand to gain significant financial returns from advertising revenues. Efforts to reconstruct the communication structures of the Web have been limited to date. As one industry commentator recently noted,

“Portals largely remain a buzzword. Making them a reality will involve employing a lot of people to sift through material, categorise it and present it in an easy fashion on the respective databases” (quoted in Windrum, 1998).

Discussions with web designers and other practitioners suggest that hopes are currently pinned on the development of improved search engines - the 'technology fix' option - rather than a restructuring of social communications (Windrum, 1998). Thus, instead of seeking to improve the efficiency of communication under a given threshold constraint, innovative effort will focus on ways of raising the threshold constraint under which the network operates. In order to assess the relative strength and weaknesses of this particular means of sustaining Metcalfe's Law, we consider the basic characteristics of the popular search engines in section 2.

2. Search Engines

The vast majority of people surfing the Web will use one or other well-known, commercially-backed search engines. These generally give more dependable results because there is a commercial interest in maintaining and upgrading them to keep pace with the growing Web. Consequently a good listing in a major search engine is more likely to bring traffic to the site than a lesser-known search engine. Search engines fall into one of two generic types; those that are automatically generated by a computer and those that are indexed by a human on a directory⁶. With the exception of the Yahoo and LookSmart search engines, all other commercial engines (AltaVista, Excite, Lycos etc.) are of the first type.

⁶To confuse matters slightly, some companies providing automated indexes have begun to offer an associated directory in which a selection of sites are reviewed or rated. For the most part these reviewed sites do not appear as the default when a query is made. Instead, a user must consciously choose to see the reviews.

Automated engines comprise three key elements. The first is the spider (or crawler). It visits a webpage, 'reads' it, and then follows links to other pages within the website. Page titles, body copy and other elements are recorded (AltaVista claims to read each word of text). A spider will regularly return to the sites it finds to check for changes. Everything the spider finds goes into the second element of the search engine, the index. The index (sometimes called the catalogue) is like a giant book containing a copy of every webpage found by the spider. If a webpage changes, then this book is updated with the new information. Search engine software is the third element of the search engine. This programme sifts through the millions of pages recorded in the index to find matches to a search. It then ranks them according to some criteria of relevancy and lists them in order.

The most frequently discussed search engine problem is the *data glut* generated by automated engines. These typically generate concordances on far more links than the user has time to process, with little or no indication as to the nature of the pages and, hence, their suitability. Suggestions on how to improve matters abound. As well as being widely discussed on the Web, articles on the subject regularly feature in computer magazines, and there are even entire books dedicated to the subject (e.g. Hurst, 1997; Kennedy, 1998; and Pfaffenberger, 1996 respectively). However, even with using the advice provided on Boolean operators and the various advanced search options provided by different search engines, the data glut problem remains. Kennedy (1998) notes how, when searching for information on the ozone layer via the UK Lycos service, the initial search for 'ozone layer' produced 5998 matches (one of which is the name of a CD shop). By adding further keywords - 'aerosol' and 'cfc' - he was able to hone in on ever more relevant information sources, stripping away 5868 of the original matches. While impressive, he was still left with the task of downloading and assessing 130 pages - no mean feat! The serious question that arises is whether, even when using current best practice, the cost/benefit ratio of exhaustive searches on such long lists is appropriate for the typical user. Even if the 'noise-to-information' ratio does not worsen as the Web continues to grow, the cost of performing exhaustive searches on ever longer lists of webpage links *is* likely to worsen.

Yahoo advocates the use of human-indexed catalogues as a means of solving the data glut problem. Here a human physically reads each webpage and adds it to a catalogue after manually determining what index category it fits under. According to Yahoo, a hierarchically structured library of categories eliminates spurious links and avoids problems of multiple listings to the same website, thus honing the search to a limited number of relevant references. However, putting the theory to the test can reveal a number of shortcomings. Repeating the test conducted by Kennedy, searching for the keyword 'cfc' yields a total of 74 identified links⁷. These are filed under two separate categories: 'Business and Economy: Companies: Industrial Supplies: Materials: Plastics: Coatings: Manufacturers: CFC International' and 'Business and Economy: Companies: Environment: Waste Management: Recycling: Refrigerants and CFCs'.

⁷ The tests reported here were conducted on 11th August 1998.

Conducting a comprehensive search of 16 pages containing 74 links is not a trivial task. What is more, despite the claims, there is still evidence of spurious links entering the Yahoo listings. For example, there are two references to the 'CFC Youth for Christ Toronto - composed of energetic young Catholic men and women' on the first reported page. If the term 'ozone layer' is added to hone down the search then Yahoo reports *zero* matches, and suggests the user manually search the category 'Business and Economy: Companies: Industrial Supplies: Materials: Plastics: Coatings: Manufacturers: CFC International'. Turning to the links contained in the category 'Business and Economy: Companies: Environment: Waste Management: Recycling: Refrigerants and CFCs', searching with just the keywords 'ozone layer' proves an equally interesting experience. On this occasion the Yahoo search engine reports 9 links but that there is 'no index category'. Of these 9 links listed, 4 are the same website link, while one is apparently a broken link to a song titled 'Destruction of Ozone Layer Should Be Stopped'.

This appears to lend some support to the advocates of robot-indexed engines who argue that human-indexed catalogues only 'appear' to solve the data glut problem. Lower numbers of initial reported links are not due to an effective screening process but due to incomplete catalogues. Humans, it is argued, cannot work fast enough to keep up with millions of new webpages being created, and so the index will always be incomplete. Certainly the claim that human-indexing eliminates spurious links and avoids problems of multiple listings to the same website does not appear to stand up. If these experiences reported are at all representative, then there is a danger that the data glut problem is simply replaced by data impoverishment through *incomplete listings*. While different, this is equally pernicious to the efficient search of the Web.

John Pike (1997) has persuasively argued that the problem of incomplete listings is not unique to human-indexed engines but is shared by robot-indexed search engines. Pike noted that AltaVista claims to be the largest Web index with 31 million pages found on 476,000 servers. Given the existing estimate of around 50 million webpages, this led Pike to ask the question 'Where are the missing pages?'. The problem, Pike argues, is not just due to the time it takes for crawlers to trawl the Web but also due to AltaVista *sampling* sites rather than indexing every page. Pike's analysis of the automated search engine technology used by AltaVista led him to conclude that the problem of picking up the entire Web is actually shared by AltaVista and all the other robot-indexed search engines.

Further problems for search engines are being created through the increasing interactivity of websites. Web authors are now experimenting with the relatively static form of the website front-page, linking it with databases, graphics facilities and other text sources whose content of frequently changes. Clearly such developments pose a serious challenge to the text-based technology employed in today's search engines. For example, the link to a site can be set up in such a way that when someone accesses the site they go to a script for an Active serverpage, or something similar, which automatically injects a database containing the latest information. In

practice this involves creating a *specific* page for that date and time. A key attraction of this procedure is that registered users can be provided with specific information for which they have a known interest. Unfortunately, the current generation of search engines are not equipped to deal with such interactivity or the tailoring of data sources. The problem of interactive webpages can, to a certain extent, be offset by 'good practice'. Webpage authors can inform the various search engines of the existence of interactive pages by providing keywords on the front-page of the site (which does the dynamic generation) describing the type of content that 'can' be generated. However this only partially helps to resolve the problem because search engines will only see a link. They cannot actually crawl through a linked database in order to identify its content.

We end this section by observing that the limitations of the current generation of text-based search engines are likely to become even more exposed as the Web continues to evolve. With regards to dynamic webpages, it is hoped that matters will improve as XML (Extensible Markup Language) gradually replaces HTML over the next 5 years. XML is more structured, and much closer to a computer language, than HTML. By facilitating structured Web delivery, XML could prove an important stimulus for the development of search engines that can process words, pictures and numerical data. Still, the diffusion of XML alone will not solve problems of incomplete listings or data glut.

3. What are the Likely Impacts on Future Web Growth?

What are the implications for the continued growth of the Web? Will Metcalfe's Law continue to hold? When considering such questions it is useful to consider the case of *Moore's Law*. This law for semiconductor performance dates back to 1964 when Gordon E. Moore (now chairman of Intel Corp.) observed that the complexity, and hence the capacity, of state-of-the-art chips had doubled every 1.5 years since the introduction of the first planar-process transistor in 1959. Extrapolating the trend, Moore suggested that this doubling would continue in the future. Indeed this 'Law' has held remarkably well to the present day. Moore's Law is essentially a rule of thumb: it doesn't explain in any detail why this doubling of capacity should occur. By contrast, Metcalfe's Law is both an observational and a theoretical law because it provides an explanation for the observed behaviour.

Moore's Law is not a 'law of nature' but the outcome of a particular social process. What makes Moore's Law so interesting, as MacKenzie (1992), v. Lente (1993) and Swann and Gill (1993) have pointed out, is the extent to which it is sustained by the self-fulfilling beliefs of those involved in the development of semiconductors. Moore's 'law' is an accepted yardstick by which different players judge their performance. In order to meet the expected doubling of performance (fearful that if they do not achieve the necessary technological improvement within the given time frame, then their competitors will), ever greater levels of financial investment and R&D effort are associated with each successive generation of semi-conductors.

Returning to Metcalfe's Law we see that it, like Moore's Law, is not an exogenously given 'natural law' but the outcome of a process which involves several interest groups - hardware providers, software providers, ISPs, and users. The set of social relationships which underpin this particular technology network must be continually constructed and reconstructed if the Web is to develop and grow. In this paper we have identified one problem threatening this development and growth; there is a distinct danger that increasing problems of 'noise' will undermine the attraction of the Web to its end users, seriously threatening the technology in its current form.

We have also noted two possible courses of action that could improve Web navigation and help sustain Metcalfe's Law. One is a reconstruction of communications within a given technical constraint. An alternative is to raise the technological constraint via a 'technology fix'. This is the course of action which has been pursued thus far. Of course one should not preclude some combination of the two options, with a degree of communication restructuring of the Web going hand-in-hand with the development of new technological solutions. However, given the past history of the Web and evidence from other sectors (such as semi-conductors), one can feel fairly confident in predicting that increasing levels of effort and financial investment will be ploughed into the development of more powerful and sophisticated search technologies.

4. The Relationship Between Web Growth and the Value of the Web

In sections 5 and 6 we present an analytical treatment of the relationship between noise, effective search and the continued growth of the Web. This builds on the economics literature on network externalities. As well as facilitating further consideration of this relationship for the Web, we are able to generate a set of results which have import for a broader class of networks. As we shall see, the issue of increasing noise has important implications for current theoretical understanding of a variety of network systems, and for best practice in the use and design of such systems.

In this section, we make some preliminary observations about the functional form of this relationship. A key starting point for an analytical model is the relationship between the value of a network and the number of users. The choice of functional form will have an important bearing on the behaviour of the model. The economics literature on network externalities has traditionally assumed that network externalities are proportional to the size of the user base (e.g. Farrell and Saloner, 1985; Katz and Shapiro, 1985). However, as Swann (1998) observes, a linear relationship has little theoretical or empirical support.

By positing a square relationship, Metcalfe's Law differs markedly to the traditional linear assumption. In discussing this particular functional form, we draw attention to

an apparent disparity between the observations that are said to support the Law, and its actual mathematical form. Piecing together some of Bob Metcalfe’s own thoughts on his ‘law’ proves an insightful exercise. Metcalfe states that the law “implies that when networks reach critical mass, like Ethernets and the Internet have, they go BOOM!” (Metcalfe, 1996). Then, because a “user equipped to communicate with 50 million other users might not have all that much to talk about with each of them... systemic network value rolls off after some N” (Metcalfe, 1995). This description is instantly familiar to students of technological innovation as the classic S-shaped diffusion curve. However, this curve is the product of a sigmoid function rather than a square polynomial function. There are two ingredients necessary to generate an S-shaped diffusion curve. The first is a positive feedback mechanism - here the existence of network externalities - that initially generates an exponential rate of growth. The second is some process that places an upper bound on the growth process. For example, the positive feedback associated with network externalities may dissipate, either because the degree of common interest between users starts to decline [assumption (i)] or because users find it increasingly difficult to identify, and therefore communicate, with others sharing a common interest [assumption (ii)]. In either event, there will be an upper bound on network size and an S-shaped diffusion curve will be generated.

5. A Formal Model of Network Externalities and the Size of a Network

The following model builds on Swann’s (1998) analysis of the functional form of network externalities. This section focuses on the value associated with being part of a large network. The following section (6) then takes into account the various types of congestion cost that can limit or detract from this value.

Suppose that users can be differentiated along a single⁸ axis $0 \leq \beta \leq 1$, and that the density of each type of user β is given by $m(\beta)$. Suppose also that the diffusion of the Web amongst the user community is summarised by the density function $f(\beta, t)$, giving the probability that a user of type β will start to use the Web at time t . Then the conditional diffusion curve, describing the proportion of users of type β that have started to use the Web *at or before* time t^* is given by:

$$F(\beta, t^*) = \int_{-\infty}^{t^*} f(\beta, t) dt \quad \mathbf{1}$$

while the aggregate diffusion curve, describing the *total* number who have adopted *at or before* time t^* is given by:

⁸ While this assumption simplifies the mathematics there is no significant loss of generality in describing users by one characteristic.

$$N(t^*) = \int_{-\infty}^{t^*} \int_0^1 m(\beta) f(\beta, t) d\beta dt \quad 2$$

As we observed earlier, it is unlikely that a user is going to visit every page and so assumption (i) of Metcalfe's Law will not hold. Therefore we let the function $g(\alpha, \beta)$ describe the probability that a user of type α will wish to examine the webpages generated by a type β user. Assuming the total utility to a user of type α is proportional to the number of pages (s)he visits⁹, the utility of α at time t^* is given by:

$$u(\alpha, t^*) = \int_{-\infty}^{t^*} \int_0^1 m(\beta) g(\alpha, \beta) f(\beta, t) d\beta dt \quad 3$$

Inspection of equations (2) and (3) confirms that in general $u(\alpha, t^*)/N(t^*)$ is not independent of t^* . This means that in general the utility of an individual user, $u(\alpha, t^*)$, is not proportional to the size of the network, $N(t^*)$.

Swann (1998) argues that, in general, utility is not proportional to the installed user base. He shows there are two special cases when this relationship will be proportional, but neither of these is particularly plausible in the present context. One condition would be that all Web users are equally likely to access any page. The other condition would be that the diffusion of Web usage is identical across all sectors and all types of users. Swann (1998) also shows that a weaker condition exists under which the average utility to all possible users is proportional to the size of the network. Again this assumption is rather implausible since it would imply that all webpages attract the same number of visitors.

In seeking to specifying a *plausible* set of assumptions about diffusion and the probability of one type of user wanting to inspect the webpages of another type, Swann finds that individual user utility is likely to be an s-shaped function of total network size. Assume, for example, that $F(\beta, t^*)$ is given by a logistic curve:

$$F(\beta, t^*) = \frac{\exp\{\beta + t^*\}}{1 + \exp\{\beta + t^*\}} \quad 4$$

and that $g(\alpha, \beta)$ declines exponentially with the square of the difference between α and β :

⁹ This neglects the diminishing returns that would set in if people get tired of reading webpages.

$$g(\alpha, \beta) = g_0 \cdot \exp\{-(\alpha - \beta)^2\}$$

5

The parameter g_0 represents the probability that someone would wish to consult the webpage of another, almost identical user. This is, so to speak, the probability that one would wish to consult the webpage designed by a twin brother or sister.

While the expression for u cannot be integrated explicitly when $F(\cdot)$ and $g(\cdot)$ take these particular forms, it can be solved numerically. Figure 2 plots the function $u(\alpha, t^*)$ against $N(t^*)$, for three values of α : 1, 0.5 and 0. Recall that $0 \leq \alpha \leq 1$, so these represent (respectively) a pioneering user, a median user and a late adopter.

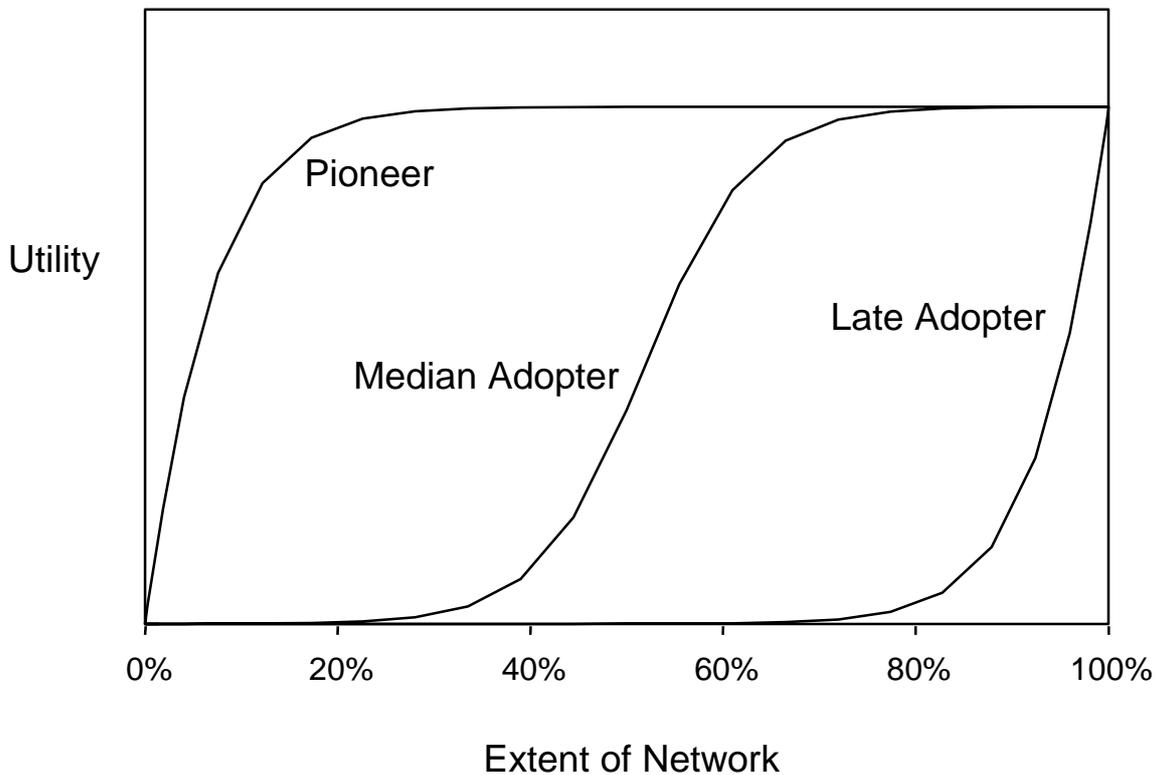


Figure 2. The relationship between utility and the extent of the network

For the median user, utility is an s-shaped function of the size of the network. When the network is small, it mainly contains members of the pioneering ‘class’ (high α) so this is of little use to the median user. It is when median users start to adopt that utility starts to take off. But when the network gets very large, the new members are typically late adopters, and the median user gains little from their membership, so

marginal utility falls off steadily.

This shape (in particular the point of inflexion) is unfamiliar in the analysis of network externalities which, as previously noted, tends to assume a linear or logarithmic functional form. The reason is that the user base variable disguises the diversity of users, and it is the order in which users appear that determines the shape of the function. The pioneering user experiences diminishing marginal utility. From his/her perspective, the best users are the first to adopt, and subsequent growth in the network is of diminishing marginal value. This shape is not so unfamiliar in the discussion of network externalities. Conversely, the late adopter experiences increasing marginal utility. The early adopters are not valued, while the last adopters are the most valuable. This last shape is perhaps the most unexpected.

What of the other intermediate consumers between $\alpha = 0$ and $\alpha = 1$? They would also have s-shaped curves, but these would not be symmetric. Indeed, it could be said that all of these functional forms are s-shaped, apart from the extremes at each end of the user spectrum¹⁰.

6. The Costs of Web Congestion

The previous section described the benefit that a user can - in principle - obtain from the Web as a network resource. We stress the words *in principle* because, as the previous sections have shown, the user incurs certain costs in accessing this information. The aim of this section is to explore how the form of the relationships in Figure 1 may alter as we take account of these costs.

In what follows, we shall compute a net benefit of using the Web, defined as:

$$B(\alpha, N) = u(\alpha, N) - C(N) \quad 6$$

where $u(\cdot)$ is the (money-metric) utility function of section 5 while $C(\cdot)$ describes how the costs of using the Web change as the size of the network increases. In particular we shall explore two types of costs, both arising from the congestion that occurs when the Web is heavily used, but of rather different types.

¹⁰ In some contexts it would be reasonable to expect that, beyond a certain point, studying additional webpages yields diminishing marginal utility. This would introduce an additional source of non-linearity in the relationship between utility and network size. However we have assumed that each additional and relevant webpage yields constant marginal utility. The non-linearities in the relationship between individual utility and total network size arise because, from α 's perspective, the number of other *relevant* users does not increase at the same rate as the aggregate measure of network size.

The first type of cost arises because (as discussed previously) the number of search ‘hits’ grows rapidly as the Web grows in size, while the relevance of many of these hits to a particular individual user is limited. In short, searching the Web involves dealing with ‘unwanted noise’. Selecting between sites containing useful information and those containing unwanted noise is a costly exercise. One can identify three distinct aspects to this particular search cost;

- (a) As the number of adopters N increases, so the total number of websites that need to be searched increases. Furthermore, the number of pages offered on existing websites tends to increase over time as users expand and develop their sites. Assume for simplicity that the number of pages offered on each new website is θN .
- (b) As the number of new websites and new webpages increases over time, so the number of internal and external site links which an automated spider needs to crawl increases. Suppose this additional condition is described by ϕN .
- (c) Good website promotion becomes necessary as N increases. Since potential visitors are unlikely to trawl through more than a couple of pages generated by a keyword search on a search engine, the pressure is for website managers to get high rankings on the popular search engines. For the existing text-based search technology, promotion requires (amongst other things) the careful selection of keywords which a potential visitor is likely to put into a search engine. These keywords are placed in the header of each webpage and for search engine spiders to read. Unfortunately useful keywords are heavily subscribed. As the total number of contributions to the Web increases, so the number of hits for a popular keyword increases. Suppose the number of hits per keyword is described by ψN .

The combined effect of (a), (b) and (c) means the total number of pages a search engine needs to visit and process increases with N at the rate of $\theta\phi\psi N^3$. This cubic form results because each increases with N ; the total number of potentially relevant sites, the number of pages and links per relevant site, and the proportion of all sites activated by a particular keyword. One immediately starts to appreciate the problems faced by search engine companies in trawling and cataloguing the Web. Indeed there may be further search costs associated with an increase in N not considered here. Still, these neatly capture the issues raised in section 2. While the form of each effect may not necessarily be proportional, it marks a reasonable starting point.

Let us assume that search costs are given by:

$$Cost\ of\ Search = C_1 \frac{\theta\phi\psi N^3}{S} \quad 7$$

where θ , ϕ and ψ are defined as above and S represents the ability of the current

search technology to identify relevant hits such that the user can extract useful information. C_1 represents the cost per unit of search time.

A second type of cost is that associated with delays in downloading material when the Web is heavily used. All Web users have experienced these delays. While the cost in time for an individual transaction may be limited, the cost can mount up rapidly as use increases. In what follows we shall use some simple principles from computer system measurement to give form to these costs. These principles recognise that as the load on a computer system increases towards its saturation level, then delays in processing any individual transaction start to increase. If the load on the Web (or throughput) is proportional to the total number of users (N), then we can write this cost as a function of N . Here we assume, following Beizer (1978) (also see Swann; 1986), that the ‘throughput-delay’ curve has the form:

$$Cost\ of\ Delay = C_2(N^{sat} - N)^{-k} \quad 8$$

where N^{sat} is the saturation load on the Web. At this saturation point the system ‘clogs up’ and $k \geq 1$ is a constant. As N approaches N^{sat} costs start to increase, without limit, because this form has an asymptote as $N \rightarrow N^{sat}$.¹¹

Combining these two cost functions with the original net benefit function, we derive:

$$B(\alpha, N) = u(\alpha, N) - C_1 \frac{\theta\phi\psi N^3}{S} - C_2(N^{sat} - N)^{-k} \quad 9$$

A model based on equation 9 is sufficiently rich to explore a range of Web-related issues. Here we consider two simple experiments that examine the relationship when, *ceteris paribus*, the first and second types of cost respectively dominate.

Figure 3 shows the relationship for different values of the parameter S , i.e. the ability of the search technology to identify websites containing potentially relevant information. Except for very large of S , the net benefit to the user reaches a maximum, and then starts to decline. This decline happens because the user cannot easily sort through the large number of ‘hits’ obtained. When S is small ($S=1$), the maximum net benefit is reached when the network is not especially large, and then the net benefit declines as sorting costs increase. As S increases, however, the improved

¹¹ This may somewhat overstate the problem on the Web. This equation was originally derived for a centralised mainframe computer system. While traffic on the Web is highly concentrated it is probably not as extreme as this. However we shall use this functional form as a useful first approximation.

filtering capacity of new, more sophisticated search technologies reduces costs and increases net benefit to the user.

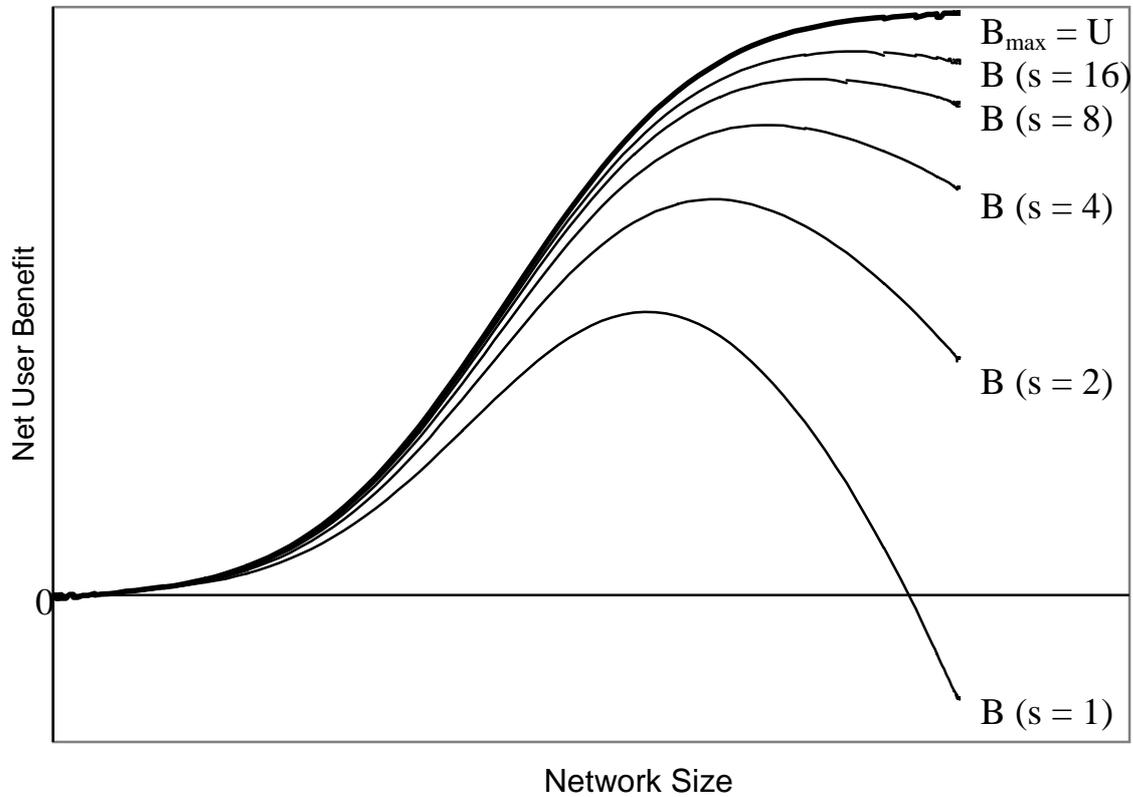


Figure 3. The effect of improvements in the quality of search engines (S)

By contrast, Figure 4 shows the relationship for different values of N^{sat} , when the second type of cost is dominant. Once again, net benefit reaches a maximum and then begins to tail off - unless N^{sat} is very large. It shows that when N^{sat} is small, the net benefit to the user tails off rapidly as the network size reaches saturation level. This will be the case even when, as here, the parameter k is set equal to 1. If $k > 1$ then the curves would tail off even more rapidly. The larger the value of N^{sat} , the more moderate are the congestion costs imposed on the Web user.

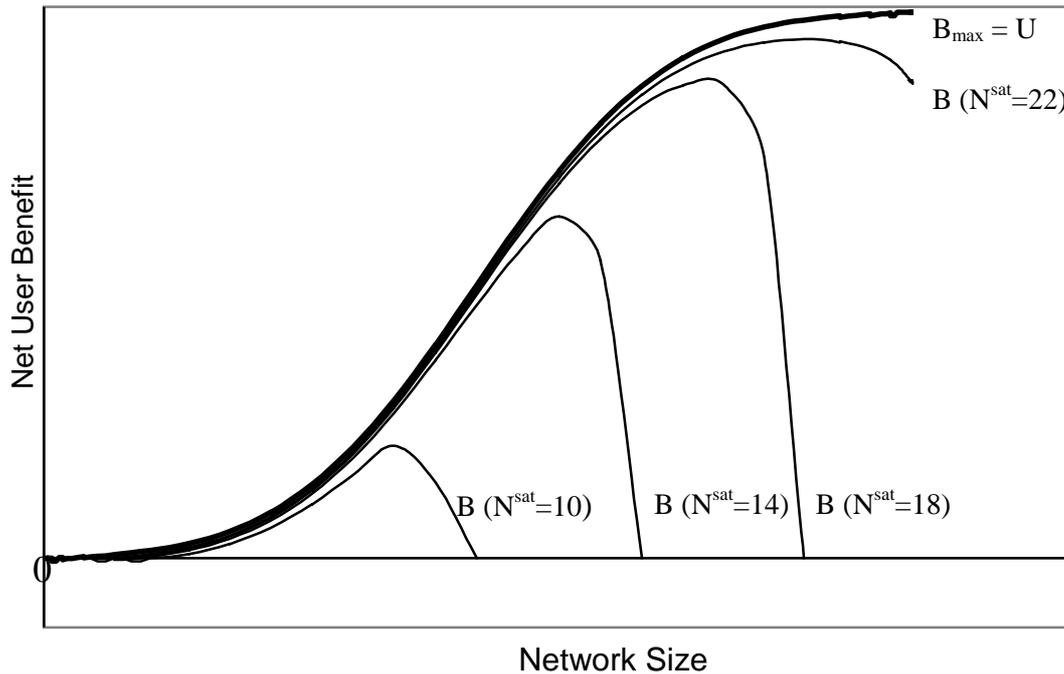


Figure 4. The effect of improvements in network capacity (N^{sat})

There are two main differences between Figures 3 and 4. First, in Figure 4, a doubling in N^{sat} has a large impact on the relationship: to see this, compare the curves for $N^{\text{sat}}=10$ and $N^{\text{sat}}=22$. By contrast, doubling the value of S in Figure 3 does not appear to have such a large impact on the net benefit relationship. Second, the rate at which net benefits start to decline after the maximum point on each curve is markedly different in Figures 3 and 4. In Figure 4, net benefit declines very rapidly as the Web approaches saturation. The performance of the Web at a particular date may appear satisfactory, but it may only need to expand a little more before serious congestion costs materialise. In Figure 3, by contrast, net benefit declines far more gradually as search costs increase.

7. Policy Conclusions

The foregoing analysis has a number of interesting policy implications. The analytical model developed in sections 5 and 6 highlights three generic areas that could be targeted in order to improve congestion on the Web;

- (a) Invest in developing improved Web navigation software, thereby raising S .
- (b) Invest in hardware and communications to increase N^{sat} , and perhaps reduce k .

(c) Create incentives for website designers to adopt a code of best practice such that the parameters θ , ϕ and ψ are kept as low as possible without compromising the quality of information placed on the Web.

In this paper we have concentrated on option (a) because this is where most R&D effort is currently being focused. It is therefore interestingly to observe that Figure 2 suggests quite large (proportionate) changes in S are required before realised net benefits approach their maximum potential. However this does not, by itself, imply that such investments are costly relative to options (b) and (c). Turning to option (b), figure 3 suggests that increasing N^{sat} is a relatively effective means of alleviating congestion costs. It may therefore be better to focus on option (b) rather than option (a) in order to sustain Metcalfe's Law. Having said this, we have noted how these results are sensitive to the functional forms chosen in equations 7 and 8. Furthermore, decisions on whether to invest effort in raising N^{sat} depends on whether users perceive delays in downloading pages as a major problem. If not, then such investments will be less valued than improvements in S .

A third course of action is to focus on (c). Codes of best practice amongst website designers can ease problems of Web congestion. For example, one well-known problem is the misuse of popular keywords in order to improve a site's search engine ranking. *Ceteris paribus*, such activities adversely affect parameter ψ . In a similar fashion, the misuse of site links to improve rankings, will also, *ceteris paribus*, raise the value of ϕ , adversely affecting user benefit. Good housekeeping, such as deleting dead links can also help matters greatly. Finally, codes of practice regarding the efficient use of webpages can assist in controlling the value of θ . The model helps demonstrate just how important these factors may be in determining the overall value of the Web for the typical user. Yet, rather surprisingly, little research has been conducted in this policy area. In order to redress this imbalance we are examining the importance of standards (broadly defined) and best practice to the continued growth of the Web, the results of which will be reported in future papers.

References

- Arthur, B.W., 1989, 'Competing Technologies, Increasing Returns, and Lock-in by Historical Events', *Economic Journal*, 99 (394), 116-131.
- Beizer, B., 1978, *Micro-Analysis of Computer System Performance*, Van Nostrand Reinhold: New York.
- Coffman, K.G. and Odlyzko, A.M., 1998, *The Size And Growth Of The Internet*, 1st July 1998, available at <amo@research.att.com>.
- David, P.A., 1985, 'Clio and the Economics of QWERTY', *American Economic Review*, 75, 332-336.
- David, P.A. and Greenstein, S., 1990, 'The Economics of Compatibility Standards: In Introduction to Recent Research', *Economics of Innovation and New Technology*, 1 (1/2), 3-41.
- Economides, N. and Encaoua, D., eds., 1996, 'Special Issue on The Economics of Networks', *International Journal of Industrial Organisation*, 14, 673-886.
- Farrell, J. and Saloner, G., 1985, 'Standardisation, Compatibility and Innovation', *RAND Journal of Economics*, 16, 70-82.
- Gilbert, R., ed., 1992, 'Symposium on Compatibility', *Journal of Industrial Economics*, 40, 1-123.
- Gilder, G., 1993, 'Metcalf's Law and Legacy', *Forbes ASAP Magazine*, 13 September 1993, also available at <<http://www.forbes.com/asap/gilder/telecosm4a.htm>>.
- Gray, M., 1997, *Web Growth Report*, available at <<http://www.mit.edu/people/mkgray/>>.
- Hurst, M., 1997, *The Problem With Web Searching*, 23 January 1997, available at <<http://www.creativegood.com/help/index.html>>.
- Katz, M., and Shapiro, C., 1985, Network externalities, competition and compatibility, *American Economic Review*, 75, 424-440.
- Kennedy, S., 1998, 'Needle in a Haystack', *EasyLife*, Issue 2, U-Net customer magazine, Summer 1998.
- Kelly, F.P., 1997, 'Charging and Accounting for Bursty Connections', in L.W. McKight and J.P. Bailey (eds.), *Internet Economics*, MIT Press: Cambridge, Mass.
- v. Lente, H., 1993, *Promising Technology: The Dynamics Of Expectations In Technological Developments*, PhD Dissertation, University of Twente, Netherlands.
- Luhmann, N., 1997, *Die Gesellschaft der Gesellschaft*, Zweiter Teilband (Part II), Suhrkamp: Frankfurt.
- MacKenzie, D., 1992, 'Economic And Sociological Explanations Of Technological Change', in R. Coombs, P. Saviotti and V. Walsh (eds.), *Technological Change and Company Strategies: Economic and Sociological Perspectives*, Academic Press: London.
- MacKie, J.K. and Varian, H.R., 1995, 'Pricing the Internet', in B. Kahim and J. Keller (eds.), *Public Access to the Internet*, MIT Press: Cambridge, Mass.
- Metcalf, R.M., 1996, 'Computer Laws Galore, But One Is Holding Back The Information Age', *InfoWorld*, 6th May 1996, available at <<http://www.infoworld.com/cgi-bin/displayNew.pl?metcalfe/bm050696.htm>>.
- Metcalf, R.M., 1995, 'Metcalf's Law', *InfoWorld*, 2nd October 1995.
- Odlyzko, A.M., 1997, *A Modest Proposal For Preventing Internet Congestion*, 3rd September 1997, available at <amo@research.att.com>.
- Pike, J., 1997, *Flawed AltaVista Internet Search Engine*, article distributed via Red Rock Eater News Service, 26th March 1997.

- Pfaffenberger, B., 1996, *Web Search Strategies*, IDG Books Worldwide.
- Swann, G. M. P., 1990, 'The Growth of an IT Network: A Case Study of PC Applications Software', in Berg J. and H. Schumny (eds.), *An Analysis of the IT Standardisation Process*, North Holland: Amsterdam.
- Swann, G. M. P., 1986, *Quality Innovation and Demand: An Economic Analysis of Rapid Improvements in Electronic Components*, Frances Pinter Publishers: London.
- Swann, G. M. P., 1998, *Are Network Externalities Linear?*, Unpublished paper, Manchester Business School.
- Swann, G. M. P. and Gill, J., 1993, *Corporate Vision and Rapid Technological Change*, Routledge: London.
- Swann, G. M. P. and Shurmer, M., 1995, 'An Analysis of the Process Generating De Facto Standards in the PC Spreadsheet Software Market', *Journal of Evolutionary Economics*, 5, 119-132.
- Windrum, P., 1998, *The Collective Invention of the World-Wide Web*, Report to DGXII, Colline Project, December 1998.
- Windrum., P., and Flanagan, K., 1997, *Recent Trends And Patterns Of Innovation In UK Web Publishing And Design*, Report to DGXII, SI4S Project, December 1997, PREST, Manchester.