

A refutation of Metcalfe's Law and a better estimate for the value of networks and network interconnections

Andrew Odlyzko¹ and Benjamin Tilly²

¹ Digital Technology Center, University of Minnesota
499 Walter Library, 117 Pleasant St. SE
Minneapolis, MN 55455, USA

odlyzko@umn.edu

<http://www.dtc.umn.edu/~odlyzko>

² ben.tilly@operamail.com

Preliminary version, March 2, 2005

Abstract. Metcalfe's Law states that the value of a communications network is proportional to the square of the size of the network. It is widely accepted and frequently cited. However, there are several arguments that this rule is a significant overestimate. (Therefore Reed's Law is even more of an overestimate, since it says that the value of a network grows exponentially, in the mathematical sense, in network size.) This note presents several quantitative arguments that suggest the value of a general communication network of size n grows like $n \log(n)$. This growth rate is faster than the linear growth, of order n , that, according to Sarnoff's Law, governs the value of a broadcast network. On the other hand, it is much slower than the quadratic growth of Metcalfe's Law, and helps explain the failure of the dot-com and telecom booms, as well as why network interconnection (such as peering on the Internet) remains a controversial issue.

1 Introduction

The value of a broadcast network is almost universally agreed to be proportional to the number of users. This rule is called Sarnoff's Law. For general communication networks, in which users can freely interact with each other, it has become widely accepted that Metcalfe's Law applies, and value is proportional to the square of the number of users. This "law" (which is just a general rule-of-thumb, not a physical law) has attained a sufficiently exalted status that in a recent article it was classed with Moore's Law as one of the five basic "rule-of-thumb 'laws' [that] have stood out" and passed the test of time [26]. Even as far back as 1996, Reed Hundt, the then-chairman of the Federal Communications Commission, claimed that Metcalfe's Law and Moore's Law "give us the best foundation for understanding the Internet" [9]. Metcalfe's Law was frequently cited during the dot-com and telecom booms to justify business plans that had the infamous "hockey-stick" revenue and profit projections. The hopes and promises were that once a service or network attained sufficient size, the non-linear growth of Metcalfe's Law would kick in, and network

and bandwagon effects would start to operate to bring great riches to the venture’s backers. This led to the argument that there was little need to worry about current disappointing financial results, and all efforts should be devoted to growth.

The fundamental insight, that the value of a general communication network grows faster than linearly, appears sound. Although it had not been enunciated explicitly until a few decades ago, it had motivated decision makers in the past. For example, it affected the development of the phone system a century ago, as shown by the following quote [19]:

The president of AT&T, Frederick Fish, believed that customers valued access and that charging a low fee for network membership would maximize the number of subscribers. According to Fish, the number of users was an important determinant of the value of telephony to individual subscribers. His desire to maximize network connections led the firm to adopt a pricing structure in which prices to residential customers were actually set below the marginal cost of service in order to encourage subscriptions. These losses were made up through increased charges to business customers.

Serious quantitative modeling of network effects dates back to the 1974 paper of Jeffrey Rohlfs [24] (see also [25], and some of the references in [24] to earlier work in the area), which was motivated by the lack of success of the AT&T PicturePhone™ videotelephony service. (For a survey of economics literature in this area, see [7].) Metcalfe’s Law itself dates back to a slide that Bob Metcalfe created around 1980, when he was running 3Com, to sell the Ethernet standard. It was dubbed “Metcalfe’s Law” by George Gilder [8] in the 1990s, and citations to it started to proliferate [9, 15, 16].

The foundation of Metcalfe’s Law is the observation that in a general communication network with n members, there are $n(n - 1)/2$ connections that can be made between pairs of participants. If all those connections are equally valuable (and this is the big “if” that we will discuss in more detail below), the total value of the network is proportional to $n(n - 1)/2$, which, since we are dealing with rough approximations, grows like n^2 .

Metcalfe’s Law is intuitively appealing, since our personal estimate of the size of a network is based on the uptake of that network among friends and family. Our derived value also varies directly with that metric. We therefore see a linear relationship between the perceived size and value of that network.

Reed’s Law [23] is based on the further insight that in a communication network as flexible as the Internet, in addition to linking pairs of members, one can form groups. With n participants, there are 2^n possible groups, and if they are all equally valuable, the value of the network grows like 2^n .

The fundamental fallacy underlying Metcalfe’s and Reed’s laws is in the assumption that all connections or all groups are equally valuable. The defect in this assumption was pointed out a century and a half ago by Henry David Thoreau. In *Walden*, he wrote [28]:

We are in great haste to construct a magnetic telegraph from Maine to Texas; but Maine and Texas, it may be, have nothing important to communicate.

Now Thoreau was wrong. Maine and Texas did and do have a lot to communicate. Some was (and is) important, and some not, but all of sufficient value for people to pay for. Still,

Thoreau's insight is valid, and Maine does not have as much to communicate with Texas as it does with Massachusetts or New York, say.

In general, connections are not used with the same intensity (and most are not used at all in large networks, such as the Internet), so assigning equal value to them is not justified. This is the basic objection to Metcalfe's Law, and it has been stated explicitly in many places, for example [12, 14, 20, 25, 26]. This observation was apparently first made by Metcalfe himself, in his first formal publication devoted to Metcalfe's Law [15]. Some users (those who generate spam, worms, and viruses, for example) actually subtract from the value of a network. Even if we disregard the clearly objectionable aspects of communication, such as spam, many users do complain about the problem of dealing with too many options.

There are additional scaling arguments that suggest Metcalfe's and Reed's laws are incorrect. For example, Reed's Law is implausible because of its exponential (in the precise mathematical sense of the term) nature [10]. If a network's value were proportional to 2^n , then there would be a threshold value m such that for n below $m - 50$, the value of the network would be no more than 0.0001% of the value of the whole economy, but once n exceeded m , the value of the network would be more than 99.9999% of the value of all assets. Beyond that stage, the addition of a single member to the network would have the effect of almost doubling the total economic value of the world. This does not fit general expectations of network values and thus also suggests that Reed's Law is not correct.

Metcalf's Law is slightly more plausible than Reed's Law, as its quadratic growth does not lead to the extreme threshold effect noted above, but it is still improbable. The problem is that Metcalfe's Law provides irresistible incentives for all networks relying on the same technology to merge or at least interconnect. To see this, consider two networks, each with n members. By Metcalfe's Law, each one is (disregarding the constant of proportionality) worth n^2 , so the total value of both is $2n^2$. But suppose these two networks merge, or one acquires the other, or they come to an agreement to interconnect. Then we will effectively have a single network with $2n$ members, which, by Metcalfe's Law, will be worth $4n^2$, or twice as much as the two separate networks. Surely it would require a combination of singularly obtuse management and singularly inefficient financial markets not to seize this obvious opportunity to double total network value by a simple combination. Yet historically there have been many cases of networks that resisted interconnection for a long time. For example, a century ago in the U.S., the Bell System and the independent phone companies often competed in the same neighborhood, with subscribers to one being unable to call subscribers to the other (see [20], or [17] for much more detail). Eventually interconnection was achieved (through a combination of financial maneuvers and political pressure), but it took two decades. In the late 1980s and early 1990s, the commercial online companies such as CompuServe, Prodigy, AOL, and MCIMail, provided email to subscribers, but only within their own systems, and it was only in the mid-1990s that full interconnection was achieved. More recently yet, AOL for many years resisted interconnecting its IM system with those of its competitors, and it is only in 2004 that an agreement, at least in principle, was reached to carry this out. In addition, short messaging services of some wireless carriers do not interoperate. And there has been a long series of controversies about interconnection policies of ISPs (Internet Service Providers). Thus the general conclusion (drawn first in

[20]) is that the incentives to interconnect cannot be too strong, and so Metcalfe's Law cannot be valid.

If we reject Metcalfe's and Reed's laws, can we replace them with a more accurate but still simple estimate of the value of a network? We propose $n \log(n)$ as an alternate rule-of-thumb valuation of a general communication network of size n . Metcalfe's law would hold if the value an individual personally gets from a network is directly proportional to the number of people in that network. This doesn't seem to hold, there is some law of diminishing returns that applies. Several arguments that we will present suggest that the value that a single user gets from being in a network of n people scales as $\log(n)$, leading to a rule-of-thumb valuation of $n \log(n)$ for a network of size n . This growth rate is faster than the linear growth of Sarnoff's Law, and so explains why connectivity and not content is king [21], and why various network and bandwagon effects do operate. On the other hand, this growth rate is only slightly faster than linear, and this helps explain why interconnection often requires time, effort, and in many cases regulatory pressure to achieve. If we have two networks, each with $n = 2^{20} = 1,048,576$ members, then (assuming the logarithm in our formula is to base 2, and the constant of proportionality is 1) each is valued at $20n$, for total value of $40n$. If these two networks interconnect, the resulting single network will have size $2n = 2^{21} = 2,097,152$, and its value by our suggested formula will be $42n$, only a 5% gain over the $40n$ total valuation of the two separate networks. The modest size of such gains helps explain why network and bandwagon effects, although definitely present and non-trivial, are often small, and therefore why so many dot-com and telecom ventures came to grief.

There have been and continue to be controversies about interconnection policies of ISPs. A particularly sensitive issue is the frequent refusal of large ISPs to peer (roughly speaking, exchange traffic freely without payment) with smaller carriers. (The refusal of AOL to interconnect instant messenger systems is very similar.) This behavior has often been attributed to abusive exploitation of market power. But there may be a more innocent explanation, based on the economic value that interconnection generates. As we show in Section 2, if Metcalfe's Law held, then interconnection would produce equal value for any two network, irrespective of their relative sizes. Hence refusal to interconnect without payment would have to be due to either obtuseness on the part of management or strategic gaming. However, if network value scales like $n \log(n)$, as we argue (or by most other rules of this type, the quadratic growth of Metcalfe's Law is very unusual in this regard) then relative gains from interconnection depend on the sizes of the networks. In this case the smaller network gains considerably more than the larger one. This produces an incentive for larger networks to refuse to interconnect without payment, a very common phenomenon in the real economy.

Our proposed $n \log(n)$ rule for network valuation does have the advantage of being consistent with observed behavior of actual communication networks, at least when it comes to rough implications of scaling. In sections 3, 4, and 5, we will present some quantitative heuristics that also lead to this same $n \log(n)$ valuation of a network of size n , although in both cases one could modify those heuristics to obtain somewhat different estimates. Obviously any rule of this type cannot be very accurate. Real networks vary wildly in

their behavior. It is even conceivable (as has been proposed in [14,25]) that their value may decrease as they grow (say as a result of spam). And there have been communication technologies that almost disappeared after a meteoric rise (as with citizens' band radios). But the overwhelmingly dominant pattern has been for usage and revenues of general communication networks to grow, often even when competing technologies appear [20]. As just one example, ordinary mail usage continued growing until just a few years ago, in spite of complaints about junk mail and the availability of numerous other services. Thus the increase in value of a network with increasing size appears to be well supported by evidence.

Over the last decade a rich literature has developed on the structure of networks. (A detailed technical survey is available in [18], an account aimed at a general audience in [29], and many references to recent literature can be found in [22].) But it appears hard to obtain any clear implications of this work for values of networks, especially since such values depend not only on structure, but also on usage. Furthermore, there are time dependencies in network valuations, as people learn to use them more intensively and in new ways, even in the absence of new members. Reed's insight about value of group formation plays a key role here, so it is necessary to study not just pair interactions, but group ones as well. (Kilkki and Kalervo have made a first step towards modeling such phenomena in [10].)

The general conclusion is that accurate valuation of networks is complicated, and no simple rule will apply universally. But if we have to select a parsimonious rule, one that is concise and yet captures many of the key features of communication networks, then we feel that our $n \log(n)$ formula fits the available evidence and is supported by reasonable heuristics.

2 Value of interconnection

When Sarnoff's Law holds, and the value of a network grows linearly in its size, as it does for a broadcast network, there is no net gain in value from combining two networks. Mergers and acquisitions in such situations are likely driven by other factors, such as scale efficiencies in operations, or attempts to increase bargaining power in purchases of content. On the other hand, when the value of a network grows faster than linearly in its size, then generally (subject to some smoothness assumptions we will not discuss), there is a net gain from a merger or interconnection, as the value of the larger network that results is greater than the sum of the values of the two constituent pieces. We next consider how that gain is distributed among the customers of the two networks.

Let us assume that Metcalfe's Law holds, and we have two ISPs, call them A and B , with m and n customers, respectively, and that on average the customers are comparable. Then interconnection would provide each of the m customers of A with additional value n (assuming the constant of proportionality in Metcalfe's Law is 1), or a total added value of

$$mn$$

for all the customers of A . Similarly, each member of B would gain m in value, so all the customers of B would gain total value of

$$nm$$

from interconnection. Thus aggregate gains to customers of A and B would be equal, and the two ISPs should peer, if they are rational. However, the incentives are different if our $n \log(n)$ rule for network valuation holds. In that case, each of the m customers of A would gain value $\log(m+n) - \log(m)$ from interconnection, and so all the customers of A would gain in total

$$m(\log(m+n) - \log(m)).$$

On the other hand, the total gain to to customers of B would now be

$$n(\log(m+n) - \log(n)).$$

If m and n are not equal, this would no longer be the same as the total gain to to customers of A . As a simple example, if $m = 2^{20} = 1,048,576$, and B has 8 times as many customers as B , so $n = 2^{23} = 8,388,608$, then (again taking logarithms to base 2) we find that interconnection would increase the value of the service to A 's customers by about 3,323,907, while B 's customers would gain about 1,425,434. Thus the smaller network would gain more than twice as much as the larger one. This clearly reduces the incentive for the latter to interconnect without compensation. This is a very simplistic model of ISP interconnection, of course, and it does not deal with other important aspects that enter into actual negotiations, such as geographical spans of networks, and balance of outgoing and incoming traffic. All we are trying to do is show that is that there may be sound economic reasons for larger networks to demand payment for interconnection from smaller one, a very common phenomenon in real life.

3 Gravity laws and the value of locality

Thoreau's comments in *Walden* about Maine and Texas not having much to communicate reflect the phenomenon that most traffic, whether in physical goods, or in information, has historically been local. Although distance is in many aspects becoming less important, as described by the "death of distance" phrase [3] and some jobs are being outsourced halfway around the world, locality is still important. As was noted in [20], an investment bank relocated its high-tech branch from San Francisco to Menlo Park, to be even closer to Silicon Valley. A recent report on New York City's recovery from the Sept. 11, 2001 disaster noted that businesses displaced by that tragedy "tended to relocate in areas where other businesses in their line of work were already clustered" [11]. Some careful quantitative studies [6] (with an account for a general audience in [30]) show that researchers working far apart lose productivity. Thus the value of geographical locality is still important.

Furthermore, even in cyberspace, one can use geographical modeling to represent other types of distances that exist. Although we are all supposedly connected through a chain of at most five acquaintances with each other [13], most of our interactions are with smaller, well-defined groups. A telephone-based language translation service that AT&T had started had disappointing results, since it was primarily emergency and legal service professionals who used it. What was discovered is that people speaking Tibetan, say, had more interest

in speaking with others who spoke Tibetan, even when those lived far away, and less with their next-door neighbors who spoke Samoan, say.

The studies of the effects of distance on interaction have found that the intensity of traffic tends to follow so-called “gravity laws”: If two cities at distance d apart have populations A and B , traffic between them is usually proportional to AB/d^α , where α is a constant, typically between 1 and 2. This observation, often ascribed to Zipf [31], actually goes back to the 19th century. (For extensive references, see [4, 20].) The degree of locality is not constant in time. Regular mail in many instances evolved towards increasing locality. On the other hand, in telephony, long distance calling has grown much faster than local (see [5] for some growth rate estimates), although local calling still dominates by far. Although this is still controversial, it appears that traffic on the Internet is also becoming more local, as had been predicted a long time ago.

The “gravity law” rule is not just an artifact of distance-sensitive pricing, since it also applies to regular mail. There is no fully satisfactory explanation for why it holds, but it is valid for a remarkably wide range of interactions, and is often used in transportation and communication facility planning. If we assume that on average the value of being able to communicate with someone at distance d does drop off as $1/d^\alpha$, then, for a uniform distribution of populations in a large disk of radius r , we find (after some calculus that we suppress) total value grows like $r^{3-\alpha}$ for $1 < \alpha < 2$, and like $r^2 \log(r)$ for $\alpha = 2$. Since the population is proportional to r^2 , the area of the disk, we deduce that a population of size n with this type of spacial distribution, the value is proportional to $n \log(n)$ for $\alpha = 2$ and to $n^{2-\alpha/2}$ for $1 < \alpha < 2$. (For $\alpha > 2$, representing extreme locality, value is proportional to n .)

What value of α is most appropriate? Different values apply for different types of traffic. Also, we are interested in traffic value (as well as the option value of being able to engage in a transaction), and that is not the same as intensity of traffic. As one example of the contrast between intensity and value (or, to be more precise, the price paid), in 1910, only about 3% of Bell System calls were toll (and even the vast majority of those were to nearby locations), but they brought in 27% of the revenues. Thus one can argue for many different values of α . However, $\alpha = 2$ is the only nice round value, occurs (at least approximately) very frequently, and produces a result that is simple to state, so we adopt it, and the resulting value of $n \log(n)$ for the valuation of a network.

4 Zipf's Law

Zipf's Law [32] (which again is just a rough empirical rule) says that if we order some large collection by size or popularity, the 2nd one will be about half of the first one in the measure we are using, the 3rd one will be about one third of the first one, and in general the one ranked k -th will be about $1/k$ of the first one. This rule has been found to apply in such widely varying places as the wealth of individuals, the size of cities, and the amount of traffic on web servers. (As Clay Shirky [27] has pointed out, it can be seen in operation today, producing inequality in very egalitarian settings, in the weblog and other spheres, even when there are no economic or political forces impelling this.)

Zipf’s Law is behind phenomena such as “content is not king” [21], and “long tails” [1], which argue that it is the huge volumes of small items or interactions, not the few huge hits, that produce the most value. It even helps explain why both the public and decision makers so often are preoccupied with the “hits,” since, especially when the total number of items available is relatively small, they can dominate. By Zipf’s Law, if value follows popularity, then the value of a collection of n items is proportional to $\log(n)$. If we have a billion items, then the most popular one thousand will contribute a third of the total value, the next million another third, and the remaining almost a billion the remaining third. But if we have online music stores such as Rhapsody or iTunes that carry 735,000 titles while the traditional brick-and-mortar record store carries 20,000 titles, then the additional value of the “long tails” of the download services is only about 33% larger than that of record stores.

Now let us suppose that the incremental value that a person gets from other people being part of a network varies as Zipf’s Law predicts. Let’s further assume that for most people their most valuable communications are with friends and family, and the value of those communications is relatively fixed - it is set by the medium and our makeup as social beings. Then each member of a network with n participants derives value proportional to $\log(n)$, for $n \log(n)$ total value.

5 Information locality

The arguments in the previous section were based on geographic locality. But cyberspace has other forms of locality, for instance locality of interests. Academia has long been structured this way, and we believe that the structure of cyberspace will be at least somewhat similar.

One can tell something about the structure of academic communication by the distribution of research articles of interest to a given researcher. This distribution is described by Bradford’s Law of scattering. In its original form [2] it said that for a typical scientist, there was a core collection A of journals that contained about a third of all articles that were of interest, another collection B of journals that was about β times as large as A that contained another third, and a final collection C of about β^2 journals that contained the remaining third, with β depending on the scientist. Since then librarians have realized that if the search continued on through sets of β^3 , β^4 , β^5 , etc., times as many journals, we would continue producing articles of interest at about equal numbers per set. So in the m most interesting journals, the number of articles of interest tends to scale like $\log(m)$, with the constant of proportionality the inverse of $\log(\beta)$. It seems reasonable to assume that the number of journals scales linearly with the number of researchers, so the total amount of potentially interesting communications among n researchers scales like $n \log(n)$.

There are many possible objections to this reasoning:

- It is obviously impossible to search all journals for all articles of possible interest. With perhaps 200,000 serials, of which about 20,000 are peer reviewed, the task is far beyond what is humanly possible. But we suspect that researchers come reasonably close to ideal results through a combination of research tools like keyword searches, citation

indexes, review articles, and peer recommendations. (Conversely the potential value of such tools can be estimated by comparing a feasible direct search versus what is likely to be out there.)

- We ignored the fact that different researchers have different constants β . We believe that individual differences average out so this does not matter.
- Bradford's Law of scattering is meant to be a rule of thumb, and not a precise quantitative law. True, but then again our $n \log(n)$ estimate is likewise meant as more of a rule of thumb than a precise quantitative law. Furthermore simple power laws are ubiquitous in self-organized systems (including social networks), and $1/n$ scaling factors are particularly common. In Bradford's law the the number of interesting articles in the n -th journal averages $1/n$ of those in the best journal, consistent with Zipf's Law.
- Not all communications of interest are of equal value. This is both obvious and difficult to measure.
- Researchers cannot read an unlimited number of interesting articles. However, articles do not have to be read to be valuable, as their availability allows quick pursuit of various research threads as they are encountered. Moreover, as the universe of articles available expands, the most valuable ones, the ones that do get read, likely grow more valuable.
- This calculation depends on the journals being ordered from most valuable to least. That is not exactly how networks grow. Communication networks do not grow independently of social relations. When people are added, they induce those close to them to join. Therefore in a mature network, those who are most important to people already in the network are likely to also be members. So additional growth is likely to occur at the boundaries of what existing people care about.

Therefore the structure of academia suggests that an $n \log(n)$ scaling rule is at least plausible for communities that self-organize on the basis of interest.

6 Conclusions

Metcalf's Law and Reed's Law both significantly overstate the value of a communication network. In their place we propose another rough rule, that the value of a network of size n grows like $n \log(n)$. This rule, while not meant to be exact, does appear to be consistent with historical behavior of networks with regard to interconnection, and it captures the advantage that general connectivity offers over broadcast networks that deliver content. It also helps explain the failure of the dot-com and telecom ventures, since it implies network effects are not as strong as had been hoped for.

References

1. C. Anderson, "The long tail," *Wired*, Oct. 2004. Available at <http://www.wired.com/wired/archive/12.10/tail.html>.
2. S. C. Bradford, "Sources of information on specific subjects," *Engineering*, vol. 137, 1934, pp. 85–86.

3. F. Cairncross, *The Death of Distance: How the Communications Revolution is Changing Our Lives*, 2nd ed., Harvard Business School Press, 2001. (1st ed. published 1997.)
4. K. G. Coffman and A. M. Odlyzko, "The size and growth rate of the Internet," *First Monday*, vol. 3, no. 10, Oct. 1998, http://firstmonday.org/issues/issue3_10/coffman/. Also available at <http://www.dtc.umn.edu/~odlyzko/doc/recent.html>.
5. K. G. Coffman and A. M. Odlyzko, "Growth of the Internet," *Optical Fiber Telecommunications IV B: Systems and Impairments*, I. P. Kaminow and T. Li, eds., pp. 17–56, Academic Press, 2002. Available at <http://www.dtc.umn.edu/~odlyzko/doc/recent.html>.
6. J. N. Cummings and S. Kiesler, "Collaborative research across disciplinary and institutional boundaries," *Social Studies in Science*, 2005, to appear.
7. N. Economides, "The economics of networks," *International Journal of Industrial Organization*, Oct. 1996. Available at <http://www.stern.nyu.edu/networks/site.html>.
8. G. Gilder, "Metcalfe's Law and legacy," *Forbes ASAP*, Sept. 13, 1993.
9. G. Kaprowski, "FCC Chair wants universal Net access - and he's serious," *Wired News*, Nov. 25, 1996. Available at <http://www.wired.com/news/politics/0,1283,645,00.html>.
10. K. Kilkki and M. Kalervo, "KK-Law for group forming services," presented at XVth International Symposium on Services and Local Access, Edinburgh Conference centre, Heriot-Watt University, Scotland, 21-26 March 2004. For information about the conference, see <http://conferences.iee.org/ISSLS2004/>. Paper available at <http://holistic.nokia.com>.
11. A. B. Krueger, "New York's economic resilience," *New York Times*, Sept. 16, 2004.
12. P. Krugman, "Networks and increasing returns: A cautionary tale," available at <http://web.mit.edu/krugman/www/metcalfe.htm>.
13. R. Matthews, "Six degrees of separation," online material, <http://www.fortunecity.com/emachines/e11/86/sixdeg.html>.
14. A. McAfee and F.-X. Oliveau, "Confronting the limits of networks," *MIT Sloan Management Review*, Summer 2002, pp. 85–87.
15. B. Metcalfe, "Metcalfe's Law: A network becomes more valuable as it reaches more users," *Infoworld*, Oct. 2, 1995. See also the May 6, 1996 column, <http://www.infoworld.com/cgi-bin/displayNew.pl?/metcalfe/bm050696.htm>.
16. B. Metcalfe, "There oughta be a law," *New York Times*, July 15, 1996.
17. M. L. Mueller, Jr., *Universal Service: Competition, Interconnection, and Monopoly in the Making of the American Telephone System*, MIT Press and AEI Press, 1997.
18. M. E. J. Newman, "The structure and function of complex networks," *SIAM Review*, vol. 45, no. 2, 2003, pp. 167–256. Available online at <http://epubs.siam.org/sam-bin/getfile/SIREV/articles/42480.pdf>.
19. J. Nix and D. Gabel, "AT&T's strategic response to competition: Why not preempt entry?," *J. Economic History*, vol. 53, no. 2, 1993, pp. 377–387.
20. A. M. Odlyzko, "The history of communications and its implications for the Internet," 2000 unpublished manuscript, available at <http://www.dtc.umn.edu/~odlyzko/doc/recent.html>.
21. A. M. Odlyzko, "Content is not king," *First Monday*, vol. 6, no. 2, Feb. 2001, http://firstmonday.org/issues/issue6_2/odlyzko/. Also available at <http://www.dtc.umn.edu/~odlyzko/doc/recent.html>.

22. G. Ravid and S. Rafaeli, "Asynchronous discussion groups as Small Worlds and Scale Free Networks," *First Monday*, vol. 9, no. 10, Oct. 2004, http://firstmonday.org/issues/issue9_9/ravid/.
23. D. P. Reed, "Weapon of math destruction: A simple formula explains why the Internet is wreaking havoc on business models," *Context Magazine*, Spring 1999, <http://www.contextmag.com/setFrameRedirect.asp?src=/archives/199903/DigitalStrategy.asp>.
24. J. Rohlfs, "A theory of interdependent demand for a communications service," *Bell Journal of Economics and Management Science*, vol. 5, 1974, pp. 16–37.
25. J. Rohlfs, *Bandwagon Effects in High-Technology Industries*, MIT Press, 2001.
26. P. E. Ross, "5 commandments: The rules engineers live by weren't always set in stone," *IEEE Spectrum*, Dec. 2003, pp. 30–35.
27. C. Shirky, Power laws, weblogs, and inequality, first published on the "Networks, Economics, and Culture" mailing list on Feb. 8, 2003. Available at http://www.shirky.com/writings/powerlaw_weblog.html.
28. H. D. Thoreau, *Walden, or, Life in the Woods*, Ticknor and Fields, 1854. Available online through Project Gutenberg, <http://www.gutenberg.org/dirs/etext95/waldn10.txt>.
29. D. J. Watts, *Six Degrees: The Science of a Connected Age*, Norton, 2003.
30. J. R. Young, "Does e-science work? Research collaborations involving multiple institutions produce fewer results, study finds," *Chronicle of Higher Education*, Dec. 10, 2004.
31. G. K. Zipf, "Some determinants of the circulation of information," *Amer. J. Psychology*, vol. 59, 1946, pp. 401–421.
32. G. K. Zipf, *Human Behavior and the Principle of Least Effort*, Addison-Westey, 1949.