

# Surfing as a Real Option

Rajan M. Lukose

Bernardo A. Huberman

Internet Ecologies Group  
Xerox Palo Alto Research Center  
Palo Alto, CA 94304  
{lukose, huberman}@parc.xerox.com

## Abstract

One of the predominant modes of accessing information in the World Wide Web consists in surfing from one document to another along hypermedia links. We have studied the dynamics of Web surfing within an economics context by considering that there is value in each page that an individual visits, and that clicking on the next page assumes that the information will continue to have some value. Within this formulation an individual will continue to surf until the expected cost of continuing is perceived to be larger than the expected value of the information to be found in the future. This problem is similar to that of a real option in financial economics.

We consider the options viewpoint as a descriptive theory of information foraging by Internet users, and we show how it leads to a kind of “law of surfing” which has been verified experimentally in several large independent datasets. But the real options perspective, which is by now a well-established field in financial economics, may also provide a rich normative model for designing rational Internet agents.

## 1. Introduction

The remarkable growth and rapidly developing complexity of the Internet and the World Wide Web has made clearer the relevance of ideas from economics to the design and understanding of the Web. For example, economics addresses congestion with usage-based pricing and auctions as a general solution to the bandwidth resource allocation problem [1–4]. The agent paradigm in artificial intelligence imagines cyberspace to be populated with rational artificial agents searching for relevant information and participating strategically in virtual auctions on behalf of users [5–7]. Here, the notion of rationality and strategy are naturally consistent with economics and game theory since they generally work from a bottom-up perspective. The models employ utility-maximizing agents, usually bounded-rational, in an environment characterized by finite resources and uncertainty. This is precisely the situation of the artificial agent. It is also the situation of the millions of people who use the Web today. Thus, to the extent that economic models are descriptive, they may be used to understand the patterns

found in the behavior of Web users, in addition to providing prescriptions for the design of artificial agents [8, 9].

Here we consider an economic model of Web surfing in which an agent must decide at each page whether to stop or continue surfing. This option to continue is treated as a “real option”, in analogy with the problem in financial economics of valuing the flexibility a firm has in making investment decisions under uncertainty [10]. In this framework, the option to continue is valued under general assumptions about the utility of a sequence of pages to a user, and a probability distribution for the number of pages a user will visit before stopping is derived. The distribution constitutes a kind of “law of surfing” which describes very well empirical measurements of user behavior [11].

Thus, a real options approach has descriptive validity, but it can also provide a rich normative model for designing rational Internet agents. The approach is by now a very well-developed area in financial economics with a large body of literature exploring various extensions of the basic theme, and the applicability to the Internet suggests new problems.

## 2. The Economics of Surfing


### 2.1 The Model

Consider an agent surfing for information by moving from page to page along hyperlinks of the Web. This clicking action assumes that there is value to be found in each page that an individual visits, and that clicking on the next page will lead to information that continues to have some value. Since the value of the next page is not certain, one can assume that it is stochastically related to the previous one, and different in value for different surfers.

Therefore, as an agent clicks on further links it faces a sequential decision problem: whether to stop surfing at a given page or to go on to the next one. Fig. 1 shows the structure of the problem in a schematic way. An agent starts at node 1 and proceeds until a decision is made to stop. At each node, the agent accrues a value  $x_n$  which is a random variable. In this treatment, the value is taken to be a scalar which summarizes various relevant factors, such as the utility to the user of the information revealed at the page. It might also include costs such as the latency-dependent congestion cost associated with accessing the node. Thus, the value is allowed to be negative. Since the agent wishes

---

To be published in the proceedings of the First International Conference on Information and Computation Economics, ICE-98, by ACM press.



**Fig. 1.** The sequence of nodes at which a surfer must decide whether to stop or pursue the option to continue. The value of node  $n$  is a random variable denoted by  $x_n$ .

to maximize the accrued utility by following an optimal stopping policy, it must decide on the value of the option to continue surfing.

In order to derive the optimal stopping policy, the values of the pages are modelled as a random walk,

$$x_n = x_{n-1} + z_n \quad (1)$$

where the value increments,  $z_n$ , are independent and identically distributed with a density  $p(z)$  having mean  $m$  and variance  $s^2$ .

The optimal stopping policy can be found working backwards from a given node  $N$  by applying a dynamic programming procedure [12]. If  $x_N$  is the value of the final page, then an agent at node  $N - 1$ , having just accrued the value  $x_{N-1}$ , will continue only if the expected value of the final node is greater than zero. Let  $V_n$  denote the value of the  $n$ th node plus the maximum value of the two possibilities, stopping and continuing. If one assumes that users discount the value of future pages at a rate  $\rho$ , the dynamic programming approach then yields the following backwards recursion relations,

$$V_N(x_N) = x_N \quad (2)$$


$$V_n(x_n) = \max \left[ 0, \frac{1}{1 + \rho} E[V_{n+1}(x_n + z_{n+1})] \right]$$

where the expectation is taken with respect to the independent random variable  $z_{n+1}$ .

From this it follows that the value of the option to continue from node  $n$  is a function of the value of the node just revealed and given by,

$$F_n(x_n) = \max \left[ 0, \frac{1}{1 + \rho} E[V_{n+1}(x_n + z_{n+1})] \right] \quad (3)$$

These results can be understood intuitively by considering that under the assumption of a random walk in page values, when a particular page is very valuable it is likely that



**Fig. 2.** The threshold value  $x_n^*$  versus the node index  $n$  for  $N = 25$  nodes and  $m = 0$ ,  $s^2 = 1$ , and  $\rho = .3$ . As long as the value of the  $n$ th page is greater than the corresponding  $x_n^*$ , the option to continue has value and it is optimal to keep surfing.


the next page will be valuable as well, and so pursuing the option to continue will tend to be profitable. On the other hand, if the page is costly, the next one will probably be costly as well, and so, when the value of the current page falls below a certain threshold, it will no longer be optimal to continue. Notice that when the expected value of continuing is greater than zero, it is optimal to do so. Thus, the threshold value, denoted by  $x_n^*$ , is the value of  $x_n$  which solves the equation

$$E[V_{n+1}(x_n^* + z_{n+1})] = 0 \quad (4)$$

which implies that after the value of a node is revealed, it is optimal to continue when  $x_n > x_n^*$  and optimal to stop when  $x_n < x_n^*$ .

These results can be made more concrete by considering the situation in which  $p(z)$  is normally distributed with  $m = 0$ ,  $s^2 = 1$ , and the discount rate  $\rho = .3$ . The optimal policy is then completely specified by the curve of thresholds  $x_n^*$  plotted versus the node index  $n$  as in Fig. 2 where the total number of nodes is set at  $N = 25$ . Notice that the value of the option to continue is a function of the value of the node just revealed, and that the function is different for each node. The functions for several nodes are plotted in Fig. 3 using the same parameter values as Fig. 2. Consider one of these curves. As described above, when the value of the node just revealed  $x_n$  is large, future pages are expected to be valuable as well, so the option to continue is valuable, and  $F(x)$  increases with  $x$ . As  $x$  decreases, the expected value of the future pages decreases, the option to continue becomes less valuable, and  $F(x)$  decreases with decreasing  $x$ . At some value  $x = x^*$ ,  $F(x) = 0$  which defines the threshold for continuing, since stopping is preferable to accepting the negative expected value of continuing.

Thus, the zeroes of  $F_n(x)$  are just the thresholds  $x_n^*$  plotted in Fig. 2. Notice also that the curves  $F_n(x)$  become flatter (monotonically) as  $n$  gets larger, or the agent progresses



**Fig. 3.** The value  $F_n(x)$  of the option to continue as a function of the value  $x$  of the page just revealed. The functions are shown for the indicated nodes, and for the same parameter values as in Fig. 2.

towards the final page. This reflects the fact that the option has less value as the end is near.

## 2.2 Discussion

This model of optimal surfing can be used for designing Internet agents. It may be possible, for example, to design an agent which forages for information in such a way that the value it accrues can be described by a model like the one given here. That is, if its decisions are such that the values of the pages it finds tend to follow, say, a random walk, then the above model provides the optimal policy for stopping.

Indeed, the model presented here can be easily put into the simple framework presented in the standard artificial intelligence textbook [13] by Russell and Norvig. On p. 490 they formalize an INFORMATION-GATHERING-AGENT who “works by repeatedly selecting the observation with the highest information value, until the costs of observing are greater than the benefits.”

Computational tractability may be an issue in designing such agents. In the case of a random walk, the number,  $N$ , of possible nodes (which could be an externally imposed “node budget”), the discount rate  $\rho$ , and the parameters  $m$  and  $s$  are required for the calculation of the optimal policy. Because analytical results are not available and numerical calculation can be computationally intensive, the policies could be calculated off-line and interpolated. Furthermore, several simplifications may be appropriate and useful in designing such agents. For example, if no node budget is imposed and the node horizon is infinite, then the threshold is simply a flat threshold at some level. In many other cases the “curse of dimensionality” which plagues dynamic programming could be overcome by the use of machine learning techniques such as neural networks and reinforcement learning [14].

We stress the fact that this economic design framework is quite flexible. For one can assume an arbitrary drift for

the random walk, which might indicate the quality of the agent. That is, an agent able to find pages which are on average increasingly valuable can be modelled by a process which has a positive drift. Also, other underlying stochastic processes can be used as well [10]. Therefore, as long as the process can be characterized, the kind of model presented here can be used to design a rational agent.

As a concrete example of these considerations, we can consider Menczer’s [15] distributed algorithm for information discovery on the World Wide Web using an adaptive population of intelligent agents making local decisions. A key issue of this work is whether the Web environment has a “semantic topology” that agents can exploit, using concepts such as “relevance autocorrelation,” which essentially encodes the notion that if the current document is relevant, one may expect that some of its neighbors will be as well. Actual Web data corroborates this view [16, 15]. The algorithm explicitly includes relevance estimation (to an original query), and possible latency costs due to network congestion as components of the “energy” (corresponding to our accrued “value”) that determines the survival of an agent. However, the actual mechanism by which the agents try to maximize the accrued “energy” is not very sophisticated and may be sub-optimal. The approach suggested here can provide a more rigorous normative model for the rationality with which these intelligent agents are endowed. In fact, it would amount to a model of decision-making very close to that advocated by financial economists in other contexts, and so is an attractive methodology for the design of such multiagent systems.

## 3. A Law of Surfing

### 3.1 Theory

The model of surfing as a real option that we just presented can also be used to elucidate the surfing patterns of individuals searching for information in the Web. In fact, under certain assumptions, we will show that a “law of surfing” holds. It specifies the probability distribution for the number of links that a surfer follows before stopping. Interestingly, this distribution was found to describe well the data obtained from several large-scale empirical studies of user behavior.

To derive the law, consider the threshold curve of Fig. 2. Clearly, because of the existence of a discount factor the threshold is asymptotically flat with the distance in nodes from the endpoint, since without discounting  $\rho = 0$  and the threshold diverges to  $-\infty$  as  $N^{1/2}$  [17]. Given the fact that the Web is so large and interconnected, and that it seems safe to assume that users discount the future while surfing, it is reasonable that a flat threshold is the appropriate model for actual users. Then, from an initial page value, a user continues surfing, with page values following a random walk as specified by Eq. 1, until the value of a page first

visits the threshold value. Thus, the number of clicks until arrival to the threshold is a random variable. In the limit of true Brownian motion, the first passage times are distributed according to the inverse Gaussian distribution [18].

Thus, the probability density of  $L$ , the length of the sequence of page visitations until stopping (determined by first passage to the fixed threshold), is given by

$$p(L) = \sqrt{\frac{\lambda}{2\pi}} L^{-3/2} e^{-\frac{\lambda}{2\mu^2 L}(L-\mu)^2} \quad (5)$$

whose mean and variance are given in terms of the parameters  $\mu$  and  $\lambda$  according to  $E[L] = \mu$  and  $Var[L] = \mu^3/\lambda$ . We discuss more fully the properties of this distribution below.

### 3.2 Measurements

To test the validity of Eq. 5, we analyzed data collected from a representative sample of America Online (AOL) Web users. For each of 5 days (29 and 30 November and 1, 3, and 5 December 1997), the entire activity of one of AOL's caching proxies was instrumented to record an anonymous but unique user identifier, the time of each URL (uniform resource locator) request, and the requested URL. For comparison with the predicted distribution, a user who starts surfing at a particular site, is said to have stopped surfing after  $L$  links as soon as he or she requests a page from a different Web site. For this analysis, if the user later returned to that site, a new length count  $L$  was started. Requests for embedded media (such as images) were not counted. On 5 December 1997, the 23,692 AOL users in our sample made 3,247,054 page requests from 1,090,168 Web sites. The measured cumulative distribution function (CDF) of the depth  $L$  for that day is shown in Fig. 4. Superimposed is the predicted function from the inverse Gaussian distribution fitted by the method of moments [19]. To test the quality of the fit, we analyzed a quantile-quantile against the fitted distribution. Both techniques, along with a study of the regression residuals, confirmed the strong fit of the empirical data to the theoretical distribution. The fit was significant at the  $P < 0.001$  level and accounted for 99% of the variance. Although the average number of pages surfed at a site was almost three, users typically requested only one page. Other AOL data from different dates showed the same strength of fit to the inverse Gaussian with nearly the same parameters.

We also examined the navigational patterns of the Web user population at Georgia Institute of Technology for a period of 3 weeks, starting on 3 August 1994. The data were collected from an instrumented version of the National Center for Supercomputing Applications' Xmosaic that was deployed across the students, faculty, and staff of the College of Computing [20]. One hundred and seven users (67% of


those invited) chose to participate in the experiment. The instrumentation of Xmosaic recorded all user interface events. Of all the collected events, 73% were navigational, resulting in 31,134 page requests. As with the AOL experiment, the surfing depth of users was calculated across all visits to each site for the duration of the study. For the combined data, the mean number of clicks was 8.32 and the variance was 2.77. Comparison of the quantile-quantile, the CDF, and a regression analysis of the observed data against an inverse Gaussian distribution of same mean and variance confirmed the ability of the law of surfing to fit the data ( $R^2 = 0.95$ ,  $P < 0.001$ ). Hence, the model was able to fit surfing behavior with data sets from diverse communities of users, several years apart, who used different browsers and connection speeds.

The previous data validated the law of surfing for a population of users who had no constraints on the Web sites they visited. We also considered the case of surfing within a single large Web site, which is important from the point of view of site design. The site used was the Xerox Corporation's external Web site ([www.xerox.com](http://www.xerox.com)). During the period 23 to 30 August 1997, the Xerox site consisted of 8432 HTML documents and received an average of 165,922 requests per day. The paths of individual users were reconstructed by a set of heuristics that used unique identifiers ("cookies"), when present, or otherwise used the topology of the site along with other information to disambiguate users behind proxies. Automatic programs that request the entire contents of the site ("spiders") were removed from the analysis. Additionally, a stack-based history mechanism was used to infer pages cached either by the client or by intermediary caches. This resulted in a data set consisting of the full path of users and the number of clicks performed at the Xerox Web site. Fig. 5 shows the CDF plot of the Xerox Web site for 26 August 1997 against the fitted inverse Gaussian defined by Eq. 5. The mean number of clicks was 3.86, with a variance of 6.08 and a maximum of 95 clicks. As with the client path distributions, both the quantile-quantile and the CDF plots of the site data showed a strong fit to Eq. 5. Moreover, these results were very consistent across all the days in the study.


For further confirmation of the model, we considered the simplest alternative hypothesis, in which a user at each page conducts an independent Bernoulli trial to make a stopping decision. This model leads to a geometric distribution of depths, which was found to be a poor fit to the data.

### 3.3 Discussion

An important property of the function  $p(L)$  is its long tail. This means that concentrating on the average value of the depth of surfing, given by  $L$ , obscures a great deal of information. In the case of the Xerox data just analyzed, for example, while the mean length was 4.2, the mode of



**Fig. 4.** The points are the empirical cumulative distribution function of the length  $L$ , the number of pages visited before stopping, for AOL users. The data were collected on 5 December 1997 from a representative sample of 23,692 AOL users who made 3,247,054 clicks. The curve is the CDF of a fitted inverse Gaussian distribution with a mean of  $\mu = 2.98$  and  $\lambda = 6.24$ .



**Fig. 5.** The points are the empirical cumulative distribution function of the length  $L$ , the number of pages visited before stopping, for measurements made on August 26, 1997 at the Xerox Corporation's Web site (<http://www.xerox.com>). The curve is the cumulative distribution function of the maximum likelihood inverse Gaussian. The mean length for this data set was 4.2 with a variance of 8.9, and a maximum of 110.

the distribution was located near one, which implies that typically most users don't go beyond one click.

Another interesting property of the law of surfing can be obtained by taking logarithms of both sides of Eq. 5. One obtains


$$\log p(L) = -\frac{3}{2} \log L - \frac{\lambda(L - \mu)^2}{2\mu^2 L} + \log \sqrt{\frac{\lambda}{2\pi}} \quad (6)$$

When  $E[L] < \text{Var}[L]$ ,  $\frac{\lambda}{\mu^2} < 1$ , which implies that there is a range of values of  $L$  such that ,

$$p(L) \propto L^{-1.5} \quad (7)$$

i.e. it scales like a Pareto distribution. This scaling behavior can be tested by noticing that on a log-log plot, the

probability density is well-approximated for large variance by a straight line with slope  $-3/2$  and that for larger values of  $L$ , the second term in Eq. 6 begins to dominate and provides a downward sloping correction. Fig. 6 shows that the approximation holds over a range of depths.



**Fig. 6.** The frequency distribution of surfing clicks on log-log scales. Data collected from the Georgia Institute of Technology, August 1994.

This scaling relation is similar to the well-known Zipf's Law [21] found in an extremely wide variety of data, which says that up to a constant, the approximate frequency of a word in a given body of text (or database, etc.) is related to its *rank* in frequency by an inverse power law. In order to make the analogy apparent in Eq. 7, rank is taken to be the depth in pages from a user's current position in a hyper-linked environment. It thus gives a description, appropriate for each user, of the likelihood that the user will arrive at a particular page which is some link distance  $L$  away. Thus, it may have implications for the valuation of advertising space on a Web page which is a given link distance away from typical entry points into Web sites.

Random variables distributed according to the inverse Gaussian also have the interesting property that they are *stable* [22]. That is, the sum of a set of independent random variables distributed according to the inverse Gaussian is itself an inverse Gaussian. This can be seen easily by considering the decomposition of a random walk to a fixed threshold into first passage time problems to intermediate thresholds. This relatively special property adds to the robustness of the distribution when comparing with measurements, because it allows the threshold to be more loosely defined (as long as it is defined fairly consistently over the samples in the dataset). For example, the Georgia Tech data in Fig. 6 is made up of length observations, one for each user, over a *length of time*. As long as users on average visit the same number of sites during that time, each observation is then the sum of inverse Gaussian random variables, and their stability ensures the validity of this kind of measurement process.

Finally, we note that the inverse Gaussian distribution is, strictly speaking, the first passage time distribution for

Brownian motion [23, 24]. Here we have a random walk which is not Brownian, but rather is discrete in “time” (links) but with continuous increments. This form of the stopping problem was studied by Wald [25] who showed that the discrete form converges to the result for Brownian motion. Furthermore, the independent increments  $z_n$  in Eq. 6 need only have finite variance for the first passage times to converge to the inverse Gaussian [26]. In fact, the two parameters of the inverse Gaussian can be written in terms of the three parameters that define the stopping problem:  $\mu = a/m$  and  $\lambda = a^2/s^2$  where  $a$  is the threshold and  $m$  and  $s$  are the mean and standard deviation of  $z_n$  as was previously defined. It is a simple consequence of these relations and the definition of the variance that

$$\sqrt{\frac{\text{Var}[L]}{\mathbb{E}[L]}} = \frac{s}{m} \quad (8)$$

which gives, independent of the threshold, some information about the “noisiness” of the underlying random walk. This dimensionless ratio, yields an answer  $\approx 1$  for most of the datasets, indicating that the underlying walk is quite noisy.

#### 4. Conclusion

In this paper we presented a real options approach to the design of agents that could forage the World Wide Web in search of information for their sponsors. In this context, the option to continue was treated as a “real option”, in analogy with the problem in financial economics of valuing the flexibility a firm has in making investment decisions under uncertainty. We then obtained expressions for the optimal stopping policy which can be incorporated into the agent’s design and also derived a law of surfing for human users of the Web. Remarkably, several independent large data sets of user surfing patterns corroborate this law [11]. The most important modelling simplification we have made is to abstract away from the topology of real Web spaces. A real Web surfer can be modelled in greater detail by taking explicit account of the graph and tree-like structure of the available pages. Our “one-dimensional” model abstracts from the topology and considers only the path actually taken. It therefore neglects the possibility of backtracking along the nodes in order to pursue another path. Such a situation can be modelled, in principle, by the methods used here, but is considerably more complicated. Constructing more detailed models may be worthwhile, however, especially since they can be tested by comparison with measurements such as those described by Cunha, et. al. in [27], which show in nice graphical form the actual traces of users’ paths, taking account of backtracks. We note however, that in spite of these simplifications the model appears to provide a good description of empirical surfing patterns.

We have made a distinction between the normative and descriptive implications of our treatment of surfing as a real option. On the one hand, since the construction of artificial Internet agents is now an important topic of research, the formal option model seems an attractive approach for designing agents that behave in a manner consistent with basic economic principles. The approach may thus have interesting normative implications for the design of agents, a great deal of which is now going on.

On the descriptive side, modelling real surfing as a real option leads to simple predictions which can actually be tested experimentally. The success of the “law of surfing” in describing actual user data is encouraging in two senses. First, it helps to fix to what degree formal economic models are actually descriptive of real users of the Web. Second, given the descriptive accuracy, it allows the construction of probabilistic algorithms which can be used to predict the activity of users [11].

Finally, these models can be used to build better browsers and interfaces, as well as provide ways to price advertising space on Web pages.

**Acknowledgments.** We thank J. Pitkow and P. Pirolli for providing the data as well as for useful discussions. R. Lukose was supported by a Citibank fellowship at Stanford University.

#### References

- [1] Jeffrey K. MacKie-Mason and Hal R. Varian. Some economics of the internet. In *Proc. of 10th Michigan Public Utility Conf.* Michigan University, November 1993.
- [2] Ron Cocchi, Scott Shenker, Deborah Estrin, and Lixia Zhang. Pricing in computer networks: Motivation, formulation, and example. *IEEE Trans. on Networking*, 1(6):614–627, 1993.
- [3] D. O. Stahl A. Gupta and A. B. Whinston. Pricing of services in the internet. Technical report, University of Texas at Austin, 1995.
- [4] C. A. Waldspurger, T. Hogg, B. A. Huberman, J. O. Kephart, and S. Stornetta. Spawn: A distributed computational economy. *IEEE Transactions on Software Engineering*, 18(2):103–117, February 1992.
- [5] Oren Etzioni and Daniel Weld. A softbot-based interface to the internet. *Communications of the ACM*, 37:72–76, July 1994.
- [6] Pattie Maes. Agents that reduce work and information overload. *Communications of the ACM*, 37:31–40, July 1994.
- [7] J. S. Rosenschein and G. Zlotkin. *Rules of Encounter*. MIT Press, Cambridge, MA, 1994.
- [8] Bernardo A. Huberman. Computation as economics. *Journal of Economic Dynamics and Control*, 22(8-9):1169–1186, July 1998.

- [9] Bernardo A. Huberman and Tad Hogg. Distributed computation as an economic system. *Journal of Economic Perspectives*, 9:141–152, 1994.
- [10] Avinash K. Dixit and Robert S. Pindyck. *Investment Under Uncertainty*. Princeton University Press, 1994.
- [11] Bernardo A. Huberman, Peter L. T. Pirolli, James E. Pitkow, and Rajan M. Lukose. Strong regularities in world wide web surfing. *Science*, 280(5360):95–97, April 3 1998.
- [12] Dimitri P. Bertsekas. *Dynamic Programming: Deterministic and Stochastic Models*. Prentice-Hall, 1987.
- [13] Stuart Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 1995.
- [14] Dimitri P. Bertsekas and John Tsitsiklis. *Neuro-Dynamic Programming*. Athena Scientific, 1996.
- [15] F. Menczer. Arachnid: Adaptive retrieval agents choosing heuristic neighborhoods for information discovery. In D. Fisher, editor, *Machine Learning: Proceedings of the 14th International Conference (ICML97)*. Morgan Kaufman, 1997.
- [16] R. Larson. Bibliometrics of the world wide web: An exploratory analysis of the intellectual structure of cyberspace. In *Proceedings 1996 Annual ASIS Meeting*.
- [17] L. A. Shepp. Explicit solutions to some problems of optimal stopping. *The Annals of Mathematical Statistics*, 40(3):993–1010, 1969.
- [18] V. Seshadri. *The Inverse Gaussian Distribution: A Case Study in Exponential Families*. Clarendon Press, 1993.
- [19] Raj S. Chhikara and J. Leroy Folks. *The Inverse Gaussian Distribution: Theory, Methodology, and Applications*. Marcel Dekker, 1989.
- [20] L. Catledge and J. Pitkow. *Comput. Networks ISDN Syst.*, 27:1065, 1995.
- [21] G. K. Zipf. *Human Behavior and the Principle of Least Effort*. Addison-Wesley, 1949.
- [22] William Feller. *An Introduction to Probability Theory and Its Applications*, volume 2. John Wiley and Sons, 1971.
- [23] Erwin Schrodinger. *Phys. Ze.*, 16:289–295, 1915.
- [24] M. V. Smoluchowski. *Phys. Ze.*, 16:318–321, 1915.
- [25] A. Wald. On cumulative sums of random variables. *Ann. Math. Statist.*, 15:283–296, 1944.
- [26] N. L. Johnson and S. Kotz. *Distributions in Statistics: Continuous Distributions 1*. Houghton Mifflin, Boston, 1970.
- [27] Carlos R. Cunha and Azer Bestavros Mark E. Crovella. Characteristics of www client-based traces. Technical report, Computer Science Department, Boston University, 1995.