Social influence and the evolution of long-tailed distributions

R. Alexander Bentley[1] Paul Ormerod[2] and Michael Batty[3]

Summary

Studies of collective human behaviour in the social sciences, often grounded in details of actions by individuals, have much to offer 'social' models from the physical sciences concerning elegant statistical regularities. Drawing on behavioural studies of social influence, we present a parsimonious, stochastic model, which generates an entire family of real-world right-skew socio-economic distributions, including exponential, winner-take-all, power law tails of varying exponents and power laws across the whole data. The widely used Albert-Barabási model of preferential attachment is simply a special case of this much more general model. In addition, the model produces the continuous turnover observed empirically within those distributions. Previous preferential attachment models have generated specific distributions with turnover using arbitrary add-on rules, but turnover is an inherent feature of our model. The model also replicates an intriguing new relationship, observed across a range of empirical studies, between the power law exponent and the proportion of data represented.

Keywords: human dynamics, scaling, culture evolution

1 Introduction

Since Pareto, the right-skew nature of income distribution has been known, while similar skewness in the frequencies of words, scientific papers, and city sizes has been recognised for decades [1 - 5]. In statistical physics, a recent explosion of interest in such distributions for social phenomena includes internet links [6, 7], author citations [8], sexual partners [9], and firm sizes and their extinctions [10, 11] amongst many others.

With socio-economic phenomena, the detailed debate over the exact form of these distributions -- for example, power laws versus similar fat-tailed functions such as the stretched exponential [12, 13] -- often involves the characterisation of the distribution at a point in time, and neglects the importance of dynamics and the underlying

[1] Corresponding Author, Department of Anthropology, Durham University, 43 Old Elvet Road, Durham, DH1 3HN, UK, Email: r.a.bentley@durham.ac.uk

[2] Volterra Consulting, Sheen Elms, 135c Sheen Lane, London SW14 8AE, UK, Email: pormerod@volterra.co.uk

[3] Centre for Advanced Spatial Analysis, University College London, 1-19 Torrington Place, London WC1E 6BT, UK. Email: m.batty@ucl.ac.uk

behaviour [12, 14] which gives rise to changes over time within any given distribution.

2. Overview of the model

Simon [3] argued that right-skew distributions were so widespread that their key similarity was likely to be 'in the underlying probability mechanisms' that led to their generation. This is clearly the case but, as noted in the social sciences for over a century [12], it is inherently a description of macro phenomena, without an explanation for the individual behaviour that gives rise to emergent properties.

We thus propose a model based upon individual agents who are boundedly rational and are influenced by the behaviour of other agents in terms of their decision-making. In other words, the agents act with social purpose, which is fundamentally different from physical or biological phenomena where the agents (or particles) are incapable of intent. The model provides four advances on previous models:

(a) It can generate a wide range of the right-skew distributions observed in cultural, economic and social situations from different combinations of its two parameters.

(b) The widely used Albert-Barabási (B-A) model [7] of preferential attachment is simply a special case of this much more general model.

(c) In terms of power law fits, there are two essential statistics, the exponent α and the fraction f of the total observations over which the power law is believed to hold. The model can replicate both observed exponents α and the fraction f from real-world observations [1, 2].

(d) Many real-world right-skew distributions exhibit constant turnover in the rankings of their constituents even if their functional form is time-invariant [14, 15]. Unlike the B-A model [7], our model is capable of generating such turnover without recourse to self-fulfilling rules such as 'aging' or variable 'fitness' of the individual elements [16].

We stress that our model is not a network model, and therefore *not* a modified B-A model. Whenever a network model is applied to social dynamics, two analogies are possible:

- The nodes represent the agents (e.g. academic papers), and the links are the relationships between them (citations)
- The links represent agents (e.g. consumers), and the nodes represent the choices the agents make (e.g., purchased brands)

Whichever the representation, new nodes in B-A models link with probability $p(k)$ to existing nodes, where $p(k)$ is a function of the degree k of the existing node (and possibly time t as well, in aging models). Agents and nodes are inseparable, connected by the network. Our evolutionary model is quite different, because the agents and their ideas are *separable*: each new agent adopts its idea either by copying another agent, or by inventing an idea of its own. This provides substantial advantages over the either network analogy.

3. The social influence model in more detail

Consider a model populated initially by N agents located in some abstract space such as a sequence of index numbers. Depending on the phenomenon, each location is an abstract representation; it could refer to the city where a firm chooses to locate itself, but it could equally well refer to the product a consumer chooses, or the idea or fashion that a person follows.

We define the *size* of a location as the number of agents at that location.

The model proceeds in a series of steps. In each step, n new agents enter the model, where the number n is fixed as a parameter in each solution of the model. With probability $1-\mu$, an agent copies the choice of location from that of an existing agent within the previous m time steps, or else with probability μ, the agent innovates by choosing a unique new location at random. In other words, the agent either copies an existing agent from the last m steps, or chooses a new location.

Here we restrict our exploration to two key parameters of the model, m and μ, by choosing convenient values for N and n. The 'memory' parameter m determines the number of steps of the previous decisions of other agents over which an agent looks when making its decision. The 'innovation' parameter μ determines the fraction of the agents who decide to take a completely new decision rather than replicating one of the decisions made by other agents.

The 'memory' parameter is one which we believe is an innovation in the literature, but it is one which is clearly realistic. A teenager contemplating downloading popular music, for example, is likely to take into account only those decisions made by other agents in the very recent past. On contrast, a firm thinking about locating in a city may consider the location decisions made by other firms over a period of years in the past.

The model can be contrasted with that of Dorogovtsev et.al. [17], which introduces an extra parameter into the preferential attachment model. This extra parameter is the exponent to which the degree of each node is powered, to generate the scores that new nodes use to link to existing nodes with probability proportional to these scores. For the value of this parameter equal to 1, the model generates the pure preferential attachment model (i.e. power-law degree distribution). For the value of this parameter equal to 0, the model generates an exponential degree distribution. For the value of this parameter tending to infinity, the distribution is winner-take-all.

We noted above some fundamental differences between our model and that of preferential attachment. These also apply to the latter with the extra parameter. In addition, we note that the modified preferential attachment model is a form of *selection* model, because longer-lasting choices are generally those with greater 'fitness' ('attractiveness'), and that variations in the distribution of popularity come from differences in intrinsic fitness.

This is quite different to our own model demonstrated here below, which achieves a range of popularity (cf. degree) distributions, range life-spans, and continual turnover without assuming *any* differences in intrinsic attractiveness/fitness of the choices. These effects arise purely through stochastic means, rather than any differences in 'fitness' assumed *a priori*.

4. The variety of distributional outcomes

A specific version of the model, with $m = 1$ (i.e., memory only of the immediately preceding step), is known in population genetics and physics [18, 19]. For the special case of $n = N$ and $m = 1$, analytical solutions demonstrate a power-law distribution [18] for $N\mu$ equal to or slightly greater than 1. For $m = 1$ and $N\mu << 1$, this gradually converges on a winner-take-all distribution as $N\mu$ approaches zero.

The case where m = *all* is a further special category of the model, where extinction or obsolescence does not occur. In this case, we can achieve different power law slopes by varying n and μ. Figure 1 shows, for example, that we can match the B-A preferential attachment model [7], obtaining a power law exponent $\alpha \sim 3$ over the entire distribution, by using m = all with $N = 1$, $n = 10$, $t = 20{,}000$, and $\mu = 0.6$.

FIGURE 1 ABOUT HERE

For socio-cultural phenomena, however, we expect memory to be limited, and thus m in general to take values below the special case of 'all'. So while we define the model to allow m to take any value between 1 and *all*, we explore here a limited range, from $m = 1$ to $m = 100$ time steps of limited memory. The combined effect of varying m along with varying the innovation parameter μ generates both a wide range of right-skew distributional forms and turnover of rankings of locations within those distributions.

Considerable anthropological and socio-economic evidence exists [20 - 24] on the plausible values for μ being no greater than 0.1.

Figure 2a plots typical solutions of the model using acceptable values of μ, while varying m (holding N = 1000, n = 100 and showing the results at time step 1000). Aside from the selected results shown in this figure, the model produces additional results ranging from a winner-take-all outcome, to a power law over the entire distribution (exponent $\alpha \sim 1.5$), to a power law fitted to the tail of varying exponent.

Figure 2b illustrates how the model parameters can be selected so that the results match real-world right-skew distributions, such as religions, website subscriptions, word use, names, and author citations.

FIGURE 2 ABOUT HERE

5. Regularity in the long tail

Table 1 lists power law tail exponents α for various recently collated social data sets [1,2] along with the fraction f (= n_{tail}/n) of total observations in the tail. A striking, and previously unreported, feature of these estimates is the relationship between α and f, where these data reveal a clear inverse correlation. The smaller the fraction f of the distribution best-fit to a power law tail: the larger the exponent α of that tail. The least-squares fit is $\alpha = 1.54 f^{-0.156}$ (r^2 = 0.952).

Figure 3 plots this relationship in the empirical data along with the least squares fit using the model, as solved 100 times, for each of μ = 0.05, 0.06 and 0.07, with m = 30 in each case (and N = 1000, n = 100, t = 1,000). The results show $\alpha = 1.56 f^{-0.155}$ (r^2 = 0.975), very similar to the data-based relationship.

FIGURE 3 ABOUT HERE

We conjecture that the continuous relationship observed in Figure 3 suggests that socio-economic power law distributions may form a continuum resulting from a generalised process with limited memory. In contrast to the special m = *all* case (Figure 1), when model runs with limited memory yield a power law over the entire distribution (f = 1), it is only with exponent α close to 1.5 (Figure 3).

6. The distribution of turnover

The model also produces continual turnover through time for any given distribution, as demonstrated by the distributions of life-spans within ranked lists (life-span being the number of time steps a location spent on the list) as in Figure 4a. This resembles the life-spans of real world social and economic fat-tail distributions in Figure 4b. The memory parameter m again expands the power of the model. Although turnover has already been demonstrated [18] for the special case $m = 1$, different values of m are needed to account for empirically observed turnover.

FIGURE 4 ABOUT HERE

Figure 5 shows distributions generated of life-spans in the top 5 (i.e., number of time steps location spent in the top 5 most popular) in tests with increasing memory m and invention rate μ (the plots in Figure 5a - 5d show progressively larger values of *m, while* each plot shows five different orders of magnitudes for μ). The effect of increasing memory m is to reduce the effect of innovation μ such that when m = all (Figure 5d), the distributions are all virtually the same form (albeit with different finite limits to the cumulative number of locations).

FIGURE 5 ABOUT HERE

7. Discussion and conclusions

We have proposed a substantial alternative to the thousands of modified Barabási-Albert (B-A) models in the last ten years. Our model derives from evolutionary models that predate B-A by many decades, and differs substantially from it and all of its subsequent modifications. With this fresh start rooted in evolutionary theory, we achieve several effects that the B-A models do not, in matching a range of real-world popularity distributions, their turnover through time, and the empirical match to the α-f relationship in Figure 3.

These real-world effects have never been matched by a single, simple model. As a result, our behavioural model has much insight to offer about the dynamics of the socio-cultural spread of ideas. The model we have presented can generate not only a wide range of long-tailed distributions but a constant turnover of the constituent agents within any given overall rank-size distribution. It is also able to replicate a newly-identified empirical relationship whereby the power law exponent increases as the proportion of data in the tail falls.

The model is quite general, despite using only two key parameters. Varying the parameter values can yield a range of distributions, such as a power law over the whole sample, a power law only in the tail, and a winner-take-all outcome. Since the parameter m represents memory and μ represents innovation in modelled decision-making, the real-world relationship between α and f in Figure 3 may result from a variation in related parameters among the different contexts of human decision-making.

This combination of results makes this model unique among the many alternatives that can produce power laws. The most commonly proposed processes such as preferential attachment, proportionate effect based on Gibrat's principle, the 'Matthew effect' and the Yule process [1,2,3,14,25], produce power laws from the positive feedback introduced by interactions between individual agents. But these 'rich get richer' models have not been able to account for flux in the constituents of the ranked distribution [26], either when growth is one of strict preferential attachment or even when growth is proportionate to a stochastic rate independent of size [27].

Even though ''dynamical problems lie at the forefront" of network science [16], in most network models, existing connections affect future connections such that change does not occur naturally, but only with imposed modifications. The reason for this is the restrictiveness of the network analogy for non-network phenomena. Let's say we represented our agents as network nodes, so the links would be instances of an agent copying the behavior of another agent. The degree distribution would thus represent how often agents were copied by other agents. This works only in the special instance where the idea and the agent are the same, such as with academic papers and their authors citing each other. But in the majority of real-world cases, the idea and the agent are *separable* - when we buy a book or use a trendy word, it never refers to a single unique user.

In the other network representation, the links represent the choosers, and the nodes represent the choices. In B-A models with aging [*e.g.* 28], the probability of the choice itself (node) diminishes with its age. This is wholly inappropriate for our own interest in the adoption of ideas, where an idea does not go extinct because it is old, but because no one uses it anymore. In many cases the opposite is true; the oldest English words, for example, are actually the ones most frequently spoken today [29].

It is much more appropriate and effective to model the limited memory of the choosers, as our model does, rather than the aging of the choices. A recent network model [30] does, in fact, briefly explore limited memory, but again equating nodes with ideas may work for the specific application to authors and their citations of other authors, but does not work at all for the much more general cases that we have explored here, where the ideas and their choosers are entirely separable.

If we are really adamant about applying a network model for behavioural choice, a network rewiring representation can be equivalent to the m =1 case for our model [17]. In these network-rewiring models, however, the links represent only the ideas currently held, not those previously held. In order to represent our model with memory beyond the last time step (m > 1), a network-rewiring model would have to represent all the previously-existing networks from the past m time steps. This would be unnecessarily complicated and a great challenge to simulate and/or solve analytically , and consequently has not been done. Our evolutionary model represents all this much more simply: each new agent adopts its idea either by copying another agent, or by inventing an idea of its own. We thus present this model as a new start, with more long-term potential for development than the previous generation of B-A models. By using an evolutionary model rather than a network model we are able to replicate real-world effects that no other model has shown.

Social scientists have been critical of modelling social and economic data by mapping onto known phenomena in physics without considering realistic behavioural motivations of the agents [12, 31, 32]. As a step in this direction, our model captures two fundamental motivations, the imitation of others and novelty in invention.

Compared to similar, less flexible versions of this model [15], a crucial new variable appears to be the memory m, which reflects different time frames to which agents will refer in different contexts.

Generating a range of long-tailed distributions with dynamic turnover, these features distinguish this model from the standard socio-economic science model of individual rational behaviour where social influence is the exception to the rule (as in, for example, 'irrational' stock market bubbles or real estate crises). With its unrealistic psychological assumptions [33] and inconsistencies with experimental results [34], the standard model suffers from a neglect of social influence, even in its modern form which permits, for example, asymmetry in the amount of information possessed by different agents [35, 36], the cost of gathering information [37], and imperfections in gathering and processing information [38].

Social influence is arguably ubiquitous among the human species [39]. In fact, rather than the agent's cost-benefit analysis that has served as a null hypothesis for rationality for over a century, an alternative is that each agent uses (consciously or not) the decisions of others as a basis for his or her own decisions.

The social-influence model we have presented allows choices among multiple possible alternatives, which rise and fall in relative popularity over time, rather than binary, 'either-or' decisions. This is truly reflective of human interactions such as the choice of a popular name for a child, the citation of an academic paper, or movement to a city where others have chosen to live. Indeed, these phenomena are inherently defined by the past decisions of others, without which there would be no cities, familiar names, or popular culture.

Acknowledgments

Batty was partially supported by the EPSRC Spatially Embedded Complex Systems Engineering Consortium (EP/C513703/1). Amy Heineike of Volterra provided programming support.

References

1. Newman MEJ (2005) Power laws, Pareto distributions and Zipf's law. *Contemporary Physics* 46: 323-351.
2. Clauset A, Shalizi CR, Newman MEJ (2007) Power-law distributions in empirical data. arXiv 07061062v1.
3. Simon HA (1955), On a class of skew distribution functions. *Biometrika* 42, 425-440
4. G.K. Zipf (1949) *Human Behavior and the Principle of Least Effort*. (Addison Wesley, Cambridge, MA)
5. Price DJD (1965) Networks of scientific papers. *Science* 149: 510-512.
6. Huberman BA, Adamic AL (1999) Growth dynamics of the World-Wide Web. *Nature* 401: 131.
7. Barabási AL, Albert R (1999) Emergence of scaling in random networks. *Science* 286: 509-512.
8. Redner, S (1998) How popular is your paper? An empirical study of the citation distribution. *European Physical Journal B* 4: 131-134.
9. Liljeros, F, Edling, CR, Amaral LAN, Stanley HE, Aberg Y (2001) The web of human sexual contacts. *Nature* 411: 907-908.
10. Axtell RL (2001) Zipf distribution of U.S. firm sizes. *Science* 293: 1818-1820.
11. Ormerod P (2006) *Why Most Things Fail: Evolution, Extinction and Economics*. (Pantheon Books, New York).
12. Borgatti SP, Mehra A, Brass DJ (2009) Network analysis in the social sciences. *Science* 323: 892-895.
13. Laherrère J, Sornette D (1998) Stretched exponential distributions in nature and economy: 'fat tails' with characteristic scales. *European Physical Journal B* 2: 525-539.
14. Batty M (2006) Rank clocks. *Nature*. 44: 592-596.
15. Bentley RA, Lipo CP, Herzog HA, Hahn MW (2007) Regular rates of popular culture change reflect random copying. *Evolution and Human Behaviour* 28: 151-158.
16. Newman MEJ, Barabási A-L, Watts DJ (2005) *The Structure and Dynamics of Networks*. (Princeton University Press, Princeton, NJ).
17. Dorogovtsev SN, Mendes JFF, Samukhin AN (2000) Structure of growing networks with preferential linking, *Phys. Review Lett.,* 85, 4633-4636

18. Evans TS (2007) Exact solutions for network rewiring models. *European Physical Journal B* 56: 65-69.
19. Gillespie JH (2004) *Population Genetics*. (Johns Hopkins University Press, Baltimore, MD)
20. Eerkens JW (2000) Practice makes within 5% of perfect. *Current Anthropology* 41: 663-668.
21. Diederen P, van Meijl H, Wolters A, Bijak K (2003) Innovation adoption in agriculture. *Cahiers d'Économie et Sociologie Rurales* 67: 30-50.
22. Srinivasan V, Mason CH (1986) Nonlinear least squares estimation of new product diffusion models. *Marketing Science* 5: 169-178.
23. Larsen ON (1961) Innovators and early adopters of television. *Sociological Inquiry* 32: 16-33.
24. Rogers EM (1962) *Diffusion of Innovations*. (Free Press, New York)
25. Yule GU (1925) A mathematical theory of evolution based on the conclusions of Dr. J. C. Willis. *Phil. Trans. R. Soc. London B* 213: 21-87.
26. Batty M (2008) The size, scale, and shape of cities. *Science* 319: 769-771.
27. Gibrat R (1931) *Les Inégalités Économiques*. (Librarie du Recueil Sirey, Paris.
28. Dorogovtsev, S.N. and J.F.F. Mendes (2000). Evolution of networks with aging of sites. *Physical Review E* 62: 1842.
29. Pagel, M, et al. (2007). Frequency of word-use predicts rates of lexical evolution throughout Indo-European history. *Nature* 449: 717.
30. Hajra, B. & P. Sen (2006). Modelling aging characteristics in citation networks *Physica A* 368: 575–582.
31. Gallegati M, Keen S, Lux T, Ormerod P (2006) Worrying trends in econophysics. *Physica A* 370: 1-6.
32. Bentley RA, Shennan SJ (2005) Random copying and cultural evolution. *Science* 309: 877-879.
33. Kahneman D (2003) Maps of bounded rationality: psychology for behavioral economics. *American Economic Review* 93: 1449-1475.
34. Smith VL (2003) Constructivist and ecological rationality in economics. *American Economic Review* 93: 465-508.
35. Akerlof GA (1970) The market for 'lemons': Quality uncertainty and the market mechanism. *Quarterly Journal of Economics* 84: 488-500.
36. Stiglitz JE (2002) Information and the change in the paradigm in economics. *American Economic Review* 92: 460-501.
37. Stigler GJ (1961) The economics of information. *Journal of Political Economy* 69: 213-225.
38. Simon HA(1955) A behavioral model of rational choice. *Quarterly J. Economics* 59: 99-118.
39. Dunbar RIM, Shultz S (2007) Evolution in the social brain. *Science* 317: 1344-1347.

Figure 1. The power law generated from the preferential attachment version of the model. As the probability distribution for a typical model run using $N = 1$, $n = 10$, $t = 20{,}000$, $\mu = 0.6$, and $m = all$ (where the generated sizes are logarithmically binned). The exponent for the power law is –2.9 ($r^2 = 0.996$), matching that reported (also by least-squares regression) for preferential attachment models [7].

Figure 2. Log-log plots of rank and size, **(a)** for typical model solutions with $N = 1000$, $n = 100$, $t = 1000$ and: $\mu = 0.01$, $m = 1$ (black); $\mu = 0.01$, $m = 100$ (red); $\mu = 0.08$, $m = 100$ (white); $\mu = 0.0001$, $m = 2$ (green). **(b)** for real-world top 100 ranked lists (dots) versus model results (lines). Top 100 lists include [41]: male baby name frequency (per million) in the 1990 US census (blue), RSS feed subscriptions 2001-2008 (orange), English words (red), cited economists 1993-2003 (purple), and religions in thousands of adherents (green). With $N = 1000$, the model fits were made with $\mu = 0.001$, $m = 50$, $n = 200$, $t = 4000$ for names, $\mu = 0.00002$, $m = 6$, $n = 2500$, $t = 10000$ for RSS feeds, $\mu = 0.00025$, $m = 85$, $n = 100$, $t = 1100$ for cited economists, $\mu = 0.004$, $m = 4$, $n = 450$, $t = 8000$ for words, and $\mu = 0.0007$, $m = 2$, $n = 100$, $t = 4000$ for religions.

Figure 3. The power law tail exponent α versus the fraction f of total observations represented by the tail. The dots show power law tails calculated for various real-world socio-cultural data sets (see Table 1 for values and errors), whose relationship (dashed grey curve) can be approximated by $\alpha = 1.54f^{-0.156}$ ($r^2 = 0.952$ except for the outlier -- the open circle -- from email lists). The thin red curve shows the least squares fit from 300 runs of our theoretical model which gives $\alpha = 1.56/f^{0.155}$ ($r^2 = 0.975$). Exponents have been estimated using maximum likelihood [1,2].

Figure 4. Life-spans of individual locations. **(a)** Typical model runs, showing the cumulative distribution of number of time steps spent in the top 5 for model runs of 1000 time steps with $N = 1000$, $n = 100$, and $m = 1$. **(b)** Life-spans of UK Number One Hits [42] for 1956-2007 (open circles), versus the model, $m = 1$, $\mu = 0.1$ (blue line), and t years in the Top 5 US boys' names [43], 1907-2006 (filled circles) versus the model, $m = 10$, $\mu = 0.001$ (red line). Since the temporal units are arbitrary, the modelled life-spans were divided by 2 to match the albums, and divided by 10 to match the names (which on the log-log plot slides the distribution to the left).

Figure 5: Modelled life-spans of individual locations, showing the cumulative distribution of number of time steps spent in the top 5 for model runs of 1000 time steps with $N = 1000$, $n = 100$ for (a) $m = 1$, (b) $m = 10$, (c) $m = 100$, (d) $m = all$. The values of μ shown are 0.1 (black), 0.01 (dark blue), 0.001 (light blue), 0.0001 (purple) and 0 (red).

Table 1. Power-law fits determined by Clauset *et al*. (2) among socio-cultural data sets.

Parameters include the number of observations *n,* the maximum observed value x_{max}, the number of observations in the tail n_{tail} and the minimum value in the tail x_{min}.

quantity	n	x_{max}	x_{min}	α	n_{tail}	$f = n_{tail}/n$
intensity of wars	115	382	2.1 ± 3.5	1.7 ± 0.2	70 ± 14	0.609
religious followers (x 10^6)	103	1050	3.85 ± 1.60	1.8 ± 0.1	39 ± 26	0.379
word count	18855	14086	7 ± 2	1.95 ± 0.02	2958 ± 987	0.157
city population (x 10^3)	19447	8009	52.5 ± 11.9	2.37 ± 0.08	580 ± 177	0.030
terrorist attack severity	9101	2749	12 ± 4	2.4 ± 0.2	547 ± 1663	0.060
surname frequency (x 10^3)	2753	2502	112 ± 41	2.5 ± 0.2	239 ± 215	0.087
paper citations	415229	8904	160 ± 35	3.16 ± 0.06	3455 ± 1859	0.008
email address books	4581	333	57 ± 21	3.5 ± 0.6	196 ± 449	0.043
papers authored	401455	1416	133 ± 13	4.3 ± 0.1	988 ± 377	0.002

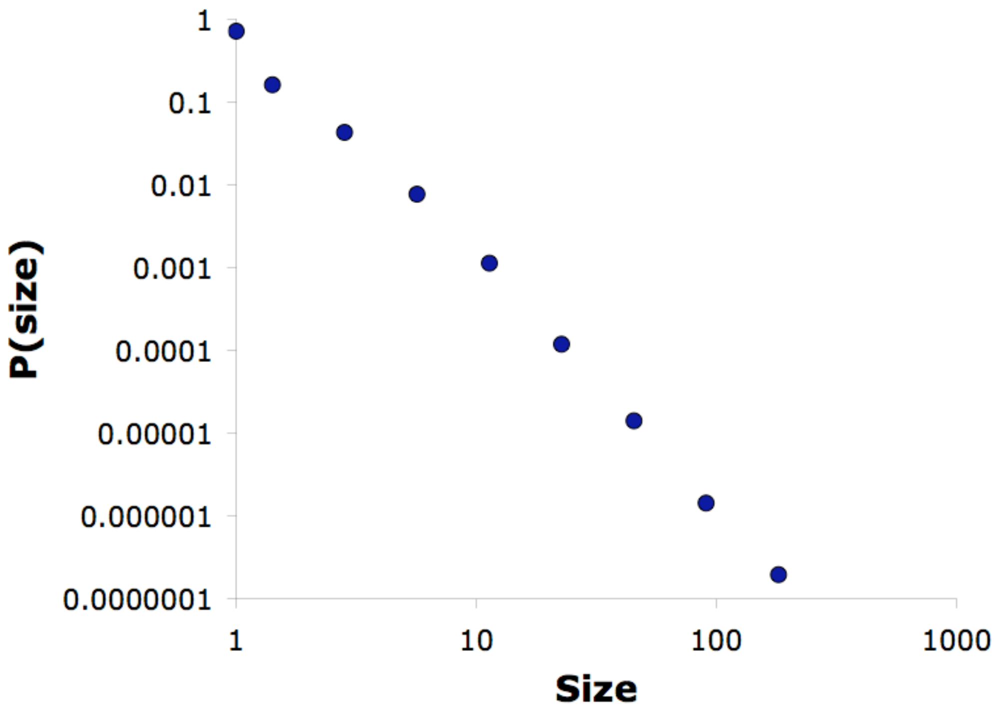

Figure 1.

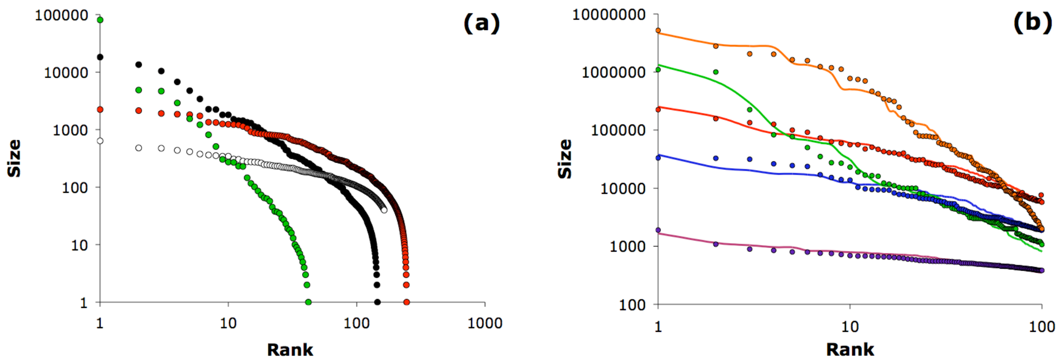

Figure 2.

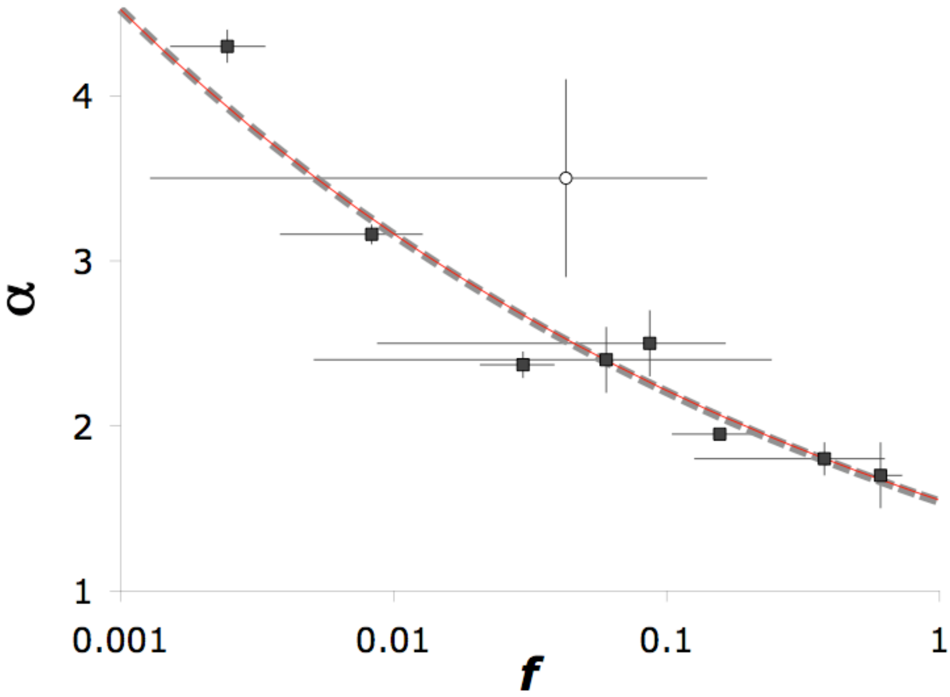

Figure 3.

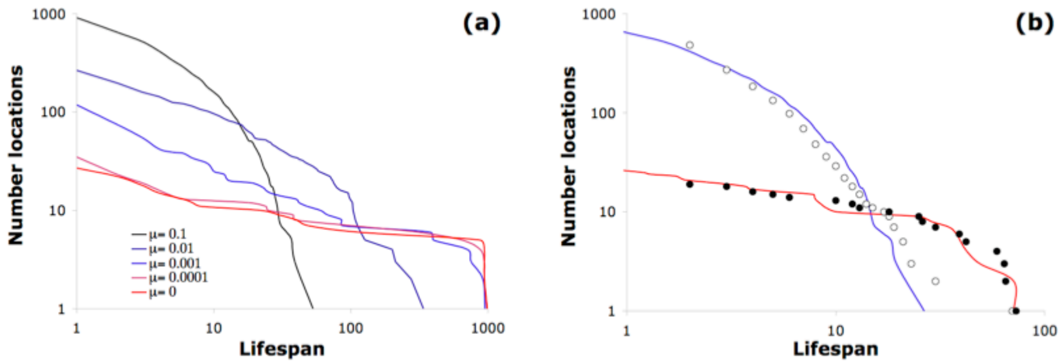

Figure 4.

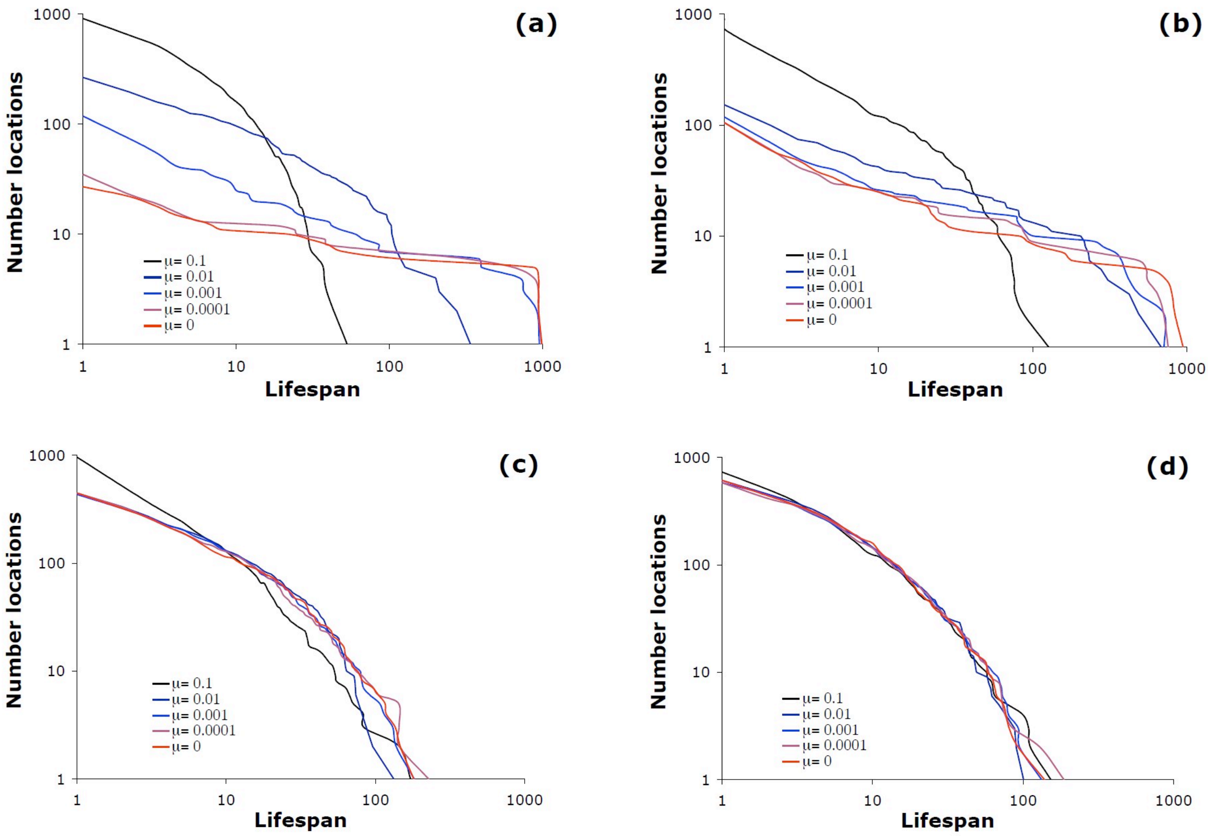
(a)
(b)
(c)
(d)
Number locations
Lifespan
1000
100
10
1
μ= 0.1
μ= 0.01
μ= 0.001
μ= 0.0001
μ= 0

Figure 5.