POWER-LAWS AND SCALING IN THE GENERALIZED LOTKA-VOLTERRA (GLV) MODEL

SORIN SOLOMON AND ABRAHAM AGAY
sorin@vms.huji.ac.il
Givat Ram, Hebrew University of Jerusalem

1. Introduction

Many natural and man-made phenomena are known to involve power-law probability distributions (e.g. Zipf 1949; Zanette and Manrubia 1997; Pareto 1897; Mandelbrot 1961, 1951, 1963; Atkinson and Harrison 1978; Cahalan and Joseph 1989).

Power-laws are common in systems composed of units that have no characteristic size, and in systems made up of auto-catalytic elements (Yule 1924; Champernowne 1953; Simon and Bonini 1958; Anderson 1995). Moreover, it has been shown that in systems which cannot be separated into dynamically independent parts, the power laws naturally take the place of the usual exponential (Boltzmann) distribution (Tsallis 1988).

It has been shown both theoretically (Solomon 1998; Solomon and Levy 1996) and using numerical simulations (Biham et al. 1998) that the Generalized Lotka-Volterra (GLV) model produces power-law distributions, but doesn't suffer from the problems and limitations of previous similar models. Since the GLV requires only a few, not too restrictive, pre-conditions, it may be expected to be widely applicable.

The GLV involves three scalar parameters and a probability distribution, each having a clear role in the model's evolution. In fact, certain model properties are "universal", i.e. independent of some of the model's parameters. This conceptual simplicity makes the GLV easily adaptable to different systems, and makes it possible to use it as a tool for description and comparison.

The features listed above make the GLV a candidate for a general method with which to simulate, analyze and understand a wide class of phenomena which are characterized by power-law probability distributions.

The aim of this paper is to provide a practical introduction to GLV modeling. We describe the basic theory, how to construct a numerical simulation of the GLV, and an application to an interesting system.

For didactic reasons we will present two previous and simpler models before introducing the GLV. This will nicely partition the topics involved into more manageable parts. The three models described in this paper are:

– Multiplicative process with a barrier
– Kesten equation
– Generalized Lotka-Volterra (GLV)

2. Some Basic Concepts

Microscopic or agent-oriented simulation is a relatively new way of generating and expressing knowledge about complex systems (Solomon 1995).

One idea this document tries to convey is the basic difference between standard methods of explanation, which are based on a global parameterization of a macroscopic dynamics and the "microscopic" ("agent-oriented") simulation method of explanation.

To illustrate the older methods, think of the high-school textbook problem of computing the trajectory of a stone under the influence of gravitation. The stone is treated as a point mass — that is the global parameterization, and we use Newton's Laws — that's the macroscopic dynamics.

This kind of explanation is not always useful. In many physical and economic situations etc., we find it is most natural to view the system under study as a collection of similar units (or agents), which interact among themselves in some well-defined way, and evolve in time.

Experience has shown that even a simple dynamics applied to a system consisting of many similar interacting agents may produce interesting macroscopic dynamics. The properties which are created by the collective behaviour of many similar agents are called emergent properties.

Some examples of systems with emergent properties are:

– Animal (and human) populations, which are composed of individuals that are constantly born (and die), and compete with other members of their species (for food, mates etc).

  The classical (scalar) Lotka-Volterra equation system was invented specifically in order to model such populations.
– The stock market, which is composed of many investors, each having a certain personal wealth. The investors buy and sell stocks, using more or less complicated strategies designed to maximize their wealth.

  Another way to characterize financial markets is to consider individual stocks (the capitalization of the individual companies traded in the market) as elementary degrees of freedom.

  We are going to use the Generalized Lotka-Volterra equation system to construct a very simple model of the stock market. In spite of being very simple, the model yields (and in doing so explains) the Pareto distribution of wealth, an economic fact known for a long time (100 years!) though never satisfactorily explained.

– Clouds, which are collections of water droplet that grow by taking in water vapor from the surrounding air. Agent-oriented models may explain the observed distribution of cloud sizes without going into all the very complicated physical details.

To make things more concrete, in the following discussion we will often use terminology borrowed from the GLV model of the stock market. However, our results and insights will be quite generic and applicable to other systems.

For example, we will often use the term 'traders' instead of agents, and talk about the 'wealth' of each trader, or the 'total wealth'. All these terms should be suitably renamed in other contexts.

3. Probability Distributions as Predictive Output of Modeling

Modeling the evolution in time of one or more traders in the stock market produces a lot of raw data. It is not very useful to produce more and more sequences of numbers without a way to analyse this data.

A natural way is to collect a lot of such sequences and analyze them statistically. For example, assuming each data sequence describes the evolution of one trader, we may compute the probability distribution over a large set of instances (they can belong to a long time segment of one simulation run or to many time sequences belonging to different runs).

Another way to look at it is to imagine a large set of (uncoupled) traders, subject to the same dynamics, and compute their distribution. Since the traders are identical, uncoupled, and subject to the same dynamics, their distribution will be the same as above. However, this way of looking at the problem leads to a model in which the various traders **do** interact. As we will see, this is a key ingredient in solving some of the problems with one-agent models and a way of making the models more realistic.

4. The Shape of the Probability Distributions

It may be a little surprising, but in certain conditions, we may predict (without performing any computer runs) how the wealth distributions produced by our models will look.

In fact, all the three models we are going to study produce probability distributions which have the following general characteristics:

– The distributions vanish for negative values.

This property follows from the way we choose the initial wealth values, and the nature of the dynamics. In practice it is found that the wealth assumes only positive values.

- The value zero usually has vanishing distribution.
  This follows from the continuity of the distribution, and the previous characteristic.
- The distribution has a long tail (which may obey a power-law or not).
  Distributions must have this property so they can be normalized.
- The distribution has a maximum.
  A "nice" continuous and non-negative function will have a maximum between two zeroes. Here we have one at zero, and one at infinity.

Some questions still remain, for example:

- The behaviour of the tail, in particular whether it decreases as a power-law, log-normal or some exponential.
- The location and height of the maximum
- The behaviour of the rising part

In this article we will be interested mainly in the shape of the tail, in particular whether it's a power-law or not. Our power-laws will be expressed in the form:

$$P(w)dw \propto w^{-1-\alpha}dw \quad (\alpha \geq 0) \tag{1}$$

5. Mechanisms for Generating Power Laws

Power-law probability distributions with power $\alpha = 0$ are obtained analytically for a multiplicative process:

$$w(t+1) = \lambda w(t) \tag{2}$$

Taking logarithms on both sides we find that $\log w(t)$ is the sum of the constant $\log w(0)$, and random variables $\log \lambda$ having the same distribution. Under the general conditions of the Central Limit Theorem (CLT) we would then get the log-normal distribution:

$$\frac{1}{\sqrt{2\pi Dt}}\frac{1}{w_t}\exp -\frac{(\log w_t - vt)^2}{2Dt} \tag{3}$$

$$v = \langle \log \lambda \rangle, \quad D = \langle (\log \lambda)^2 \rangle - v^2 \tag{4}$$

Note that the log-normal becomes an ever "flattening" w^{-1} distribution as the time goes to infinity.

Power-law probability distributions with power $\alpha > 0$ are obtained analytically for a multiplicative process modified in such a way that $w(t)$ will have a lower bound (or barrier) w_{min}. Such a modification may be accomplished in two ways:

– Making the distribution of λ dependent on $w(t)$, the outcome of the random process. In order to repel the wealth from w_{min}, the distribution is skewed to the right whenever $w(t)$ gets near the barrier value.
– A "manual reseting" of $w(t)$ whenever it drops below the "barrier" w_{min}:

These modifications are obviously not allowed by the CLT, and the derivation above which led to the log-normal distribution cannot be applied.

The consequences of such modifications are far reaching and it is plausible that they change the physics in an essential way, as they are equivalent to a localized drift towards higher values of wealth. The somewhat unexpected result is that the tail of the log-normal distribution is "pushed downwards" to yield power-laws with positive α.

The main problem with this power-low generating mechanism based on single-agent dynamics is that it works only for small values of λ which typically make $w(t+1)$ smaller than $w(t)$. This is not the case in nature, where populations and economies expand (at least for certain time periods). Moreover, the exponent of the power law is highly unstable to fluctuations in the parameters of the system. We will see later how the multi-agent GLV solves these problems.

There is an easy way to check if our data obeys a power-law or a log-normal distribution. On a log-log graph, a power-law yields a straight line, and the log-normal yields an inverted parabola. Of course, a simple exponential will appear as a straight line on a linear-log graph. To prove these observations, take the logarithm of the defining equations, and remember that on the log-log plot, distances are proportional to the logarithm.

The fine point here is that a power-law tail is difficult to distinguish from a log-normal tail, especially if the latter is very flat.

6. An Example: The Kesten Equation

The discrete Kesten equation may serve as an example of a simple single-agent stochastic model, and as a "building block" for the GLV model:

$$w(t+1) = \lambda w(t) + \nu \tag{5}$$

The Kesten equation may (crudely) describe the time-evolution of the wealth of one trader in the stock market. In this context λ may be called the "success factor", and ν is a "restocking term", which takes into account the wealth acquired from external sources (e.g. using the state built infra-structure, subsidies etc).

Of course, w represents the wealth of our lone trader, who at each time-step performs a stock transaction. $w(t)$ is the wealth at time t, and $w(t+1)$ is the wealth at time $t+1$. In practice t will assume positive integer values.

It's clear that λ and ν should be random variables obeying some given probability distribution Π; in the Kesten system both are taken to have positive distributions. We may take As a first approximation, we may take a uniform distribution for λ with $\langle\lambda\rangle < 1$.

The Kesten model produces power-law distributions for w if λ is predominantly smaller than 1, and λw is smaller than w on the average, i.e. in a shrinking economy. In this case w is bounded, and the ν coefficient is always large enough to provide the required drift to the right.

7. The Generalized Lotka-Volterra Equation System

The GLV implements a more complex dynamics than the Kesten equation. The advantages, however, outweigh the difficulties. In particular, the model ensures a stable exponent of the power law even in the presence of large fluctuations in the parameters. In particular, the value of $\langle\lambda\rangle$ can vary during the run (and between the runs) without affecting the exponent of the power law distribution. In particular, λ can take typical values both larger or smaller than 1.

We have N traders, each having wealth w_i, and each w_i is evolving in time according to:

$$w_i(t+1) = \lambda w_i(t) + aw - cww_i(t) \tag{6}$$

Here w without index is the total wealth, which supplies the only coupling between the traders:

$$w = w_1 + w_2 + ... + w_n \tag{7}$$

λ is positive and random with a probability distribution similar to the λ success factor we used in the previous section.

We will take the coefficients a and c to be constant although in general they may be functions of time, reflecting the changing conditions in the environment.

The coefficient a expresses the auto-catalytic property of wealth at the social level, i.e. it represents the wealth individuals receive in subsidies, services and social benefits as members of society. That is the reason that it is proportional to the total wealth.

The coefficient c originates in the competition between each individual and the rest of society. It has the effect of limiting the growth of w to values sustainable for the current conditions and resources.

Consider for a moment the differential version of the equation, with λ a non-stochastic variable having value 1. The steady state solution is $w_i = a/c$, independent of the w coupling, and traders with approximately this wealth will tend to keep it.

Another interpretation is to consider w_i, the capitalization of company i. In this case, λ represents the fluctuations in the market value of the company, while a represents the correlation between the value of each company and the market index and c represents the competition between the companies for the finite amount of money in the market. In this case, the model will predict the emergence of a power-law in the probability distribution of company sizes (capitalization). In particular, this would imply that the weights of the various companies composing the S&P 500 are distributed by a power law. In turn this would imply that the S&P fluctuations follow a truncated Levy distribution of corresponding index.

Yet another interpretation is to consider w_i as the size of coordinated trader sets (traders adopting similar investment policies), assuming that the sizes of these sets vary self-catalytically according to the random factor λ, while the a term represents the diffusion of traders between sets. The nonlinear term c then represents the competition for individual traders membership between these investing schools. In this case, the model would predict a power-law in the distribution of trader set sizes. Again, such a distribution would account for the truncated Levy distribution of market fluctuations (with index equal to the exponent of the power-law of trader set sizes).

It would be interesting to discriminate between the validity of the three interpretations by studying experimentally market fluctuations and the distributions of individual wealth, companies capitalization and correlated investors sets, respectively.

8. Correspondence with the Scalar LV Equation

There is a simple connection between the GLV and the conventional scalar Lotka-Volterra equation, which describe the time evolution of homogeneous populations:

$$w(t+1) = (1-d+b)w(t) - cw^2(t) \tag{8}$$

where b is the birth rate, d the death rate, and c describes the competition between individuals.

The GLV was designed to produce the scalar equation upon summation over all traders. w in the GLV should be equivalent to w in the scalar equation:

$$\sum \lambda w_i + aw - cww_i = \lambda w + aNw - cw^2 \tag{9}$$

Equating the corresponding coefficient gives:

$$\lambda + Na = 1 - d + b \tag{10}$$

Examining the order of the various terms suggests that a and c have to be of order $1/N$ or smaller in realistic systems.

9. Wealth Scalability of the GLV Equation System

Power-laws are scalable in the sense they retain their exponent upon scaling the argument by some factor. An equation producing power-laws is also expected to be scale invariant in this sense, so it's interesting to see what happens to the GLV upon the scaling of w_i.

Scaling w_i by γ means:

$$w_i \longrightarrow \gamma w_i \quad w \longrightarrow \gamma w \tag{11}$$

Applying this scaling to the GLV we get:

$$\gamma w_i(t+1) = \gamma \lambda w(t) + \gamma^2 c w w_i(t) \tag{12}$$

It is useful to define a new variable $u_i = w - w_i$ and write:

$$\gamma w_i(t+1) = \gamma \lambda w_i(t) + \gamma a w_i(t) - \gamma c u_i w_i + a u_i - \gamma^2 c w_i^2 \tag{13}$$

The last two terms "break" the scaling, but we are going to show that in the range where a power-law is expected they are small. In principle we should compare the terms we want to drop to all the terms we retain (or to the smaller one among them).

Ignoring for the moment the scaling factor, and taking u_i to be of same order as w, which is certainly plausible for large N, we find that the terms are of order: w/N, w/N^2, w^2/N^2, w/N, and w^2/N^3, respectively.

Dropping the "good" second term will make our job much easier. The fourth term is the toughest, comparing it with the first and third terms, we find that it can be dropped if:

$$w_i > w/N, \quad w_i > a/\gamma c \tag{14}$$

The first condition states that the power-law is valid for the above-average traders. The second condition is very interesting, it seems to require that a/c should be small.

Since w_i is always less than w, we can write:

$$w/N < w_i < w \tag{15}$$

10. Scaling the Number of Traders in the GLV

The GLV is a system of equations coupled by the total wealth w, when scaling the number of traders N, w is also scaled, and each individual equation changes.

If we want to preserve the pre-scaled behaviour of the traders, we should add a $1/N$ factor to the terms containing w, or equivalently, use the average wealth.

The $1/N$ factor can be attached to the a and c coefficients:

$$a = a'/N \quad a' = Na \tag{16}$$

$$c = c'/N \quad c' = Nc \tag{17}$$

The new primed coefficients are less sensitive to changes in the number of traders, and so are better suited to comparing similar systems with different sizes.

11. A Useful Formula for the Exponent

Having found a power-law distribution $w^{-1-\alpha}$ for the GLV, we would like to have a formula for the exponent α. Such a formula can be derived using the instantaneous ideal value of w in the steady state.

Summing the GLV over i, and substituting the instantaneous ideal value $I(t) = \langle w(t) \rangle = N\langle w_i(t) \rangle$, gives (after neglecting the fluctuations and relaxation time):

$$I(t+1) = \langle \lambda \rangle I(t) + \langle a(t) \rangle N I(t) - \langle c(t) \rangle I^2(t) \tag{18}$$

In the steady state $I(t+1) = I(t)$, and we can write I instead of $I(t)$ and rearrange:

$$I = \frac{\langle \lambda \rangle - 1 + N\langle a(t) \rangle}{\langle c(t) \rangle} \tag{19}$$

Substituting back into the GLV:

$$w_i(t+1) = [\lambda(t) - \langle \lambda(t) \rangle + 1 - N\langle a(t) \rangle] w_i(t) + a(t) w(t) \tag{20}$$

The last term can be neglected, since it is small in the range in which we are interested: $w/N < w_i < w$.

It is convenient to define a new variable:

$$z = \lambda(t) - \langle \lambda(t) \rangle + 1 - N\langle a(t) \rangle \tag{21}$$

It is easy to see that the last approximation produced a multiplicative process:

$$w_i(t+1) = z w_i(t) \tag{22}$$

Updating the traders in the steady state with the process equation shouldn't change the distribution:

$$w^{-1-\alpha} dw = \int_\lambda \Pi(\lambda) \left[\frac{w}{z}\right]^{-1-\alpha} \frac{dw}{z} d\lambda \tag{23}$$

The equation above involves a few subtle points, so a few remarks are in order:

- The equation describes the changes in the distribution due to one iteration of the multiplicative process.
- The left hand side represents the flow of traders leaving a certain wealth level w, and the right side represents the total flow of traders arriving at the wealth level w.
- w on the left hand side is the value before the updating, but w on the right hand side is after the updating. That's why we divide dw on the right side by z.
- We don't integrate over w, only over λ, and after the integration has been performed both sides will stay multiplied by dw.

Dividing both sides by their common factors gives:

$$\int \Pi(\lambda) z^{\alpha} d\lambda = 1 \tag{24}$$

This equation can be solved numerically for α.

12. The Nature of the Fluctuations

Stock prices are known to have a power-law distribution, and to fluctuate with Levy-stable (Mandelbrot 1963).

13. Implementing the GLV

To become a useful simulation a GLV implementation should have the following features:

- We should compute the w_i probability distribution, and either display the results in a table, plot it on the screen, or create a file for a plotting utility.
- We can't produce reasonable probability distributions with too few agents. Values of N of order 1000 were used in real work.
- We need to "smooth" the probability distributions by averaging over successive distributions.
- A good random number generator (preferably more than one) should be used.
- It is desirable to have some methods for updating the agent values, e.g. updating the value of w after each trader updating.
- A method to produce λ values with a specified (e.g. non-uniform) distribution should be provided.
- A flexible way of controlling the values of a and c should be provided.

14. Modeling with the GLV

Having gained some experience with the GLV, we will try to identify when it is possible to model a system with it. The characteristics of the model are of course necessary conditions:

– Having an extensive property (in the sense of statistical physics), assuming positive values.
– A power-law probability distribution of the said extensive property. The exponent is not critical, but should be less than -2 ($\alpha > 1$).
– The number of agents should be fixed (or a modification of the GLV should be used). All agents should have the same status, since the model treats them equally.

Modeling a phenomenon with the GLV is quite straightforward:

– Choose an extensive property which is known to obey a power-law distribution.
– Identify the constants a and c, and the distribution of λ in the problem.
– Get an approximation for the distribution of λ, and find the number of agents N.
– Use the approximation formula for α, and a simulation program to explore and find values of Na and Nc (these are more significant than a and c) which produce results similar to experimental data.

References

P. W. Anderson, J. Arrow and D. Pines, eds. The Economy as an Evolving Complex System (Redwood City, Calif.: Addison-Wesley, 1988)
A.B. Atkinson and A.J. Harrison, Distribution of Total Wealth in Britain (Cambridge University Press, Cambridge, 1978).
O. Biham, O. Malcai, M. Levy, S. Solomon, Phys. Rev. E **58**, 1352 (1998)
R. Cahalan and J. Joseph, Monthly Weather Review **117**, 261 (1989)
D. G. Champernowne, Econometrica **63**, 318 (1953)
Y. Ijiri and H. A. Simon, *Skew Distributions and the Sizes of Business Firms* (North-Holland, Amsterdam, 1977)
S. Solomon and M. Levy, J. Mod. Phys. C **7**, 745 (1996)
B. Mandelbrot, Comptes Randus **232**, 1638 (1951)
B. Mandelbrot, Econometrica **29**, 517 (1961)
B. Mandelbrot, J. Business **36**, 394 (1963)
V. Pareto, Cours d'Economique Politique, Vol 2 (1897)
H. A. Simon and C. P. Bonini, Am. Econ. Rev. **48**, 607 (1958)
S. Solomon, Ann. Rev. Comp. Phys. (1995)
S. Solomon, Computational Finance 97 (Kluwer Academic Publishers 1998)
C. Tsallis, J. Stat. Phys. **52**, 479 (1988).
U. G. Yule, Phil. Trans. B. **213**, 21 (1924)
D. H. Zanette and S. C. Manrubia, Phys. Rev. Lett. **79**, 523 (1997)
G. K. Zipf, Human Behaviour and the Principle of Least Effort (Addison-Wesley Press, Cambridge, MA, 1949)