

Bounded geometric growth: motivation for the logistic function

William S. Harlan

August 2007

INTRODUCTION

The logistic function appears often in simple physical and probabilistic experiments. A normalized logistic is also known as an S-curve or sigmoid function. The first derivative of this function has a familiar bell-like shape, but it is not a Gaussian distribution. Many use a Gaussian to describe data when a logistic would be more appropriate. The tails of a logistic are exponential, whereas the tails of a Gaussian die off very quickly. To decide which distribution makes more sense, we must be aware of the conceptual model for the underlying phenomena.

In biology, the logistic describes population growth in a bounded environment, such as bacteria in a petri dish. In business, a logistic describes the successful growth of market saturation. In engineering, the logistic describes the production of a finite resource such as an oilfield or a collection of oilfields.

After discussing examples, we will see how a bound to exponential growth leads to logistic behavior. There are other forms of the logistic function with extra variables that allow more arbitrary shifts and scaling. First, I limit myself to the form derived most naturally from the Verhulst equation. Normalizations clarify the behavior without any loss of generality. Finally, I use a change of variables for fitting recorded data in physical units.

EXAMPLES

Exponential (geometric) growth is a widely appreciated phenomenon for which we already have familiar mental models. Investments and populations grow exponentially (geometrically) when their rate of growth is proportional to their present size. You can take almost any example of exponential growth and turn it into logistic growth by putting a maximum limit on its size. Just make the rate of growth also proportional to the remaining room left for growth. Why is this such a natural assumption?

Growth in a petri dish

Let us consider the bacteria in a petri dish. This is an easy way to create a logistic curve in nature, and the mental model is a simple one.

A petri dish contains a finite amount of food and space. Into this dish we add a few microscopic bits of bacteria (or mold, if you prefer). Each bacterium lives for a certain amount of time, eats a certain amount of food during that time, and breeds a certain number of new bacteria. We can count the total number of bacteria that have lived and died so far, as a cumulative sum; or more easily, we can count the amount of food consumed so far. The two numbers should be directly proportional.

At the beginning these bacteria see an vast expanse of food, essentially unlimited given their current size. Their rate of growth is directly proportional to their current population, so we expect to see them begin with exponential growth. At some point, sooner or later, these bacteria will have grown to such a size that they have eaten half the food available. At this point clearly the rate of growth can no longer be exponential. In fact, the rate of consumption of food is now at its maximum possible rate. If half the food is gone, then the total cumulative population over time has also reached its halfway point. As many bacteria can be expected to live and die after this point as have gone before. Food is now the limiting factor, and not the size of the existing population. The rate of consumption of food and the population at any moment are in fact symmetric over time. Both decline and eventually approach zero exponentially, at the same rate at which they originally increased. After most of the food has disappeared, the population growth is directly proportional to the amount of remaining food. As there are fewer places for bacteria to find food, then fewer bacteria will survive and consume a lifetime of food. Although the population size is no longer a limit, their individual rates of reproduction still matter.

The logistic function can be used to describe either the fraction of the food consumed, or the accumulated population of bacteria that have lived and died. The first derivative of the logistic function describes the rate at which the food is being consumed, and also the living population of bacteria at any given moment. (If you have twice as many bacteria, then they are consuming food at twice the rate.) This derivative has an intuitive bell shape, up and down symmetrically, with exponential tails. The logistic is the integral of the bell shape: it rises exponentially from 0 at the beginning, grows steepest at the half-way point, then asymptotically approaches 1 (or 100%) at later times. The time scale is rather arbitrary. We can adjust the units of time or the rates of growth and fit different populations with the same curve.

Let us quickly examine two slightly messier examples, to see the analogies.

Market share

The market share of a given product can be expressed as a fraction, from 0% to 100%. All markets have a maximum size of some kind, at least the one imposed by a finite number of people with money. Let's assume someone begins with a superior product and that the relative quality of this product to its rivals does not change over time. The early days of this product on the market should experience exponential growth, for several reasons. The number of new people exposed to this new product depends on the number who already have it. The ability of a business to grow, advertise, and increase production is proportional to the current cash-flow. An exponential is an excellent default choice, in the absence of other special circumstances (which always exist).

Clearly, when you have a certain fraction of the market, geometric growth is no longer possible. Peter Norvig coined this as Norvig's Law: "Any technology that surpasses 50% penetration will never double again (in any number of months)." But let's also assume we have no regulatory limits and no one abusing a larger market share (bear with me). This product should still naturally tend to a saturated monopoly of the market. Such market saturation is typically drawn as a sigmoid much like a logistic. In fact it is a logistic, given no other mechanisms. After saturation, the rate of change of market share is proportional to the declining number of new customers. In any given month, a consistent fraction of the remaining unconverted customers will convert to the superior product. That is, we have a geometric or exponential decline in new customers for each reporting period.

Mining and oil

Finally, let's examine the discovery and exhaustion of a physical resource, such as mining a mountain range, or exploration and production of oil in an field. The logistic has long been used to predict the production history, the number of barrels of oil produced a day, in any oil field. The curve also accurately handles a collection of oilfields, including all the oil fields in a given country. Such a calculation was first used by King Hubbert in 1957 to predict correctly the peak of total US oil production in the early 1970's.

Earliest oil production is easily exponential, like many business ventures. As long as there is vastly more oil to be produced than available, then previously produced oil can proportionally fund the exploration and production of new oil wells. Success also increases our understanding of an area and improves our ability to recognize and exploit new prospects, so long as there is no noticeable limit to those prospects. At some point though, the amount of oil in a given field becomes the limiting factor. Like bacteria in a petri dish, fewer oil wells find a viable spot in the oil field in order to produce a full lifetime. The maximum rate of production is achieved, very observably, when half of the oil has been produced that will ever be produced. (That is

not to say that oil does not remain in the ground, but it cannot be produced economically, using less energy than obtained from the new oil.)

Oil production from individual oil fields do often show asymmetry, falling more rapidly or more gently after a peak than expected from the rise. Petroleum engineers have learned that deliberately slowing production increases the ultimate recoverable oil from a field. Gas production of a single field tends to maintain a more constant rate of production until the pressure abruptly fails, dropping production to nothing. But while individual fields may have unique production curves, collections of fields in a region or country tend to follow a more predictable logistic trend, with the expected symmetry.

A contrived game of darts

I also find it useful to think of an entirely artificial game of throwing darts. The dartboard is finite in size and can only hold a certain number of darts. Our aim is good enough to hit the target every time, but bad enough that we have no control over where the dart lands in the target. If a dart hits an empty spot, then it always sticks. If it hits too close to an existing dart, it will bounce off. With me so far? Good.

Now we add a silly rule to get geometric growth. We will measure time in “turns.” For each turn, we are allowed to throw more darts. For the first turn, we are allowed to throw one dart, which is guaranteed to stick. For the second turn we are allowed to throw two darts, which are very likely to stick. For each turn after that, we are allowed to throw as many darts as are currently attached to the board. Our growth in the beginning should be almost geometric, doubling with each turn. But as more darts fill the board, we will see more of our throws bounce off. By the time the board is half full, we will be able to throw enough darts with each turn to fill all remaining space. But each dart now has only fifty percent chance or less of sticking to the board. We get a logistic curve if we plot the fraction of the dartboard that has been filled as a function of turns. Since the number of darts is finite, we don’t get an exact fit, but the expected number of darts on the board should describe a logistic curve.

Similarly, you can imagine bacteria spores landing in a petri dish as darts landing on a dartboard. If food is available, a bacterium survives and breeds. If the food was already consumed by a prior bacterium, then the new one will die.

Some people have compared marketing and oil and oil exploration to throwing darts, though probably not with this model in mind. In real life, our aim is not random. But, in effect, an improved aim does not change the game fundamentally. Some of the dartboard could be marked with customers or oilfields, and the rest be considered a miss. All we have done is reduce the size of the dartboard and provide more ways for the dart to miss its target.

The number of darts for the next turn should still depend on the total number of previous successes. The final stages of the game should not change either. When most of the targets are gone, our progress will be proportional to the fraction of remaining opportunities.

Bacteria, as well, probably do not fill a petri dish uniformly, but spread from a center. Most living populations will have some evolved ability to find food. Yet, such changes to the rules should accelerate or decelerate population growth, without changing the underlying geometric limits.

EQUATIONS

The scenarios described above do not come close to representing all problems that can be modeled as a logistic curve. The function solves certain estimation problems involving the parameters of a Gaussian random variable. Such an S-curve is also convenient for signal processing applications such as neural networks. To help our intuition, I will nevertheless explain the notation with the previous examples in mind.

Keep in mind that these distributions also represent expectations or probabilities. Imagine that each limited resource is composed of a finite number of unique identities, such as an individual customer, a certain barrel of oil, a particular bit of food, or an empty spot on the dartboard. The logistic represents the probability that a unique quantum of a resource will be consumed by a particular moment in time. Since the same probability distribution applies to all quanta, you expect an actual realization to resemble a histogram with roughly the same shape. Thinking of the logistic as a probability distribution will help when we try fit actual data.

The Verhulst equation and the logistic function

Let us use $f(t)$ to represent a fraction of some quantity limited to values between 0 and 1. This fraction is a function of time t .

We expect this fraction to increase over time. The rate of increase, the first derivative, will always be positive:

$$\frac{df(t)}{dt} > 0.$$

Units of time are fairly arbitrary for such problems. For the function to approach a value of 1 asymptotically, time must continue to positive infinity. To avoid a small non-zero value to begin growth, we can allow the function to begin arbitrarily early at negative infinity, where it can approach 0.

The scale of time units, whether seconds or days, is also arbitrary. We'll choose a scale that most conveniently measures a consistent change in the

function. Let us put the halfway point, at zero time so that

$$f(0) = 1/2. \tag{1}$$

For earliest values of t , we expect $f(t)$ to increase geometrically. That is, we expect the rate of increase to be proportional to the current value:

$$\begin{aligned} f(t) &\rightarrow 0, \text{ and} \\ \frac{df(t)}{dt} &\propto f(t), \\ \text{as } t &\rightarrow -\infty. \end{aligned} \tag{2}$$

Similarly, as time increases and our function approaches unity, we expect the rate of growth to be proportional to the remaining fractional capacity.

$$\begin{aligned} f(t) &\rightarrow 1, \text{ and} \\ \frac{df(t)}{dt} &\propto 1 - f(t), \\ \text{as } t &\rightarrow \infty. \end{aligned} \tag{3}$$

This assumption is worth dwelling upon in light of our previous examples. Given an almost complete saturation of our available capacity, growth cannot be limited any longer by the existing population. The only remaining limitation to continued growth is the size of the remaining opportunities for growth. If the remaining opportunities shrink by half, then the chance of our getting one of those opportunities must also decline by half. Here, I find the dartboard analogy very helpful.

Let us combine these two proportions (2) and (3) into a single equation that respects both:

$$\frac{df(t)}{dt} \propto f(t)[1 - f(t)].$$

For appropriate time units, we can avoid any scale factors and write

$$\frac{df(t)}{dt} = f(t)[1 - f(t)]. \tag{4}$$

This is slightly simplified version of the Verhulst equation, which originated in studies of populations.

The rate of growth at any time is proportional to the population and to the remaining available fraction. Both factors are always in play, though one factor dominates when the value of the function approaches either 0 or 1.

By centering this equation at zero time with (1), we can rearrange the Verhulst equation (4) and integrate for $f(x)$ with

$$\left[\frac{1}{f(t)} + \frac{1}{1 - f(t)} \right] df(t)/dt = 1,$$

$$\begin{aligned} \frac{d}{dt}\{\log f(t) - \log[1 - f(t)]\} &= 1, \\ \log f(t) - \log[1 - f(t)] &= t, \\ \log \left[\frac{f(t)}{1 - f(t)} \right] &= t, \text{ and} \end{aligned} \tag{5}$$

$$\frac{f(t)}{1 - f(t)} = \exp(t). \tag{6}$$

Finally, we arrive at the simplest form of a logistic equation:

$$f(t) = \frac{\exp(t)}{1 + \exp(t)} = \frac{1}{1 + \exp(-t)}. \tag{7}$$

Some versions include include arbitrary scale factors for time or for the fraction itself. We've avoided those by normalization to fractions and convenient time units. Later we will use a change of variables useful for fitting physical data.

First notice that this equation is anti-symmetric, with an additive constant:

$$\begin{aligned} 1 - f(t) &= 1/[1 + \exp(t)] = f(-t); \\ f(t) + f(-t) &= 1. \end{aligned} \tag{8}$$

The asymptotic growth at the beginning mirrors the asymptotic limit at the end. We can think of the used capacity or remaining capacity as mirror images of each other. This is particularly striking because our rate of uncontrolled growth in the beginning also determines our rate of diminishing returns in the end. To lose this antisymmetry, we would need to introduce different (fractional) powers in our original proportions (2) and (3).

As a curiosity, this derivation also shows that the ratio (6) of used capacity to the remaining capacity increases exponentially for all times. This suggests an alternative derivation. The ability to improve this ratio is proportional to the ratio itself. This might explain the unjustified optimism that sometimes accompanies the exhaustion of a depleting resource. The ratio of success to failure is still growing geometrically! This might also motivate somewhat the use of the S-curve in neural networks. The ratio (odds) of certainty to uncertainty is allowed to grow exponentially with new information. Overall, however, I don't find this behavior very helpful to intuition.

(Neural networks, and logistic regression prefer a different motivation for their S-curve. Notice that the logarithm of this ratio (5) is called the logit function. This particular logit function is a simple linearly increasing function of time. The logit function is also equal to the negative derivative of the binary entropy function, which measures the uncertainty of two possible outcomes. Integrating shows that the binary entropy is changing as a negative square of time, a parabola convex down, centered at zero time.)

The derivative $df(t)/dt$ is often a more interesting quantity than $f(t)$ itself. For example, in oil production, this might be the number of barrels produced a

day (with an appropriate scale factor). It could be the annual growth in market share, the rate at which a population grows, or the rate of consumption of food.

$$\begin{aligned}\frac{df(t)}{dt} &= \frac{\exp(-t)}{[1 + \exp(-t)]^2} = \frac{1}{[\exp(t/2) + \exp(-t/2)]^2}, \\ \frac{df(0)}{dt} &= 1/4, \text{ and } \frac{df(\pm\infty)}{dt} = 0.\end{aligned}\tag{9}$$

The maximum rate of increase, by design, occurs at time zero. It is also a perfectly symmetric bell-shape, rising from zero to a maximum value of $1/4$, then declining again, with exponential tails. In this form you can see more clearly how the exponential on one side eventually overwhelms the one on the other.

In this form, the derivative (9) has unit area, integrating to 1. The equation is also useful as the probability distribution function (pdf) that a given resource (food, oil, or customer) will be used at a particular moment in time.

Fitting real-world data

Assume you have some data that you think might be described by a logistic curve. You have the data up to a certain point in time. You might not be halfway yet. Can you see how well the data are described by a logistic? Can you predict the area under the curve, or the halfway point?

From a partial dataset, we do not yet know the ultimate true capacity, and we use real time units. Let's use another form of the Verhulst equation more useful for real-world measurements.

To get a form similar to that used by Verhulst for his population model, we can replace

$$f(t) \equiv Q(t)/k,\tag{10}$$

where $Q(t)$ is a measurable capacity or or population, and k is an unknown upper limit, called the "carrying capacity."

We also substitute

$$t \equiv r(t_r - \bar{t}_r),\tag{11}$$

with t_r for measurable time units, with r for an unknown constant growth rate, and with \bar{t}_r for an unknown reference time. The reference time \bar{t}_r is when we expect to reach half of the maximum capacity:

$$Q(\bar{t}_r) = k/2.\tag{12}$$

With these substitutions, we rewrite the Verhulst equation (4) as

$$\begin{aligned}\frac{dQ(t_r - \bar{t}_r)}{dt_r} &= r[1 - Q(t_r - \bar{t}_r)/k]Q(t_r - \bar{t}_r); \\ \frac{dQ(t_r)}{dt_r}/Q(t_r) &= r - (r/k)Q(t_r).\end{aligned}\tag{13}$$

Notice that the measurable quantities on the left of (13) are a linear function of the measurable quantities on the right. The slope of the line is r/k , and the vertical intercept of the line is r .

The quantity $\frac{dQ(t_r)}{dt_r}/Q(t_r)$ in (13) could be called the fractional rate of growth. It is the current rate of growth divided by the cumulative value so far. We do not need to know ultimate rates, capacities, or reference times to calculate this quantity. At earliest times, when $Q(t_r)$ is small relative to k , the fractional rate of growth (13) achieves a maximum value of r .

We can make a graph with this fractional rate of growth on the vertical axis, and with the cumulative value $Q(t_r)$ on the horizontal. For every time at which we measure these two quantities, we can place a point on the graph. All values are positive and fall inside the upper-right quadrant.

If the data fit a logistic curve, then we should be able to draw a straight line through them. The slope and vertical intercept of the line allow us to estimate the unknown constants r and k . The vertical intercept, where $Q(-\infty) = 0$, is the rate r , and the horizontal intercept is the maximum carrying capacity $Q(\infty) = k$.

So what about the reference time, \bar{t}_r ? As time increases our data points move along this line, but not uniformly. Time units do not appear explicitly, except as a sampling parameter. The time \bar{t}_r corresponds to the data point with half of the ultimate capacity, as in (12). We may not have enough data to identify this point from this graph.

Another drawback to this particular way of graphing data is that early times will show much greater scatter than later times. When $dQ(t_r)/dt_r$ and $Q(t_r)$ are small, their ratio will show greater variation for small variations in either. This particular linearization is more suitable for an age of graph paper. I prefer to fit the logistic more directly.

Using the \bar{t}_r definition (12) as a boundary condition, we can also rewrite the logistic function (7) in measurable units:

$$Q(t_r) = \frac{k}{1 + \exp[-r(t_r - \bar{t}_r)]}. \quad (14)$$

Here we can see more clearly that k is the ultimate maximum value of $Q(t_r)$.

If we fit $Q(t_r)$ directly, our fit should improve with time. The value is a cumulative one, integrating measurements over longer periods of time. Again, we can expect more variation at earlier times.

Instead, let us examine an absolute rate of increase $P(t_r)$ that we can also measure:

$$P(t_r) \equiv \frac{dQ(t_r)}{dt_r} = \frac{kr}{\{\exp[r(t_r - \bar{t}_r)/2] + \exp[-r(t_r - \bar{t}_r)/2]\}^2}. \quad (15)$$

Note the peak value is $P(t_r) = dQ(\bar{t}_r)/dt_r = kr/4$.

Now we have a function with more consistent variations over time. The incremental change during a short interval of time will tend to follow the underlying distribution, with greater deviations as we shorten the interval.

Actually, it isn't difficult simply to scan reasonable values for all three parameters k , r , and \bar{t}_r and minimize some misfit to $P(t_r)$. You can also plot the misfit as contours of multiple parameters and get a better idea of your sensitivity to each.

Choosing a best measure of misfit is still necessary. Least-squares, the default choice for many, makes sense only if you think that errors in your measurements are Gaussian and consistent over time. This seems unlikely. Lower magnitudes have less potential for absolute variation than larger ones. We could instead minimize errors in the ratio of a measured magnitude of $P(t_r)$ to the expected magnitude. Or equivalently, we can minimize errors in the logarithm of $P(t_r)$. If we minimize the square of those errors, then we are assuming that variations in our measurements are multiplicative, following a log Gaussian distribution. This is much better, but I think still not optimum.

Another way to think of the problem is that the logistic derivative $P(t_r)$ in (16) describes a probability of a particular quantity being exploited or consumed at a particular point in time. A given customer, bacterium, or barrel of oil, is most to appear near the peak time \bar{t}_r rather than near the tails. Given a certain realization of that probability, our recorded data, what parameters maximize the probability of that data? It turns out that this likelihood is maximized by a minimum cross-entropy.

Let our recorded data be pairs of samples $\{P^i, t_r^i\}$ indexed by i . Then the best distribution $P(t)$ should minimize

$$\min_{k, r, \bar{t}_r} \sum_i \left\{ P^i \log \left[P^i / P(t_r^i) \right] \right\}. \quad (16)$$

$P(t_r)$ is a function of these three unknown parameters (k, r, \bar{t}_r) .

The $P(t_r)$ that minimizes this cross-entropy is the one that makes the actually recorded data most probable.

It is not difficult to set physical limits on each parameter and exhaustively try all combinations of the three, with a dense sampling. You can speed convergence with Newton linearizations, but it is not necessary.