

BAK-SNEPPEN MODEL OF EVOLUTION

APPLICATIONS AND IMPLICATIONS

Bak, Per and Kim Sneppen, 1993, Punctuated Equilibrium and Criticality
in a Simple Model of Evolution: Physical Review Letters, Vol. 71, No. 24, p 4083-4086

Programmed by Steve Baedke
Documentation by Lynn S. Fichter

BIOLOGICAL (MIS)-ANALOGIES: (Help File Three)

Because the B-S model is very stripped down, one must be cautious making analogies with biology. Some critics have argued that the model is so stripped down that it is not an accurate reflection of real evolution. This becomes especially true when trying to use biological terms or concepts to describe the model. Making connections between the terminology of the model and biological terminology is important since we do want to make comparisons, but it must be done judiciously. For example:

1. The evolutionary units in the model are called “species”, and the species undergo selection, but the analogy is not exact. Traditional evolutionary theory states that selection does not operate on species, it operates on individuals, and to a biologist the distinction is important.

But, the B-S model is trying to capture something more abstract, how a system containing agents that undergo coupled selection behaves. Bak and Sneppen refer to their model as “coarse grained” meaning that an entire “species” is represented by a single fitness, and indeed when only the last individual of a species exists, just before the species goes extinct, that last individual contains all the fitness remaining in the species. Or, as Bak and Sneppen say it (1993, p 4086), “*Since all species belonging to a single avalanche become extinct together, they might well be viewed as a single organism.*” Still, saying the species have fitness and undergo selection violates many biologists’ sensibilities.

On the other hand, some evolutionary biologists such as Steven Stanley (1975, 1979 - p 35) argue that not only are individuals selected, but at a larger evolutionary scale species are also selected, as well as clades of species, and ecosystems. Conversely, Richard Dawkins (1990) argues that it is genes that are selected, rather than organisms (or, organisms are the genes’ survival mechanisms). Evolutionary biologists still argue over what constitutes a selection agent, but in the B-S model it does not matter whether the agents being selected are individuals, species, or ecosystems. The universal principles of complex systems apply to all scales, and calling these agents “species” is just a convenient handle.

2. Sometimes we use the terms *extinction* and *mutation* to describe the changes taking place in the model. For example, when the lowest fitness species changes fitness, the species is said to go extinct and be immediately replaced by a new species of different fitness. Likewise, when the species to the left and right of a replaced species change fitness they are said to mutate. But, in fact, neither of these changes are analogous to extinction or mutation. Rather, the fitness of all three “species” is simply changed up or down at random.

On the other hand, those species above the threshold, and not involved in an avalanche, are undergoing no change at all. Thus, by comparison those dropping below the threshold during an avalanche are dropping out of the stable ecosystem, and are going “extinct”, to be replaced by new species whose fitness has a good chance of having a fitness above the threshold allowing them to become a member of a stable part of the ecosystem.

The further analogy is that an avalanche destabilizes an ecosystem causing a rapid turnover in species that, by Red Queen evolution, raises their mutual fitnesses above the threshold to establish a new stable ecosystem. In the B-S model this occurs when an avalanche shifts sideways to involve a different set of species.

Because we are using these models to understand the behavior of biological systems it is easy to use biological terms to refer to parts of the model. Analogy makes the model easier to understand, but the analogies must be examined for just how parallel they are. Nonetheless, the B-S model exhibits behavior that mimics very closely what we observe in large scale patterns of biological evolution.

THEORETICAL FRAMEWORK: (Help File Four)

Prior to the early 1970's evolutionary theory stated that evolution took place by slow, steady changes over long periods of time - *gradualism*. Biological evolution is an exceedingly complicated process involving the interactions among uncountable numbers of individual organisms, and between those organisms and their environment. It is easy to get lost in the complicated and messy details while trying to understand the mechanisms behind the behavior we observe. Swirling around any historical event are so many facts that it can be hard to determine just what is important.

Per Bak and Kim Sneppen developed their extremely stripped down evolutionary model in the early 1990's to test some observations made by Niles Eldredge and Steven Jay Gould (1972) that the fossil record shows *punctuated equilibrium*, spurts of rapid change followed by long periods of stasis (little or no change), as well as observations by David Raup (1992) who noted that extinctions are episodic and take place at all scales of observation. Since both of these behaviors are similar to the sand pile behavior that Bak studied, the B-S model was designed to see if an evolutionary model could be designed that would self-organizes to a critical state, and exhibit avalanche behavior that follows a power law distribution.

This strategy is a traditional one in science, build a model that contains only the most basic elements, and see if its behavior has the same features as real systems. If there are similarities between the model and real world then perhaps we can get some insight into the evolutionary dynamics of the real system.

When a model as simple as this one, just by the working out of the mathematics (computational viewpoint) can illustrate such complex behavior, as well as have power-law relationships similar to those found in the natural world, it points to some universal phenomena. In this, the B-S model is like X_{next} , exceedingly simple in appearance, but amazingly deep in behavior and meaning. It is not a guarantee evolution works this way, but it does provide a new way of looking at the evidence.

Some of the theoretical concepts behind the model are:

1. Punctuated Equilibrium

From the time of Darwin evolution has been seen as gradual change taking place over very long periods of time. Gradualism is also the mantra of the Modern Synthesis, the basis of all modern theories of evolution. This changed in 1972 when Eldredge and Gould published their paper "Punctuated Equilibria: An Alternative to Phyletic Gradualism." The paper produced a rage of controversy generally split along discipline lines. Biologists, who were used to studying evolution over short time spans (years to hundreds of years), and using population genetics to study shifts in gene frequency, were adamant gradualists. On the other hand, paleontologists, or at least some of them, who observed evolution over very long spans of time (tens of thousands to billions of years) were just as adamant that the fossil record recorded long periods when no

evolutionary change took place - *equilibrium*, or stasis - *punctuated* by short intervals of very fast change - thus *punctuated equilibrium*.

Both gradualism and punctuated equilibrium were based on empirical evidence; it is just that the scales of observation were different. This controversy has begun to die down a little; each side has declared victory, and gone on about their business.

Yet, it would be remiss to not mention that neither biologists nor paleontologists would likely find the B-S model palatable. For one, the B-S model is the work of physicists, and evolutionary biologists do not take kindly to them messing in their domain. But also, the B-S model argues that large scale patterns of evolution are neither the result of all those messy details the biologists love, nor the result of changing environmental conditions or catastrophic events like meteorite impacts that the paleontologists love (even if the originator of the asteroid impact theory of extinction was the physicist Louis Alvarez). In fact, both paleontologists and biologists would tend to find the B-S model loathsome.

But, of course, in the B-S model small scale, gradual changes are not insignificant; it is they that lead the evolution to the critical point. And, environmental changes are not insignificant. They might well be the sensitive-dependent trigger that sets off an extinction avalanche. But, in both cases the system is most susceptible to change at the critical state regardless of what finally precipitates the change. The B-S model in no way negates any mechanisms the biologists or paleontologists postulate. It just says that if biological evolution is a self-organized critical system then there are additional mechanisms that we need to understand.

2. *Red Queen Evolution*

Prior to 1990 the idea of biological fitness was summarized in Sewall Wright's (1932) Shifting Balance Theory of a fitness landscape. Different fitnesses were seen as peaks on a landscape, and the height and positioning of these peaks were fixed. In this model an organism increased its fitness by climbing up the side of its peak until it reached the top, which is the maximum fitness it can achieve, and then it is stuck. It is a great theoretical problem how an organism can get off its present peak to find and climb an even higher fitness peak and continue evolving since to do so it has to traverse the valleys between the peaks, requiring it to temporarily lower its fitness - not something natural selection does naturally. Up, or extinct, are the only options. Several solutions have been proposed to solve this problem, including:

1. Wright's shifting balance theory which suggests that drift in local populations can allow a shift from one adaptive peak to another,
2. The evolution of sexual reproduction, and
3. Red Queen evolution.

Red Queen evolution states that species do not undergo evolutionary change in isolation but are driven by changes in the species they interact with. For example, a prey species runs fast to avoid being eaten, which spurs the predator to evolve to run faster to catch them, which causes the prey to evolve to run even faster to get away.

The Red Queen comes from Lewis Carol's "Through the Looking Glass" (Alice in Wonderland) where there is a scene in which Alice is running alongside the Red Queen but they are not getting anywhere. Alice asks why they do not seem to be moving, and the Red Queen replies that in Wonderland people need to run fast just to stay in the same place. If you want to get ahead you must run even faster. Red Queen change is a form of evolution, or learning, where two interacting entities are continuously forced to improve as a result of the interaction. Each is driven to evolve faster by the interactions between them. Indeed, without this mechanism they may not evolve at all. Witness the species above the threshold, and not involved in an avalanche - they are static.

Red Queen behavior in the B-S model occurs when the adjacent species (left and right) also change along with lowest fitness species. The rationale is that the adjacent species, lying closest to the changing species, will be the most affected by its changes. The fitness of the adjacent species may go up or down, and this is captured by allowing these fitnesses to change at random. In biological evolution the changes will not always be random since natural selection is driving them, but if even random changes can lead evolution to the critical state, then directed evolution will just be more efficient at accomplishing it. And the B-S model shows that even with mutations taking place at random, the system is very efficient not only at raising the threshold to the critical level, but also keeping the system continuously poised there.

2. *Dancing Fitness Landscapes.*

Closely related to the Red Queen theory is the idea of “dancing fitness landscapes.” In 1991 Kauffman and Johnson proposed the idea that neither the fitness of an organism, nor the fitness landscape are fixed. An organism’s fitness is always related to the fitness of the organisms it interacts with. If one organism changes, then the fitness of the organisms it relates to has also changed through no fault or effort of their own. After all, they are a *system* related by positive and negative feedback.

The image of a “dancing fitness landscapes” captures the idea that life is a dance where you cannot climb to some optimal fitness and just sit there, but must continuously “dance” to keep your fitness maximized in a continuously changing landscape - and that landscape is changing in part because other species you interact with are changing.

Dancing fitness behavior is observed in two places in the B-S model. It occurs first at the beginning of the run. The lifting of the threshold line occurs through Red Queen behavior. The second occurrence is during the avalanches when the threshold is near 0.66. Species not involved in the avalanche are, of course, not changing, and their fitness is fixed. But if an avalanche moves into their region then Red Queen behavior is operating to raise the average fitness of avalanching species up to the threshold.

3. *Self-Evolved Criticality and Avalanche behavior*

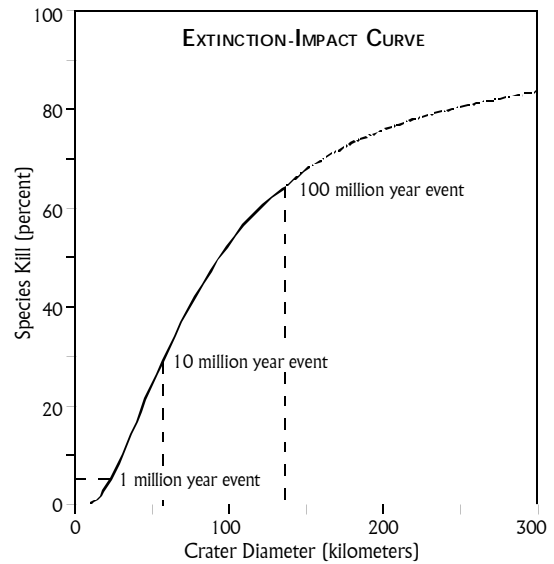
Per Bak in his sand pile experiments developed the idea that any open system that is capable of self-evolution will develop to the critical state, and undergo avalanche behavior. The Bak-Sneppen model applies these concepts to evolutionary systems. By coupling species together in a Red Queen relationship, a change in one species is analogous to dropping a grain of sand onto the sand pile. Changing the fitness of one species destabilizes the two adjacent species causing their fitness to change, which may in turn destabilize species adjacent to them causing their fitness to change, and so on. This avalanche of changes keeps avalanching until at some point the lowest fitness generated rises equal to or above the threshold, at which point the avalanche is over. However, because the threshold is so high it is very easy for Rule one to change the next lowest fitness species to something lower than the threshold, which then affects the two next to it by Rule Two, . . . and we are off on another avalanche.

The deeper significance of this avalanche behavior is that it has no identifiable cause. Yes, one mutation ends up below the threshold at random to set off the avalanche, but that is just the point. It is at random, with no identifiable cause, just one of many minor changes, and the result, at high criticality, is an avalanche, a wave of extinctions and new species generation that may go on for thousands or millions of events before the avalanche is over.

We are always looking for the cause of an extinction, the smoking gun that killed off all those species. But the B-S model demonstrates that because the system self-evolves to the critical state it is inherent just within the dynamics of the system for avalanches to occur. Extinctions are not necessarily caused by some external event, such as an asteroid impact or a dramatic climatic change. Such events may, of course, cause

an extinction, or may destabilize a critical system enough to precipitate an extinction, but the B-S model implies that such external events are not necessary.

Some idea of the complexity of the issue Raup (1992), using empirical data, devised an “extinction impact curve” that showed that the sizes of extinction events were correlated with the sizes of corresponding meteorite impacts (measured in Mean Waiting Time, larger impacts requiring exponentially longer waiting times). As can be seen, the resulting graph is a straight forward, convincing curve. From that data alone one might argue that all extinctions are the result of meteorite impacts. But, it is not that simple. Thirty-five million years ago a 2-3 mile diameter meteorite called the Chesapeake Invader, struck in the vicinity of the modern Chesapeake Bay (Poag 1999) excavating a 50-mile-wide, mile deep crater. An impact this size is impressive; the initial shock blast would spread out in a 600 mile diameter zone, and an enormous volume of dust ejected into the atmosphere, obscuring the sun and lowering temperatures for weeks to months. Based on Raup’s kill curve this impact should correspond with a sizable extinction, but in fact, there is no recorded extinction at this time.



The implications for the record of life are several. First, the system must be at the critical state before some external event can trigger an avalanche of extinction, and, conversely, if a system is not at the critical state the external trigger may strike without effect. Second, because a system at the critical state is sensitive dependent some significant external trigger may not be necessary to trigger an extinction. The trigger may simply be the flap of a butterfly wing, or the extinction of a single species that results in an avalanche that propagates through the entire system. Because the system at the critical state is sensitive dependent it may be futile looking for a specific trigger. Or, as Stuart Kauffman said, *the critical point is not a “nice place to be.”*

4. Power-law Relationships.

Finally, Bak has observed that self-evolving critical systems show power-law distributions. Indeed, power-law relationships are so characteristic of self-evolved critical systems that we use the discovery of power-law relationships to infer that a system must be a self-evolved critical system, even if other evidence is absent. The Bak-Sneppen model was created to not only test for Red Queen evolution, self-evolved criticality, and avalanche behavior, but also to demonstrate that avalanches within the system do indeed follow a power-law distribution. This, of course, is a computational viewpoint question - the only way to find out is to run the model. And, indeed, the computer model plots out in the upper right hand window the size of each avalanche and the threshold fitness at which it occurs. If you plot these on a log-log graph what you get is a power law distribution.

The B-S model is self-critical, and by implication any evolutionary system that incorporates the rules of the model will also exhibit self-organized criticality and avalanche behavior. It is Bak and Sneppen’s contention that this model does mimic some fundamental processes in biological evolution since both exhibit power-law distributions in the avalanche (extinction) sizes.

It is not perfect, however. It is not enough for both the model and life to exhibit a power-law distribution, the slope of the power-law line should be the same in both. They are close, but not exact, and so while the

Bak-Sneppen model captures some of what goes on in biological evolution, there must still be other “tuning controls” that will result in power-law distributions that are nearly identical in the model, and in life.

Nonetheless, the extremely close correspondence between the behavior and output of the model and the large scale patterns in the record of life argue that the B-S model has captures something essentially important about biological evolution, and it is not to be found in all the messy details of biology. Real life is a self-organized critical phenomena.

Larger Implications: (Help file 5)

Bak, Flyvbjerg, and Sneppen in a 1994 paper in the New Scientist (March 12) titled “Can We Model Darwin” explore some of the larger implications of B-S model of evolution.

1. *“We have studied many different versions of this model, and in all cases we found self-organized criticality. The behavior is robust - as it must be to represent real evolution since our models will certainly differ from the real thing when it comes to details.”*

Yes, indeed. Life is robust. Life always finds a way. Through 4 billion years of trials and tribulations life has always made it through, and come back even stronger. A model like the B-S gives us some insight into why this is true.

2. *“But, the most interesting feature of the model is its extreme simplicity and robustness; it has no subtle structure. We believe that only a few conditions need to be satisfied for the model to work. The picture which seems to emerge from our model is that while evolution does take ecology to the critical point by its self-organizing dynamics, the fitness of that point is not particularly high.”*

Fitness in the model optimizes at 0.66, and although it can climb higher it is exceedingly difficult, and takes exponentially larger periods of time to creep up more, 0.66 is not a particularly high fitness.

3. *“The critical point is not, as (Stuart) Kauffman once described it, “a nice place to be.” So “survival of the fittest” does not imply evolution to a state where everybody is well off. On the contrary, individual species are barely able to hang on - like the grains of sand in the critical sand pile.”*

And, we might add, the same is true for all self-evolving complex systems. It becomes most pertinent when we think about individuals in a society. All of us are jockeying for position, for power, for wealth. We live in a social structure where our “fitness” relative to all the other people we interact with is constantly changing. A close friend moves away, or gets a promotion where they move to a new social class, and our relative fitness has changed.

Of course, there is always that urge to find a safe niche to rest in, where maybe the vicissitudes of life will not affect us. Some people retreat and become hermit-like, some claw their way to wealth and/or power in the belief that such positions will make them safe. The B-S model by implication tells us no such safety is possible. And we have stories and movies that explore this theme.

But, perhaps if there is one positive lesson we might derive from all this, it is to remain as flexible and unflappable as possible. There will be ups and downs, most will be small, some will be big, but they are all just part of the system. Learn to go with the flow. Every avalanche is followed by a period of rebuilding. Give a prayerful bow to the inevitable, and get on with building a new life.

4. And finally, Gaia. *“In the critical state the species are connected at all scales, as illustrated by the power law distribution of avalanches. Since all species belonging to a single avalanche become*

extinct together, they might as well be viewed as a single organism. We thus have a hierarchical organization of organisms, up to and including the total ecology; thus one may speculate that the whole system in the self-organizing critical state acts as one interconnected organism, as suggested by Lovelock's Gaia hypothesis." (Bak, Flyvbjerg, and Sneppen, 1994, p 4086).

Like punctuated equilibrium, the Gaia Theory of James Lovelock (1979, 1988) raised a fury of controversy when it first appeared. There were scientists who hated it with a passion, and a smaller number who were intrigued by the idea, and willing to explore it. The history of the controversy is interesting and instructive, but again not our purpose here. The controversy has died down a little, and some of the critics are not so critical, even if not accepting.

What the Gaia theory has done is forced us to examine the earth and the life on it from a fundamentally new perspective, and frankly when Lovelock proposed the idea he was forced to speak largely theoretically and in metaphorical terms since there was very little hard data to support it. To be fair, that data may have been absent because no one was looking for it. This is not unusual, when Einstein proposed relativity he argued for the existence of phenomena no one had ever seen or looked for, and it took decades to test some of Einstein's predictions.

The Bak-Sneppen evolutionary model casts interesting light on the Gaia Theory. Lovelock, and Lynn Margulis, proposed that not only was all life on the earth interconnected, it also was connected with, controlled to some extent, the physical and chemical state of the earth in a state favorable for life. Or, as they say themselves.

"By looking at life through Gaia's telescope, we see it as a planetary-scale phenomenon with a cosmological life span. Gaia as the largest manifestation of life differs from other living organisms of Earth in the way that you or I differ from our population of living cells. At some time early in the Earth's history before life existed, the solid Earth, the atmosphere, and oceans were still evolving by the laws of physics and chemistry alone. It was careering (sic), downhill, to the lifeless steady state of a planet almost at equilibrium. At that instant the living things, the rocks, the air, and the oceans merged to form the new entity, Gaia. Just as when the sperm merges with the egg, new life was conceived.

"The Gaia hypothesis, when we introduced it in the 1970's supposed that the atmosphere, the oceans, the climate, and the crust of the Earth are regulated at a state comfortable for life because of the behavior of living organisms. Specifically, the Gaia hypothesis said that the temperature, oxidation state, acidity, and certain aspects of the rocks and waters are at any time kept constant, and that this homeostasis is maintained by activity feedback processes operated automatically and unconsciously by the biota. Solar energy sustains comfortable conditions for life. The conditions are only constant in the short term and evolve in synchrony with the changing needs of the biota as it evolves. Life and its environment are so closely coupled that evolution concerns Gaia, not the organisms or the environment taken separately."[The Ages of Gaia, 1988, p 19]

In evolutionary systems terms Lovelock and Margulis are saying that the earth and the life on it are all part of the same system, and because they are part of the system they are all connected by positive and negative feedback. Furthermore, as part of the same system they all self evolve to the critical state and undergo avalanche behavior. Now, this argument will not persuade anyone who does not accept Gaia in principle, but for someone who accepts the idea that the physical, chemical, and biological parts of the earth are all part of the same system, and accepts universality, then the Bak-Sneppen model is no-brainer support for the existence of Gaia.

Literature Cited

Bak, Per and Kim Sneppen, 1993, Punctuated Equilibrium and Criticality in a Simple Model of Evolution: Physical Review Letters, Vol. 71, No. 24, p 4083-4086

Bak, Per, Flyvbjerg, Henrik, and Sneppen, Kim, 1994, Can We Model Darwin: New Scientist, March 12, pp. 36-39.

Dawkins, Richard, 1990, The Selfish Gene: Oxford Univ Pr (Trade), 352 pages, ISBN: 0192860925

Eldredge and Gould (1972). Punctuated Equilibria: An Alternative to Phyletic Gradualism. Models in Paleobiology. San Francisco: Freeman, pp. 88-102

Kaufmann, S. A. and Johnsen, S.J., 1991, "Dancing Fitness Landscapes and Coevolution to the Edge of Chaos," preprint

Lovelock, James, 1979, Gaia: A New Look at Life on Earth. Oxford: Oxford University Press

Lovelock, James, 1988, The Ages of Gaia. New York: Bantam

Raup, David M., 1992, Extinction : Bad Genes or Bad Luck?: W.W. Norton & Company; ISBN: 0393309274

Stanley, Steven, 1975, "A theory of evolution above the species level": Proceedings of the National Academy of Sciences, USA, vol. 72(2) pg. 646.

Steven Stanley, 1979, Macroevolution: Pattern and Process, 1979.

Wright, Sewell, 1932, The Roles of Mutation, Inbreeding, Crossbreeding and Selection in Evolution: in, W. B. Provine (Ed.), Sewall Wright Evolution; Selected Papers (pp. 161-171. Chicago, Ill.: University of Chicago Press, 1986