


Centre for Philosophy of Natural and Social Science

Measurement in Physics and Economics

Technical Report 19/02

*Self-Organised Criticality - What It Is and
What It Isn't*

Roman Frigg



Editor: Peter Dietsch

Self-Organised Criticality – What It Is and What It Isn't.

Roman Frigg

Department of Philosophy, Logic and Scientific Method

London School of Economics

Houghton Street

London WC2A 2AE, England

r.p.frigg@lse.ac.uk

Abstract

The last decade and a half has seen an ardent development of Self-Organised Criticality (SOC), a new approach to complex systems, which has become important in many domains of natural as well as social science such as geology, biology, astronomy and economics, to mention just a few. This has led many to adopt a generalist stance towards SOC, which is now repeatedly claimed to be a universal theory of complex behaviour.

The aim of this paper is twofold. First, I provide a brief and non-technical introduction to SOC. Second, I critically discuss the various bold claims that have been made in connection with it. Throughout, I will adopt a rather sober attitude and argue that some people have gotten carried away by fancy contentions far too readily. My overall conclusion will be that none of these bold claims can be maintained. Nevertheless, stripped of exaggerated expectations and daring assertions, many SOC models are interesting vehicles for promising scientific research.

Keywords

Self-organised criticality, scaling law, models, formal analogy.

1. Introduction

Unifying theories have a long history. Ever since antiquity attempts have been made to explain a seemingly heterogeneous variety of phenomena by a few general principles. One of the most recent attempts in this vein is Self-Organised Criticality (SOC), a new approach to complex systems, which has become important in many domains of natural as well as social science.¹ SOC has been used to model phenomena as diverse as the dynamics of granular materials, evolution, earthquakes, forest fires, landscape formation, solar flares, river networks, mountain ranges, volcanic activity, traffic jams, plasmas, superconductors, fluctuations in stock markets, brain function and the spreading of epidemics, to mention just a few.

This has led many to adopt a generalist stance towards SOC, which is now repeatedly said to be a new paradigm for the explanation of complex phenomena or even to instruct us about how nature works. Although not all scientists explicitly endorse views similar to this, the absence of critical voices and the prevailing enthusiasm in the scientific community to come up with ever more applications of SOC and indicates that this indeed seems to be a widespread conviction.

This should make us suspicious. Does SOC really provide us with a general account of complex phenomena and if yes what kind of account is it? Despite its prominence within science, SOC has received no attention from philosophers so far. The aim of this paper is to get philosophers of science acquainted with SOC and to present a critical assessment of the theory as well as the bold claims that have been voiced in connection with it. To this end, I first provide a basic introduction to SOC (Sec. 2), and then discuss in what sense SOC is said to be a general theory that instructs us about the working of nature (Sec. 3). In section 4, I criticise these views and argue that, even on the most modest reading, the claim is untenable: SOC is not even close to

providing anything like an ‘essential dynamical process for everything that evolves and becomes complex’ (Paczuski and Bak, 1999, p. 2). This, however, does not prevent SOC from being a useful tool for promising research – false models may be valuable vehicles for promising research. Why this is so and how false models are used is discussed in section 5.

As this short survey indicates, the thrust of my discussion is deflationary. The bold claims put forward in connection with SOC cannot be sustained. This raises the questions whether SOC is just another item in a long line of inflated claims about the scope of a particular theory – as recent examples catastrophe theory, chaos theory or fractal geometry readily come to mind. The answer is yes and no.

The affirmative part of the answer asserts that the claims put forward in connection with SOC are indeed exaggerated in a similar way in which claims about fractal geometry, say, are exaggerated. SOC is not the essential dynamical process for everything that evolves and becomes complex just as fractal geometry is not the geometry of nature.² But why then, might one ask, should we spill any ink on one more of these unsustained pretensions? The answer to this question is that though, given past experiences with bold claims of that sort, a deflationary verdict may not strike us as too surprising, it needs argument – there is no *prima facie* reason why *every* bold theory *must* be mistaken. However, while fractals and chaos, for instance, have received a fair amount of attention in the past, not a single philosophical discussion of SOC has been published yet. This is rather surprising given that SOC is just the matter that is spreading at the moment and that people are interested in. The present paper aims at filling this gap.

The negative part of the answer points towards a disanalogy between the role SOC and the other aforementioned theories play outside the realm of mathematics and theoretical physics. SOC has diffused into applied disciplines to a much greater extent than any other of the recently fashionable theories. While catastrophe theory, for instance, by and large only paid lip service to

applications and remained a mathematician's concern, SOC is widely applied outside theoretical physics, where it originated. In particular in the earth sciences, from which many examples discussed in the sequel are drawn, SOC gains more and more momentum. Geologists write papers on SOC, publish them in geological journals and discuss the issue at conferences.³ This is, I take it, another reason why we need a critical examination of what SOC is and of what it achieves. Given that the generalist claims about SOC cannot be held up, what are we to make of these uses? This is the question I address in the second part of the paper.

2. A Primer on SOC

At first glance, there is a bedazzling multitude of SOC models.⁴ However, a closer look at the literature reveals that most of these fall into one of two main groups (Vespignani and Zapperi, 1995, p. 6347). The first group, the so-called 'stochastic models', consists of models with a stochastic dynamics operating in a deterministic environment. The paradigm example for such a system is the sand-pile model; further important examples include earthquake models and forest fire models. The second group, the so-called 'extremal' or 'quenched' models, consists of models with a deterministic dynamics operating in a random environment. The most important example of this type is the Bak-Sneppen model of evolution. In this section I will introduce both types of models and discuss important examples.

(a) Stochastic models

To become familiar with the kind of situations the first kind of SOC deals with, just remember some early childhood experiences in the sandbox (some people might be acquainted with similar scenarios from what occasionally happens on their desks). We set out to build up a pile of sand from scratch. We add sand in small portions (e.g. grain by grain) at random locations. At the beginning nothing happens. But as the pile grows, now and then the addition of a further grain triggers a local motion whereby the sand rearranges itself slightly. If, at that stage, we suddenly stopped adding sand, the pile would almost immediately come to rest in a ‘metastable’ state. As the process continues, the pile grows, its slope becomes steeper, and small avalanches start spilling over the sides. After a while, the pile reaches a critical size and large avalanches will occur that may span the entire pile. At this point the pile will grow no more, the amount of sand which slides off is equal to the amount of sand which is added. This is the self-organised critical state. In this state, avalanches of all sizes may occur; from the toppling of one single grain to ‘local’ landslides and up to huge avalanches tearing down good parts of the pile. This state is critical in the sense that a single local event such as the falling of a grain at a certain position can lead to effects affecting the system in its entirety; it is self-organised because this pattern occurs spontaneously, there is no outside agent or ‘invisible hand’ who imposes order. Hence, what has happened is that the addition of sand has turned the pile from a stable system, in which the dynamics of individual grains is local, into a critical system whose emergent dynamics is global.

After having laid down the leading idea, we have to give it a precise formulation in order to make it mathematically tractable. To this end, we model the sandbox by a two-dimensional quadratic lattice and introduce co-ordinates such that the tuple (n_x, n_y) denotes the cell whose co-ordinates are $x=n_x$ and $y=n_y$ respectively, where n_x and n_y are natural numbers ranging over $1, \dots,$

N ; thus the lattice has N^2 cells. Furthermore, we assume that our ‘grains’ are regular cubes, whose side-length is just the length of a cell of the lattice, and which can be stacked on top of each other like chocolate bars. Then we assign to every cell (n_x, n_y) an integer $Z(n_x, n_y)$ which represents the number of grains on the cell. What we get looks very much like a chessboard with dice piled up on it (see Fig. 1). Within this model, adding a grain of sand to the pile amounts to first choosing a cell (n_x, n_y) at random and then increasing the number $Z(n_x, n_y)$ by one: $Z(n_x, n_y) \rightarrow Z(n_x, n_y) + 1$. This process is now repeated again and again in regular discrete time steps (for instance one grain a second). But since we don’t want to pile up cubes ad infinitum, we have to introduce a ‘toppling rule’. This is done as follows: We stipulate that particles only flow in axial directions and can leave the system only at the margin. The ‘toppling dynamics’ is implemented by setting a threshold Z_c above which any given $Z(n_x, n_y)$ is unstable. Once $Z(n_x, n_y)$ exceeds Z_c the unstable site ‘tumbles’ by dropping four grains off and adding one of them to each of its nearest neighbours. Thus, for $Z(n_x, n_y)$ greater or equal to Z_c we have:

$$Z(n_x, n_y) \rightarrow Z(n_x, n_y) - 4$$

$$Z(n_x \pm 1, n_y) \rightarrow Z(n_x \pm 1, n_y) + 1$$

$$Z(n_x, n_y \pm 1) \rightarrow Z(n_x, n_y \pm 1) + 1$$

For convenience, Z_c is usually chosen to be equal to four, but in our model any number equal to or greater than four will do. It is easy to see that if the lattice is ‘crowded enough’, the addition of one more grain at a particular site may trigger a ‘domino effect’: The neighbours of the unstable site receive one more grain each of which makes (some of) them unstable as well, and from here the perturbation spreads to the neighbours’ neighbours and so on. The system has reached the SOC state. This ‘chain reaction’ results in an avalanche which keeps going until all

sites are stable again. Hence, minor perturbations may lead to effects of all sizes, and the perturbation may eventually propagate throughout the entire lattice.⁵ For convenience we also stipulate that while an avalanche is going down, no further grains are added to the system, the addition of sand stops when the first grain topples and just starts again when the system has reached a stable state.

Extensive numerical calculations have been carried out on this model and have yielded an astonishing result: The distribution of avalanche size and duration displays a power law behaviour, with all sizes appearing. That is, when we let the process go on for a long time and count the number s of grains that topple in an entire avalanche, and measure its duration d , and then count how many avalanches of size s and duration d have occurred, we find that these numbers, $N(s)$ and $N(d)$ respectively, obey a power law:^{6,7}

$$N(s) = s^{-a} \text{ and } N(d) = d^{-b}.$$

For the above model, numerical calculations yield that the exponents approximately equal 1.1.⁸ This power law is the hallmark of SOC and turns out to be crucial in future applications.




Fig. 1: There are two grains on cell (3, 3), hence $Z(3, 3) = 2$.

But what actually *is* SOC? Intuitively, the idea of a SOC state is quite clear, yet is there a rigorous definition? This turns out to be a hard question to answer. In the many pages written on SOC one hardly ever encounters a definition of what SOC really is. One finds vague phrases like ‘power law behaviour without parameter fine tuning’ or some loose talk about ‘scale invariance’ or ‘power law distributed events’. To improve the situation, Grinstein (1995, p. 270) suggests defining the SOC state as a ‘statistical steady state that is produced by processes with an infinite separation [...] of time scales [...], and exhibits scale invariance without further tuning of parameters or long-range attractions.’ He notes that this definition is not unproblematic and by no means uncontroversial. Be that as it may, I don’t want to dwell on this issue here since nothing of what follows hinges on it. For my purposes an intuitive understanding of SOC is sufficient.⁹

What is the bearing, one might ask now, of this fancy game on ‘real’ science? What do we learn about the world by tampering with an idealised sand-pile consisting of cubic grains obeying oversimplified laws of motion such as the toppling rule? Surprisingly enough, it has turned out that stochastic SOC models are much more than just an enjoyable pastime. Models of this kind have been constructed for all kinds of systems and its proponents claim they successfully explain the behaviour of such diverse phenomena as earthquakes or forest fires.¹⁰

To get a flavour of how this is achieved, I will present a brief discussion of the earthquake case.¹¹ We are all familiar with the phenomenon: Most of the time the surface of the earth is at rest, but every now and then the ground shakes. There are a few very large earthquakes with catastrophic consequences and many more smaller ones. But what does ‘many more’ mean? Earth scientists have pursued this matter and have come up with an interesting result: If one counts the number $N(m)$ of earthquakes of a given magnitude m one finds that they obey a

glaringly simple exponential distribution now known as the Gutenberg-Richter law:

$$N(m) = k * m^{-(1+b)},$$

where $b \approx 0.5-0.6$ and k is some constant (Chen *et al.*, 1991). But for the scientist with some knowledge about SOC this rings a bell: Though this law is a purely empirical finding and much older than SOC, it is precisely the power-law distribution for the avalanche size. This immediately raises the question: Are earthquakes like the sand slides in the pile model? Is the Gutenberg-Richter law an indication that the crust of the earth is in a SOC state?

Proponents of SOC have answered this question with an emphatic ‘yes’. A preliminary justification for this ‘yes’ points to the similarities between the crust of the earth and the sand-pile model: Avalanches correspond to earthquakes, adding a grain corresponds to tectonic plate motion providing the energy for the earthquakes, and so on (see for instance Bak, 1997, p. 89).

Admittedly, this is a bit too rough and ready. For this reason researchers have come up with different models of the crust of the earth, of which the so-called block spring model is the most popular. In this model the surface of the earth is represented by a two-dimensional array of blocks connected to each other by springs which represent the pressure on the material near a fault. The tectonic motion is simulated by fixing all blocks on a plate which is dragged along another (fixed) plate (see Fig. 2). Some tampering with this model soon reveals that the mathematics describing the motions of the blocks is very much the same as in the sand-pile model. Hence, the motion of the blocks exhibits SOC and the Gutenberg-Richter law holds true. Since the details of all this do not matter for my purposes, I omit them here. Moreover, this is not merely lucky coincidence; all earthquake models geophysicists have come up with univocally

exhibit SOC behaviour. For this reason, the proponents of SOC conclude, one can legitimately claim that the crust of the earth has organised itself into a critical state.




Fig. 2: The block-spring model.

(b) Extremal models

The second group of SOC models deals with processes operating in a random environment whose dynamics is deterministic. The most important model in this group is the so-called ‘Bak-Sneppen model’ of evolution. In this subsection I shall provide a brief discussion of this model and the heuristic ideas that lead to its construction.¹²

It is a widely (though not univocally) held conviction among Darwinians that evolution is a continuous and gradual process. On this gradualist view, biological change is driven by selection for small mutations, operating more or less uniformly at all times and all places. Unfortunately, this picture cannot explain mass extinctions like the vanishing of the dinosaurs about 65 million years ago, though this is a well-established fact. The usual move to salvage the concept of a

smooth evolution then is to look for extrinsic causes like volcanic eruptions, climatic changes or extraterrestrial objects hitting the earth, which are then supposed to account for the ‘irregularities’ in this otherwise smooth process.

This picture does not square with the following. The late Jack Sepkoski did extended research on thousands of marine fossil records and found that these large extinction events are just a part of a general pattern: If one divides time in consecutive units of four million years, then counts the fraction of species that disappeared since the previous period (i.e. if one calculates the extinction rate), and finally draws a histogram showing the number of periods with an extinction rate between 0% and 10%, 10% and 20%, and so on, then one finds that this histogram can be very well approximated by a power law of the same type as we have seen above. The mass extinctions then are at the extreme end of that curve.

As in the earthquake case, this raises the question whether this power-law distribution is a hint that evolution is a SOC process. After all, why not? Bak (1997, p. 114) writes: ‘Biology involves interactions among millions of species, each with numerous individuals. One can speculate that the dynamics could be similar to that of sand-piles with millions of interacting grains of sand.’ In this picture, an extinction can be viewed as an avalanche, where the death of a species causes the death or birth of other species in the same way that the addition of a grain of sand leads to toppling of other grains. However, not even real ‘aficionados’ of SOC pretend that speculations are enough. To substantiate this claim, various mathematical models of evolution have been constructed, most notably the so-called Bak-Sneppen model (Bak and Sneppen, 1993). In this model a continuous fitness number between zero and one is attributed to every species and selection is taken to act on the level of species. Then consider a given number S of species. To model the interaction between these species, think of them as placed on a circle, where each species only interacts with its two neighbours in the following way (in reality this might

correspond to a predator-prey chain): At each (discrete) time step, the species with the lowest fitness goes extinct and is replaced by a new species with a random fitness. However, this affects the neighbours as well (e.g. if a bird goes extinct the fitness of the beetles who were the prey of the bird as well as the fitness of the predators who hunted the bird are affected). For this reason, the fitness levels of the two neighbours of the extinct species also get replaced by a new random fitness number.

This is the set-up of the model. If now, starting from a random initial distribution of fitness, one calculates how species evolve, one finds that the model exhibits SOC. There are periods where almost nothing happens, and suddenly an avalanche of extinctions occurs. The size as well as the duration of the avalanches obey the by now well-known power-law distribution and thus correspond very well to Sepkoski's empirical findings. This suggests that evolution indeed is a SOC process.

Needless to say, this model was just a starting point, numerous variants, ramifications, and improvements have been suggested. All these, however, remain faithful, not to the detail, but to the spirit of the Bak-Sneppen model.

3. In what sense is SOC claimed to be ubiquitous?¹³

The examples discussed in the previous section by no means exhaust the class of successful applications of SOC. As already noted in the introduction, SOC has been applied to problems like forest fires, landscape formation, solar flares, brain function, and many more. Given this empirical success, some have claimed that SOC really is *the* crucial concept for understanding

the working of the world. To mention just a few, Per Bak (1997, title) claims to have discovered ‘how nature works’, Creutz (1997, p. 147) thinks that ‘[t]he idea provides a possible “explanation” of the omnipresent multi-scale structures *throughout the natural world*’ (my emphasis), Paczuski and Bak (1999, p. 2) assert that ‘punctuated equilibrium [i.e. SOC] dynamics is the essential dynamical process for *everything* that evolves and becomes complex’ (my emphasis), Blanchard and co-workers (1999, p. 375) observe that SOC has become a ‘new paradigm for the explanation of a huge variety of phenomena in nature as well as social science’, and Buchanan (2000, title) thinks that SOC is ‘ubiquitous’ and has discovered ‘why the world is simpler than we think’.¹⁴ Although this sample is certainly not representative and not all scientists explicitly endorse bold claims of this sort, the prevailing and wide-spread enthusiasm in the physics community to come up with ever more applications of SOC indicates that something along these lines is what a good many scientists at least implicitly subscribe to.

This gives rise to two questions: First, what is really claimed when SOC is said to be ubiquitous? In what sense is SOC a general feature of nature? Second, once we have settled for a specific reading of the claim, is the claim true or not? Is SOC really ubiquitous? In this section I address the first question; the second will be discussed in the next. My conclusion will be rather sober. I will argue that a closer look at how SOC is being put to use renders the ubiquity claim untenable.

What does it mean to say that SOC is ubiquitous? The most obvious interpretation of this claim is to think about SOC as a general theory (like thermodynamics, for instance) at the heart of which lie powerful general principles which turn out to be applicable across different and apparently unrelated fields. This seems to conform well with the views just mentioned.

Can this be true? Is SOC a general theory in (roughly) the same way as thermodynamics is? A closer look at how SOC is applied reveals that this cannot possibly be the case. A remark of

Bak's highlights the problem. While writing about earthquake models he notes that the model of evolution can be thought of as an earthquake model 'simply by a change in terminology' (1997, p. 172). This phrase nicely summarises the way in which SOC is put to use. Roughly speaking, the general strategy is this: take the sand-pile model (or another well understood SOC model) and reinterpret the terms in its mathematical description so that they refer to the new system. In the case of the earthquake model just mentioned, one reinterprets the mathematical term which stands for fitness in the evolution model so that it refers to the barrier distribution over a fault plane that generates earthquakes (cf. *ibid.*). In this way, we obtain an earthquake model from an evolutionary model almost for free. And this is by no means an atypical example. At the heart of SOC lie a few 'core models', which are 'applied' to different domains by reinterpreting the elements of one of the models. It is this strategy that generates the overwhelming number of SOC applications.

But do we apply a theory to a new domain by merely reinterpreting the terms of its calculus? I do not think so. Applying a theory amounts to bringing its general principles to bear on a particular situation. To stick with the example, thermodynamics is applied to a concrete situation by showing that the system's behaviour is an instance of its two general physical principles (which, very roughly speaking, state that energy is conserved and that entropy cannot decrease). In the case of SOC, however, things are different. SOC is applied by taking one of the core models whose mathematics is then reinterpreted so that its terms refer to features of the new target system. Hence, what lies at the heart of SOC are not physical principles but a few core models or, to be more precise, their mathematical skeleton which is then 'applied' by giving its terms different interpretations.

For this reason, SOC does not seem to be a scientific theory. Whatever stance one might take towards scientific theories, it seems to be uncontroversial that an uninterpreted calculus is not

part of empirical science at all. A scientific theory is *about something*, be it earthquakes, forest fires or electrons. But a calculus by itself is not about anything (at least not about anything in the physical realm). It becomes a scientific theory just by qualifying its elements with non-formal concepts like mass or energy.

To drive the point home, consider the example of linearity. Linear equations are arguably the most widely used mathematical structure in science (they are excessively applied in physics, economics, psychology, biology, and many other disciplines), but clearly they are not a *theory of everything* they describe; nobody would call linearity a theory of elasticity, the natural rate of unemployment (NAIRU), and learning (linear response theory), just because linear equations are used to describe all these phenomena. Similarly, we would not say that Bohr's theory of the atom and Newton's theory of the solar system are identical because they are based on the same calculus; and it would strike us as rather outlandish, to say the least, if someone claimed that partial differential equations were a theory of diffusion, heat-conduction, quantum motion, and the vibration of a membrane just because this type of equation is used in all these cases.¹⁵ Hence, I conclude that SOC is not a scientific theory at all, at least not in the same sense thermodynamics, quantum mechanics, or electrodynamics, for instance, are theories.¹⁶

But if SOC is not a theory, what is it then? As I see it, the above discussion of the way in which SOC is applied suggests that it is best considered as a group of models which are connected to each other by what Mary Hesse (1963, pp. 75f.) has called 'formal analogy'. Two objects are related by formal analogy if they bear no similarity other than that of both being interpretations of the same mathematical calculus. This is the situation we find in SOC. What sand-piles, earthquakes, and traffic jams (as construed by SOC) have in common is precisely that they are differing interpretations of the same calculus. At the core of SOC lie a few mathematical

models which are interpreted in ever different ways in every application, whereby they turn out to be a powerful mathematical tool to tackle a wide range of problems.

On this reading, the claim that SOC is ubiquitous amounts to saying that a variety of different phenomena from various fields can successfully be modelled within the SOC-framework. SOC is seen to provide tools, which can be used to build adequate models of vastly different systems. This seems to be a fair construal of what SOC does, and also one with which its proponents can live. Hence the question arises: Is it true?

4. Is SOC really ubiquitous?

It is undoubtedly the case that SOC models have been constructed for an impressive variety of phenomena; hence the ‘numerical’ part of the claim is beyond discussion. However, in order to claim that SOC *is* ubiquitous one not only has to come up with a myriad of models for all kinds of things. One also has to show that these models really capture the essential features of the systems modelled, and hence provide a solid account of the processes at work. In what follows I shall argue that this is the problematic bit. Most of the SOC models are so highly idealised that it is, to say the least, an open question whether they really instruct us about what is going on in nature.

As an illustration consider the paradigm sand-pile. In a by now well-known experiment, Frette and his co-workers (1996) set out to check whether a pile of rice exhibits SOC or not. What they found is that the dynamics of the pile reproduces the power-law distribution of the avalanche-size typical for SOC only in the case of very elongated grains, but not in the case of round-ish ones.

The more spherical the grains become, the less SOC-like behaviour the pile exhibits. Hence, SOC is by no means a general pattern of the behaviour of granular piles. And worse, this failure to account for how a host of real piles behave is not just due to ‘experimental noise’; rather it is rooted in the simplifying assumptions of the model. The movement of the cubes in the model (the sand-pile model is a cellular automaton model) does not at all correspond to the movement of real grains: cubes move just from one square to the next in axial directions, and they can even jump upwards, whereas the behaviour of real sand shows a variety of slip, slide, and roll motions the model cannot account for (Kardar, 1996, p. 22).

A further problem is that the model works under the assumption that the grains move instantaneously from one field to the next and then come to rest. Real grains, however, show inertia (roughly speaking, the tendency to keep moving once set in motion). It takes some time to accelerate them and then it takes again some time until they eventually come to rest. But the SOC model does not take the inertia of the grains into account and this is another reason why the model fails to capture the dynamics of real piles. And the worst is yet to come: As Dickman *et al.* (2000, p. 38) point out, the neglect of inertia is necessary to obtain SOC in the first place. When we try to augment the model and build inertia into it, SOC vanishes; and the reason why the model gets things more or less right in the case of elongated grains is that their inertia is extremely small. Hence, there seems to be no way that this model can ever account for how most real piles behave, or in other words, it is not even approximately true. To conclude, the model gets neither the kinematics nor the dynamics of most real piles right, and the latter do not even reproduce the crucial power-law distribution of the event-size. For this reason, it is at best unclear what it tells us about how a real pile works. Frette *et al.* (1996, p. 51) draw the conclusion that ‘SOC behaviour is not a universal characteristic of slowly driven granular media.’ In view of all this, I don’t think we are entitled to say that real sand-piles do exhibit SOC.

However, other models seem to be better off since they, unlike the sand-pile model, do yield the right predictions. The evolutionary model, for instance, posits a power law distribution for the extinction of individual species, and this is exactly what palaeontologists found when they went through their fossil extinction records.¹⁷ The latter is certainly a remarkable finding, but is it sufficient to establish that evolution exhibits SOC? I don't think that this is correct. Power laws and SOC states are two different things. The finding of a power law can, but need not be a footprint of an underlying SOC structure; or in the jargon of confirmation theory, we face the problem of alternative hypotheses. There might be other, yet unknown, mechanisms for generating the same power law.¹⁸ Therefore, further evidence is needed to establish the claim that evolution really is a SOC process; in particular, one has to show that the model captures the essence of the phenomenon.

At this point, however, one may have one's doubts. One cannot help having the feeling that the evolutionary model is constructed simply to get the statistics right. 'Real' evolution does not have much in common with a game based on the static replacement of one species by another whose fitness is chosen randomly. The basic assumption of the model, that the species with the lowest fitness goes extinct and is replaced by a new one with random fitness, has some plausibility only if the extinction is due to an external cause (such as the impact of an asteroid). It is just false if the extinction is due to natural selection, as premised in the model. So in order to be realistic the model has to assume what it was built to do away with: external causes. More generally, a closer look at the set-up of the model reveals that it does not seem to capture what is really going on in this enormously complex process we call evolution. Philip W. Anderson, in his review of Bak's 1997, puts this proviso as follows: '[...] but the question here becomes whether the model really captures the essence of the phenomenon or is constructed simply to do what it does. An evolving ecology [...] has more to it than the right statistical distribution of extinctions

[...]. Contrary to Stephen J. Gould's prejudice, evolution does have a direction (we need not call it "progress"), and it is not a static replacement of one species by another. So Bak's "life" is not life as we know it.' (1996, p. 773)

Hence, neither in the case of sand-piles nor in evolution is it clear that the present models warrant the conclusion that the real systems exhibit SOC. And these examples are by no means exceptions. Although there are cases where a SOC-model can be considered an adequate description of the system's behaviour (elongated rice, or the movement of domain walls in type II superconductors, for instance; for the latter see Dickmann *et al.*, 2000, p. 38), similar provisos apply in almost all cases. Faithful representation is a necessary condition to sustain the claim that things in reality *are* like in the model. But SOC-models meet this condition only in the most exceptional cases. For these reasons I conclude that the currently available models do not provide sufficient grounds to sustain the claim that a vast class of real-world systems actually *do* exhibit SOC.

5. The usefulness of false models

My discussion so far has been rather critical. I have argued that SOC is not a theory but a group of models. Moreover, these models are often such gross oversimplifications of real situations that one cannot legitimately claim that the latter really do exhibit SOC. Taken at face value, what SOC tells us about a real system is patently false, and therefore the claim that SOC tells us how nature works cannot be sustained.

Where does this leave us? Should we conclude that SOC should be assigned to the waste paper basket sooner rather than later? I think that this would be too hasty for two reasons. First, truthful representation of facts is not the only route to scientific success. As Morgan and Morrison (1999, 28) point out: ‘We do not assess each model based on its ability to accurately mirror the system, rather the legitimacy of each different representation is a function of the model’s performance in specific contexts.’ Also false models (and theories) can be a vehicle for successful research and finally lead to an increase in knowledge. Second, SOC is a research programme (in a Lakatosian sense) rather than a ‘finished’ scientific edifice. The ‘hard core’ of the programme is the belief that earthquakes, evolution, traffic jams and so on are essentially SOC processes in the way specified by the paradigm models discussed in the first section.¹⁹ But research programmes are almost always ‘born refuted’ and an immediate fit between the conjecture and experience is never achieved. One must treat budding programmes leniently; they may take decades before they achieve reasonable empirical success.²⁰ For this reason, the question is not so much whether the currently available SOC models give faithful accounts of certain processes, but rather whether SOC is a promising research program or not. This, however, is a question we cannot settle now: One of the punch lines of Lakatos’ account of research programmes is that the methodology is backward looking. A programme can be evaluated only *ex post facto*. There is no way to foresee how successful it will be before things have actually been done. But SOC has hardly started (the key paper was published in 1987), so it is definitely too early to assess its fruitfulness as a program and to either praise or dismiss it.²¹

The remainder of the paper is devoted to a discussion of the first of the above points. Why deal with false models? *Prima facie*, the observation that most SOC models are oversimplified, incomplete, and deficient in many other ways leaves us with little reason for dealing with them. If we want to do so nevertheless we have to answer the question of how these models tell us

something about the real world. The block spring model, for instance, is not even remotely accurate as a description of an actual geological situation. But how then can it be of any use in understanding real tremblings? In what ways can a model lead to the comprehension of a situation in the world when, as applied to this situation, it is just plainly false?

To answer these questions I will distinguish between the intrinsic and the extrinsic usage of a model. I call the usage ‘intrinsic’ when the model is employed as a tool for enquiry, i.e. to find out something about the world. I term the usage of a model ‘extrinsic’ when its purpose is strategic rather than investigative, that is, when its importance does not stem from its success as a tool of enquiry but, but from its pointing in new directions, articulating new ideas or ‘opening doors’.

(a) Intrinsic usages

In this subsection, I discuss different ways in which SOC models can be useful despite their falsity.

Learning from a model by building and using it: As Mary Morgan has pointed out (1999), learning from a model happens in at least two places: in building it and in using it. (1) *Building models:* Constructing a model is essentially a process of conceptualisation. Begin with some visual and experimental evidence about the behaviour of the system (in some cases also some background theory) and then interpret, conceptualise, and categorise this evidence in a certain way. Or, in Nancy Cartwright’s terms (1983, pp. 133-4), we produce a prepared description on the basis of the information gathered in an unprepared description. This involves applying a certain conceptual machinery to the phenomena and, more generally, looking at the system under investigation from a particular point of view. It is this process of interpreting and conceptualising

that takes place in the development of a model, which involves learning about the problem at hand.

This is what happens when one sets out to build a SOC model. As an example take the formation of landscapes. Why do they look the way they do? To answer this question we intend to think about landscapes in terms of SOC. This involves, among other things, to consider them as something dynamic rather than static (contrary to our common experience where landscapes appear static), to view them as the products of small outside influences, to focus on the breaking down of avalanches of different sizes, and to pay attention to the (possible) emergence of fractal patterns. By doing so (or by trying to do so) we learn a great deal about landscapes, even if the model we end up with is not a realistic representation of the real system.

In the example given, this conceptualisation even gave rise to empirical work. Based on these ideas, Tamas Vicsek and his co-workers built an experiment in which an ‘initial landscape’, a granular pile which had the shape of a ridge, was exposed to the influence of water. This set-up was then used to gather information about how landslides occur, what effects they have and how they, in the end, produce a relief (see Bak, 1997, pp. 75-7).

(2) *Using models*: Models are instruments for investigating the world rather than passive representations and we learn from them by using them. There are various different ways in which learning by using can take place. Among these, one is particularly important in connection with SOC, namely simulation. Roughly speaking, a simulation imitates one process by dint of another one (Hartmann, 1996, p. 83). However, this does not imply that one process is a perfect replica of the other. As Morgan and Morrison (1999, p. 29) observe, ‘we may be able to simulate the behaviour of phenomena without necessarily knowing that the simulated behaviour was produced in the same way as it occurred in nature.’ In other words, what is needed in order to make a simulation successful is only that we can map the model’s predictions onto empirical facts in a

direct way, i.e. that the model can reproduce the concrete details of the behaviour of a system, but this without necessarily representing the ‘higher level’ structures and properties adequately as well.

This notion of simulation is of great use in the case of SOC. Due to the discreteness of the models, computers can easily handle them and since we do not require that the process in the model (e.g. the dynamics of the cubes in the stylised sand pile) be an exact representation of the real dynamics of the system we can use these models to simulate all kinds of systems on a computer. A concrete instance of this use of SOC models is the study of river networks (see Bak 1997, pp. 80-4). Simple SOC models for the formation of landscapes have been developed which can simulate the process of the creation of a river network on a computer. This simulation is successful in the sense that many aspects of the computed river network agree with empirical observations. Nevertheless, and this is the salient point, the use of this model for the purpose of simulation does not commit one to view that SOC *is* the underlying mechanism that creates these networks in nature.

*Models as sketches:*²² Another way to make sense of SOC models is to understand them as sketches. While realistic models aim to depict the main characteristics of a system at least more or less accurately, sketch models intend to highlight particular features by isolating and exaggerating them. The assumptions of the model are chosen in such a way that they ‘give an impression’ of some selected aspect of reality, not by depicting it accurately, but by ‘overstating’, mis-describing, and deliberately distorting it. For this reason, the degree of approximation is irrelevant in the design of the model. The objective is to tell a simple story that captures some of the main features of the phenomenon, and the best way to do so is to disfigure these features into an extreme case (compare Gibbard and Varian, 1978, p. 665 and pp. 673-676). This, however, should not result in mere science fiction. The sketch must be such that, though false as a

description, it illuminates selected aspects of the system under investigation. Concretely, this can among other things mean that the model serves one of the following purposes:²³ The sketch can help to answer questions about more realistic models, provide an arena for answering questions about properties of more complex models, ‘factor out’ phenomena that would not otherwise be seen, serve as a limiting case of a more general model (or two false models may define the extreme of a continuum of cases in which the real case is supposed to lie), or it can lead to the detection and estimation of relevant variables.

This makes a good match with SOC. In fact, my discussion of SOC so far strongly suggests that many SOC models are best understood as sketches: They do not depict a system accurately but try to ‘give an impression’ of selected parts of reality by isolating, mis-describing, and exaggerating them (earthquakes and evolution, for instance), or they merely bear similarities of some kind to certain situations (SOC as applied to economic, social, and historical settings).²⁴

Most SOC models are sketches, but *what* exactly do they sketch and *how* do they do so? In what way do we learn something about real earthquakes by studying the block-spring model, for instance? What do we learn about evolution by tampering with the respective SOC model? Unfortunately, this is by no means clear. Gibbard and Varian hit the nail on the head when they remark that ‘[o]ften when a model is presented, only the briefest suggestive remarks are made about its bearing on the world [...]’ (1978, p. 676). Apart from the usual observation that the model reproduces the power law behaviour known from experience, by and large one looks in vain for further comments about what we learn from these models about actual systems and how they relate to real things. For this reason I think that further research on SOC will have to lay more stress on this issue and think in more detail about how these models relate to the real world.

(b) Extrinsic usages

False models as a starting point: It is a methodological truism that ‘one has to start somewhere’, and quite often a false model may provide a good point of departure. The strategy then is to ameliorate the false model and to build a better one on its basis. Iterating this procedure leads to a series of models in which every one is an improvement of its predecessor. In this way, a simple and false model can initiate a series of models of ever increasing complexity and accuracy.²⁵

This is quite common for SOC models. Many of them do indeed serve as a starting point for the construction of more complex models. Looking at the literature, one often comes across attempts to improve existing models by either altering their original set-up or building more details into the picture. For sure, there is still a long way to go until accurate models will emerge (if at all, this is unforeseeable at the moment), but scientists definitely do use these models as a starting point.

Models as heuristic devices: False models can serve a heuristic purpose in that they point in new directions or help to articulate new ideas, without yielding any correct results. Among the many heuristic functions a model can have, the most important one in connection with SOC is to *think differently*: A model, even though clearly false, may lead us to think differently about certain problems, motivate new questions, shed a different light on some issues, and finally make it easier to adopt an altogether different point of view. In doing so, the model acts as an antipode to stagnant assumptions, undercuts too readily accepted hypotheses, and helps to ‘de-familiarise’ deeply entrenched styles of reasoning. In short, a false model can indicate alternative ways to deal with a phenomenon.

Different SOC models have been immensely fruitful in this respect. Consider again the case of earthquakes: Before the onset of SOC, the aim of earthquake science was to detect patterns in large quakes and to predict them effectively. To this end, people were discussing the ‘characteristic size’ of a trembling, all kinds of statistics about sizes and frequencies of their

occurrence were made, and various indicators for quakes were investigated, such as the level of ground water or even the behaviour of animals. This has been an unsuccessful approach. No typical size has been identified and earthquakes are still as unpredictable as they have ever been. It is the merit of SOC to have initiated a shift of our focus: There is no difference in principle between small and big tremors, and the laws governing big quakes (if there are any at all) are not different from those governing small ones. Big and small events are both part of one and the same picture and by focusing on just one category we miss out on essential features of the whole. Moreover, due to their stochastic nature, earthquakes (small and big) are not predictable at all in the same way that we can predict other physical events. Having initiated this shift of focus is one of the enduring contributions of SOC to geophysics. Even if at the end of the day all current SOC models have found their space on the scrap yard of science, earthquake research will never be the same again.

As another example, consider evolution. One of the enduring contributions of SOC models to evolutionary biology is to have demonstrated that there is no contradiction (in principle) between Darwin's account which views evolution as a slowly and steadily advancing process, and the punctuated equilibrium theory suggested by Gould and Eldridge, according to which evolution is a stop-and-go progression (see Bak, 1997, p. 141). Although the model does not teach us much about evolution, it shows that these seemingly contradictory accounts can be compatible, and this is an important result.

Similar remarks, I think, apply to almost all disciplines in which SOC models have been used. As a heuristic tool, the sand-pile can suggest a new way of thinking about certain processes, and it can be employed as an antipode to widely held but unfruitful assumptions.

Summary and Conclusion

I have argued that SOC cannot possibly be a general theory in the same way thermodynamics or Newtonian mechanics are general theories. Rather, SOC is a group of models, which are related to each other via formal analogy. But even on this modest reading, the claim that SOC is ubiquitous is impossible to maintain, since most SOC models are gross oversimplifications and cannot in any way be considered realistic descriptions of their target system. Nevertheless, a lot can be learned from the construction and the use of SOC models, or when understood as sketches they can be vehicles for successful research. Moreover, they have a great heuristic value since they help ‘open new doors’, articulate innovative ideas, and at the same they ‘de-familiarise’ deeply entrenched styles of reasoning that have become stumbling blocks for research.

Acknowledgements

I would like to thank Mo Abed, Nancy Cartwright, Peter Dietsch, Carl Hoefer, Samuel Kutter, Adrian Lister, Mary Morgan, Julian Reiss, Christoph Schmidt-Petri and Philip Thonemann for comments on earlier drafts and/or for helpful discussions. Research for this paper was supported by the ‘Measurement in Physics and Economics’ project at the LSE.

References

- Achinstein, P. (1968). Concepts of Science. A Philosophical Analysis. Baltimore: Johns Hopkins University Press.
- Anderson, P.W. (1996). From the Top Down. (Review of Per Bak: How Nature Works: The Science of Self-Organised Criticality.) Nature, 383, 772-3.
- Bak, P. (1997). How Nature Works: The Science of Self-Organised Criticality. Oxford: Oxford University Press.
- Bak, P., & Sneppen, K. (1993). Punctuated equilibrium and Criticality in a Simple Model of Evolution. Physical review Letters, 71, 4083-86.
- Bak, P., Tang, C., & Wiesenfeld, K. (1987). Self-Organised Criticality: An Explanation of 1/f Noise. Physical review Letters, 59, 381-84.
- (1988a). Self-Organised Criticality. Physical Review, 38A, 364-74.
- (1988b). Self-Organised Critical Phenomena. In: H. Bai-Lin (Ed.), Directions in Chaos, Vol. 2 (pp. 16-45). Singapore, New Jersey & Hong Kong: World Scientific.
- Blanchard, Ph., Cessac, B., & Krüger, T. (2000). What Can We Learn about Self-Organised Criticality from Dynamical Systems Theory? Journal of Statistical Physics, 98, 375-404.
- Blanter, E.M., Shnirman, M.G., Le Mouel, J.L., & Allegre, C.J. (1997). Scaling Laws in Blocks Dynamic and Dynamic Self-Organized Criticality. Physics of the Earth and Planetary Interiors, 99, 295-307.
- Bons, P.D., & van Milligen, B.P. (2001). A New Experiment to Model Self-Organized Critical Transport and Accumulation of Melt and Hydrocarbons from Their Source Rocks. Geology 29, 919-929.

- Buchanan, M. (2000). Ubiquity. The Science of History ... Or why the World is Simpler than We Think. London: Weidenfeld & Nicolson.
- Cartwright, N. (1983). *How the Laws of Physics Lie*. Oxford: Oxford University Press.
- Chen, K., Bak, P., & Obukhov, S.P. (1991). Self-Organised Criticality in a Crack-Propagation Model of Earthquakes. Physical Review, 43A, 625-30.
- Clar, S., Drossel, B., & Schwabl, F. (1996). Forest Fires and other Examples of Self-Organized Criticality. Journal of Physics, Condensed Matter, 8, 6803-24.
- Clar, S., Drossel, B., Schenk, K., & Schwabl, F. (1999). Self-Organised Criticality in Forest-Fire Models. Physica, 266A, 153-159.
- Creutz, M. (1997). Cellular Automata and Self-Organised Criticality. In G. Bhanot, S. Chen, & P. Seiden (Eds.), Some New Directions in Science on Computers (pp. 147-172). Singapore, New Jersey & Hong Kong: World Scientific.
- Dickman, R., Muñoz, M.A., Vespignani, A., & Zapperi, S. (2000) Paths to Self-Organised Criticality. Brazilian Journal of Physics, 30, 27-41.
- Doi, M., & Edwards, S.F. (1986). The Theory of Polymer Dynamics. Oxford: Oxford University Press.
- Frette, V., Christensen, K., Malthe-Sørensen, A., Feder, J., Jøssang, T., & Meakin, P. (1996). Avalanche Dynamics in a Pile of Rice. Nature, 379, 49-52.
- Frigg, R. (2002). Review of Mark Buchanan: Ubiquity. The Science of History ... Or Why the World is Simpler than We Think. British Journal for the Philosophy of Science, 53, 585-591.
- Gibbard, A., & Varian, H.R (1978). Economic Models. Journal of Philosophy, 75, 664-77.
- Grinstein, G. (1995). Generic Scale Invariance and Self-Organised Criticality. In A. McKane *et al.* (Eds.), Scale Invariance, Interfaces, and Non-Equilibrium Dynamics (pp. 261-93). New York & London: Nato ASI Series.

- Hartmann, S. (1995). Models as a Tool for Theory Construction: Some Strategies of Preliminary Physics. In W.E. Herfel, W. Krajewski, I. Niiniluoto, & R. Wojcicki (Eds.), *Theories and Models in Scientific Process* (pp. 49-67). Poznan Studies in the Philosophy of Science and the Humanities, 44.
- (1996). The World as a Process. In R. Hegelsmann et al. (Eds.), Modelling and Simulation in the Social Sciences from the Philosophy of Science Point of View. Dordrecht: Reidel.
- Hesse, M.B. (1963). Models and Analogies in Science. London & New York: Sheed & Ward.
- Kardar, M. (1996). Avalanche Theory in Rice. *Nature*, 379, 22.
- Kirchner, J.W., & Weil, A. (1998). No Fractals in Fossil Extinction Statistics. *Nature*, 395, 337-8.
- Lakatos, I. (1970). Falsification and the Methodology of Scientific Research Programmes. In I. Lakatos & A. Musgrave, Criticism and the Growth of Knowledge (pp. 91-196). Cambridge: Cambridge University Press.
- Malamud, B.D., Morein, G., & Turcotte, D.L. (1998). Forest Fires: An example of Self-Organised Critical Behaviour. *Science*, 281, 1840-42.
- Malamud, B.D. & Turcotte, D.L. (1999). Self-Organised Criticality Applied to Natural Hazards. Natural Hazards: Journal of the International Society for the Prevention and Mitigation of Natural Hazards, 20, 93-116.
- Morgan, M.S., & Morrison, M. (Eds.) (1999). Models as Mediators. Perspectives on Natural and Social Science. Cambridge: Cambridge University Press.
- Morgan, M.S. (1999). Learning from Models. In Morgan and Morrison 1999, 347-88.
- Paczuski, M. & Bak, P. (1999). Self-Organisation of Complex Systems. e-print cond-mat/9906077.

- Psillos, S. (1995). The Cognitive Interplay between Theories and Models: The Case of 19th Century Physics. In W. Herfelm *et al.* (Eds.), *Theories and Models in Scientific Process* (pp. 105-33). Poznan Studies in the Philosophy of Science and the Humanities, 44.
- Rundle, J.B., Klein, W., & Gross, S. (1996). Dynamics of a Travelling Wave Model for Earthquakes. Physical Review Letters, 76, 4285-90.
- Schuster, H.G. (1988). Deterministic Chaos : An Introduction. Weinheim & Cambridge: Cambridge University Press.
- Shenker, O. (1994). Fractal Geometry is not the Geometry of Nature. Studies in History and Philosophy of Science, 25, 967-981.
- Solé, R.V., Manrubia, S.C., Benton, M., & Bak, P. (1997). Self-Similarity of Extinction Statistics in the Fossil Record. Nature, 388, 764-67.
- Stanley, M.H.R. *et al.* (1996). Scaling behaviour in the Growth of Companies. Nature, 379, 804-6.
- Stacy, S., McCloskey, J, Bean, C.J., & Ren, J. (1996). Heterogeneity in a Self-Organised Critical Earthquake Model. Geophysical Research Letters, 23, 383-386.
- Sutton, J. (2001). The Variance of Firm Growth Rates: The ‘Scaling’ Puzzle. Unpublished manuscript. London School of Economics.
- Vespignani, A., & Zapperi, S. (1998). How Self-Organised Criticality Works: A Unified Mean field Picture. Physical Review, 57E, 6345-62.
- Wimsatt, W. (1987). False Models as Means to Truer Theories. In N. Nitecki & A. Hoffman (Eds.). Neutral Models in Biology (pp. 23-55). Oxford: Oxford University Press.

Footnotes

¹ SOC is independent of earlier approaches to complexity and self-organisation such as synergetics, for instance. Therefore, I won't say anything about them in this paper.

² For a critical discussion of fractals see Shenker (1994).

³ See for instance Blanter et al. (1997), Bons et al. (2001), Malamud et al. (1999), Steacy et al. (1996), to mention just a few.

⁴ The two classical references are Bak, Tang, and Wiesenfeld (1987) and (1988a). Introductions may be found in Bak (1997); Bak, Tang, and Wiesenfeld (1988b); Buchanan (2000); Paczuski and Bak (1999). More advanced presentations of the subject matter include: Dickman et al. (2000); Creutz (1996); Blanchard et al. (1999), to mention just a few.

⁵ I should note that this effect has nothing to do with *sensitive dependence on initial conditions* which is a characteristic of chaotic systems.

⁶ Note that this is not literally true. Although the law holds true over a wide range there is an upper bound (the total number grains in the system) and a lower bound (one grain) to the size of avalanches.

⁷ Laws of this kind are often referred to as 'scaling laws'.

⁸ Considerable efforts have been made to calculate the exponents of these power laws analytically. Unfortunately all these attempts have so far been unsuccessful (Bak, 1997, pp. 62-64; Creuz, 1996, p. 3).

⁹ I should also mention that originally SOC has been suggested as an explanation of the occurrence of $1/f$ noise and fractal pattern formation in the dynamical evolution of certain systems. However, recent research has revealed that the connection between these and SOC is rather loose (Dickman *et al.*, 2000, p. 37). For this reason I will not discuss them any further.

¹⁰ See Bak (1997) and references therein; Paczuski/Bak (1999), pp. 2-3.

¹¹ The second important family of stochastic SOC models are the forest fire models. Space constraints prevent me from discussing them here, good surveys include: Clar *et al.* (1996) and (1999) and Malamud *et al.* (1998).

¹² The classical reference is Bak and Sneppen (1993). Accessible surveys can be found in Bak (1997), pp. 16-19, 105-174; Paczuski and Bak (1999), pp. 7-11.

¹³ This and the following Section are based on Frigg (2002).

¹⁴ In what follows I shall adopt Buchanan's terminology and use the phrase 'SOC is ubiquitous' to express the view that SOC is a general feature of nature.

¹⁵ I should note that my attack against the reinterpretation of a calculus is only directed against its use to produce an (allegedly) unified theory. As a heuristic move it has proved immensely fruitful and as such there is nothing wrong with it. For more details see Achinstein (1968), chs. 7 and 8; Hesse (1963), ch. 2 and Psillos (1995), sec. 3.

¹⁶ Some may find this way of looking at the issue too reminiscent of the (no longer fashionable) syntactic view of theories. However, my use of 'syntactic' jargon is merely a matter of convenience and nothing in the argument hinges on it. One may make the same points within the framework of the semantic approach to theories, though in this case things get a bit more verbose.

¹⁷ I should mention that Kirchner and Weil (1998) have pointed out the statistically self-similar structure Solé *et al.* (1997) purport to have found in the extinction time-series is an artefact of the interpolation methods used by these authors. However, nothing in my argument depends on how this question is settled.

¹⁸ And there actually are: Intermittent mappings produce the same $1/f$ noise pattern as SOC (Schuster 1988, 91-97). Or consider the following case from economics to further illustrate the

point (thanks to Nancy Cartwright for having drawn my attention to this example). Stanley et al. (1996) observed that for firms with similar sales the distribution of annual growth rates has an exponential form. Exploiting an analogy with statistical physics, the authors explained this empirical finding by assuming the presence of certain correlations across the firm's units. However, in a recent paper, Sutton (2001) has pointed out that the occurrence of these power laws is not necessarily indicative of some form of correlation between the units of the firm but can just as well be explained within a simple model that allows for the different businesses of a firm to vary in size. Hence, two totally different set-ups can equally well account for the empirically observed power laws.

¹⁹ Compare Lakatos' characterisation of the Newtonian programme as the 'metaphysical principle' that 'the planets are essentially gravitating spinning-tops of roughly spherical shape' (1970, p. 136).

²⁰ See Lakatos' lucid illustration of these features of research programmes in his discussion of Newton (1970, pp. 133-138).

²¹ It would be interesting to discuss SOC as a Lakatosian research programme since there are many attention-grabbing peculiarities to its 'hard core' and its heuristics. Space constraints prevent me from doing so here.

²² This group of models coincides with what Gibbard and Varian (1978) call 'caricatures' and with what is also known as 'toy models' (Hartmann, 1995). Though both terms quite adequately mirror the character of these models, they have a pejorative connotation I would like to avoid. For this reason I prefer the neutral term 'sketch'.

²³ These points are mentioned in Wimsatt (1987, pp. 28-31).

²⁴ Though most of the time the rhetoric in the writings on SOC is realist, there are sometimes remarks suggesting that scientists actually think about their models in this way: Bak occasionally calls SOC models ‘metaphoric pictures’ (1997, p. 50 and p. 62), or ‘toy models’ (1997, p. 187); and in the introduction to their paper on evolution Bak and Sneppen remark that ‘[a]lthough the model is very abstract, and necessarily oversimplified, we believe that it contains elements representing the mechanism of real evolution’ (1993, p. 4083).

²⁵ Wimsatt (1987) provides a detailed and instructive case study of how biologists, departing from simple and patently false genetic models, have constructed a series of ever better pictures. Thereby he convincingly shows how successful false models can be as a starting point.