

Extracted chapter from:

This Impermanent Eternity:

Scrutinizing the sustainability paradigm

J. Kasmire

December 2, 2014

Chapter 1

Experimenting

Science is a way of thinking much more than it is a body of knowledge.

Carl Sagan Sagan (1974)

This chapter continues by focussing on the third TM cycle step, that of creating and executing transition experiments. As before, this chapter begins with some known knowns drawn from the sustainability paradigm map that frame a known unknown. Transitions theory makes another appearance, this time accompanied by chaos theory as both theories disagree about the differences between radical and incremental innovations and about the relationship between those two types of innovation and a system-wide emergent transition. This disagreement forms the basis of the unknown to be explored in this chapter.

This particular unknown is very important because innovation, and specifically radical innovation, is thought to be the key to successful transitions and therefore key to a successful transition to sustainability. For this reason, transition experiments are typically geared toward increasing innovation, increasing the radicalness of innovations, protecting innovations with radical potential, or all of these at once. Thus, this chapter not only scrutinises the conflicting views on innovation as a key to transitions but also the way that experiments based on those views work as part of a TM cycle.

1.1 What is known and suspected but not known

This chapter begins with myriad knowns plucked from the sustainability paradigm map before proceeding to explain how those knowns encircle a known unknown that matters for sustainability and that relates to the third TM cycle step. The chapter then runs through an entire mini-TM cycle, ending with first order learning that expands or improves the knowns on the sustainability paradigm map and second order learning that reflects on the validity of the entire map in light of the experimental results.

1.1.1 The knowns

Complex causes, complex effects Although the classical paradigm sought, and often found, seemingly simple cause-effect relationships, the CAS paradigm does not look for or expect to find such simple, linear relationships. Instead, complex effects and phenomena, like transitions or radical innovations, are understood to be the result of multiple interacting factors at multiple system levels. Historical analyses and transition experiments work to tease apart these factors, which so far appear to fall into three main divisions. Each of these divisions then lends itself to a category of transition experiments.

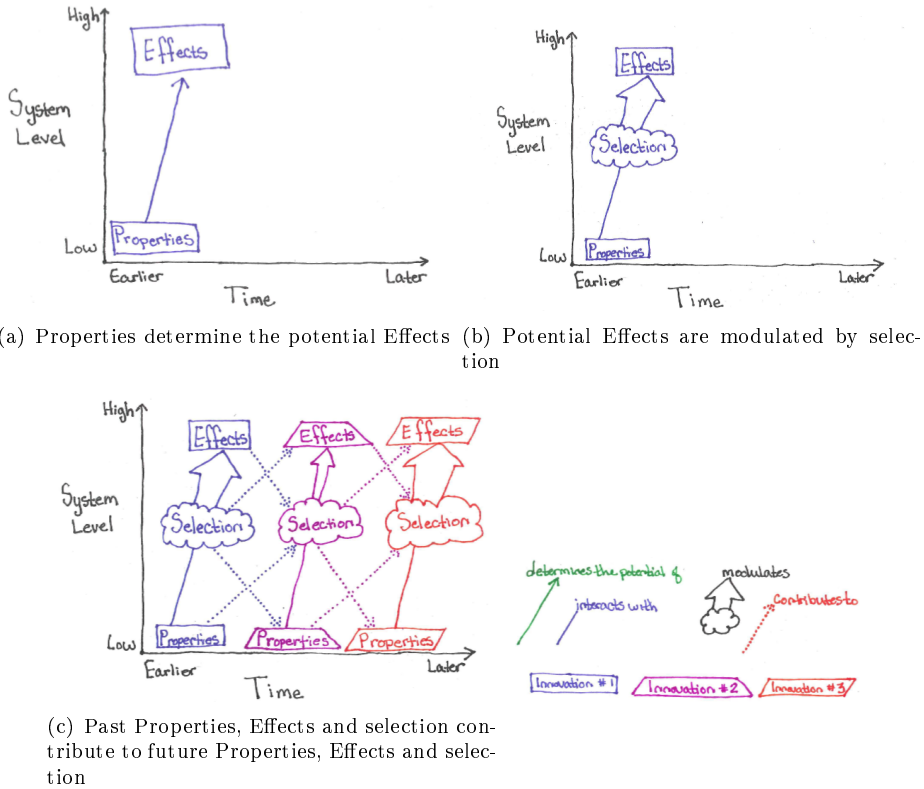


Figure 1.1: The three main complex and interacting factors governing how an innovation's Properties relate to its Effects in a CAS.

First, innovations are emergent phenomena, meaning that their Effects are visible at one level, while the Properties belong to a lower level. This makes an innovation's Properties difficult to examine for the radicalness that causes the Effect of 'jumping' across the adaptive fitness landscape. Transitions are also innovations (Chappin, 2011) but their Properties are high-level enough to be examined. Such examinations suggest that all radical Effects are linked to at least one radical Property (although

radical Properties do not guarantee radical Effects) or that Properties at any system level determine the potential Effects at higher levels (See Figure 1.1(a)). Thus, some transition experiments want to increase the number of radical Properties in order to increase the number of radical Effects, including transitions. These transition experiments are equivalent to producing more ‘innovation seeds’ or more ‘radical innovation seeds’, typically by boosting innovation, boosting the radical Properties of innovations, or recognising radical Properties before their Effects are known.

Next, Properties and Effects are not linked together in a vacuum. Selection pressure acting on innovations is understood to balance variation so that strong selection means less radical innovation while weak selection means more radical innovation. When they first appear, radical innovations are usually ill-adapted, uncompetitive ‘hopeful monstrosities’ because jumping across the fitness landscape sometimes means arriving at a lower elevation. Some of these ill-adapted new radical innovations are quickly wiped out by strong and consistent high-level selection pressures. Others find themselves protected in niches with weaker, inconsistent or at least differently oriented selection pressures. Developing within that protection offers radical innovations a chance to reach the full radical potential of their Properties so that they produce radical Effects. Thus, while an innovation’s Properties determine its potential Effects, those Properties interact with selection to modulate the potential for Effects (See Figure 1.1(b)). Along these lines, many transition experiments use Strategic Niche Management (SNM) or similar models to reduce or alter the selection pressures within niches or sub-systems. These transition experiments are equivalent to observing or creating ideal growing conditions into which the ‘radical innovation seeds’ can be planted, typically by creating or altering subsidies, exemptions or delayed constraints in promising areas or by reducing pressure on front-runners and blue-sky thinkers.

Finally, the selection-variation balance is not fixed or externally imposed but the product of previous innovations-selection interactions. Radical innovations appear to disrupt selection and create windows of opportunity for more radical Effects in the future while incremental innovations appear to bolster selection, strengthening the dynamic equilibria and smothering radical Effects. This makes the innovation-selection balance self-reinforcing and means that transitions or other radical Effects are rarely the result of single innovation with especially radical Properties or specific, especially favourable selection environments. More commonly, radical Effects at any system level are the emergent outcome of multiple, interacting, mutually supportive lower level innovations, diffusions, behaviours, or changes that create feedback loops and other non-trivial behaviour. As such, they only occur when an innovation has sufficiently radical Properties *and* enjoys a sufficiently low selection environment *and* can ride a self-reinforcing wave of disrupted selection pressure, which could explain why radical innovations and transitions are so rare. Thus, an innovation’s Properties and selection environment interact to determine its own Effects, but they are also determined by those of past innovations and in turn determine those of future innovations (See Figure 1.1(c)). Therefore, some transition experiments set out to link up innovations, niches and emergent patterns to create self-reinforcing behaviours. These

experiments might be equivalent to repeatedly planting the best innovation seeds in the best growing conditions to observe how each planting influences the following. Ideally, these transition experiments involve multiple component experiments with different and potentially aligned aims.

Unfortunately, transition experiments have so far failed to produce a transition. The failures are especially clear for experiments reducing or altering selection pressures through SNM which “were certainly over-optimistic” (Grin et al., 2010, p. 104) as they have mostly slowed innovation development, produced uncompetitive innovations, and created subsidy addictions (Moor, 2001; Needham and Faludi, 1999). Although not successful in producing a transition, those experiments have arguably produced the positive outcomes of showing what doesn’t work.

Complex cause, various effects CAS are not only complex, but also chaotic. One important chaotic behaviour is the self-organising criticality (SOC), classically depicted with a sand pile. Roughly equivalent grains of sand are iteratively added to the top of a pile, where most leave the pile apparently unchanged, although rarely some precipitate an avalanche. Avalanche frequency varies as a power of the size, so that there are many small avalanches but very few enormous ones in a power law distribution. There are no general rules about the properties or features of the grains of sand that cause big or small avalanches since most grains of sand cause no avalanche at all, whatever their size. Instead, the outcome of each additional grain of sand depends on the specifics of that particular grain of sand, where and when it falls, and its relation to every other grain of sand that has already fallen.

Power law distributions are seen in other natural phenomena such as earthquakes, solar flares, meteor strikes, forest fires, landslides, and cracks in pavement. These too would seem to be the product of SOC in chaotic systems. More importantly, power law distributions are also visible in many socio-technical phenomena, including patent citations, scientific journal citations, stock market fluctuations, the size of cities, return on investments, firm size, trading volume, executive pay and income distribution (Arthur and Polak, 2006; Bak et al., 1987; Gabaix, 2008; Gabaix et al., 2003). These power law distributed measures are almost precisely the same ones used to identify radical innovations, front-runners, niches ripe for protection or innovators with radical potential, which suggests that radicalness may be a SOC in a chaotic system. Consequently, researchers have begun to investigate radical innovations and transitions in socio-technical CAS as responses to a SOC driver (Arthur and Polak, 2006; Silverberg and Verspagen, 2005).

Observers may detect patterns or relationships in Effects or other radical system outputs, and they may translate those observations into hypothesised causal relationships for predicting future Effects, just as avalanches can be investigated and interpreted to produce causal explanations or predictive models of sorts (Mudge, 1965). Unfortunately, the hypothesised ‘causes’ of recognised SOC outcomes tend to be frustratingly dissimilar, unique, contingent and path dependent while conditions that seem quite similar produce no avalanches or SOC effects. This doubly dissociates the suggested ‘causes’ from the effects and leaving the observed effects a general

explanation and rendering the pattern stubbornly unpredictable (Hough, 2009). The Properties and Effects of innovations are also doubly dissociated, with radical Effects coming from innovations that show no clear radical Properties (Bradshaw, 1992; Bradshaw and Lienert, 1991; Kasmire et al., 2012; Korhonen and Välikangas, 2014; Sarkikoski, 1999) while innovations with apparently very radical Properties producing no Effects at all (Johnson, 2010). Predictions of ‘the next big thing’ in innovation are often equally, but more hilariously, inaccurate (Benford, 2010).

The patterns and hypothetical causal relationships they suggest are observer dependent, which may say more about the observer and the system boundaries they are wont to impose than about the system upon which those boundaries are imposed. With this in mind, it is relatively easy to see radical innovations as jumping across a fitness landscape by simply leaving the connection between the origin and destination outside the system description. Similarly, it is easy to see a sand pile as the *same* sand pile after some grains of sand are added and as a *different* sand pile only after an avalanche, even though the sand pile can also be seen as different after every single addition. The before and after differences when there is no avalanche will seem very trivial to most observers, but chaotic systems are sensitive to ostensibly trivial differences because of the butterfly effect. Furthermore, the system boundaries of irreversible and path dependent systems cannot be compared over time because the exact conditions before an interaction can never be recreated to disentangle its contribution. Imagine trying to exactly reset an entire sand pile to the way it was before an avalanche to see if a different grain of sand would cause a similar slide. It is equally ridiculous trying to remove all traces of a radical innovation, reverse all ageing and learning, and introduce a different innovation to see if it produces the same Effects as the first.

1.1.2 The unknowns

The gap of unknown now begins to materialise amongst the knowns. The continued debate about what are the necessary and sufficient conditions for the rare, rapid and widespread advances that come from radical innovations and transitions and how they differ from the necessary and sufficient conditions for the common, slow moving, incremental change within a dynamic equilibrium reveals a very common and often shared assumption. That assumption is “that these two types of evolutionary change must arise through very different internal mechanisms and from very different underlying principles” rather than both being ‘special cases of a common theory . . . one system with 2 different types of response’ ” (Cohen and Stewart, 2000, p. 334).

This means there are two possible explanations for radical innovations or transitions, but each has very different consequences for transition experiments and other attempts to understand, interpret, predict, influence, or reproduce the phenomenon of interest. If innovations are like seeds, then their radical or incremental potential is already contained within them, even though their ability to reach that potential is modulated by the environment, which is in turn modulated by previous innovations. On this basis, historical analyses and transition experiments are part of the “ana-

lytical component of research into transitions [that] focuses on tracing, recognising and measuring transition patterns-not in the classical, deterministic sense, but in the co-evolutionary sense, making use of recent insights derived from complexity theory” (Grin et al., 2010, p. 121). These pattern-centric analyses hope to untangle the complex causal relationships, at least to some extent, in order to improve the ability to create, shape or manage the Effects, including radical innovations and transitions.

If innovations are like grains of sand, then their potential to produce radical or incremental Effects is not within themselves, nor their immediate environment, nor even nearby interactions ¹. Instead, an innovation’s Effects would be the emergent outcomes of the entire history of the entire system, with no general or meaningful causes in a chaotic, non-periodic, evolving system and no possibility of predicting, influencing, managing or reproducing specific Effects. What’s more, if the SOC and resultant power law distributions are any indication that the system is operating at the edge of chaos, then that system may already be precariously poised at or near the point at which all coupled sub-systems are tuned to each other (Kauffman and Johnsen, 1991), meaning that any interference can only be detrimental. A pile of sand will suffer very little from such detrimental meddling, but a CAS populated by real people has much more to lose, such as its ability to adapt or to maintain advantageous adaptations.

The knowledge gap to explore therefore relates to whether innovations are more like seeds, packed with potential, or like grains of sand, with their potential resting in a system that never repeats. And yet, that knowledge gap goes a bit further than just asking which explanation is more appropriate because each explanation entails different ideas about how to best interact with the system, with one advocating intervention and the other saying that intervention can only make things worse.

1.2 Step One - Transition arena, problem definition and system description

1.2.1 The transition arena

Transitions theory and TM are once again part of the transition arena, this time representing a transition arena participant who detects meaningful patterns in the various historical analyses of past innovations and transitions. This transition arena participant believes that those patterns hint at causal relationships and that transition experiments are a way to examine, explore and better understand those patterns and relationships in order to predict, influence or manage the desired Effects. This transition arena participant endorses the three factors described and depicted above (See Figure 1.1) as the necessary and sufficient conditions for producing a radical innovation or transition and uses these as the basis for designing transition experiments.

In general, transition experiments are learning how to manage the selection-variation balance in a CAS. Tilting the balance toward variation at the expense of

¹Remember that not interacting is a kind of interaction, although it is harder to recognise.

selection is intended to temporarily destabilise and instigate a transition but the balance must be shifted back toward selection once the system has been caught in a sustainable attractor. As part of learning how to manage the selection-variation balance, transition experiments tend to fall into various categories including increasing innovation (encompassing boosting total innovation and boosting radical innovation), protecting promising innovations in appropriate niches, and making connections between transition experiments to build up self-reinforcing behaviours. Given that transition experiments have had quite limited success so far, this transition arena participant would conclude that there must be a quite complex relationship between the various factors and between the conditions and outcomes of transition experiments.

Chaos theory The other transition arena participant is chaos theory, which is related to but not the same as complexity theory (See ??). This participant believes that, contrary to the impression given by history books, innovations Properties are not meaningfully or causally linked to Effects and that the links found in historical analyses are the result of cherry picking, mental gymnastics, subjectivity and a “natural tendency to romanticise breakthrough innovations, imagining momentous ideas transcending their surroundings, a gifted mind somehow seeing over the detritus of old ideas and ossified tradition.” (Johnson, 2010). Since innovation, CAS and SOC are inherently messy, non-intuitive and unwieldy, this transition arena participant puts more store in the results coming from ABM and other ‘in silico’ investigations than in historical analyses, which allow researchers to observe the entire unpredictable, chaotic, evolutionary, and dynamic process without the same biases and distortions that arise from after the fact analyses (Almirall and Casadesus-Masanell, 2010; Arthur and Polak, 2006; Auerswald et al., 2000; Ethiraj et al., 2008; Frenken, 2001, 2006; Gavetti and Levinthal, 2000; Goldenberg et al., 2001; Ma and Nakamori, 2005; Rivkin, 2000; Silverberg and Verspagen, 2005). In these models, incremental changes are not dismissed as unimportant, excluded from the model, or idealised as some sort of equilibrium but are allowed to accumulate, interact, serve as building blocks and foundations for future innovations, and contribute to SOC that occur at some distant point in time, space or system level within the system. Notably, models have found innovations that qualify as radical coming from purely incremental technological development (Arthur and Polak, 2006; Silverberg and Verspagen, 2005), meaning that novelty, change and radical Effects do not need a separate or external cause.

This transition arena member sees transition experiments as trying to create the conditions that preceded a past SOC according to subjective impressions of similar conditions. Pointing to the lack of transition experiment success, this transition arena participant responds with:

Increasing innovation Evolving systems are already working at the limits of their capacity to innovate (Kauffman, 2002), although most observers will cheerfully ignore much of that innovation by labelling it a ‘dynamic equilibrium’. This is not to say that innovation cannot be manipulated, just that manipulations are

unlikely to work as intended could, at least temporarily, reduce innovation².

Increasing radical innovation The power law distributions governing radical Effects are the result of SOC, which is self-organising and scale-invariant, meaning that fiddling about with rewards and punishments will not make a blind bit of difference to the overall ratio of radical to incremental innovations.

Protecting niches Fully closed systems usually go to equilibria and semi-closed systems to dynamic equilibria, so artificially isolated sub-systems are unlikely to produce the desired non-equilibrium behaviour. If variation and selection are allied rather than opposing, so that more innovation entails more selection and vice versa, then semi-isolated niches will only have *different* selection pressures rather than weaker selection pressures. Thus, changing a niche will either suffocate an innovation's development or optimise it to conditions that appear nowhere else outside the niche.

Linking to create emergence Creating links between niches or niche-innovations is contradictory to the idea of protecting them in semi-isolation, but it at its best this idea demands that the emergent outcome of the links and interactions can be anticipated or managed, even though emergence is by definition intractable, unpredictable and surprising.

1.2.2 The problem definition

Now, the transition arena has arrived at a bit of an impasse. Both participants agree that there are two system responses but disagree on number of distinct causes for those responses. Models capable of producing the two distinct responses can be built with one cause or two causes, but the two viewpoints cannot be reconciled into a single system description since they (more obviously than most) necessarily exclude each other. Therefore, instead of blending the system descriptions as a transition arena normally does, I blend the two perspectives in a different way. The model itself is built along the lines of the chaos theory participant and draws on past models based on the NK model underlying the concept of fitness landscapes (Kauffman, 1993) or on lattice percolation models (Silverberg and Verspagen, 2005). Both of these types of models have limitations in the way changes in fitness over time are represented, the way innovations relate to each other, and the representations of complexity and co-evolution. A less common model improves on these by incorporating the cumulative nature of innovations and a variable concept of fitness (Arthur and Polak, 2006), but is inflexible, difficult to program or work with, and complicated to interpret (Arthur, 2009; Arthur and Polak, 2006; Scott, 2010). The ABM system description emerging from this transition arena takes the important features of these models, but with some simplifications to make it flexible, easy to work with and interpretable.

²Don't worry, as Dr. Ian Malcolm in Jurassic Park said, "Life, uh... finds a way."

While the model fits the chaos perspective, the experiments performed on it match the perspective of transitions theory, specifically replicating the various kinds of transition experiments. The simplifications and abstractions of this model mean that the innovation rates, innovation's Properties, or selection pressures can be directly manipulated rather than influenced in the circumspect ways that real-world transition experiments are restricted to. The perspectives of each transition arena participant are therefore accommodated together, limiting the tendency for models to confirm the perspective according to which they were built³ which is a persistent problem of model building. More importantly, this blend of perspectives moves the problem to be addressed away from explaining an observed phenomena (which both perspectives are equally capable of doing) and toward an exploration of whether the actions advocated by one perspective can influence system outcomes in a model that matches another perspective. Essentially, the problem is now about determining the role for transition experiments in a system where innovation is shaped by SOC.

1.2.3 The system description

Environment description Perhaps unusually, I begin with the environment, which is a 'landscape' made of 'peaks' and 'valleys', both of which are lists of integers defined at the initialisation of the model. The peaks represent the needs that drive technology evolution, such as the need for heating, transport, fashionable shoes or mousetraps, but are simpler than the logical operator needs used by Arthur and Polak Arthur and Polak (2006). In the model, as in the real world, these needs represented by these peaks can be satisfied in many ways with new innovations continually trying different and potentially better ways to do so until the need has been well and truly satisfied, which is possible with logical operators or integers, but arguably not possible for fashionable shoes. The valleys represent dead ends, non-functional, problematic or physically impossible technologies. These might include an innovative combinations of two individually useful technologies that work at cross purposes, such as chocolate teapots, glass hammers or tissue paper roofs⁴ or they might include imaginable but physically impossible technologies like perpetual motion machines.

Before initialisation, the modeller specifies a maximum integer, the density of peaks and the density of valleys (See Figure 1.2). When the model is initialised, the peaks and valleys are randomly created according to their respective densities between one and the maximum integer. The same number can be both a peak and a valley, representing a technology that would solve a societal need perfectly but which is not practically possible⁵ and there are no restrictions on whether adjacent numbers can

³Garbage in, garbage out. This famous modelling maxim reminds modellers that building a model that behaves the way you expect doesn't mean that the world works like the model, only that the model represents one possible way that the world could work.

⁴Someone who celebrates Easter might well enjoy a chocolate teapot with a little chocolate bunny peaking out and the winner of the Builder of The Year might get a trophy made of glass or crystal in the shape of a hammer. However, specialised, absurd or gag-gift purposes aside, these are all pretty useless innovations, so are simplified in this model as total dead-ends.

⁵Cold fusion, or communism, perhaps?

both be peaks, valleys, or one of each. Although this model uses the fitness landscape metaphor, the usual fitness landscape concepts of local and global optima do not apply. The landscape peaks can be understood as being equally high, although the way they are clustered or scattered across the landscape will mean that their slopes are not equally steep.




Figure 1.2: A typical landscape set up for these Adder model experiments, with the maximum integer set to 1000, peak density set to 70% and valley density set to 25%, creating 700 peaks and 250 valleys.

Agent description The agents in this model represent innovations or technologies with an arrangement of component parts. Each technology agent has a *structure* made of a fixed number of *components* separated by the arithmetic *operators* $+$ and $-$. The very first technology, called the primitive, has only a single defined component, a one, and a single defined operator, a $+$. During the model operation, which I explain fully below, the primitive’s structure is cloned and then modified by randomly defining (or redefining) one of the operators and one of the components. When an operator is (re)defined, that operator is replaced by either a $+$ or a $-$, but when a component is (re)defined, that component is replaced by the entire structure of any technology currently marked as available for use. This allows new technologies to be built out of previously existing technologies, and for two technologies with all the same component parts to have those parts arranged differently, with consequences for fitness.

Summing over a technology’s structure determines its *product*. The primitive always has a product of one, but later technologies can have higher or lower (even negative) products. A technology’s product is compared to the predetermined peaks in the landscape to find its *nearest peak* and the absolute difference between the product and that nearest peak is recorded as its *distance-to-nearest-peak*. The number of non-zeros that appear in its recursive structure determines its *cost*, which can never be negative. Thus, the primitive always has a cost of one, but other technologies can have much higher costs. Cloned technologies are only incrementally modified versions of their parent, but the small change between a $+$ and a $-$ or between one component and another means that products and costs can differ widely between a parent and offspring (See Table 1.1).

A technology agent’s fitness can also be very different to that of its parent because

Technology agent	Structure	Product	Cost	Fitness
Primitive	+1	1	1	0.25
X	+1-1+1+1	2	4	0.125
Y	+1-1+1+(+1-1+1+1)	3	7	0.107
Z	+1-(+1-1+1+1)+1+(+1-1+1+1)	2	10	0.05
Q	+1+1+1+1+1+1+1	7	7	1

Table 1.1: The structure, product, cost and fitness of example technology agents in a run with fitness landscape peaks at 3 and 7.

fitness is calculated as:

$$\frac{1}{1 + \text{distance-to-nearest-peak}} * \frac{\text{product}}{\text{cost}}$$

This means that fitness can be negative, but never higher than 1. Through the cloning process, the population of technologies can adapt over time, growing progressively fitter by either achieving a product closer to a peak or by reducing cost relative to their product, with the ultimate fitness possible only both product and cost exactly equal a peak.

Technology agents also record whether or not they are in the *active-repertoire*, which can be understood to mean they are currently available for use. At any point in time, the top ten fittest technologies per peak are automatically in the active repertoire because they represent the best available solutions for each societal need. They also record all the technologies used as a component in their own structure in their *contributor-tech-whos*, and when a technology agent is put into the active repertoire, all of its contributor-tech-whos are also put into the active repertoire. This is meant to account for the way technologies do not disappear when a better or more efficient one comes along, but continue to be used, even in new innovations if they are cheap, reliable or ubiquitous⁶. Finally, technology agents also record which others use their structure as a component in their *who-uses-me*, who is their *parent*, and who are their *offspring*.

Standard model operation Now that we know about the environment and agents, let’s take a look at what happens in the model when it operates in standard mode. The landscape of peaks and valleys and the primitive agent are created during initialisation. Then, and at the beginning of every time step, the active repertoire is cleared. Following this, the ten fittest technologies per peak are placed in the active repertoire. If two or more tie for a position in the fittest, the tie is broken randomly. Then, all technologies that appear in the contributor-tech-whos of any technology

⁶Consider how people continue to drive older car models, even when better ones are available, how the parts for those old cars are still produced, and how an inventor might use those old parts when tinkering around in the garage instead of always choosing to experiment with brand new, state-of-the-art parts.

already in active repertoire are also placed in the active repertoire, recursively, until no more that qualify. At the first time step, the only technology agent in the active repertoire is the primitive.

Next, the fittest technology per peak is cloned six times. If there is a second fittest, it is cloned three times, and if there is a third fittest, it is cloned once. Again, any ties are broken randomly. Clones have an exact copy of their parent’s structure, with one randomly selected operator replaced with either + or - and one randomly selected component replaced with the entire structure of another technology drawn randomly from the active repertoire. Although every clone changes one operator and one component, the operator might replace a + with a + or a - with a - for no net change, but there is only a vanishingly rare chance that a component will be replaced by the same component. At the first time step, the only available technology is the primitive, so it is cloned six times. Most of the operator modifications will have no effect on the product or the cost, even if they do result in a change of sign, since most of the operators precede zeros. At this first step, the primitive is the only technology in the active repertoire, so the component change of all offspring will involve inserting the entire structure of the primitive into one of the component positions.

At every subsequent time step, the active repertoire is once again cleared and then filled, this time with some of the offspring as well as the primitive. Now multiple technologies are cloned and there are multiple options in the active repertoire to be inserted as components, meaning that the second generation of clones will be less homogeneous than the first and could approach new peaks as their products diverge. Clones whose product exactly equals a valley are killed off, never getting a chance to be in the active repertoire, to clone themselves or to become a component, but still count as one of their parent’s clones for that time step.

A peak is approached when any technology has that peak as his nearest peak but a peak is only reached when a technology’s product exactly equals the peak. The fittest technology agents for reached peaks continue to produce clones and technologies whose product exactly equals a peak can still be bumped out of the fitness rankings by subsequent technologies whose product also exactly equals the peak, but whose cost is lower. As technologies move down the fitness rankings they stop qualifying for the active repertoire on their own merits and will only remain active if some much fitter technology agent uses their structure as a component part. Occasionally, a technology bumped out of the active repertoire will trigger a wave of obsolescence as many other technologies suddenly no longer qualify.

This standard mode of operation is useful for creating a control experiment that lays out the baseline metrics against which the transition experiments with non-standard operation can be compared. The standard mode also allows a large parameter sweep to explore the general model behaviour and to settle on the final set of parameters that are most interesting for the remaining experiments.

Non-standard model operation The standard model is important, and preliminary testing showed that it produced the desired SOC behaviours for a wide range of initial parameter settings. However, various non-standard modes of operation are

needed to create the experimental cases that mimic the typical transition experiments. Transition experiments use indirect methods of influence but this model allows direct manipulation. In the cases with the **Extra Radical** mode of operation, approximately ten percent of clones have two operator changes and two component changes while the remaining ninety percent having only one of each change as in the standard model operation. Clones with twice as much change will be less like their parents and so can be understood as travelling further across the landscape in a single generation than normal. These ‘jumps’ can be interpreted as more radical Properties which should, according to the transition arena participant representing transitions theory, produce more radical Effects. In the **Extra Innovation** experimental cases, the number of clones produced is increased so that the the fittest technology agent produces ten clones instead of the standard six, the second fittest produces five clones instead of three, and the third fittest produces three clones instead of one. This mimics the desired outcome of transition experiments that seek to increase total innovation, which could also produce more radical Effects. The experimental cases with the **Niche Protect** mode of operation identify a peaks with at least one but fewer than four technology agents that share that nearest peak and apply selection differently. The technologies so identified will all produce six clones each to represent the way reduced selection pressures within newly populated ‘niches’ could allow the better development and/or radical Effects. Finally, the **All Together** mode of operation combines all of the other modes to explore how they interact. Although already tested individually, individual transition experiments are also tested together to allow for mutually reinforcing feedback loops to produce far greater effects.

The metrics There are three main metrics of interest in this model, each of which has two subtly different sub-metrics.

Problem solving success The ultimate goal of transition experiments is to solve the problem of unsustainability. Problem solving success is usually measured in an NK or fitness landscape model by how well randomly placed agents manage to climb the highest peak in the landscape. Peaks are all equally tall in this model and agents are not randomly placed, but a useful alternative would be to measure how well the agents manage to climb all of the peaks available. The recursive nature of the technologies means this model is quite resource intensive, so rather than measure the amount of time needed to climb all peaks a better metric is the percentage of all peaks in the landscape climbed within the fairly arbitrary limit of 30 time steps.

Within this main metric, there are two relevant sub-metrics. The first measures *the percentage of peaks exactly reached* within the time limit, which measures how many of the peaks in the landscape have at least one agent with a product that equals the peak’s integer. However, since peaks can be sandwiched between valleys, or even share their integer with a valley, it is important to measure how many peaks are satisfied in some way, even if not exactly reached. Thus, the second sub-metric is *the percentage of peaks satisfied* within the time limit, which

measures how many peaks are the nearest peak to at least one technology agent, regardless of how close that agent’s product is to the peak integer.

Technology creation Despite the main goal of solving the problem of unsustainability, a common immediate and short term goal of transition experiments is to boost innovation, which is especially relevant for the Extra Innovation cases. The first and most obvious sub-metric here is *the total number of technologies created*. However, since the number of agents created is highly related to the number of peaks available in the landscape, other important sub-metrics are *the number of technologies per peak in the landscape* and *the number of technologies per satisfied peak*.

Radical to incremental ratios Finally, another common immediate and short term goal of transition experiments is to boost the ratio of radical to incremental innovations, regardless of whether the total number of innovations also rises, which is most obviously relevant to the Extra Radical cases. Distinguishing between radical and incremental innovations is not easy, especially as there are no meaningful differences in Properties for most modes of operation. Just as in the real world, Effects are the easiest measures of radicalness. The two sub-metrics of interest here are *the number of technology with six or more offspring* and *the number of technologies used as a component 50 or more times*.

1.3 Step Two - The visions, pathways and agenda

1.3.1 The visions

Transition theory and TM expect that some, if not all, of the experimental cases will produce significant changes in the metrics of interest (See Table 1.2) compared to the control case. Effectively, the transitions theory visions can be interpreted as the alternative hypothesis in that the experimental manipulations are expected to matter. For example, the Extra Innovation case is expected to produce a greater number of innovations while the Extra Radical case is expected to produce a higher ratio of radical to incremental innovations. The All Together case is expected to show the greatest changes, especially in the number of peaks satisfied.

Experimental case	Problem solving success	Metric	
		Technology creation	Radical/incremental ratio
Extra Radical	Slight increase	Unclear	Significant increase
Extra Innovation	Slight increase	Significant increase	Unclear
Niche Protect	Slight increase	Slight increase	Unclear
All Together	Significant increase	Significant increase	Significant increase

Table 1.2: The visions for each metric expected in each non-standard mode of operation as compared to the standard according to the transition arena participant representing transitions theory or TM.

All cases are expected to increase the number of peaks satisfied within the time limit, but direct or linear relationships are not expected. For example, a large increase in total innovation compared to the standard mode of operation might only lead to satisfying a few more peaks. Although there are good reasons to look at the number of peaks satisfied within the time limit, some experimental cases, like those created in the Niche Protect mode, are expected to show more effect on the number of peaks reached than on those satisfied. The niche protection is meant to give fledgling innovations a chance to develop in ways that allow them to competitively meet societal needs, which suggests that the Niche Protect experimental cases will be particularly good at reaching peaks once they have been approached but not necessarily superior at approaching peaks compared to other experimental cases.

Chaos theory does not have such high hopes because this vision is, more or less, aligned with the null hypothesis, where the experimental manipulations are not expected to matter significantly (See Table 1.3). For example, the Extra Radical experiment is anticipated to impact on the radical to incremental ratio for two reasons. First, technologies are unlikely to share a nearest peak with their parent after the first few time steps. Even in the purely incremental cloning of the standard model operation, offspring quickly spread across the landscape with no need for extra ‘jumping’ afforded by extra-radical Properties. Secondly, if the ratio of radical to incremental innovations is governed by a SOC, it will conform to a scale-invariant power law distribution and will remain the same for 1000 agents or 100,000 agents.

Most of the metrics are expected to be shaped in some way by SOC, meaning that they are likely to arise regardless of differences in initial conditions. The insensitivity to initial conditions of SOC may seem at odds with the sensitivity to initial conditions that chaotic systems are known for, but this conflict is resolved by considering specifics and generals. Simulation run specifics, such as the number of generations needed to reach a given peak or the timing and size of waves of obsolescence within the active repertoire, are low-level details which makes them sensitive to initial conditions and will vary greatly. Simulation run generals, such as the ratio of agents with radical and incremental Effects or the number of peaks satisfied, are higher level details which makes them insensitive initial conditions and unlikely to vary.

The number of technologies created within the time limit is the only low-level metric, making it the only one unlikely to be governed by a SOC and the only one likely to vary much at all. However, the more technologies, the more likely they are to land on a valley or exactly replicate the cost and product of another. Related to this, this transition arena participant expects that more variation will produce more selection, rather than less, and that less selection will produce less variation. Importantly, this means that neither a greater number of innovations, nor a reduction in selection for the Niche Protect cases, are expected to change higher level metrics at all.

Experimental case	Problem solving success	Metric	
		Technology creation	Radical/incremental ratio
Extra Radical	No effect	Slight increase	No effect
Extra Innovation	No effect	Slight increase	No effect
Niche Protect	No effect	Slight increase	No effect
All Together	No effect	Slight increase	No effect

Table 1.3: The visions for each metric expected in each non-standard mode of operation as compared to the standard according to the transition arena participant representing chaos theory.

1.3.2 The pathways

Specific paths cannot be prestatated because the adaptive population of technology agents is free to wander freely across the run-unique landscapes under the influence of the chaotic and random elements. However, some general pathways can be imagined because this model has a (semi-)closed and finite system state space. Within that system state space, the model is expected to moves from a State of one peaks satisfied (but not reached) to a State of almost all peaks satisfied (with most reached), which is almost surely an attractor created by model operation rules. Thus, this model is expected to follow a classic transition pathway as the system output falls into the attractors within the single state space.

That envisioned pathway can be further elaborated. The combinatorial explosion built in to this model means that initially, agents will have a competitive advantage for being the first to reach a new, unexplored peaks, regardless of how inefficient their product to cost ratio is. This would probably qualify as a first mover advantage. After discovering and populating a peak, the fitness advantage falls to more efficient agents with a lower cost for a given product. The system output is therefore expected to move from a State in which fitness is mostly a matter of first mover advantage to a State where fitness is mostly a matter of efficiency, measurable by a gradual decrease in average cost. Were fitness to be judged differently, for example with a penalty for the depth of recursion in the structure, then the pathway would have been expected to move from a State of first mover advantage to a State of shallowest possible structures. Any number of other possible ways to judge fitness would produce any number of equivalent envisioned pathways which would be shared equally by both transition arena participants.

However, the two transition arena participants are not in total agreement on all aspects of imagined pathways. The transition arena participant represented by transitions theory or TM understands incremental innovation as refinements of existing technology, adaptations that improve an innovation’s ability to solve a problem or meet a need, or efficiency gains. Thus, this transition arena participant would envisages that technologies will be nearer to the peak that their parents are near to and that parents will be booted out of position by their own offspring. If the fittest agent for a given peak is examined over time, this transition arena participant would expect the fittest technologies agents to be clones of the formerly fittest technologies except

for when a valley lies between a parent and a peak. The chaos theory transition arena participant, sees no useful distinction between radical and incremental innovations in terms of their Properties, does not expect offspring to be competing on the same peak as their parent (after the first few time steps), and does not expect the fittest agents for a given peak over time to show any family lines.

As an interesting aside, the system state space is closed and finite, with no capacity for new dimensions or selection pressures to redirect the system output. Closed systems normally go to a dynamic equilibrium, but this model might not. If left to run as long as needed to reach all peak with total efficiency⁷, then there would still be SOC in the size of the active repertoire, and maybe in other system output. Testing this is out of the question because of the heavy resource use of this model. Were it possible to show, such a result would very strongly suggest that within even very simple models, with the strongest possible competition and the least room for variation, system behaviour resists being idealised or simplified as a dynamic equilibrium, but remains strongly chaotic.

1.3.3 The agenda

The agenda is really only applicable to transition experiments that have a clear event horizon to avoid. As with most ABM or alternative transition experiment platforms, this model has no real place for an agenda, and the transition arena therefore has nothing much to say about it.

1.4 Step Three - The transition experiments

As this chapter focusses most strongly on transition experiments, it should come as no surprise that there are more of them in this chapter than in any other. The first is, of course, a control experiment run with the standard model operation over a range of parameters to create six experimental cases (See Table 1.4). This control experiment provides a baseline comparison for the remaining experiments with non-standard model operation.

1.4.1 Experiment one - Control

The parameters

The results In general, the model successfully replicates the behaviour of the logic circuit model on which it is most closely based (Kasmire et al., 2012) and displays signs of SOC. The first peaks are reached fairly quickly by technologies whose structure contains only the primitive or first generation technology components while later peaks are reached by technologies whose structure contains technologies created throughout the run. Mapping the parent-offspring network of a typical run reveals an organic

⁷One might even call such a situation peak-peak.

Parameter	Value	Justification
Time steps	30	Surprisingly short, due to combinatorial explosion and intensive resource use
Repetitions	10	Surprisingly few, also down to the resource use
Max integer	1000	A nice, round number
Density of peaks	25, 50, 75	Early tests showed that extreme values prevented successful runs
Density of valleys	25, 30	Allow one experimental case with valleys > peak
Components in structure	10	A nice, round number
Competition level	3	Standard setting, governs how many agents per peak produce offspring
Offspring production	6, 3, 1	Standard setting, governs the rankings for offspring production
Change to structure for clones	Incremental	Standard setting, all offspring have a change to operator and 1 component
Niche protection	Off	Offspring production does not depend on how many agents share a peak

Table 1.4: The parameter settings marking out six distinct experimental cases within the control experiment.

structure (See Figure 1.3), with offspring and uses per technology apparently following power law distributions (See Figures 1.4(a) and 1.4(b)).

The active repertoire grows as the simulation continues, although it does not grow in a simple or linear way and instead undergoes crashes that suggest its size could be governed by a SOC (See Figure 1.5(a)). The number of peaks develops a classic S-shaped curve over the course of the 30 time steps (See Figure 1.5(b)), meeting the expected pathways of both transition arena participants. Cost efficiency starts slowly but later also shows a classic S-shaped curve, again meeting the expected pathways of both transition arena participants, although not within the 30 time steps of a normal run (See Figure 1.5(d)).

With more valleys, the first peak integers sometimes coincided with valley integers or were blocked by valleys, which derailed the recursive technology build out process and ‘crashed’ the run. Over half of the runs with a valley density of 30% crashed, showing spectacularly poor problem solving success, very low technology creation numbers and very unusual radical to incremental ratios. Although this outcome is further confirmation that the Adder model is successfully replicating important features of the models on which it is based, the crashed runs are very disruptive to the metrics. Thus, the cases with 30% valley density will be left out of the remaining analysis and will not be used in the rest of the experimental cases.

The remaining cases with 20% valley density provide the baseline behaviour for each metric against which the remaining cases will be compared (See Table 1.5).






Figure 1.3: By the tenth time step, the agents in a normal run have formed a very organic looking networks. The primitive agent is highlighted in green and one particularly fit and prolific agent is highlighted in pink.



(a) Offspring per parent agent




(b) Uses per agent


Figure 1.4: Both offspring per parent technology (1.4(a)) and uses per component technology (1.4(b)) show power law-like distributions.




(a) Active repertoire size



(b) Number of peaks reached



(c) Average technology cost within 30 time steps



(d) Average technology cost over longer time frame

Figure 1.5: The active repertoire grows non-monotonically with some size crashes, reminiscent of SOC (1.5(a)) while the number of peaks reached grows in a classic logistic curve (1.5(b)). Average technology cost does not reduce within the 30 time step limit (1.5(c)), but does when left to run for much longer (1.5(d)).

The three cases within this control experiment lay out a range of results for each metric. Unsurprisingly, the cases with the fewest problems were the most successful at solving those problems within the 30 time step limit, although the cases with the most problems were not the least successful at solving them. The technology creation metrics and the radical to incremental ratio metrics also show a range of results, with total technologies created clearly related to the number of peaks in the landscape. The most interesting of these metrics is the exceptional radical to incremental ration according to the number of technologies with 6 or more offspring relative to those without in the case with a peak density of 50%. A closer look at the individual runs in this experimental case revealed a single run with an astonishing ratio of 4.23 prolific parents relative to all other technology agents. Although this specific run appears to be an outlier, it is quite an interesting one since some transition experiments are specifically interested in increasing increase the ratio of radical to incremental innovations.

Finally, examining the fittest technologies per peak over time in some of the specific runs that did not crash answered the disagreement between the transition arena participants regarding success within family lines. The fittest technologies per peak were the offspring of the previous fittest technologies for that same peak only for the peaks nearest the primitive agent and typically only for the first few time steps. The fittest technologies per peak for all other peaks and throughout most of the simulation were not direct clones of the previously fittest technologies for that peaks. This supports the chaos theory vision and the idea that incremental but combinatorial innovation does not necessarily entail incremental improvements toward the same purpose or goal as the original.

Based on the standard mode of operation in this control experiment, and mindful of the resource intensive nature of this model, I decided to avoid using a 30% valley density and a 75% peak density in all future experiments. The remaining experiments

Metric	Sub-metric	% peak density 25	% peak density 50	% peak density
Problem solving success	% reached	81	55	72
	% satisfied	97	63	77
Technology creation	Technologies total	39,048	68,605	98,731
	Technologies/landscape peak	156	137	132
	Technologies/satisfied peak	161	216	171
Radical/incremental ratios	Offspring	0.049	0.622	0.06
	Uses as component	0.028	0.028	0.014

Table 1.5: The mean and standard deviation of peaks reached and approached per case.

with non-standard modes of operation will thus focus on a narrower set of parameters that cover only peak densities of 25 and 50% combined with 20% valley density.

1.4.2 Experiments two through five - Manipulating innovation in a complex system

The parameters The first three non-control experiments have all but one of the standard settings used in the control experiment or motivated by the results of that control experiment. The final non-control experiment continues to use many of those same standard settings, but includes all three of the individual changes used in the other non-control experiments.

The results The non-control experiments also experienced some crashed runs. Although these crashes influence the calculated metrics, they were regularly distributed between the non-control experiments and were comparable in frequency to the crash rate of the control experimental cases with 20% valley density. This suggests that crashed runs are a part of normal model behaviour, and so were not excluded from

Parameter	Value				Justification
	Extra Innovation	Extra Radical	Niche Protect	All Together	
Time steps	30	30	30	30	Same as co
Repetitions	10	10	10	10	Same as co
Max integer	1000	1000	1000	1000	Same as co
Density of peaks	25 and 50	25 and 50	25 and 50	25 and 50	Based on r control
Density of valleys	20	20	20	20	Based on r control
Components in structure	10	10	10	10	Same as co
Competition level	3	3	3	3	Same as co
Offspring production	10, 5, 3	6, 3, 1	6, 3, 1	10, 5, 3	Increased production
Change to clone structure	Incremental	Radical and incremental	Incremental	Radical and incremental	Increased d between pa clone
Niche protection	Off	Off	On	On	Reduced for underp peaks

Table 1.6: The parameter settings for all non-control experiments, each containing two experimental cases.

the analysis. The results for all experiments, including the control, are presented here, grouped by the metric for easier interpretation.

Problem solving success metrics Success at solving (or nearly solving) societal needs and problems is the ultimate aim of transition experiments, so the number of peaks reached and satisfied is the most important set of metrics (See Table 1.7). Interestingly, there was no clear positive effect on either sub-metric from any of the non-standard modes of operation as compared to the control when the peak density was 25%. In fact, most of the transition experiment cases did worse than the control. At the same time, all of the non-standard modes of operation did better than the control when the peak density was 50%. However, there was little to choose from between them and none of them did better than the high water mark set by the 25% peak density control case.

Interestingly, the transitions theory visions expected the Niche Protect and All Together modes of operation to do particularly well at reaching more peaks compared to the control, but this was not shown to be true. The Niche Protect case with 50% peak density tied for first and the All Together case tied for third among the cases with 50% peak density, while the Niche Protect and All Together cases were the worst and second worst cases of those with 25% peak density.

Technology creation metrics Both of the transition arena participants expected that the technology creation metrics would show an effect from the various experimental cases and these expectations were certainly borne out by the results (See Table 1.8). As predicted, the Extra Innovation and All Together cases showed the largest and clearest effects, increasing drastically for all of the technology creation sub-metrics. Most of the other experimental cases performed roughly the same on these technology creation metrics when compared with the control experiment, except for the Niche Protect experiment which showed a slight increase for the number of technologies created per landscape peak. Allowing the fittest agents to produce more clones (either in general or within the protected niches) leads to greater innovation in total, even if this does not produce clear or straightforward benefits for problem solving.

Radical to incremental ratio metrics Finally, the various metrics tracking the ratio of radical to incremental innovations also failed to live up to the expectations

% peak density	Metric	Control	Extra Innovation	Extra Radical	Niche Protect	All Together
25	% reached	81	80	81	71	71
	% satisfied	97	97	97	88	88
50	% reached	55	72	80	80	80
	% satisfied	63	82	91	91	91

Table 1.7: The various metrics covering problem solving success for all experiments and cases.

% peak density	Metric	Control	Extra Innovation	Extra Radical	Niche Protect
25	Technologies total	39,048	116,602	36,902 36,626	68,705
	Technologies/ landscape peak	156	466	148	147
	Technologies/ satisfied peak	161	480	153	167
50	Technologies total	68,605	135,910	70,729	79,369
	Technologies/ landscape peak	137	272	141	159
	Technologies/ satisfied peak	216	332	155	198

Table 1.8: The various metrics covering technology creation for all experiments and cases.

of the transition arena participant representing transitions theory (See Table 1.9). This transition arena participant expected to the radical to incremental ratio to be a manipulable value and also expected to see the greatest manipulation in the Extra Radical and All Together cases where some technologies were created to be less like their parent than normal.

Remember that metric for the number of technology agents with six or more offspring compared to all other technology agents had one particularly anomalous run in the case with 50% peak density that distorted the average value. For this reason, all of the other experimental cases with 50% peak density are far lower than the control case. Setting this anomalous case aside and comparing all of the other cases to each other, there are still no clear patterns suggesting that any of the non-standard modes of operation have any significant effect on the radical to incremental ratio according to offspring. Neither were there any more anomalous runs that might hint at how an unusually large number of radical innovations might be created. The same is generally true for the other sub-metric, the one comparing the number of technology agents used as a component 50 or more times. Although there is a bit more variability for this sub-metric, there are still no clear patterns suggesting any consistent or meaningful effects from any of the transition experiments.

Another interesting result is that the two sub-metrics did not agree with each other in relation to control experiment. Many cases showed that a rise in one sub-metric could be accompanied by no change or a fall in the other sub-metric for the same case. Further, comparing the various non-standard modes of operation showed that a consistent change for one peak density was not reflected in the other peak density.

All of this suggests that the two sub-metrics are unlikely to be measuring the same thing, even if they both seem to be measures of an innovation’s radical Effect.

1.5 Step Four - Learning

1.5.1 First order learning

The new knowns show that very few of the visions and expectations of the transition arena participant representing transitions theory or TM were met. Those that were met were also shared by the transition arena participant representing chaos theory, such as the expectation that the system would transition from a State of no-problems-solved to a State of most-problems-solved as the system output discovered and fell into the attractor created by the model operation rules. However, the results generally favoured the visions and expectations of chaos theory. This does not definitively speak to the likely outcomes of transition experiments in the real world, because these results do not prove that transitions or radical innovations are actually governed by a SOC. However, they are very interesting in that they offer one possible explanation for why transition experiments have not yet produces any clear or significant signs of the desired effects.

More specifically, the first new known that supports chaos theory is that **purely incremental development does not mean incremental progress toward a solution on a given problem**. This new known comes from examining the fittest agents per peak over time in the control experiment, which showed that technologies continuously bumped each other out of the top ranks, but that the fittest agents were not replaced by their own offspring. Instead, the combinatorial nature of technological development meant that (re-)defining a single operator and component allowed technologies to take steps of variable length, in effect jumping all over the fitness landscape.

The next important new known is that **the low-level features of a system are potentially malleable**, as shown but the successful increases in the technology creation metrics for many of the experimental cases. Unfortunately, another of the new knowns is that **the emergent, mid- and high-level features of a system are resistant to change**, regardless of whether these features are targeted directly or indirectly through lower level changes, as show by the lack of changes seen in the radical to incremental ratio and problem solving success metrics. This suggests two

% peak density	Metric	Control	Extra Innovation	Extra Radical	Niche Pro
25	Offspring	0.05	0.04	0.04	0.05
	Uses ascomponents	0.03	0.02	0.02	0.06
50	Offspring	0.62	0.05	0.05	0.06
	Uses ascomponents	0.03	0.04	0.02	0.02

Table 1.9: The various metrics covering radical to incremental ratios for all experiments and cases.

things. First, that unlike the low-level features, emergent features of interest may be governed by SOC that reliably arise over a wide range of initial conditions and that consistently produce the same distributions and outcomes. Second, this also suggests that variation and selection are not balanced against each other in a zero sum game, but that an increase to one could be an increase to the other.

Yet another of the new knowns is that **changes to the emergent, mid- and high-level features of a system are often detrimental**. Not only are these system features resistant to change, but any change that does take place is only likely to push them away from the edge of chaos where they operate most effectively. This suggests that transition experiments aiming to improve on these mid- or high-level system features, such as the Extra Radical case, are more likely to disrupt the SOC governing the creation of radical innovations and make the situation worse.

Another new known to emerge from these experiments is the idea that **radical Properties are not linked in any meaningful or causative way to radical Effects**, as shown by the poor radical to incremental ratios shown in the Extra Radical and All Together cases that included some technologies with more radical Properties. Related to this is the new known that **identifying radical Effects is highly dependent on the selected metric**, which suggests that the various concepts grouped together under the category of Effects are actually independent as well as observer dependent.

The final new known to come out of these experiments is that **cooperative or mutual reinforcement between various experimental interventions cannot be predicted**, as most obviously shown in the fact that the All Together case performed worse on most of the metrics than the individual intervention most closely associated with that metric and sometimes performed worse than the control experiment with no intervention. This does not mean that interventions or system manipulations are incapable of interacting to form self-reinforcing behaviours, just that it may be much more difficult than expected to design interventions or manipulations that do so.

The new unknowns With every batch of new knows comes some new unknowns, some of which are quite problematic. For example, the idea that variation and selection are not equal and opposite forces means that a new unknown surrounds whether **attempts to boost innovation might be counterproductive by boosting selection at the same time**. If more innovation only means more competition, then it will produce the same number of winners relative to losers, but more losers in total. Related to this, another new unknown relates to whether or not **attempts to decrease, redirect or weaken selection pressures can only decrease, redirect or weaken innovations**. This might explain why SNM, subsidies, and other schemes designed to subsidise, protect or support innovation by reducing selection have mostly produced subsidy addiction and poorly developed, uncompetitive innovations.

The subjective nature of deciding whether any given innovation is radical or not is a well recognised problem, but the knowns from these experiments suggest that ‘radicalness’ is not a property of the innovation at all, nor even the system in which it is found. Instead, radicalness may only be an entirely subjective and relative im-

posed as the result of observer dependent choices of measures applied within observer dependent system boundaries, like ‘rich’ or ‘old’. This means that another unknown to explore relates to whether or not **Properties or Effects are entirely observer dependent patterns with no important relation to the systems in which they are observed.**

The one of the most important new unknowns to come out of these experiments is whether **innovations are more like seeds or snowflakes** or whether transitions theory or chaos theory has a more accurate view of the real-world. Of course, the answer does not have to be so binary. Causal relationships can co-exist alongside SOC, each of which might be more or less important in different (sub-)systems or at different points in time. However, the most important new unknown to emerge from these experiments is perhaps whether or not **we are prepared to accept explanations without clear or meaningful causal relationships and that cannot be used to improve influence or control?**

Practical application and policy recommendations for the Westland-Oostland Greenport These new knowns and unknowns have some important policy applications. The first of these is that attempts to increase the amount of innovation in a system may be moderately successful, but attempts to manage or increase the proportion of radical to incremental innovations or the number of problems that are solved are unlikely to be successful but likely to be detrimental. For example, local governments could easily spur innovation in the greenhouse horticulture sector by funding all proposed innovation projects, but could only expect a small minority of them to produce radical Effects. However, a proposal to fund only ‘radical’ innovations would have to find a way to predict radical Effects before they can be observed, when subjectively determined radical Properties are not clearly linked in any meaningful or causative way to the equally subjective radical Effects. To make matters worse, such a policy is likely to disrupt the normal mechanisms governing innovations and might end up wasting time and money by stalling development of the targeted technologies and potential competitors, fostering subsidy addiction, and creating system lock-in, as seems to be the case with CHP technologies (van der Veen, 2012). Thus, policy-makers should avoid discriminating between innovations solely on the basis of how radical their Properties seem to be and should instead embrace policies that support all innovation or innovations that have already demonstrated Effects. Echoing a practical recommendation of the last chapter, policy-makers are also advised not to assume that various measures of radical Effects are interchangeable or that they actually measure any objective quality of the innovation. Also related to this issue, policy-makers need to consider whether the costs and benefits of boosting innovation are adequately distributed. Encouraging innovation might lead to risks falling unfairly on those least able to withstand them but who had been encouraged to innovate when otherwise would not.

One problem of fitness landscape models is that imply that agents solve exogenous problems in a fixed solution space and in a static environment, while every solution in the real world brings new problems to solve (Johnson, 2010; Kelly, 2010). Like the

proportion of radical to incremental innovations, the proportion of problems solved or nearly solved to those completely unsolved seemed very inflexible, resisting all changes that were not obviously negative. Both of these issues could be governed by a SOC, already operating at the limits of their capacity, suggesting that any interference can only be negative. A Westland-Oostland Greenport would be the way that solving the problem of uncertain crop production due to temperature fluctuation has created the new problem of reliance on dwindling fossil fuels. Problems seem to just shift around within a system, from being the problems of local crop producers or consumers to become the problems to global residents. Thus, policy-makers are recommended to consider how problems and problem solving efforts are distributed, both throughout the system and over time, when deciding how to approach new problems and solutions.

More practical issues arise from the evidence that emergent phenomena cannot be predicted nor engineered. For example, one policy might require all growers in the Westland-Oostland Greenport to install water storage capacity to mitigate flood risks while another encourages those same growers to employ water storage as a heat buffer to CHP units more effective and flexible. At first glance, the two policies seem to push for the same thing, increased water storage, and so should reinforce each other and lead to much greater water storage capacity than either policy alone. However, flood risk water storage cannot be easily be used to store heat while both types of storage compete for investment and space, forcing growers to cut corners, comply according to the bare minimum, spend more money than originally considered, or to avoid meeting the policies as much as possible. It doesn't seem like they should, but the two policies conflict and work at cross purposes to reduce the effectiveness of both. Policy-makers can not predict the interactions of all their policies⁸, but they do need to be open to detecting problems in unexpected quarters and seeking advice widely to understand who is most affected and when. Importantly, policy-makers should be aware that the best intentioned interventions can still produce negative outcomes, that those negative outcomes cannot be predicted, and that the best they can hope for is to react appropriately when those outcomes become apparent.

Finally, some of the new knowns and unknowns from this mini-TM cycle do not have such clear practical applications, but are nevertheless important for policy-makers to consider. Many policies and plans are based on the idea that effects, either desirable or not, have a clear cause that can be encouraged or avoided as needed. The proposed cause-and-effect relationships seem intuitively useful and are very appealing, not least because they promise a certain degree of prediction and control. For example, governments may want to copy the success of some policy in a new context and so will look for the crucial causative factors that drove the first success. Unfortunately, wanting a given effect to have a reliable, reproducible and predictable cause does not make it so. Policy-makers are therefore advised to consider that not all effects have causes, in the way that causes are usually understood, and that finding a cause does not guarantee it can be manipulated to control the desired effects.

⁸Policy-makers surely know that two wrongs do not make a right, but they may not consider that two (or more) rights do not always make a right.

1.5.2 Second order learning

Reflecting on the third step The results of the first order learning are then applied reflexively to the entire TM cycle. As this chapter focusses on the third step of the TM cycle, so too does the second order learning. Thus, the role of transition experiments in advancing the sustainability paradigm and its paradigmatic map come under the metaphorical microscope of second order learning.

Experiments are a scientific way of figuring out the ‘hows’ and ‘whys’ behind observed phenomena, and they rely on pre-identifying and measuring, if not also controlling, the important variables thought to be at play in the phenomena of interest. Experiments also depend on that phenomenon being reproducible and on that phenomenon being influenced by or otherwise interacting with some variables, although not necessarily those identified ahead of time. For this reason, mental operations and other totally (at the time) unobservable events within people’s minds were long considered inappropriate for scientific experimentation (Skinner, 1945). If some phenomena is considered suitable for experimentation, then repetition and variable control are necessary to allow statistical analysis to eliminate the errors, imprecisions, and individual vagaries and reveal the underlying causes and regularities. Thus, transitions and other high-level innovation Effects are generally assumed to have causes that can be discovered as well as influenced, managed or reproduced, even if those causes are not easily observable or currently understood. However, second order learning demands that these underlying assumptions must be questioned.

First, the assumption that important variables should be identified ahead of time and then measured and/or controlled is problematic for SOC and other chaotic behaviours. The factors that turn out to be important for any given SOC can only be identified after the fact, at which point control is totally impossible and even precise measurement is difficult. Even this post-mortem identification of important factors is complicated by the role of observer dependence, which tends to emphasise factors that were present and noticeably changing over those that were important by their absence or constancy. This explains why innovations with radical Effects are so often assumed to have radical Properties, even when those Properties appear for all the world to be totally incremental while innovations with seemingly radical Properties that fail to produce radical Effects are generally ignored as inconvenient (Benford, 2010).

Next, the idea that experiments need to be reproducible is important for statistical analysis as well as for scientific review from peers, but SOC are rare, follow power-law distributions, and cannot be reproduced on demand because the ‘causes’ identified after the fact tend to be very specific and context dependent. This is a problem for the traditional statistical analysis which relies on normal distributions and measures of central tendency to separate out the errors, imprecisions and vagaries from the meaningful and influential factors. Without these traditional statistical analysis tools, SOC defy all description, refuse to obey any rules, and remain doggedly unpredictable. The specificity, rarity, rule defying behaviour and unpredictability is also a problem for scientific peer-review, which finds it impossible to reproduce the necessary conditions.

Those conditions associated with the SOC are as unique and unprecedented as those associated with ‘normal’ behaviour, all of which are impossible to reproduce. However, attempts to reproduce those conditions must decide what aspects are important to control and what are not, thus bringing observer dependence back in and allowing lots of differences as long as the outcomes are ‘close enough’. All of this could help explain why transitions and radical Effects occur spontaneously throughout the world and throughout history but have so far been impossible to produce through deliberate effects and why a diffusion or transition that succeeded in one sub-system seems impossible to replicate in another.

And finally, experiments are about elucidating assumed cause-effect relationships, and so are ill-suited to phenomena that do not have cause-effect relationships to divine. Chaotic systems are still deterministic, so SOC and other chaotic behaviours do have causes, but the causes of emergent effects involve every component and interaction (or lack of interaction) throughout the entire history of the entire system. Thus, the causes are, effectively, the whole of the system itself. This could explain why innovations with truly radical Properties, such as Mendel’s genetic ideas or Babbage’s proto-computer (Johnson, 2010) have proved themselves genuinely useful to modern eyes but were utterly ignored in their own time. Their value comes from the entirety of the systems in which they are embedded, which made them incapable of success in a system without the necessary supporting components (multiple observations of discontinuous inheritance for Mendel’s genetics and mass-production, generally accessible education and ready energy supply for Babbage’s engines).

Without these assumptions, many of the knowns associated with transition experiments began to fall apart. There is no longer any reason to link radical Properties to radical Effects through an innovation because those Properties can no longer be considered objective qualities of the innovation nor can they be considered meaningful ‘causes’. Indeed, the differences between radical and incremental become meaningless when Properties are nearly impossible to investigate and when Effects belong to the entire system more than they belong to the innovation. Likewise, there reproducing desirable Effects loses meaning when all Effects are totally and utterly unique or when SOC ensure that some Effect of equal magnitude will happen in due course without any effort at all.

Most challenging of all, there may be no reason to continue searching for cause-effect relationships behind innovations, diffusions, transitions and the other high-level system behaviours of interest. These behaviours may not have the kind of ‘causes’ people usually think of as causes, which seems unsettling because the alternative is generally understood to be randomness, or effect with no cause at all. However, this is a false dichotomy. There are many non-random mechanisms for change that do not have clear cause-effect relationships, such as evolution. These mechanisms are unpredictable, unreproducible and unmanageable, but that unpredictability, unreproducibility and unmanageability is not a problem to solve so much as an just the way things are. Stubbornly continuing to search for causes when there are none because of ideological biases could be preventing any progress toward actual understanding of how things are, as with those who reject evolution as an explanation for how things

are simply because it does not meet their expectations of what such an explanation should be like. More worryingly, a stubborn search for non-existent causes of a given outcome may be making that outcome less likely.

Reflecting on the map of knowns and unknowns Transition experiments set out to find cause-effect relationships and so the results are interpreted as supporting or failing to support *particular* cause-effect relationships, but not as supporting or failing to support the search for cause-effect relationships. At its worst, this is a case of the same ‘garbage in, garbage out’ principle that plagues model building whereby assumptions shape the tests and the tests confirm the assumptions. The reliance on pre-identifying important factors, reproducing desired phenomena and discovering cause-effect relationships that underlies transition experiments harks back to the classical paradigm. This might be surprising given that complexity and chaos “challenged Newtonian determinism and destroyed the beliefs in control and prediction, emphasising the end of certainty and strongly criticising the reductionism approach” (Grin et al., 2010, p. 138). CAS and its associated concepts successfully opened up new topics for study and “influenced many other research fields with insights on our limited understanding of the world and on how to deal with structural uncertainties . . . but it has not yet delivered a well-grounded and empirically tested new paradigm” (Grin et al., 2010, p. 138). For many, complex or non-linear systems must still be approached through the same basic experiments of the classical paradigm with a thin veneer of CAS applied over the top. The classical paradigm goals of total control and prediction became the CAS goals of ‘improved influence’ and ‘reduced uncertainty’. Supply networks, for example, are generally recognised as complex systems, but the literature on their management continues to emphasise increased predictability and control through the creation of negative feedback loops to reduce their dimensionality, complexity, dynamics, and emergent properties (Choi et al., 2001).

Whether or not CAS is recognised as a new paradigm at this point in time is immaterial for a couple of reasons. First, those living through a scientific revolution are not well placed to identify the paradigm shift, nor to spot which ideas and concepts will get the boot, which will remain with modifications, and which up-and-coming ideas will come out on top (Kuhn, Thomas, 1970). Second, scientific revolutions happen at all scales. Even one scientist or scientific field that feels a CAS revolution means that CAS can be called revolutionary. Exactly how many scientists or scientific disciplines have to share that view in order for it to be inarguably classed as a scientific revolution is unclear, and probably unimportant. That number will surely change over time, so again, the revolutionary nature of CAS is best left for history to judge. But the third reason is the most important. If there is no longer any meaningful distinction to be made between radical and incremental innovations, then there should be no meaningful distinction between revolutionary science and normal puzzle-solving science. All science changes the entire system and moves it irreversibly into a totally unprecedented and unique state. Any differences between one state and another are a matter of observer dependence, so that one scientific discovery can seem revolutionary to some and not to others. The many individual scientific discoveries that do not

seem to trigger a paradigm shift nevertheless contribute to every paradigm shift that happens for the rest of the future. Every paradigm shift would not have happened, or would not have happened in quite the same way, without each and every one of those seemingly non-revolutionary scientific contributions, even the ones that seem to work counter to those that are most closely linked to that eventual revolution. Therefore, it is not important whether CAS is seen as a new paradigm or not, or whether the next generally recognised paradigm is closely associated with CAS or not.

But the problem runs much deeper than whether or not transition experiments are searching for classical paradigm-esque cause-effect relationships, because almost *all* experiments are predicated on the idea that variables can be identified, measured and controlled to reveal predictable and manageable cause-effect relationships. That is just good science, according to the scientific method and, it could be argued, the definition of science. CAS may or may not become generally recognised as a paradigm changer, but even if it does there is no guarantee that it can ever change the high-level paradigm that defines the ‘proper role of science’. As the Carl Sagan quote at the beginning of the chapter says, science is not a body of knowledge but a way of thinking. Now, that way of thinking is based around the statistically derivation of cause-effect relationships in order to predict, influence, manage or control the future, all of which are hard concepts to relinquish because they offer the hope of eventual prediction and control.

Bibliography

- Esteve Almirall and Ramon Casadesus-Masanell. Open versus closed innovation: A model of discovery and divergence. *Academy of management review*, 35(1):27–47, 2010.
- W Brian Arthur. *The nature of technology: What it is and how it evolves*. Simon and Schuster, 2009.
- W Brian Arthur and Wolfgang Polak. The evolution of technology within a simple computer model. *Complexity*, 11(5):23–31, 2006.
- Philip Auerswald, Stuart Kauffman, José Lobo, and Karl Shell. The production recipes approach to modeling technological innovation: An application to learning by doing. *Journal of Economic Dynamics and Control*, 24(3):389–450, 2000.
- Per Bak, Chao Tang, and Kurt Wiesenfeld. Self-organized criticality: An explanation of the 1/f noise. *Physical Review Letters*, 59(4):381–384, 1987.
- Gregory Benford. *The Wonderful Future that Never was*. Hearst Communications, 2010.
- Gary Bradshaw. The airplane and the logic of invention. *Cognitive models of science*, 1992.
- GL Bradshaw and Marsha Lienert. The invention of the airplane. In *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society*, pages 605–610, 1991.
- É.J.L. Chappin. *Simulating Energy Transitions*. PhD thesis, PhD thesis, Delft University of Technology, Delft, the Netherlands. ISBN: 978-90-79787-30-2. URL: <http://chappin.com/ChappinEJL-PhDthesis.pdf>, 2011.
- Thomas Y Choi, Kevin J Dooley, and Manus Rungtusanatham. Supply networks and complex adaptive systems: control versus emergence. *Journal of operations management*, 19(3):351–366, 2001.
- J. Cohen and I. Stewart. *The collapse of chaos: Discovering simplicity in a complex world*. ePenguin, 2000.

- Sendil K Ethiraj, Daniel Levinthal, and Rishi R Roy. The dual role of modularity: innovation and imitation. *Management Science*, 54(5):939–955, 2008.
- Koen Frenken. *Understanding product innovation using Complex Systems Theory*. University of Amsterdam, 2001. ISBN 90-6464-05-80.
- Koen Frenken. *Innovation, evolution and complexity theory*. Edward Elgar Publishing, 2006.
- Xavier Gabaix. Power laws in economics and finance. Technical report, National Bureau of Economic Research, 2008.
- Xavier Gabaix, Parameswaran Gopikrishnan, Vasiliki Plerou, and H Eugene Stanley. A theory of power-law distributions in financial market fluctuations. *Nature*, 423(6937):267–270, 2003.
- Giovanni Gavetti and Daniel Levinthal. Looking forward and looking backward: Cognitive and experiential search. *Administrative science quarterly*, 45(1):113–137, 2000.
- Jacob Goldenberg, Donald R Lehmann, and David Mazursky. The idea itself and the circumstances of its emergence as predictors of new product success. *Management Science*, 47(1):69–84, 2001.
- John Grin, Jan Rotmans, and Johan Schot. Transitions to sustainable development. *New Directions in the Study of Long Term Transformative Change*, New York, 2010.
- Susan Hough. *Predicting the unpredictable: the tumultuous science of earthquake prediction*. Princeton University Press, 2009.
- Steven Johnson. *Where Good Ideas Come From: The Natural History of Innovation*. Penguin, 2010.
- J Kasmire, Janne M Korhonen, and Igor Nikolic. How radical is a radical innovation? An outline for a computational approach. *Energy Procedia*, 20:346–353, 2012.
- S.A. Kauffman and S. Johnsen. Coevolution to the edge of chaos: Coupled fitness landscapes, poised states, and coevolutionary avalanches*. *Journal of Theoretical Biology*, 149(4):467–505, 1991.
- Stuart A Kauffman. *The origins of order: Self-organization and selection in evolution*. Oxford University Press, USA, 1993.
- Stuart A Kauffman. *Investigations*. Oxford University Press, 2002.
- Kevin Kelly. *What technology wants*. Viking Press, 2010.
- Janne M Korhonen and Liisa Välikangas. 11 Constraints and ingenuity: the case of Outokumpu and the development of flash smelting in the copper industry. *Handbook of Organizational and Entrepreneurial Ingenuity*, page 254, 2014.

- Kuhn, Thomas. The structure of scientific revolutions. *Chicago and London*, 1970.
- Tiejun Ma and Yoshiteru Nakamori. Agent-based modeling on technological innovation as an evolutionary process. *European Journal of Operational Research*, 166(3):741–755, 2005.
- André Moor. Towards a grand deal on subsidies and climate change. In *Natural Resources Forum*, volume 25, pages 167–176. Wiley Online Library, 2001.
- MR Mudge. Rockfall-avalanche and rockslide-avalanche deposits at Sawtooth Ridge, Montana. *Geological Society of America Bulletin*, 76(9):1003–1014, 1965.
- Barrie Needham and Andreas Faludi. Dutch growth management in a changing market. *Planning practice and research*, 14(4):481–491, 1999.
- Jan W Rivkin. Imitation of complex strategies. *Management Science*, 46(6):824–844, 2000.
- Carl Sagan. Can we know the universe? Reflections on a grain of salt. *Broca’s brain*, pages 15–21, 1974.
- T Sarkikoski. A Flash of Knowledge—How an Outokumpu Innovation Became a Culture. *Outokumpu Oyj, Riihitontuntie 7 B, Espoo, PL 140, Finland, 1999. 303*, 1999.
- E. Scott. Cooption and Catalysis in a Model of Technological Evolution. September 2010.
- Gerald Silverberg and Bart Verspagen. A percolation model of innovation in complex technology spaces. *Journal of Economic Dynamics and Control*, 29(1):225–244, 2005.
- Burrhus F Skinner. The operational analysis of psychological terms. *Psychological review*, 52(5):270, 1945.
- R.A.C. van der Veen. Case study ‘CHP diffusion in Dutch greenhouses’ Final report. <http://repository.tudelft.nl/view/ir/uuid:df888718-fa67-4aba-b336-4f06e929b56c/>, 2012. The Final Report of the CHP Case Study.