Dekker, S. W. A. (2013). Drifting into failure: Complexity theory and the management of risk. In: Banerjee, S. (Ed.), *Chaos and Complexity Theory for management: Nonlinear Dynamics*, pp. 241-253. Hershey, PA: IGI Global Business Science Reference.

Drifting into failure: Complexity theory and the management of risk

Sidney W. A. Dekker

The Gaussian copula, an equation first published by David Li in 2000, was a beautiful thing—in isolation. Its intention was to price collaterized debt obligations, and work out whether they were moving in the same direction. The copula was an enabler of mortgaging the most hopeless of homeowner prospects. Millions of securities could be traded on the back of a single number (a security is something that shows ownership, or right of ownership of stocks or bonds, or the right to ownership connected with derivatives that get their value from some underlying asset). As more and more webs of interactions and relationships and interdependencies and feedback loops started growing around it, however, it became part of a complex system. The copula became the trigger of a recession that swelled the number of homeless families in the US by 30% inside of two years.[i] It ended up bringing global lending to a virtual standstill, triggering a worldwide financial crisis and a deep recession. How did a once good idea like this drift into failure, and how can such a risk of collapse be managed? That is what this chapter is about.

In a crisis, all correlations go to 1

The copula is about probabilities, and whether they are associated with other probabilities. By putting in different bonds the Gaussian copula function produced a single number that became easily manipulable and traceable by the world of quantitative finance. It could show correlations between bonds that might default and bonds that might not.

$$Pr[T_A<1,T_B<1] = \Phi_2(\Phi^{-1}(F_A(1)),\ \Phi^{-1}(F_B(1)),\gamma) \qquad (1)$$

In the Gaussian copula (see equation (1)), probability Pr is a joint default probability, the likelihood that any two members of a pool (A and B, each of which might contain hundreds or even thousands of mortgages, for example) will both default. This is what investors are looking, and the rest of the copula provides the answer. Φ_2 is the copula: it couples (hence the Latin copula) the individual probabilities associated with pools A and B to come up with a single number. Errors here can massively increase the risk of the whole equation producing garbage. Survival times (T_A and T_B) represent the amount of time between now and when A and B can be expected to default (i.e. fail to repay). Li was inspired in this by actuarial science, which has a similar concept for what happens to somebody's life expectancy when their spouse dies. Distribution functions F_A and F_B represent the probabilities of how long A and B are likely to survive. These are not certainties, so small misassessments or miscalculations here can lead to a much greater production or risk than the formula would have you believe. The idea about equality (=) between the copula and probability of default might be a dangerously misleading concept in this formula, as it suggests that there is no room for error, and is a short notation that says "is" or "equals" or "is equal to," which muffles the role of real-world uncertainty, fuzziness, and precariousness. The gamma γ at the end is the all-powerful correlation parameter, which reduces correlation to a single constant. This, in this context, is something that should be highly improbable, if not impossible. It was, however, the part that really made Li's idea irresistible and pervasive.

In isolation it looked elegant, and it worked well. But the world in which the Gaussian copula was released was the world of collaterized debt obligations, growing exponentially in size and complexity from the 1990's onward. Collaterized debt obligations, an invention from the late 1980's, are a type of synthetic asset-backed security. Some sort of corporate entity needs to be constructed to hold assets as collateral (say, somebody's house), and then collect interest which can be sold as packages of cash flow to investors. Collaterized debt obligations can either come from a special purpose entity that acquires a portfolio of underlying assets such as mortgage-backed securities, commercial real estate bonds and corporate loans. Or the special purpose entity issues bonds (those collaterized debt obligations) in what are called different tranches (levels of risk), from which the proceeds are then used to purchase a portfolio of underlying assets.

The risk and return for somebody who has invested in a debt obligation (often without really knowing it, by the way) depends on how the obligations and their tranches are defined (and this may not be communicated very clearly or get lost in the various layers of buying, selling and reselling), and only indirectly on the underlying assets. So in effect, the investment depends on the assumptions and methods used to define the risk and return of the tranches. Like all asset-backed securities, debt obligations enable the originators of the underlying assets to pass credit risk to a variety of other institutions or chop it up and distribute it to an immeasurable number of individual investors, who in turn might resell it. Risk is distributed and sold into invisibility. With the multiplying layers of players involved in the bond market in sending this money (or debts) around, risks were chopped up and scattered into oblivion.

The use and increasing opacity of financial instruments like debt obligations expanded dramatically on the back of a growing number of asset managers and investors. There was no intelligent design behind this, no single smart designer who had it all figured out beforehand, or could figure it all out even while things were playing out. The growth of asset management was in part a response to—and source for—a growing need for stock market investments and mortgages. From the 1990's onward, an increasing set of players in modern society (from private individuals to government pension funds to sovereign nations) turned to stock markets for financial returns and presumed future security. It intersected with the realization in many developed nations of impending pension fund depletions and sovereign debt obligations that would need to be met at some point. Many countries embarked an aggressive strategies to get people to invest in their own future pensions. At the same time, the call of home ownership as route to establishing independent wealth found renewed vigor in the United States and elsewhere, even as the middle class saw its gains hollowed out by wage stagnation and price inflation (particularly hard-to-prospect costs for things like college tuition and healthcare). There was enough borrowable money to go around. Everybody could get a loan.

The amount of investable money that came available each month was staggering. Attracted by all the opportunity, asset management exploded across a bewildering array of actors, from investment trusts to commercial banks, investment banks, pension funds, private banking organizations, insurance companies, mutual fund companies and unit trusts. With so many layers between them and what their money was doing out there scattered across the bond market, many people had little idea what exactly they were investing in, and their closest fund manager or financial advisor might not really have known either. And how did the asset managers make money? The issuer of a collaterized debt obligation, which might typically be an investment bank or some other entity working in their stead, earned a commission when it issued the bond, and then earned management fees during the remainder of its life. The ability to earn money from originating and securitizing loans, coupled with the absence of any residual liability (you didn't own an asset you could lose, after all, like a house), skewed the incentives of banks and financial managers in favor of loan volume rather than loan quality. More was better. And the Gaussian copula was the great enabler. Trade anything you could get your hands on, any debt you can find, run it through the copula and you'd see how it did or would possibly do. In a sense it was a replay of let's-securitize- everything-we-can, even securities themselves" enthusiasm that swept Enron in the late 1990's. Securitizations exploded, with everything from lotto winnings to proceeds from tobacco lawsuits being turned into securities that could be sold to the investing public (McLean and Elkind 2004).

A world of distant assets, future bets, and seemingly virtual debts that could be traded, hedged, securitized—and generate money for those organizations that did the trading (or at least make their accounting figures look really good). Everybody used the Gaussian copula, from bond investors to Wall Street banks to ratings agencies and financial regulators. In complexity science this is known as positive feedback: more leads to more. Everybody started using the copula, because everybody started using it. Organizations that didn't risked being left behind in the skyrocketing securities trade. Once they embraced the copula as a basic instrument to assess risk, they no longer needed to look at the quality of any underlying securities. All they had to do was look at a correlation number, and out would come a rating telling them how safe or risky a tranche was.

In finance, it is impossible to ever reduce risk outright. We can only try to set up a market in which people who don't want risk sell it to those who do. But in the collaterized debt obligations market, people used the Gaussian copula model to convince themselves they didn't have any risk at all. The Gaussian copula was, if anything, a multiplier—something that made the trading easier, and attracted more trading-on-trading (i.e. hedging other people's hedges, taking out loans to invest in other people's loans) something it could do by collapsing the complexity of the risk in these deals into a single number. It was easy to convince people that good things move together. If house prices go up, house prices go up. Bubbles grow when everybody is doing what everybody is doing. Again, that is positive feedback. But this also works the other way: bad things also tend to move together. Bubbles get punctured when everybody is doing what everybody is doing. If one mortgage defaults, then it is not unlikely that one hundred others will too. As the saying goes: in a crisis, all correlations go to 1.

When house prices stopped rising, less became less. Once again, there was positive feedback. Fewer people wanted to lend money because fewer people wanted to lend money. The huge web of interconnections and interdependencies that had grown on the back of the unfathomable amount of available bond market money now started backfiring, reverberating in the opposite direction. It was a feedback loop that made the world run dry of borrowable money in very short order. And it triggered an economic crisis and a huge recession.

After the fact ("ex post"), it all seems like a really bad idea. But that is the language of rational choice theory. After all, the badness of the idea became obvious only once we could see how it turned out. And how it turned out depended not on the Gaussian copula function per se. It depended on the world the function was released into, and the mass of relationships and interdependencies that started growing around it as people saw how it could work for them locally, relative to their goals, knowledge and focus of attention.

Drifting into failure

Managers often have trouble grasping the complexity and normality that gives rise to such large events. The growth of complexity in our organizations and their technologies has outpaced our understanding of how complex systems work and fail. There is a marked Newtonian legacy in our thinking about organizational failure: risk is located in components and in the linear progression or propagation of component failures, producing a failure trajectory in which there is a proportionality or symmetry between causes and effects. These assumptions animate most reliability and quality work. But many technologies (like the Gaussian copula) have got ahead of such theories. We are able to build things—deep-sea oil rigs, jackscrews, collaterized debt obligations—whose properties we understand in isolation. But in competitive, regulated societies, their connections proliferate, their interactions and interdependencies multiply, their complexities mushroom.

The idea put forward in this chapter is that organizations do not just fail because of component breakage or linear propagations of breakdowns. Instead, failure breeds opportunistically, non-randomly, among the very structures designed to protect an organization from disaster. It aims to understand how organizations gradually decline into breakdown despite the best intentions of many involved. Failure, like success, can in part be seen as emergent—not present in the components that make up the system, but visible only as a higher-order property. This suggests that managers see their organizations more as ecological systems, not as machines that are insulated and isolated from their environments.

A common pattern seems to be a drift into disaster—a slow, incremental decline into bad judgment by organizations that take past results as a guarantee for continued success. These were systems that drifted into failure. While pursuing success in a dynamic, complex environment with limited resources and multiple goal conflicts, a succession of small, everyday decisions eventually produced breakdowns on a massive scale. Small steps that increased risk, that took the organization away from previously accepted norms, were seen as routine and non-problematic or even necessary to achieve local gains. The focus on tweaking more return, production, or profit out of finite resources obscures a longer-term picture of drift toward disaster. It is hidden behind incrementalism, technical uncertainty and complexity, non-linear feedback loops, weak signals, personnel changes, job mandate horizons and more. Organizations drift into failure by succeeding, by doing what it takes to produce success in a world with limited resources and multiple goal conflicts in a dynamic environment.

This does not mean that failure is inevitable, or that all risk is unmanageable. It does, however, require that managers think with a new vocabulary to complement the componential discussion that keeps dominating their work today. Complexity theory is used in this chapter for precisely that purpose. This is necessary for identifying new leverage points to stop drift at levels and in ways that are not componential or structural.

Features of drift into failure

One of the greatest conundrums in these cases is to explain why a slide into disaster, easy to see and depict in retrospect, was missed by those who inflicted it on themselves, or by those responsible for preventing it. What obscures it, mostly, is normal work by normal people in what looks to everybody to be a normal organization. The decisions, trade-offs, preferences and priorities that people made, even if seemingly out of the ordinary or even immoral after the failure, were seen as normal and common sense at the time. Just as we like to believe that ours are now. Here are five concepts that together may characterize drift (Dekker 2011):

- Scarcity and competition
- Incrementalism, or small steps
- Sensitive dependence on initial conditions
- Unruly technology
- Contribution of the protective structure

Let us look at how these factors work together to produce drift. No organization operates in a vacuum. This is a fundamental feature of complex systems: they are open to influences from the outside and are in constant transaction with what happens around them. Any organization operates and must try to survive in an environment that has real constraints, such as the amount of capital available, the number of customers reachable, and the qualifications of available employees. There are also hard constraints on how fast things can be built, developed, driven. And, most importantly, there are other organizations that try to do the same thing.

Competitive pressure is important. The major engine of drift hides somewhere in the conflict between production and protection, in the tension between operating with risk under control and operating at all. This tension provides the energy behind the slow, steady disengagement of practice from earlier established norms or managerial constraints. This disengagement can eventually become drift into failure. As a system is taken into use, it learns, and as it learns, it adapts. A critical ingredient of this learning a normalization of deviance (Vaughan 1996), the apparent insensitivity to mounting evidence that, from the position of retrospective outsider, could have shown how bad managerial judgments and decisions have become.

One big problem is a managerial feedback imbalance. Information on whether a decision is cost-effective or efficient can be relatively easy to get. How much is or was borrowed from safety in order to achieve that goal, however, is much more difficult to quantify and compare. Evidence from a feedback imbalance, then, suggests that the system can operate relatively risk-free, yet more efficiently.

Although each decision is locally rational, making sense for decision makers in their time and place, the global picture can become one of drift into failure. Incrementalism, or small changes, can lead to big events. It is sometimes easy for managers to overestimate the controllability of their organizations and technologies. But many of these are quite unruly: disorderly and not amenable to discipline or control (Wynne 1988). One aspect of this, in complexity and systems thinking, is sometimes also called “sensitive dependency on initial conditions” (or butterfly effect). Whether technology takes an unruly trajectory toward failure may depend on seemingly innocuous features or infinitesimally small differences in how it all got started or influenced with small managerial decisions. In complex systems, managerial actions control very little, but influence almost everything (Page 2007).

Managers are typically not alone in driving the activities of their organizations. The operation of their technologies is surrounded by structures meant to keep its risk under control (e.g. boards, regulators, shareholders). This tends to produce an array of rules, routines, procedures, guidance material, prescriptions, expert opinions. This weaves vast webs of relationships in which the operation of the organization is suspended—a web, moreover, that has no clear boundaries, no obvious end or beginning. This creates a wonderful paradox. The whole structure that is designed (and has evolved) to keep the organization’s risk under control, can make the functioning and malfunctioning of that organization more opaque. The meaning of signals gets constructed, negotiated, and transacted through the web of relationships that is strung throughout this structure. Sometimes, the weak signals that are left over trigger only weak organizational responses, if any at all (Weick and Sutcliffe 2007). What if the protective structure itself contributes to the construction and treatment of weak signals, and by extension to drift—in ways that are inadvertent, unforeseen, and hard to detect? The organized social complexity surrounding the operation, all the committees, working groups, regulatory interventions, approvals, and manufacturer inputs, which all intended to protect the system from breakdown, could actually help set its course toward failure.

The introduction of a new piece of technology or organizational instruments (CDO’s, for example) is typically followed by negotiation, by discovery, by the creation of new relationships and rationalities. “Technical systems turn into models for themselves,” said Weingart (pp. 8-9), “the observation of their functioning, and especially their malfunctioning, on a real scale is required as a basis for further technical development.” (Weingart 1991). Rules and standards do not exist as unequivocal, aboriginal markers against a tide of incoming operational data (and if they do, they are quickly proven useless or out of date). Rather, rules and standards are the constantly updated products of the processes of conciliation, of give and take, of the detection and rationalization of new data. As Brian Wynne said (Wynne 1988):

> "Beneath a public image of rule-following behavior and the associated belief that accidents are due to deviation from those clear rules, experts are operating with far greater levels of ambiguity, needing to make expert judgments in less than clearly structured situations. The key point is that their judgments are not normally of a kind—how do we design, operate and maintain the system according to 'the' rules? Practices do not follow rules, rather, rules follow evolving practices" (p. 153).

Nor is there a one-way or unproblematic relationship between the original rules or requirements and subsequent operational data. Even if the data, in one reading, may prove the original requirements or rules wrong, this doesn't mean that complex systems, under the various normal pressures of operating economically, reject the requirements and rules and come up with new ones. Instead, the meaning of the data can get renegotiated. People may want to wait for more data. The data can be denied or downplayed in managerial or board meetings.

Drift and complexity science

Mechanistic thinking about failures, that is, the Newtonian-Cartesian approach, means going down and in. Understanding why things went wrong comes from breaking open the system, diving down, finding the parts, and identifying which ones were broken. This approach is taken even if the parts are located in different areas of the system, such as procedural control, supervisory layers, managerial levels, regulatory or board oversight. In contrast, systems thinking about failures means going up and out. Understanding comes from seeing how the system is configured in a larger network of other systems, of tracing the relationships with those, and how those spread out to affect, and be affected by, factors that lie far away in time and space from the moment things went wrong.

With his formulation of General Systems Theory, Von Bertalanffly helped establish a scientific foundation for an alternative to Newtonian-Cartesian thought in the late 1960's (Bertalanffy 1969). Today that alternative is known as complexity and systems theory. More recently, Cilliers summarized the characteristics of complex, as opposed to Newtonian, systems (Cilliers 1998):

- Complex systems are open systems—open to influences from the environment in which they operate and influencing that environment in return. Such openness means that it is difficult to frame the boundaries around a system of interest.
- In a complex system, each component is ignorant of the behavior of the system as a whole, and doesn't know the full effects of its actions either. Components respond locally to information presented by them there and then. Complexity arises from the huge, multiplied webs of relationships and interactions that result from these local actions.

- Complexity is a feature of the system, not of components inside of it. The knowledge of each component is limited and local, and there is no component that possesses enough capacity to represent the complexity of the entire system in that component itself. This is why the behavior of the system cannot be reduced to the behavior of the constituent components, but only characterized on the basis of the multitude of ever-changing relationships between them.
- Complex systems operate under conditions far from equilibrium. Inputs need to be made the whole time by its components in order to keep it functioning. Without that constant flow of actions, of inputs, it cannot survive in a changing environment. The performance of complex systems is typically optimized at the edge of chaos, just before system behavior will become unrecognizably turbulent.
- Complex systems have a history, a path-dependence. Their past is co-responsible for their present behavior, and descriptions of complexity have to take history into account.
- Interactions in complex systems are non-linear. That means that there is an asymmetry between for example input and output, and that small events can produce large results. The existence of feedback loops means that complex systems can contain multipliers (where more of one means more of the other, in turn leading to more of one, and so forth) and butterfly effects.

Open systems mean that it can be quite difficult to define the border of a system. What belongs to the system, and what doesn't? This is known as the frame problem. It is very difficult to explicitly specify which conditions are *not* affected by an action. System theory itself provides no answer to the frame problem. Where you place the frame is up to you, and up to the question you wish to examine. The adage of forensic science to "follow the money" is the same commitment. On the one hand, it leaves the observer or investigator entirely open to where the trail may take her or him and how it branches out into multiple directions, organizations, countries. That is where the system of interest is open. On the other hand, the commitment frames the system of interest as that which can be expressed monetarily.

In a complex system, each component is ignorant of the behavior of the system as a whole. This is a very important point (Cilliers 1998). If each component "knew" what effects its actions had on the entire rest of the system, then all of the system's complexity would have to be present *in that component*. It isn't. This is the whole point of complexity and systems theory. Single elements do not contain all the complexity of the system. If they did, then reductionism could work as an analytic strategy: you could explain the whole simply by looking at the part. But in complex systems, you can't, and analytic reduction doesn't work to enhance anybody's understanding of the system. Complexity is the result of a rich interaction, of constantly evolving relationships between components and the information and other exchanges that they produce.

In the mechanistic worldview, it is enough to understand the functioning or breaking of parts to explain the behavior of the system as a whole. In complexity and systems thinking, where nothing really functions in an unbroken or strictly linear fashion, it is not. Recall, from the second chapter, the outlines of drift into failure. Here is how they overlap with what complexity theory may have to offer:

- Resource scarcity and competition, which leads to a chronic need to balance cost pressures with safety. In a complex system, this means that the thousands smaller and larger decisions and trade-offs that get made throughout the system each day can generate a joint preference without central coordination, and without apparent local consequences: production and efficiency get served in people's local goal pursuits while safety gets sacrificed—but not visibly so;
- Incrementalism, where constant organizational and operational adaptation around goal conflicts and uncertainty produces small, step-wise normalization where each next decrement is only a small deviation from the previously accepted norm, and continued operational success is relied upon as a guarantee of future safety;
- Sensitive dependence on initial conditions. Because of the lack of a central designer or any part that knows the entire complex system, conditions can be changed in one of its corners for a very good reason and without any apparent implications: it's simply no big deal. This may, however, generate reverberations through the interconnected webs of relationships; it can get amplified or suppressed as it modulates through the system;
- Unruly technology, which introduces and sustains uncertainties about how and when things may fail. Complexity can be a property of the technology-in-context. Even though parts or sub-systems can be modeled exhaustively in isolation (and therefore remain merely complicated), their operation with each other in a dynamic environment generates the unforeseeabilities and uncertainties of complexity;
- Contribution of the entire protective structure (the organization itself, but also the regulator, legislation, and other forms of oversight) that is set up and maintained to ensure safety (at least in principle: some regulators would stress that all they do is ensure regulatory compliance). Protective structures themselves can consist of complex webs of players and interactions, and are exposed to an environment that influences it with societal expectations, resource constraints, and goal interactions. This affects how it condones, regulates and helps rationalize or even legalizes definitions of "acceptable" system performance.

The concern behind complexity and drifting into failure is how a large number of things and processes interact, and generate organizational trajectories when exposed to different influences. Resource scarcity and goal oppositions form one such pervasive influence. They express themselves in thousands smaller and larger trade-offs, sacrifices, budgetary decisions—some very obvious, others hardly noticed. The ripple effects of such decisions and trade-offs are sometimes easy to foresee, but often opaque and resistant to anything resembling deterministic prediction. Incrementalism shows up in all kinds of subtle ways as people in the organization adapt, rationalize and normalize their views, assessments and decisions.

The contribution of the protective structure to such adaptation and normalization, as well as exposure to its own resource constraints and goal interactions, is another influence on this. Such influences ebb and flow to different parts of the operational organization or even originate there, and are negotiated, dealt with, ignored or integrated.

Drifting into success

One key word for the current age seems to be resilience. Organizational resilience is the ability to recognize, adapt to and absorb problem disturbances without noticeable or consequential decrements in performance (Hollnagel, Woods et al. 2006). Diversity is a critical ingredient for resilience, because it gives an organization the requisite variety that allows it to respond to disturbances. With diversity, a system has a larger number of perspectives to view a problem with and a larger repertoire of possible responses. Diversity means that routine scripts and learned responses do not get over-rehearsed and over-applied, but that an organization has different ways of dealing with situations and has a rich store of perspectives and narratives to interpret those situations with. Diversity is a key property of complex systems. It is a property that can be harnessed so that complexity is given a role in preventing a drift into failure rather than creating it. Whether the complexity of drift can be managed depends very much on how managers can recognize and capitalize on diversity through the features of drift described above.

Resource scarcity and competition

Resource scarcity and competition leads to a chronic need to balance cost pressures with safety. Thousands smaller and larger managerial decisions and trade-offs that get made throughout the system each day can generate a joint preference, a common organizational direction. The decisions also have no apparent local or even knowable global consequences.

Resource scarcity and competition is itself a feature of complexity, of course. It is the natural by-product of operating in an environment with interactions and interdependencies across players, where multiple groups are after the same scarce resources and will try to deliver the same service or product faster, better and cheaper. It is impossible to make these features go away other than deciding to no longer compete, to leave the business, to go do something else. An organization cannot typically decide to operate without scarcity or competition (even monopolies often still suffer from resource scarcity in the markets where they operate).

The tensions and paradoxes that are produced by resource scarcity and competition are natural phenomena for complex systems. Managers shouldn't necessarily try to resolve them, because they probably won't be successful anyway. Complex systems interact with other complex systems, and this leads to tensions and paradoxes that can never be fully reconciled. And their may be benefits to this. In complex social organizations, the seemingly opposing forces of competition and cooperation often work together in positive ways—competition within an industry can improve the collective performance of all participants.

But there is another side to this. When a manager is convinced that the organization's business is going faster, better and cheaper, she or he should get a bit uncomfortable. Locally optimizing processes so that they become faster, better and cheaper probably makes good local sense. But it may not make good global sense. The sum of local optimizations can be global sub-optimization. Fully optimizing a complex systems is both impossible and undesirable. It is impossible because we will never know whether we have actually achieved a totally optimal solution. Even if things get tweaked optimally in a local sense, there will likely be influences of such steps in other parts of the complex system that we might not even be aware of. The cost of optimization in one locality easily gets exported elsewhere—and show up as a constraint that forces a sub-optimal solution there. Extra work may be required, or extra slack may need to be built in elsewhere to accommodate the pressures or demands or constraints imposed by another part of the system. Fully optimizing a complex system is undesirable because the lack of slack and margin can translate small perturbations into large events (Perrow 1984). Tight interdependencies or coupling means that there are more time-dependent processes (meaning they can't wait or stand by until attended to), sequences that are invariant (the order of processes cannot be changed) and little slack (e.g. things cannot be redone to get it right). It means that small failures can cascade into bigger ones more quickly than if there is margin. Diversity of opinion (e.g. on a board) can be one way to make the organization stop and think but it requires the credibility and the courage to say "no" when everybody says "yes."

Sensitive dependency and small steps

Because there is no central designer, or any part, that knows the entire complex system, local actors can change their conditions in one of its corners for a very good reason and without any apparent implications. To them, it's simply no big deal, and in fact it may bring immediate gains on some local goal trade-off that people routinely face (e.g. getting the system out the door versus checking it once more). Changing these conditions, however, may generate reverberations through the interconnected webs of relationships. The reclassification of foam debris events from flight safety to maintenance issues in the case of Space Shuttle Columbia was such a change in starting conditions that eventually linked to the accident (CAIB 2003). A difference in an initial condition like this can get amplified (or suppressed) as its influences and post-conditions meander through the organization.

These changes in condition (a redefinition or reclassification of a risk event, or a particular denotation of the status of a pre-operational system, for example) are not momentous in themselves. Which is why they are relatively easy to achieve, and hard to detect or stop. These are, in fact, often small steps. They are small steps that help the system locally optimize or rationalize a corner of its operations. This can actually be turned into a managerial advantage. Small steps happen all the time as complex organizations adapt to their environment. This means that possible opportunities for reflection on practice offer themselves up quite frequently. Calling on people to reflect on smaller steps probably does not generate as much defensive posturing as challenging their more momentous decisions. Retracting a small step, should people conclude that that is necessary, may also not be as expensive. In other words, small steps could mean that there is political and organizational space for critically reflecting on them, and perhaps ultimately space for not taking one of them.

Doing this may require some managerial action, however. Small steps and the normalization they can entail are often no longer be visible or seen as significant by insiders. Outsiders might think about this very differently, and they may represent a resource that managers could capitalize. Having people come in from other operational workplaces does at least two things. It brings fresh perspectives that can help insiders recalibrate what they consider “normal:” an invitation for critical reflection on one’s own practices.

Unruly technology

Unruly technology, which introduces and sustains uncertainties about how and when things may fail, can also be turned into a lever for managing the complexity of drift. Recall how technology can remain unruly. Even though parts or sub-systems of a technology can be modeled exhaustively in isolation, their operation with each other in a dynamic environment generates the unforeseeabilities and uncertainties typical of complexity. Technology in a complex world, like a Gaussian copula is never “finished.” Declaring it operational has much more to do with us, and with our organizational and political constraints, than with the technology and how and where it operates.

To capitalize on unruly technology, rather than be vexed by it, managers should invert their perspective. We may traditionally have seen unruly technology as a pest that needs further tweaking before it becomes optimal. This is a Newtonian commitment: we can arrive at an exhaustive description of how the technology works, and we can optimize its functioning (i.e. find one best way of deploying it). Complexity theory, in contrast, suggests that we see unruly technology not as a partially finished end-product, but as a tool for discovery. The feature of drift discussed above, small steps, is actually a wonderful property here. Small steps can mean small experiments. Many small steps means many small experiments and opportunities for discovery. These are small experiments without necessarily devastating consequences, but with the potential to create important results and insights. David Li’s surprise and dismay about the enormous popularity of a Gaussian copula he invented, for example, generated information both about our environment *and* about us that could have been used to our (and the world’s) advantage.

Contribution of the protective structure

Protective structures themselves typically consist of complex webs of players and interactions, and are exposed to an environment that influences it with societal expectations, resource constraints, and goal interactions. This affects how it condones, regulates and helps rationalize or even legalizes definitions of “acceptable” organizational performance. It means that there is often something inescapably paradoxical and corruptible about the role of a protective structure. In principle, the idea of a protective structure is the addition of more diversity to the system. Bringing in outsiders is an obviously good idea—it gets done when consultants come in, but also when a regulator comes in. But viewpoints between operator and regulator about what is risky can begin to overlap, even while promoting an image of control and diversity.

Perhaps the three terms most closely associated with protective structures—regulation, compliance, oversight and inspection—are all fundamentally mismatched to complexity. Complex systems cannot in principle be regulated. Regulation, in its bare-bones form, means controlling or maintaining particular parameters so as to keep things operating properly. The idea of regulation is locked in a machine metaphor (and therefore Newtonian assumptions) of how an organization works. This includes not only an image of an organization as a collection of parts and interconnections, but the idea that we can arrive at a complete description of how the system works versus how it is supposed to work. This, of course, is where compliance-based approaches come in. If regulators discover gaps between rules and practice, they may typically try to close those gaps by pressing for practice that is more compliant with the rules.

Complex systems, however, produce behaviors that are more akin to living systems than to machines. Self-organization and emergence, let alone creative evolution, are impossible in a machine. And these behaviors are all very difficult to hold up against a set of rules for how the system is supposed to work. Rules cannot even accommodate creative evolution and emergence, because it would mean that somebody or some agency has designed, in advance, the emergence and creative evolution and self-organization and then condensed it all into rules that reflect some ideal-type against which actual evolving practice can be matched. Emergence, creative evolution and self-organization cannot be designed beforehand. That is the whole point of complexity.

If a regulator cannot regulate a complex system, then what can it do? Will a regulator always be caught behind the curves of self-organization and emergence, holding a bag of obsolete rules that came from less evolved systems? Languages of compliance and regulation can perhaps be juxtaposed against those of co-evolution and counter-evolution. Rather than a regulator, complex systems should have a co-evolver/counter-evolver. This must be an organization that has the requisite variety not only to have an idea of the complexity of the operational organization (and thus has to co-evolve with how that organization evolves). It should also have requisite variety to counter-evolve. At least in its theories or models, it should be able to generate alternative outcomes to the small steps that get made by the operational organization.

This is also why the idea of risk oversight is problematic. Oversight implies a big picture. A big picture, in a sense of a complete description, is impossible to achieve of a complex system. Not only is that computationally intractable, complex systems (as said many times before in this book) evolve and change and adapt the whole time. Nailing its description down at one moment in time means very little for how it will look a next moment. If the big picture of oversight, however, implies a sensitivity to the features of complexity and drift, then it might work. Oversight can try to explore complex system properties such as interconnectedness, interdependence, diversity, and rates of learning and selection that go on inside of the complex system. This is perhaps the type of oversight that people inside the operational organization are not capable of because of the locality principle (actors and decision makers in the organization only see their local interactions). But for a regulator (as for managers), it means learning an entirely new repertoire of languages and countermeasures: from componential, determinist, compliance to co/counter-evolving complexity.

References

Bertalanffy, L. v. (1969). General system theory; foundations, development, applications. New York,, G. Braziller.

CAIB (2003). Report Volume 1, August 2003. Washington, DC, Columbia Accident Investigation Board.

Cilliers, P. (1998). Complexity and postmodernism: Understanding complex systems. London, Routledge.

Dekker, S. W. A. (2011). Drift into failure: From hunting broken components to understanding complex systems. Farnham, UK, Ashgate Publishing Co.

Hollnagel, E., D. D. Woods, et al. (2006). Resilience engineering: Concepts and precepts. Aldershot, UK, Ashgate Publishing Co.

McLean, B. and P. Elkind (2004). The smartest guys in the room: The amazing rise and scandalous fall of Enron. New York, Portfolio.

Page, S. E. (2007). "Making the difference: Applying a logic of diversity." Academy of Management Perspectives **11**: 6-20.

Perrow, C. (1984). Normal accidents: Living with high-risk technologies. New York, Basic Books.

Vaughan, D. (1996). The Challenger launch decision: Risky technology, culture, and deviance at NASA. Chicago, University of Chicago Press.

Weick, K. E. and K. M. Sutcliffe (2007). Managing the unexpected: Resilient performance in an age of uncertainty. San Francisco, Jossey-Bass.

Weingart, P. (1991). Large technical systems, real life experiments, and the legitimation trap of technology assessment: The contribution of science and technology to constituting risk perception. Social responses to large technial systems: Control or anticipation. T. R. LaPorte. Amsterdam, Kluwer**:** 8-9.

Wynne, B. (1988). "Unruly Technology: Practical Rules, Impractical Discourses and Public Understanding." Social Studies of Science **18**(1): 147-167.

[i] Associated Press (2010, July 18). *Number of homeless families grew in 2009, report says.* International Herald Tribune, p. 4.