

THE CYNEFIN MINI-BOOK

An Introduction to Complexity
and the Cynefin Framework

Greg Brougham

InfoQ^{live}

ENTERPRISE SOFTWARE
DEVELOPMENT SERIES

Copyrighted material

Contents

Preface	2
The Cynefin Framework	5
Introduction.....	6
Limits of rationality	6
The Cynefin framework.....	8
Considerations.....	10
Cynefin and narrative	15
Shared context and sense making	19
Introduction.....	20
Contextualisation	22
Shared context.....	23
Portfolio management	29
Introduction.....	30
Background.....	31
Complicated portfolio template.....	34
Complex portfolio template.....	36
Portfolio initiatives review	38
Portfolio management.....	39
Cynefin dynamics	41
Further reading	45
About the Author	46
References	47

Preface



This mini-book started out as a series of papers that were experiential in nature, which were intended to provide an introduction to Cynefin, so you won't find discussions about the ontology, epistemology, and phenomenology of the approach (okay, you will: ontology is mentioned once in the third paper when disorder is discussed, but it is an important part of the message). Shane Hastie suggested making them into a mini-book, which also provided the opportunity to add some text on the use of narrative. This discussion provides a bridge between the first paper, which contains the introduction, and the third paper in the series, which is about sense making. It also provided the opportunity to include some writings on Cynefin dynamics, on which there is little written. This deals with the management of groups and moving between exploration (un-order) and exploitation (order).

One of the key messages that comes from complexity is that you should work with fine-grained objects, leverage distributed cognition, and ensure disintermediation. Since this is a new management approach, there are not many stories to reference, but the article about Lotus in The Sunday Times that was published in February 2015 is interesting in that it embodies all of these principles. The new CEO asked for all three existing car models to be broken down into their parts, which were then laid out on tables for inspection. All 900 employees of the company were involved in this exercise and they were asked to tag the components using a traffic-light system. The components were either to be kept, supply renegotiated, redesigned, or discarded. This also ensured that everyone was on the same page and understood why these changes were being proposed. The exercise resulted in saving around 20 kg and £3,000. It also led to the quality of the cars being improved.

One thing that I did not also address in the papers was the derivation of the name Cynefin. It is a Welsh word and the literal English translation is "habit" or "place", but this does not convey its full meaning. Quoting Mike Pearson's *In Come I*: "It is the piece of earth where a community has lived – a community with whom we identify the places that we have lived." Dave Snowden noted in a tweet that its meaning is similar to the Maori word *tūrangawaewae*, which means "a place to stand". This again falls short in terms of the richness of the word, and Te Ara: The Encyclopedia of New Zealand expands the meaning to "places where we feel

especially empowered and connected. They are our foundation, our place in the world, our home.”

I hope you find these writings of interest and use.

What is in an InfoQ mini-book?

InfoQ mini-books are designed to be concise, intending to serve technical leaders looking to get a firm conceptual understanding of a new idea, framework, technology or technique in a quick yet in-depth fashion. You can think of these books as covering a topic strategically or essentially. After reading a mini-book, the reader should have a fundamental understanding of the concepts covered, including when and where to apply them, how they relate to other ideas and technologies, and an overall feeling that they have assimilated the combined knowledge of other professionals who have already figured out what these concepts are about. The reader will then be able to make intelligent decisions about the concepts once their projects require them, and can delve into sources of more detailed information (such as larger books or tutorials) at that time.

Who this book is for

This book is aimed specifically at architects, project managers and stakeholders who are interested in a short introduction to the subject of complexity, and Cynefin and its related practices in particular. It is not intended to be a replacement for training, but to demonstrate some of the practices and the value that they offer in dealing with an increasingly uncertain world.

Reader feedback

We always welcome feedback from our readers. Let us know what you think about this book – what you liked or disliked. Reader feedback helps us develop titles that you get the most out of.

To send us feedback, email us at feedback@infoq.com.

If you have a topic that you have expertise in and you are interested in either writing or contributing to a book, please take a look at our mini-book guidelines on <http://www.infoq.com/minibook-guidelines>.

PART **ONE**

The Cynefin
Framework

Introduction

The Cynefin framework and its practices can be used to address the uncertainty of the modern world. The practices can be used to complement traditional approaches to programme and portfolio management. They provide a more comprehensive approach that reflects the needs of management in an ever more uncertain world.

We have seen a couple of books on the issue of uncertainty appear over the last decade, but neither has a comprehensive framework that allows us to deal with the modern world's increased uncertainty.

If anyone doubts that traditional models are struggling to deal with modern market dynamics, they only need to look at the demise of the Monitor Group. This was the company of Michael Porter (the father of strategic analysis) and used his market-analysis model. He based his approach on rigorous analysis of market forces and the assumption that this leads to a rational, structured approach that would result in a competitive advantage. It became apparent that while it could “help explain excess profits in retrospect, it was almost useless in predicting them in prospect.” Matthew Stewart (2009) notes “Most successful strategies emerge through action; they become perspicuous only in hindsight.” A.G. Lafley and Roger Martin (2013) also note that market dynamics are not this simple and the world is increasingly complex, global, and competitive. We've learned that there are limits to rationality, but how can we address uncertainty and take advantage of turbulence in the market place?

Limits of rationality

When we are faced with a problem, we assume that all we need to do is elaborate the options, select one, and then execute. This assumes that causality is determinable and therefore that we have a valid means of eliminating options. What we mean by causality is that we can relate cause and effect; if we take a certain action, we know what the effect will be — or

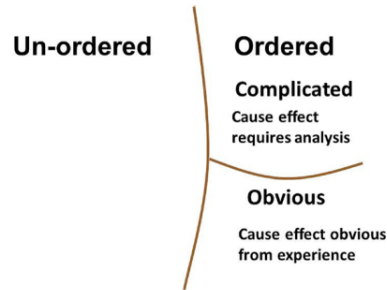
Un-ordered

Ordered

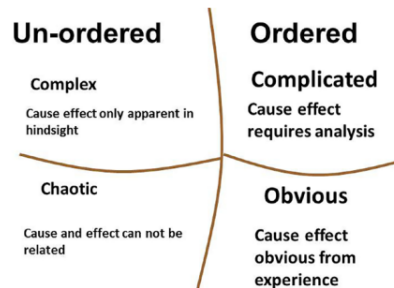


given an effect, we can determine what caused it. This is not always the case and we need to acknowledge that there are systems in which we can determine cause and effect and those in which we cannot. We call the former ordered and the later un-ordered¹ systems².

In an ordered system, the system is highly constrained, the behaviour is highly predictable, and the causality is either obvious from experience or can be determined by analysis. If the cause is obvious then we have a simple system, and if it is not obvious but can be determined by analysis, we say it is a complicated system as cause and effect (or determination of the cause) is separated by time.



For an un-ordered system, we cannot determine causality. However, we find that some of these systems are stable, and the constraints and behaviour evolve over time through the interaction of the components. They are dispositional in nature — that is, they are disposed to move or evolve in



a certain direction — but causality can only be determined in hindsight, and no amount of analysis will allow us to predict the behaviour of the system. This domain, we call complex. But there are also some systems that are not stable and which we can only described as chaotic: there are few to no constraints, and behaviour is random. There is one additional domain that needs to be considered: systems that we have not yet determined — we put these in the “disorder” bucket.

We can apply another definition to systems that are ordered as we can take them apart and put back together again — for example, a car or an aeroplane; those that are un-ordered can never be deconstructed then put back together. For example, think of making mayonnaise.

¹ Disordered is a separate state, addressed below.
² With a tip of the hat to Spencer-Brown.


We also need to acknowledge that systems are not always stable and the state of the system may change over time. A system may be stable and predictable, but its performance degrades over time or it may simply break. In the case of degradation, we may have to involve an expert (say, a mechanic in the case of a car) to analyse what is going on. In the case of breakage — for example, a car that has broken down and we have been thrown into chaos — we may need a rescue service to recover the vehicle and take it to a garage in order for the problem to be analysed. We have moved from obvious (driving the vehicle) to chaotic (the vehicle is broken) to complicated (we need a mechanic) to hopefully return to obvious (the car is again working).

The Cynefin framework

These domains of obvious, complicated, complex, and chaos, along with disorder (not yet determined), are the domains of the Cynefin complexity framework. The framework allows us to describe reality, and gives us techniques and practices that can be applied to manage in the complicated and complex domains. These practices complement the traditional approaches that are applicable where order holds. It is not the case that the old practices don't work, but we must realise that they only work within certain boundaries and that if we are not in an ordered space, they are not applicable.

Obvious³ (known knowns): Here, we know what we are doing and have seen it a thousand times before, so we sense, categorise, and respond (S-C-R). We expect to see best practices employed.

Complicated (known unknowns): We don't know what is going on but we know that we can analyse what has happened and work it out, so we sense, analyse, and respond (S-A-R). This is the domain of good practice.



³ This domain was previously called "simple".

Complex (unknown unknowns): We cannot determine what will cause a particular outcome but we can run some experiments to see if they move us in the right direction, so we probe, sense, and respond (P-S-R). This is also the domain of multiple hypotheses: there is no right or wrong answer, so we may want to run a series of experiments or run a number in parallel. This is the domain of emergent practice.

Chaotic (unknowable unknowns): Here, the system is not stable but we need to do something as it is not viable to wait, so we act (do something), sense, and respond (A-S-R). This is the domain of novel practice.

Disorder (not determined): These are the items whose domain we have yet to determine.

These practices are summarised in the following table, which also suggests what we should consider in each domain. This touches on a number of practices such as crews as an alternative to teams or social-network simulation (SNS) — all of which indicate the richness of Cynefin.

	NATURE	RESPONSE	PREPARE
CHAOS	<p>Unknowable unknowns Temporary state - no time No evidence of any constraint High turbulence no patterns Old certainties no longer apply</p>	<p>Act-Sense-Respond Speed of authoritative response vital Follow and enforce heuristics Focus on constraints not solution Use the opportunity to innovate</p>	<p>Establish heuristics & Parables Human sensor networks Crews as crisis management teams Simulation games for key managers Multi-perspective dissent feedback</p>
COMPLEX	<p>Unknown unknowns Messily coherent, patterns discernible Partial changing constraints Flux within stabilities Evidence supports contradiction</p>	<p>Probe-Sense-Respond Monitor safe-to-fail experiments All contradiction within heuristics Flex constraints to manage emergence Agility key to amplification/dampening</p>	<p>Create time and space for reflection Human sensor networks operational SNS & like to create networks Scenario planning (inc. micro) Build and monitor 'requisite diversity'</p>
COMPLICATED	<p>Known unknowns Ordered, predictable, forecastable Constraints evident and enforceable Stable with constraints Evidence susceptible to analysis</p>	<p>Sense-Analyse-Respond Determine experts or process to resolve Manage & enforce process Monitor effectiveness of constraints Focus on exploitation not exploration</p>	<p>Right people & process, right time Process engineering with feedback Diversity of experts in network Sound analytical practice Stand aside but stay in touch</p>
OBVIOUS	<p>Known unknowns Familiar, certain, well worn pathways Constraints self-evident to all Stable within universal constraints Self evident solutions</p>	<p>Sense-Categorise-Respond Ensure sound process in place Monitor for noncompliance & deviance Test for complacency Protect some pet mavericks</p>	<p>Watch for outliers Usable process Right support people for key staff Automate, but not automata Anonymous appeal/whistle blowers</p>

This is a development of the original framework that appeared in the November 2007 Harvard Business Review paper on decision making. Let's look at a situation and see how to use the framework:

The first thing we need to do is to differentiate between order and un-order: have we seen this before or have any experience that can be leveraged? If so, then the problem is ordered. If we haven't seen this before and it is truly novel, then this is something that we may be better off explor-

ing first; therefore, we treat it as un-ordered. If there is no time to consider the domain in which we are operating, then we should consider the worst case, not the most simple — which means that if it is ordered, then deal with it as complicated and engage experts and if it is un-ordered, then treat it as chaotic and act in an attempt to stabilise the situation. This avoids oversimplifying the situation and fooling ourselves that we know the answer, only to find that we don't.

Considerations

There are a couple of general points that we need to discuss before we talk specifically about Cynefin practices. These relate to complicated and complex domains.

Groupthink/naivety: The issue with the complicated domain, the domain of experts, is that there is a tendency towards groupthink. There is value in groups as evidenced by the wisdom of crowds, which is based on the average view of a group of experts/experienced people in the area with the key point that the individuals are not allowed to engage each other. The participants may not disclose their guesses to each other and each estimate must be made in isolation to avoid participants influencing one other. This is not necessarily practical on a day-to-day basis so we need an effective strategy to address this and to improve information scanning.

Alfred Sloan said, “Gentlemen, I take it that we are all in complete agreement on the decision here. Then, I propose that we postpone further discussion... to give ourselves time to develop disagreement and perhaps gain some understanding of what the decision is all about.”⁴

The issue of groupthink can be addressed by engaging diverse groups of individuals and ensuring that some of these come from other domains of expertise and therefore provide a naive view.

The mantra is that there are no dumb questions, but you need people who are willing to ask such questions. Engaging experts from diverse domains can provide this. It is worth noting that President Franklin D. Roosevelt dragged people from one meeting to another if they were bright and he thought that they would provide an interesting point of view (Chip

⁴ Summing up of a GM senior executive meeting – see <http://www.economist.com/node/13047099>

Heath and Dan Heath 2013). Dave Snowden (Fitzsimmons 2014) tells a story of the engagement of anthropologists to study the management structure of a hotel chain. This gave rise to an interesting chapter in the report on the mating habits of the staff. This should be considered when you engage the experts.

Multi-hypotheses: We have talked about there being no clear answer in the complex domain, which means that we will have competing views. Remember, in this domain, cause and effect are not directly related so we have messy coherence and may be able to ascertain patterns forming over time.

We need to understand and accept that this is not bad in itself and sometimes needs to be embraced. This means that we are not looking for the right answer but for a series of ideas that we can test to find out what works, as more than one idea may be viable. The complex domain is the domain of multi-hypotheses, so aim to run a number of parallel experiments that test these hypotheses and maximise the potential for learning.

Conflict resolution

The use of multi-hypotheses is also a useful technique for conflict resolution. The hypotheses can be outlined and experiments run to determine which are valid and should be progressed. It may be the case that opposing ideas have validity and have benefit. See the portfolio forms in the portfolio-management section for an outline of the experiments.

There is a good example in the Heath brothers' *Decisive* where they discuss the largest direct-car-sales company in America. This started as an experiment to see if people would consider buying cars online and the response was overwhelming. The experiment had to be stopped as the company was losing money because the trial had them purchase cars from traditional dealers to meet the online orders. This type of issue is going to become more common in the future and we need a more flexible and comprehensive model.

Again, naivety has a role to play as we cannot assume that we know what will happen. As long as someone can make a coherent argument, their idea has value and should be explored.

Obliquity: The last point to make is that of the need to recognise the value of obliquity, which is the practice of achieving objectives indirectly. You may know that Apple doesn't focus on traditional management

measures but uses net promoter score (NPS) to assess satisfaction from the customer's perspective.⁵ It uses NPS to guide its product development and marketing efforts. It does make use of traditional metrics internally but these are not targets in themselves.

To expand on this, I'll use an example from sailing. Whether sailing upwind or downwind, we are trying to get to the next mark as quickly as possible and the common metric used to indicate progress is VMG (velocity made good). VMG is a direct measurement that doesn't take into account all the other variables that are in play such as wind strength, point of sail, etc., so we typically use an oblique measure, which is that of target velocity. The boat's performance characteristics are used to determine the optimum speed given the wind strength and point of sail. The helm then steers the boat based on this indirect measure, which leads to optimisation of the VMG to the next mark (I would expect the navigator to have VMG on his instruments but not on the main instruments so the helm is oblivious to the direct measure).⁶

What we ideally want to use is an oblique approach, as this avoids the cognitive bias that comes from using a traditional measure that can be gamed. This is explored in John Kay's excellent *Obliquity: Why Our Goals Are Best Achieved Indirectly*.

Closing comments

Cynefin provides an approach and a set of practices for addressing the uncertainty that increasingly faces management today. It provides them with the means to realise that they are facing a messy, intractable problem and with the tools to enable them to make progress in this imperfect world.

This is a significant change from the more traditional approaches, which try to reduce a problem to a set of rational actions and acknowledges that

5 Interestingly you cannot optimise for customer satisfaction and shareholder value. See Roger Martin's "The Age of Customer Capitalism."

6 This also allows velocity headers and lifts to be handled, but a discussion of steering to targets is not the focus of this paper. You find dinghy sailors do this intuitively, hence the mantra of ease-hike-trim, and one of the issues that they have when moving to large boats is dealing with the array of information available on a typical race boat. If you are interested in this area, I suggest Will Oxley's recently published "*Modern Race Navigation*".

in some instances we cannot predict the outcomes. Instead of obsessing about predicting the future, we can move to controlling the future, and we therefore don't need to predict everything. This is the value of Cynefin. The portfolio-management section considers the practices related to the running the experiments.