UNCLASSIFIED

Australian Government
Department of Defence
Defence Science and Technology Group

Certainty and its discontents

Dr Darryn Reid
Principal Scientist
Joint & Operations Analysis Division
Defence Science & Technology Group

Certainty and its discontents

Abstract:

We find ourselves, happily, on the cusp of scientific revolution. Recent years have seen a rapidly growing recognition, across many scientific fields, of the ubiquitous presence in application domains of deep uncertainty and complexity. We often hear the problem couched in terms of a metaphoric hedge of uncertainty and complexity phenomena that has sprung up and flowered over our familiar methods, as if complexity and uncertainty were something new, at least in measure if not in essence, and as if our growing awareness were primarily an empirical matter. Yet arguably what is new is not the presence of complexity and uncertainty phenomena in the worlds of our studies, nor the availability of the empirical evidence, but rather the expanding transdisplinary scientific basis in theory by which we look for the evidence in the first place and then comprehend its meaning.

Uncertainty manifests in defence and military circles overwhelmingly in problem contexts dominated by the potential for catastrophic failure, particularly cascading systemic failure; indeed, it is difficult to imagine a defence problem domain where unpredictability and failure are not predominant - the predominant - features. Yet traditional analysis approaches often struggle to handle multiple kinds of failure and often fail to treat deep uncertainty at all. The convenience of degenerate symmetry conditions that do not hold in messy application domains is simply unsustainable. Moreover, uncertainty also creates opportunity just as readily as it does the possibility of failure; there is much more to handling failure and opportunity under uncertainty than maximising expected future utility. While defence and military matters are hardly unique in this regard, the extreme nature of the problem environments here perhaps brings manifestations of deep uncertainty through failure potential and opportunity into starkest relief.

The need for expansion of both the fundamental theory about the nature of uncertainty and how we might effectively deal with it and the application of theory to practical matters across the sciences suggests mathematics may find itself again poised in a uniquely central position. In a time when the almost ubiquitous emphasis in research resource allocation remains on demonstrating immediate application, the need and opportunity for the development of coherent new mathematical theories to support and enable the applied sciences has ironically perhaps never been greater. We find ourselves, happily, on the cusp of a mathematically enabled revolution.

Dr Darryn J Reid is Principal Scientist in the Defence Science and Technology Group, and has been with DST since 1995. He has worked in distributed systems, machine learning and artificial intelligence, interoperability, formal reasoning and logics, operations research, simulation, optimisation and optimal control, electronic warfare, intelligence analysis, missile targeting and control, command support systems, complexity, nonlinear dynamics and ergodic theory, web-based technologies, software development, functional languages, formal languages and model theory, theory of computation and algorithmic information theory, crowd modelling, economic theory and military theory. He holds the degrees of Bachelor of Science in Mathematics and Computer Science, Bachelor of Science with First Class Honours in Mathematics and Computer Science, and Doctor of Philosophy in Theoretical Computer Science from the University of Queensland. He has strong research interests in pure and applied mathematics, theoretical and applied computer science, philosophy, military theory and economics. He is also an artist. In other words, he knows just enough to understand how ignorant he is. He is currently trying to age as disgracefully as possible, with the support of his beautiful wife Julie and their son Tyler.

Cherished certainties

Ontological uncertainty Keynes 1922

Epistemological uncertainty Knight, 1922

Neutrality of Money ‘axiom’

Gross substitution ‘axiom’

Ergodic system ‘axiom’

Sampling the future in the system model is equivalent to sampling the past

The system model’s time average converges to its ensemble average

The system model is stationary and regular

Ergodic

Fully observable

The art and science of military thinking

Antoine-Henri Jomini

Dietrich von Bülow

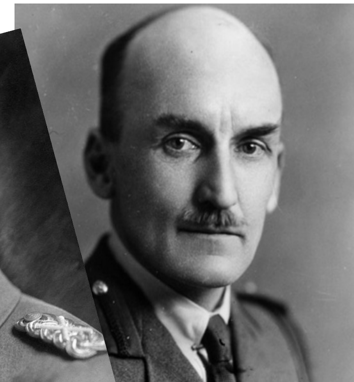

J.F.C. Fuller

B.H. Liddell Hart

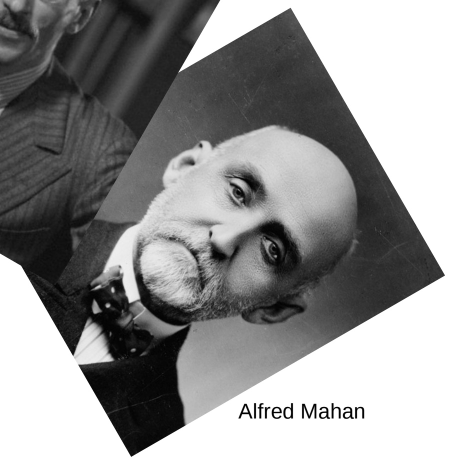

Alfred Mahan

1. **Rationality as the basis for knowledge development and decision-making**
2. **Look to the sciences for account of method**
3. **Knowledge as verified truth**

   The goal of military theory: a complete and correct theory of war and battle
   Decision-making as the making of correct decisions
   Knowledge as pushing back uncertainty as a diminishing residual

Symmetries and their invariants

Ontological uncertainty Keynes 1922

Epistemological uncertainty Knight, 1922

knowledge as verified truth

The future and past are *symmetrical*

→ We can basically ignore time!

<u>Invariant condition</u>: sampling the future in the system model is equivalent to sampling the past

The system model's time average converges to its ensemble average

The system model is stationary and regular

Ergodic

Fully observable

Third Choice: disputed from the start

Carl von Clausewitz, as a Russian Colonel

Gerhard Scharnhorst

August von Gneisenau

Helmuth von Molkte the Elder

"Efforts were therefore made to equip the conduct of war with principles, rules, or even systems. This did present a positive goal, but failed to take an adequate account of the endless complexities involved." [emphasis mine]

- Carl von Clausewitz

"No matter how obvious and palpable the difference between knowledge [science] and ability [art] may seem … it is still extremely difficult to separate them entirely … it is impossible to imagine a human capable of perception but not of judgement, or vice versa, it is likewise impossible to separate art and knowledge ..." [emphasis mine]

- Carl von Clausewitz

Clausewitz: uncertainty *and* chance

Ontological uncertainty
Clausewitz c.1820-1830, Keynes 1922

Epistemological uncertainty Knight, 1922

The future and past are *symmetrical*

→ We can basically ignore time!

<u>Invariant condition</u>: sampling the future in the system model is equivalent to sampling the past

The system model's time average converges to its ensemble average

The system model is stationary and regular

The secret life of paradoxes, part 1

Self-referential Paradox.

Gödel sentence: “This sentence is unprovable.”

$T \vdash \rho \equiv Pr(\overline{\neg\rho})$

“If this sentence is provable, then there is a shorter proof of its negation.”

$T \vdash \rho \equiv \exists y(Proof(\overline{\neg\rho},y) \wedge \forall v<y \neg Proof(\overline{\rho},v))$

$T \vdash \rho \equiv \exists y Proof(\overline{\neg\rho},y) \exists y Proof(\overline{\rho},y)$

$T \vdash \rho \equiv Pr(\overline{\neg\rho}) Pr(\overline{\rho})$

Using Provability Logic GL:

fix p ↔ ~□p

fix p ↔ ~□p □~p

f = ~□(◊T) → ◊T

Fixed-Point Paradox.

“If the system proves its own consistency, then it is inconsistent;
if it is consistent, then it cannot prove its own consistency.”

“The Halting Problem for Turing Machines is unsolvable.”

The limits of knowing!

Infinite Regress Paradox.

“1. The sentences numbered 2 or more are unprovable.”
2. The sentences numbered 3 or more are unprovable.”
…
“n. The sentences numbered n+1 or more are unprovable.”
…

The secret life of paradoxes, part 2

Self-referential Paradox.

Gödel sentence: “This sentence is unprovable.”

$$T \vdash \rho \equiv Pr(\overline{\neg\rho})$$

“If this sentence is provable, then there is a shorter proof of its negation.”

$$T \vdash \rho \equiv \exists y(Proof(\overline{\neg\rho},y) \wedge \forall v{<}y \neg Proof(\overline{\rho},v))$$

$$T \vdash \rho \equiv \exists y Proof(\overline{\neg\rho},y) \exists y Proof(\overline{\rho},y)$$

$$T \vdash \rho \equiv Pr(\overline{\neg\rho}) Pr(\overline{\rho})$$

Fixed-Point Paradox.

“If the system proves its own consistency, then it is inconsistent;
if it is consistent, then it cannot prove its own consistency.”

“The Halting Problem for Turing Machines is unsolvable.” → **Algorithmic Entropy!**

Infinite Regress Paradox.

“1. The sentences numbered 2 or more are unprovable.”
2. The sentences numbered 3 or more are unprovable.”
…
“n. The sentences numbered n+1 or more are unprovable.”
...

Entropy, reloaded

Entropies are invariants under class symmetries (isomorphisms)

Different ways of describing essentially the same system

Algorithmic entropy is about information content:

Invariant under effective procedure equivalence

K(ω) is the size of the shortest procedure yielding ω

K(ω|φ) is the size of the shortest procedure yielding ω from φ

Algorithmic entropy is about compression/decompression

We can use entropy to measure effectively the amount of information from *outside* the system we would need to solve our problem

* There are different entropy formulations dependent on context

** Entropy in economics: 'exposure'

Ergodicity revisited

Ontological uncertainty
Clausewitz c.1820-1830, Keynes 1922

Epistemological uncertainty Knight, 1922

Ergodic

Fully observable

The future and past are *symmetrical*

→ We can basically ignore time!

Invariant condition: sampling the future in the system model is equivalent to sampling the past

The system model's time average converges to its ensemble average

The system model is stationary and regular

The system model is technically *trivial*

The system is guaranteed to be paradox-free

The uncertainty of future events given past events decreases to zero

Split rationality of Operations Analysis

Operations Analysis (OA) was born in WW1 and coalesced in WW2

“The systematic study of operational application of capability.”

“The ‘quantitative’ determination of the most effective use of limited resources.”

There were many notable successes …

… but by the 1960s the limits were well recognised …

Expected Future Utility Maximisation in uncontrolled military environments?

In environments whose dominant feature is deep uncertainty?

Where the dominant manifestation is ubiquitous potential for terminal failure?

Where the the other major manifestation is opportunities of unpredictable realisation?

Where futures are acutely and unpredictably sensitive to past decisions?

Where there is no single definitive problem formulation?

Where there is a mess of interacting features and components?

So where do we go from here?

U-boat Type VIIC under air attack

The insufficiency of ergodicity

Ontological uncertainty
Clausewitz c.1820-1830, Keynes 1922

Epistemological uncertainty Knight, 1922

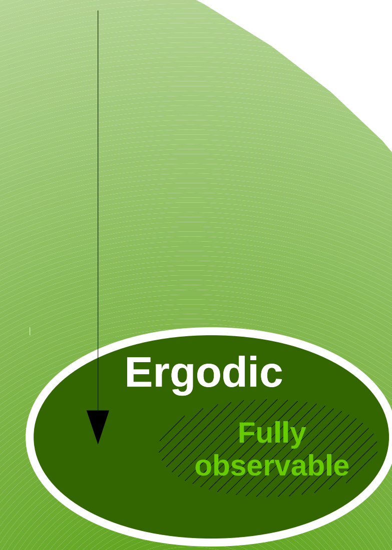

The future and past are not symmetrical

We cannot ignore time

Samples from the future is cannot be substituted with samples from the past

The system model's time average does not converge to its ensemble average

The system model is not stationary and probably not regular either

The system model admits self-reference

The system contains paradoxes that manifest in fundamental uncertainty

Modelling Complex Warfighting

Motivation: Available current approaches are inadequate for the conditions of extreme uncertainty, complexity and uncertainty and terminal failure potential of our problem domains.

Purpose: to revolutionise Defence & military Operations Analysis, around complexity and uncertainty.

Topic: determining and evaluating options for the effective allocation of limited Defence resources under Force Design and Employment, where the consequences of the decisions are realised in futures that the decision-makers cannot know nor fully control at the time of the decision.

Approach: multiple transdisciplanary teams of DST with a diverse group of academic parters who are co-investing in building national capability in Operations Analysis research and application for Defence & National Security.

The unnecessity of ergodicity

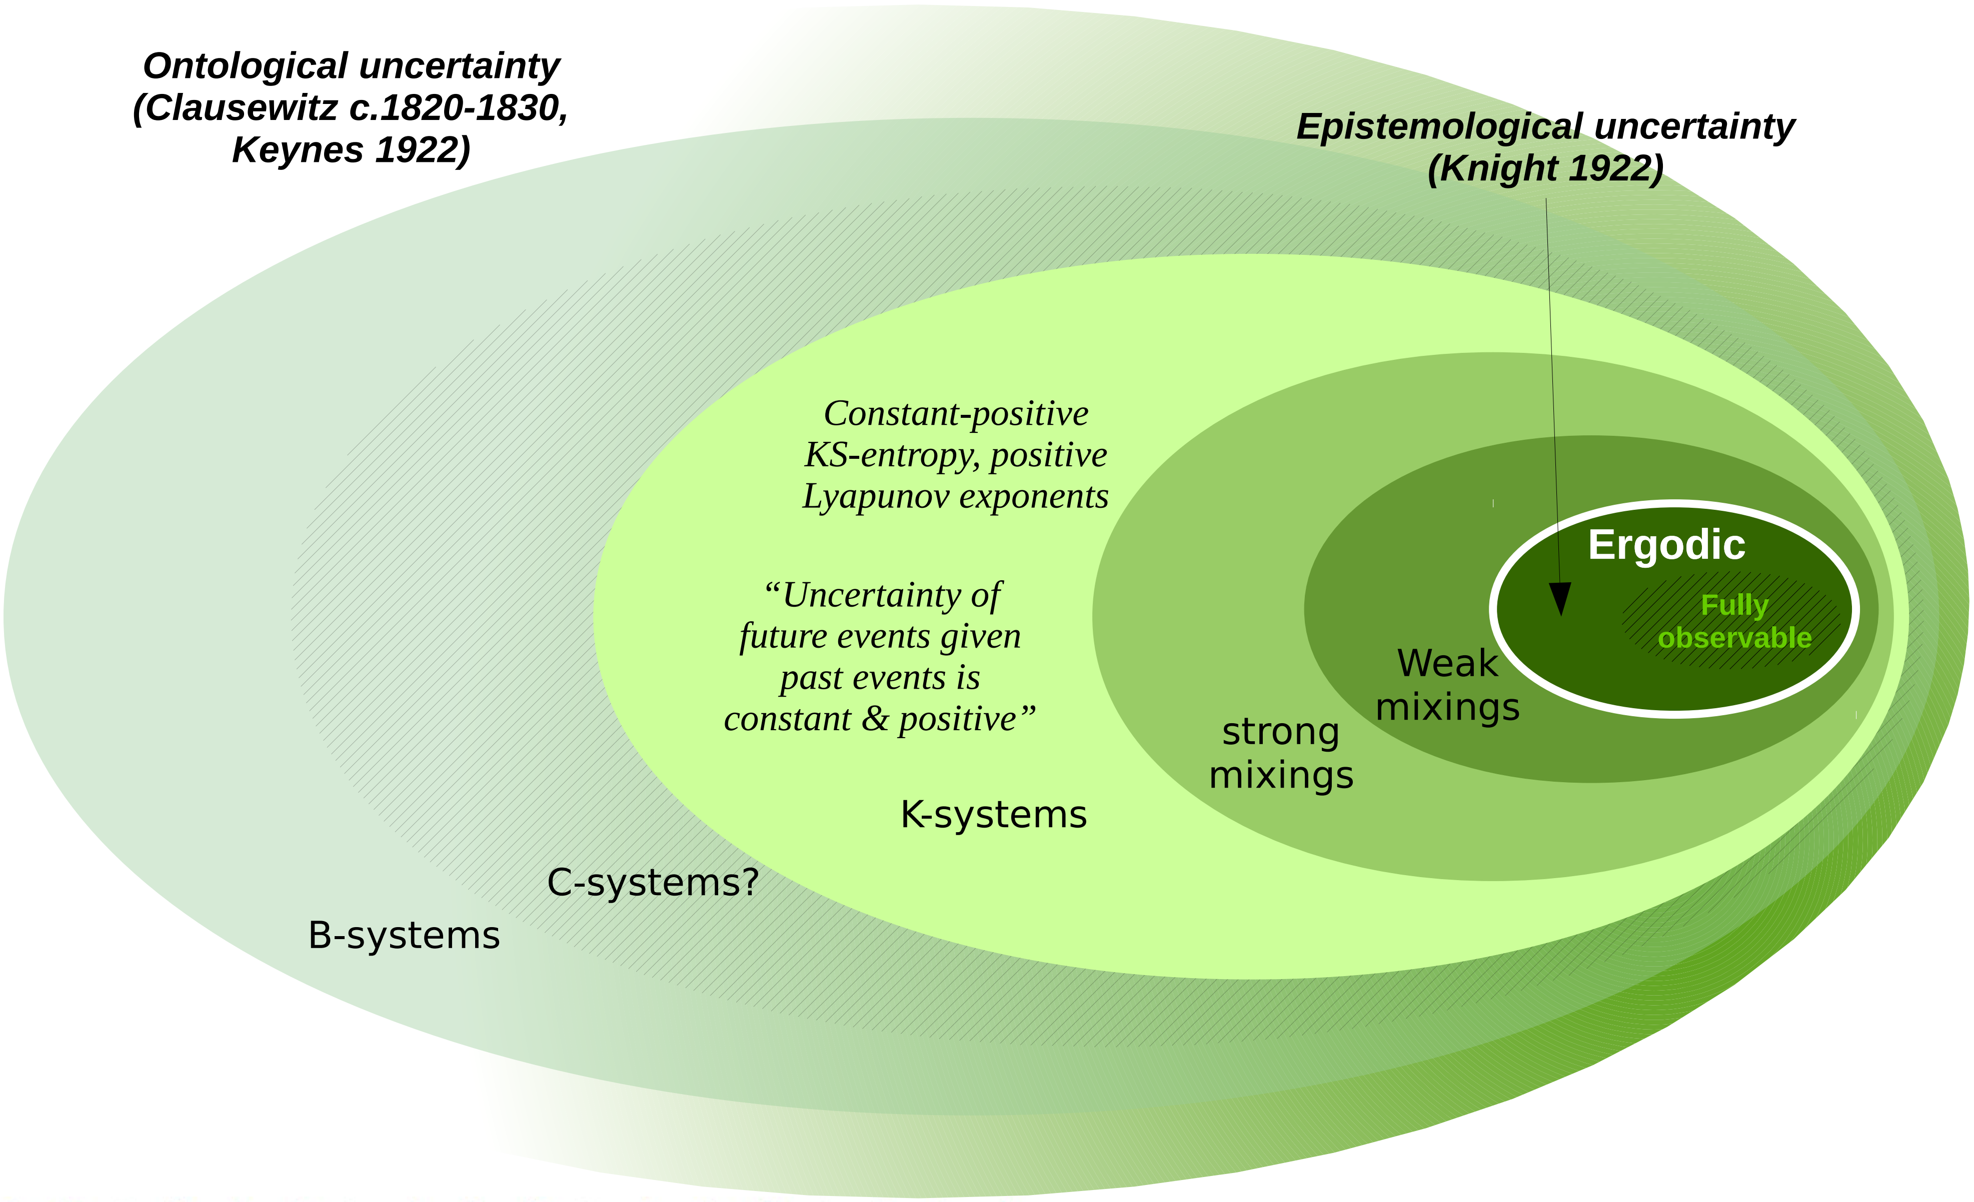

Beyond ergodicity

"Time is a mechanism for stopping everything from happening at once"

- Introducing time and allowing non-stationarity
- Identifying paradoxes directly (often self-reference involving time)
- **Identifying alternate symmetries and their corresponding invariants**

  - the formulation of entropy follows from isomorphism class
- Formulating around entropy of events rather than the events themselves
- Formulating around entropy of costs rather than costs themselves
- Presence of terminal failure states is a primary feature
- Opportunities are just affordable failures with viable potential pay-off
- Feasible scenario spaces (path dependencies) rather than tuned scenarios

Mathematical potentialities

Measures of relative utility with respect to different command functions of force design options

Determine crucial sensitivities to failure in force options (rather than maximising efficiency or expected utility)

Feasible scenario spaces and path dependencies

Generating novel events in simulations

Entropic data sets to determine
dynamic phase shifts

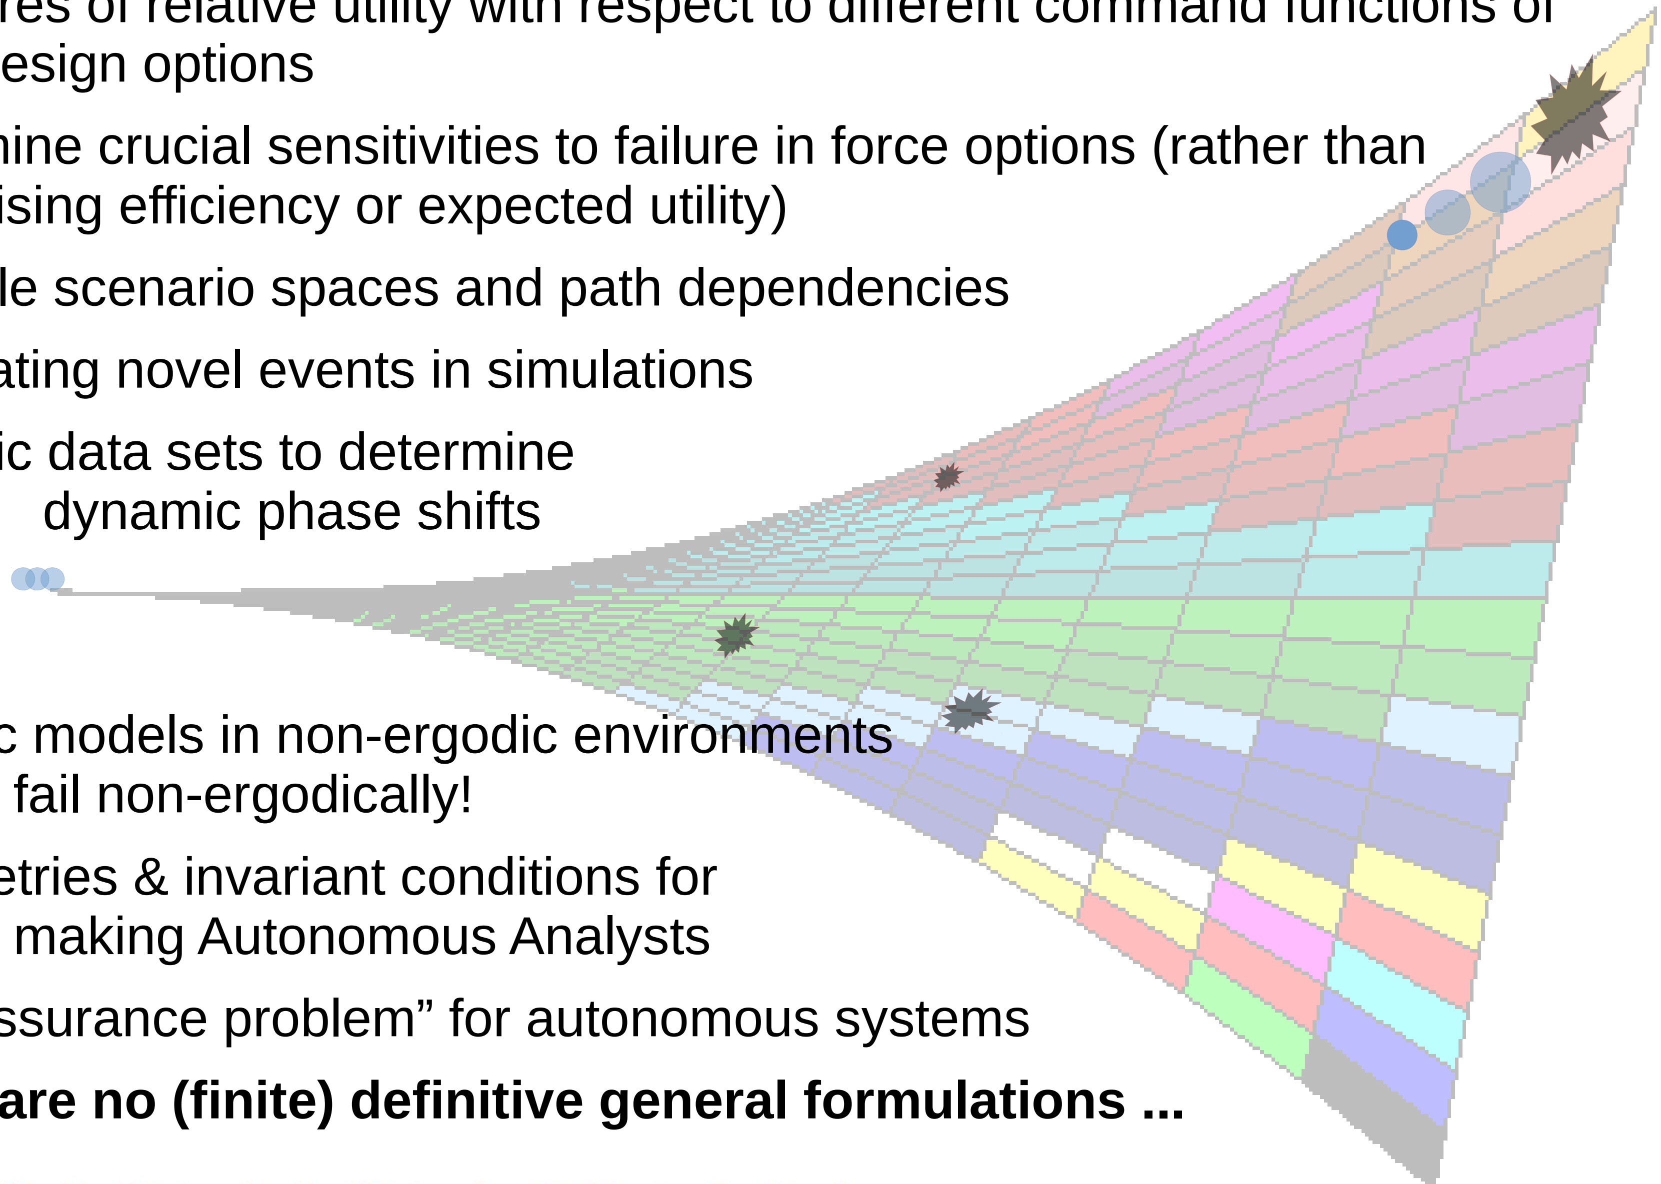

Ergodic models in non-ergodic environments
fail non-ergodically!

Symmetries & invariant conditions for
making Autonomous Analysts

The “assurance problem” for autonomous systems

There are no (finite) definitive general formulations ...

In closing

It simply does not matter how good a model is on average if we cannot bear the consequences the first time the strong assumptions we build it around slam unhappily into a nonconforming reality.

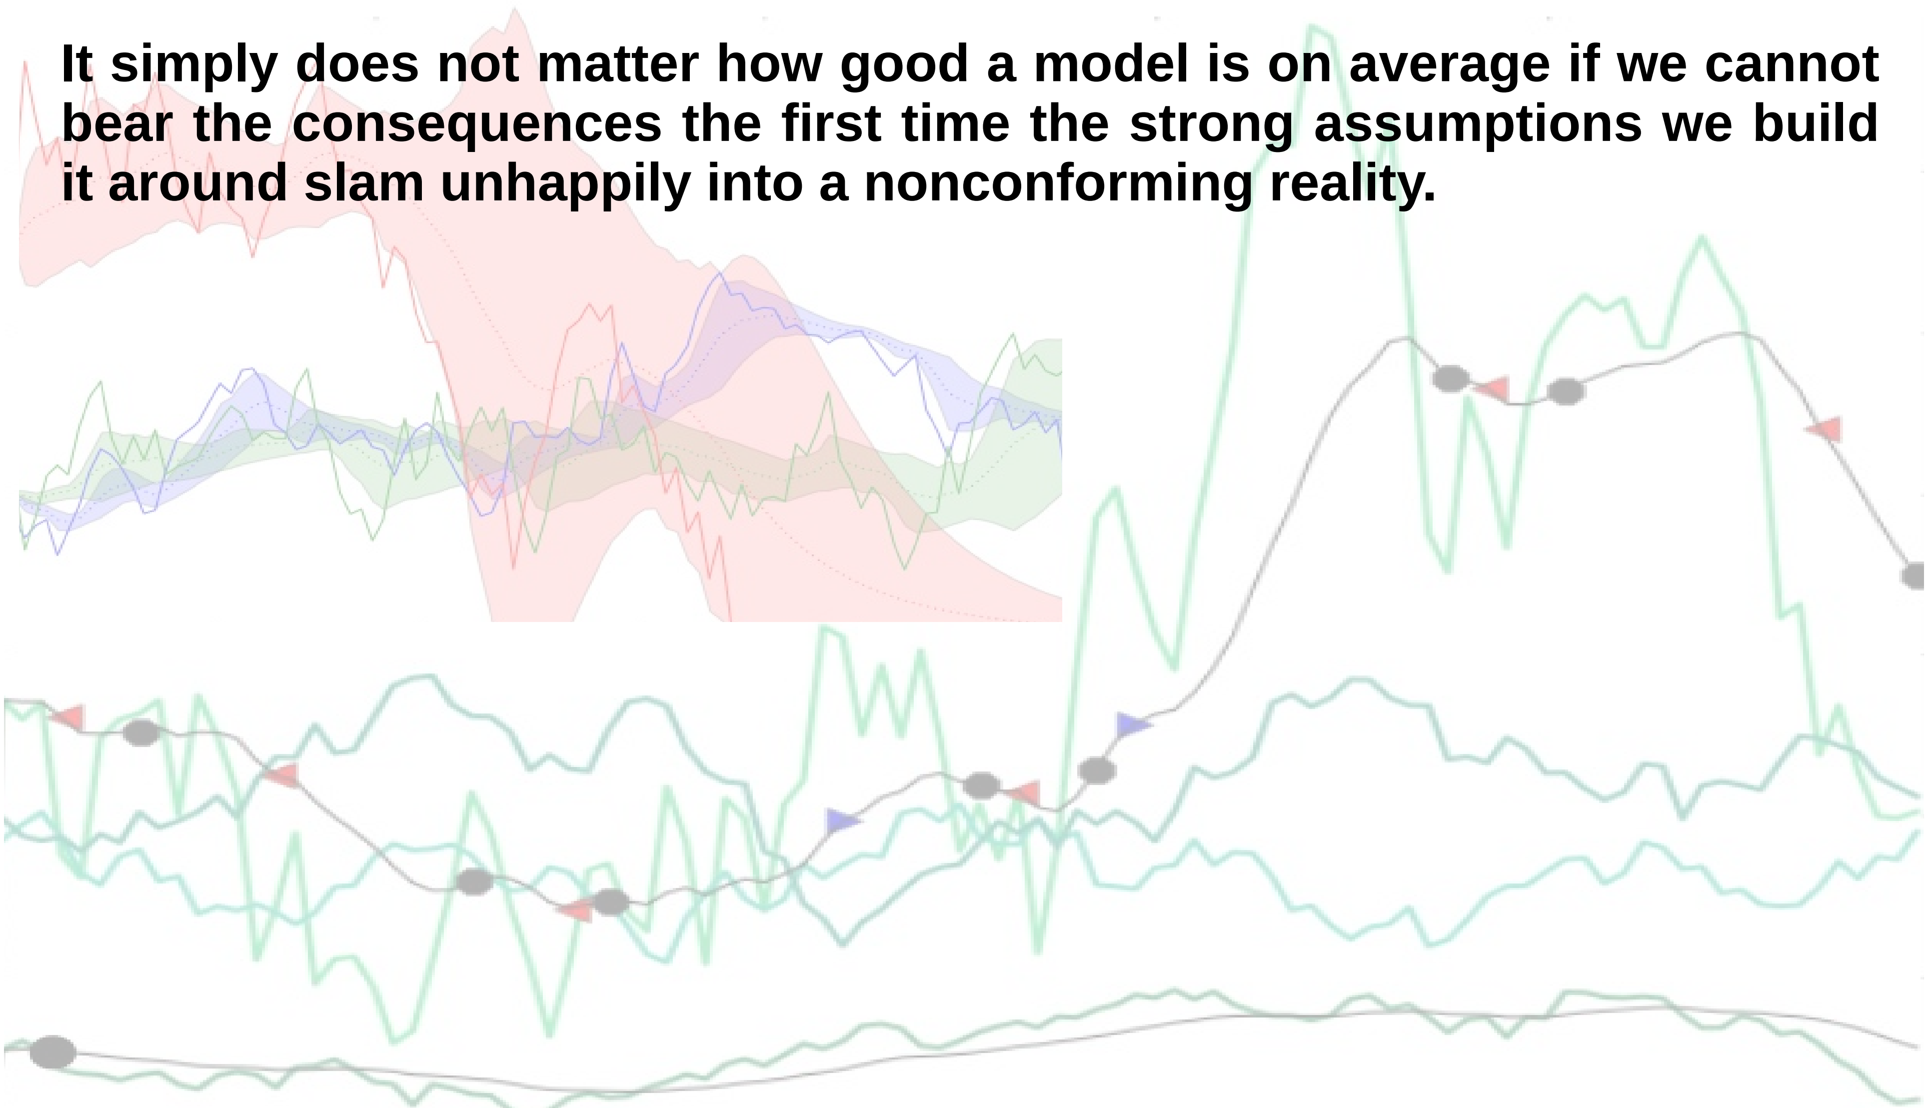

Abstract:

We find ourselves, happily, on the cusp of scientific revolution. Recent years have seen a rapidly growing recognition, across many scientific fields, of the ubiquitous presence in application domains of deep uncertainty and complexity. We often hear the problem couched in terms of a metaphoric cliff of uncertainty and complexity phenonema that has sprung up and flowered over our familiar methods, as if complexity and uncertainty were something new, at least in measure if not in essence, and as if our growing awareness were primarily an empirical matter. Yet arguably what is new is not the presence of complexity and uncertainty phenomena in the worlds of our studies, nor the availability of the empirical evidence, but rather the expanding transdisplinary scientific basis in theory by which we look for the evidence in the first place and then comprehend its meaning.

Uncertainty manifests in defence and military circles overwhelmingly in problem contexts dominated by the potential for catastrophic failure, particularly cascading systemic failure; indeed, it is difficult to imagine a defence problem domain where unpredictability and failure are not predominant - the predominant - features. Yet traditional analysis approaches often struggle to handle multiple kinds of failure and often fail to treat deep uncertainty at all. The convenience of degenerate symmetry conditions that do not hold in messy application domains is simply unsustainable. Moreover, uncertainty also creates opportunity just as readily as it does the possibility of failure; there is much more to handling failure and opportunity under uncertainty than maximising expected future utility. While defence and military matters are hardly unique in this regard, the extreme nature of the problem environments here perhaps brings manifestations of deep uncertainty through failure potential and opportunity into starkest relief.

The need for expansion of both the fundamental theory about the nature of uncertainty and how we might effectively deal with it and the application of theory to practical matters across the sciences suggests mathematics may find itself again poised in a uniquely central position. In a time when the almost ubiquitous emphasis in research resource allocation remains on demonstrating immediate application, the need and opportunity for the development of coherent new mathematical theories to support and enable the applied sciences has ironically perhaps never been greater. We find ourselves, happily, on the cusp of a mathematically enabled revolution.

Certainty and its discontents

Abstract:

We find ourselves, happily, on the cusp of scientific revolution. Recent years have seen a rapidly growing recognition, across many scientific fields, of the ubiquitous presence in application domains of deep uncertainty and complexity. We often hear the problem couched in terms of a metaphoric hedge of uncertainty and complexity phenomena that has sprung up and flowered over our familiar methods, as if complexity and uncertainty were something new, at least in measure if not in essence, and as if our growing awareness were primarily an empirical matter. Yet arguably what is new is not the presence of complexity and uncertainty phenomena in the worlds of our studies, nor the availability of the empirical evidence, but rather the expanding transdisplinary scientific basis in theory by which we look for the evidence in the first place and then comprehend its meaning.

Uncertainty manifests in defence and military circles overwhelmingly in problem contexts dominated by the potential for catastrophic failure, particularly cascading systemic failure; indeed, it is difficult to imagine a defence problem domain where unpredictability and failure are not predominant - the predominant - features. Yet traditional analysis approaches often struggle to handle multiple kinds of failure and often fail to treat deep uncertainty at all. The convenience of degenerate symmetry conditions that do not hold in messy application domains is simply unsustainable. Moreover, uncertainty also creates opportunity just as readily as it does the possibility of failure; there is much more to handling failure and opportunity under uncertainty than maximising expected future utility. While defence and military matters are hardly unique in this regard, the extreme nature of the problem environments here perhaps brings manifestations of deep uncertainty through failure potential and opportunity into starkest relief.

The need for expansion of both the fundamental theory about the nature of uncertainty and how we might effectively deal with it and the application of theory to practical matters across the sciences suggests mathematics may find itself again poised in a uniquely central position. In a time when the almost ubiquitous emphasis in research resource allocation remains on demonstrating immediate application, the need and opportunity for the development of coherent new mathematical theories to support and enable the applied sciences has ironically perhaps never been greater. We find ourselves, happily, on the cusp of a mathematically enabled revolution.

Dr Darryn J Reid is Principal Scientist in the Defence Science and Technology Group, and has been with DST since 1995. He has worked in distributed systems, machine learning and artificial intelligence, interoperability, formal reasoning and logics, operations research, simulation, optimisation and optimal control, electronic warfare, intelligence analysis, missile targeting and control, command support systems, complexity, nonlinear dynamics and ergodic theory, web-based technologies, software development, functional languages, formal languages and model theory, theory of computation and algorithmic information theory, crowd modelling, economic theory and military theory. He holds the degrees of Bachelor of Science in Mathematics and Computer Science, Bachelor of Science with First Class Honours in Mathematics and Computer Science, and Doctor of Philosophy in Theoretical Computer Science from the University of Queensland. He has strong research interests in pure and applied mathematics, theoretical and applied computer science, philosophy, military theory and economics. He is also an artist. In other words, he knows just enough to understand how ignorant he is. He is currently trying to age as disgracefully as possible, with the support of his beautiful wife Julie and their son Tyler.

I will open with something about my role being about the development of goals for research programmes that are scientifically developed, intensely scrutinised, and above all, ridiculously ambitious. So in the spirit of ambitious goals, my ambitious goal I have set for myself today to not be boring.

This sets the talk up for what I want to achieve particularly around the part about our Modelling Complex Warfighting SRI, perhaps with an implication regarding SRIs in general and what my Principal Scientist role is really all about (well, according to me as the architect of my role, at least), and JOAD's research and research capability development. I want to create a realisation or perhaps even surprise in the area about the MCW SRI concerning the ambitiousness of our research and the high tolerance for failure and how this has produced a collaborative research platform that is widely seen as immensely successful. This opening also sets up a tie-back closing to complete the point about the importance of not being boring in scientific research; while the opening is intended to be playful and flippant, it disguises the very serious and central point about not simply looking for incremental improvements with better solutions to problems we alrealy know about, but rather conceiving of completely new and better problem choices in the first place.

A bit of a pitch at the opening will be to talk a bit about the driving part in the abstract about the problem not being a dearth of data but a lack of theory by which to both drive data collection and to understand what the data means. This is to set us up for some outrageously ambitious research goals that I will talk about later, and also why although such goals are extremely ambitious are actually a lot more doable than might be manifest at first blush; the power to achieve such things lies in the deep underlying mathematical theory.

This slide is just a place to keep the abstract with the talk, and a place for a few notes.

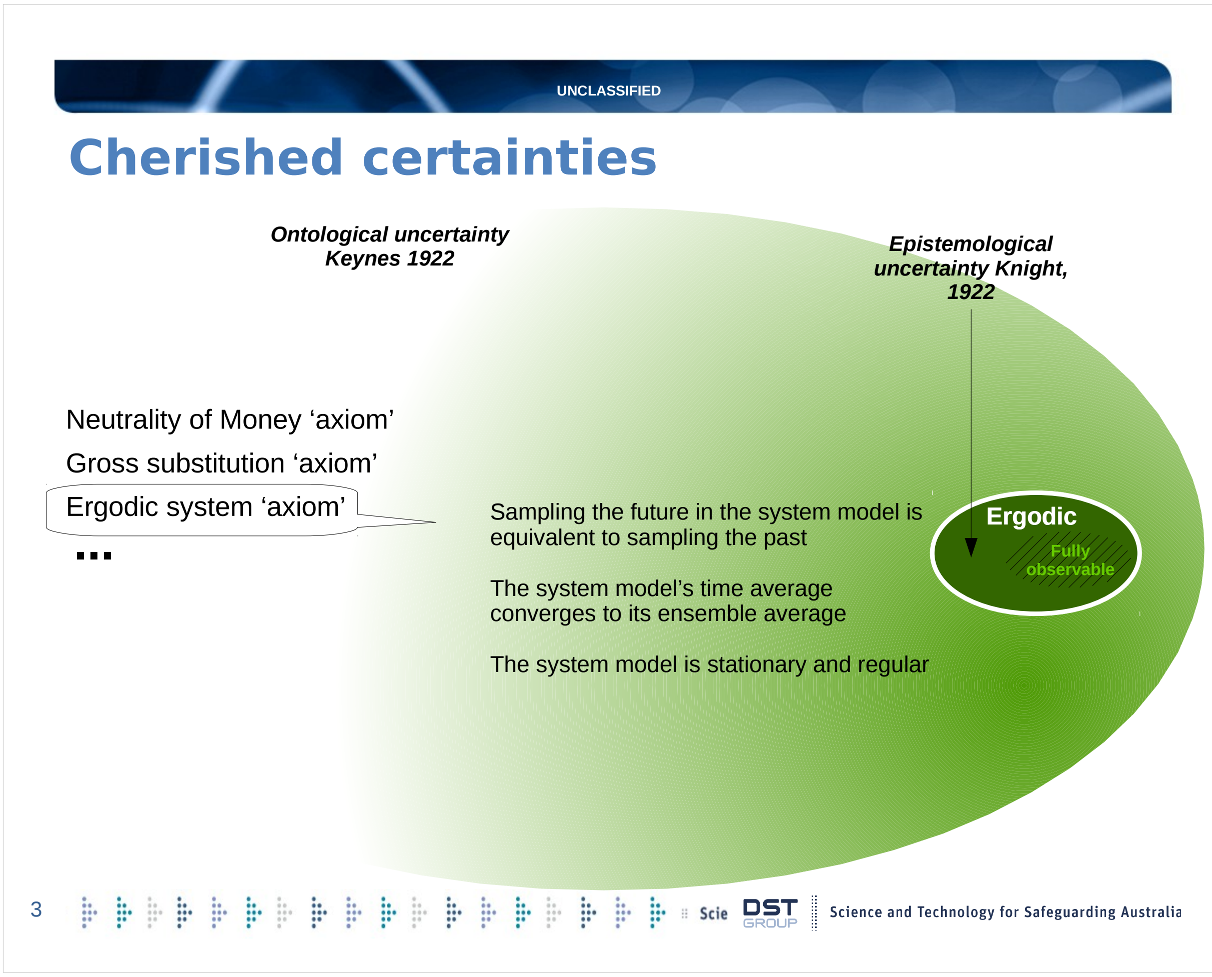

I start by touching on the foundational assumptions of economic theory: the neutrality of money axiom holds that change in monetary supply only affects nominal variables (such as prices) and not real variables (like employment and consumption); the gross subtitution says that if demand for something increases then its relative price will increase causing demand for it to spill over instead into a lower priced substitute; and, most importantly today because it will be a focal point for my talk, ergodicity is the assertion that the unrealised future can be effectively sampled by sampling instead from the past. It appears all over the place in a few different yet equivalent forms.

Ergodicity shows up all over the place: in economics, in Operations Analysis, and in Artificial Intelligence, it manifests commonly in terms of the future being statistically the same as the past; in physics, and in AI as well, it often shows up as an assumption about the time average of a system converging to the ensemble average; in mathematics, we might talk about a system being stationary and regular.

The overall diagram is intended to illustrate where two key concepts in economics about undertainty fit; both views of uncertainty were proposed in 1922 but they are very different conceptions. This is an area I think that causes a lot of confusion in economics, but the difference becomes very clear from a mathematical point of view. Knight's epistemological uncertainty – literally uncertainty arising from limits in knowledge – is about what we now often call "partial observability" in computer science, AI, and hard OA, or "state-space ignorance", or just "ignorance" in social science, behavioural economics and soft OA. Keynes' much larger and more fundamental ontological uncertainty – literally meaning uncertainty that is simply intrinsic to reality – is about what we can now describe in terms of fundamental limits to knowing. Ontological uncertainty occurs in systems where even knowing the complete rules of the system does not empower us to answer questions about the future states of that system. And expecting to know the complete rules of the system is itself unreasonable, to boot ...

The key to understanding how we can know the complete rules of the system yet be unable to answer questions about its future states is understand the difference between determinism and predictability. They are simply not interchangeable, nor are they even connected by a simple logic in the obvious manner that people tend to assume. In fact, they are only effectively equivalent under ergodicity!

UNCLASSIFIED

The art and science of military thinking

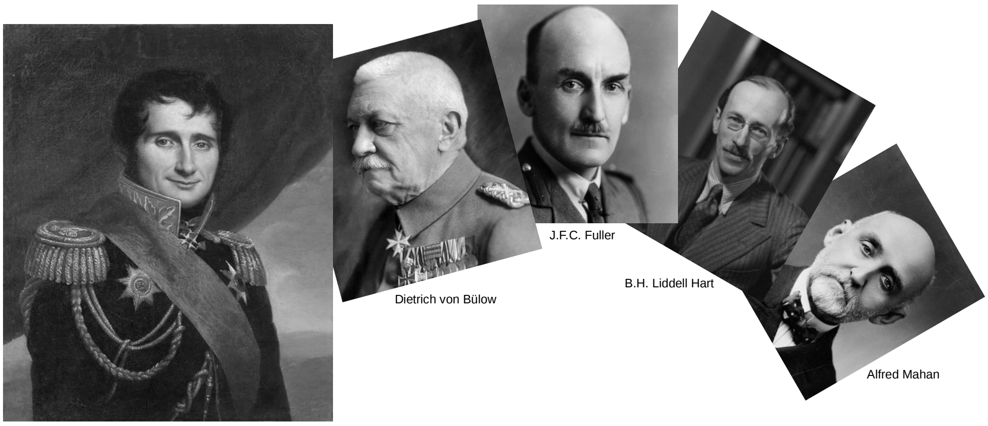

Antoine-Henri Jomini

1. Rationality as the basis for knowledge development and decision-making
2. Look to the sciences for account of method
3. Knowledge as verified truth

The goal of military theory: a complete and correct theory of war and battle
Decision-making as the making of correct decisions
Knowledge as pushing back uncertainty as a diminishing residual

We could pick up the story of how ergodicity became a ubiquitous assumption from the Enlightenment, but, to be more topical to my area of concern, we'll pick it up in the case of military theory at the Military Enlightenment. Any other field you care to name has a parallel story. There were three basic choices: the choice in favour of rationality, the choice in favour of looking to the sciences for explanation of rationality, and the Third Choice, in favour of a positivist/foundationalist account of science that has had profound implications for the field ever since.

The most influential figure here was the French-speaking Swiss Baron Antoine-Henri de Jomini, who claimed to provide the one true interpretation of Napoleon Bonaparte, arguing him to have instantiated the executive near-perfection of the conduct of war and battle. His noticeably managerially-oriented theory heavily borrowed from the mathematical terminology of the preceding Enlightenment, with geometrical schemes of effective manoeuvre; against a backdrop of widespread fascination with the rapidly emerging sciences, and where mathematics held the appearance of the embodiment of pure truth and of almost godlike power to provide absolute knowledge using the powers of analytical reason, this was, if nothing else, clever marketing; dependent on his audience for his income, Jomini was a shameless self-promoter (which annoyed people, including Clausewitz and the Duke of Wellington, who called him a "pompous charlatan").

It was more than just marketing: Jomini set in motion a programme for military theory oriented around determining for war and battle a complete universal theory; when generalised above his geometric formulations for 19th Century manoeuvre, Jomini's account of battle amounts to determining the decisive point and directing superior combat power against it. For Jomini, war and battle are essentially unchanging, tweaked around the fringes by technological and political developments, yes, but otherwise static. Professional soldiers should stick to the practical business of orchestrating the actions of forces on the battlefield according to the governing geometrical laws of manoeuvre, which he then claimed to outlay with near perfection for the benefit of an audience newly eager to develop doctrine and, in particular, professional education and training programmes.

Jomini is consequently prescriptive and instrumental – he said "instructional", which in modern terms means 'doctrinal'. He was overwhelmingly influential in military instruction: at the United States Military Academy at West Point, his works were taught exclusively up to the American Civil War, and they continue to feature prominently today, and he promoted a more scientific approach at the Russian military academy that he helped to established. Though he shows precious little intellectual development over time, this did nothing to diminish his influence; if anything, the constancy of his position served to help expand it.

Yet the goal of a universal theory sets up a conundrum, and Jomini was well aware that not everything can be reduced in what he thought a 'scientific' manner: while he overstated the proportion of relevant problems that are fully solvable with the application of the correct method, he also knew that there are problems that are not amenable to full and final resolution. In the *Art of War*, which Jomini wrote after Clausewitz' early death when he no longer posed a competitive threat, Jomini made numerous concessions to Clausewitz' influence. Most importantly here are the passages here concerning the limits of theory, limits of applicability of fixed rules beyond the tactical, and denying that war is ultimately "a positive science" capable of providing mechanistic formulations for the conduct of battle, despite his use of geometrical terminology and prescriptive intent.

Jomini's solution to the conundrum about the limits of rational enquiry was to follow what was already a well established pattern, and to cleave matters into two distinct camps: the supposed definitive certainties and analytical methods of scientific rationality on one side; on the other, the exercising of judgement, particularly judgement necessary under conditions of uncertainty, as art, with its supreme practitioner taking the form of the military genius. War and battle are thus, for Jomini, led by military heroes whose creative genius lies beyond rational penetration, while rigorous rational analysis using the methods of science provides support.

The debate over war as art or science goes back at least to the middle ages (at least to early 15th Century), but in the Military Enlightenment, what emerged was war as *both* art *and* science, with debates over the proportions generating a stragely static kind of dynamic tension ever since.

The Third Choice for a positivist notion of science manifests as a tension between art and science because of the narrowness of application that this notion of science yields. Art is allegedly about dealing with complexity because allegedly science cannot, and is presented as domain of genius.

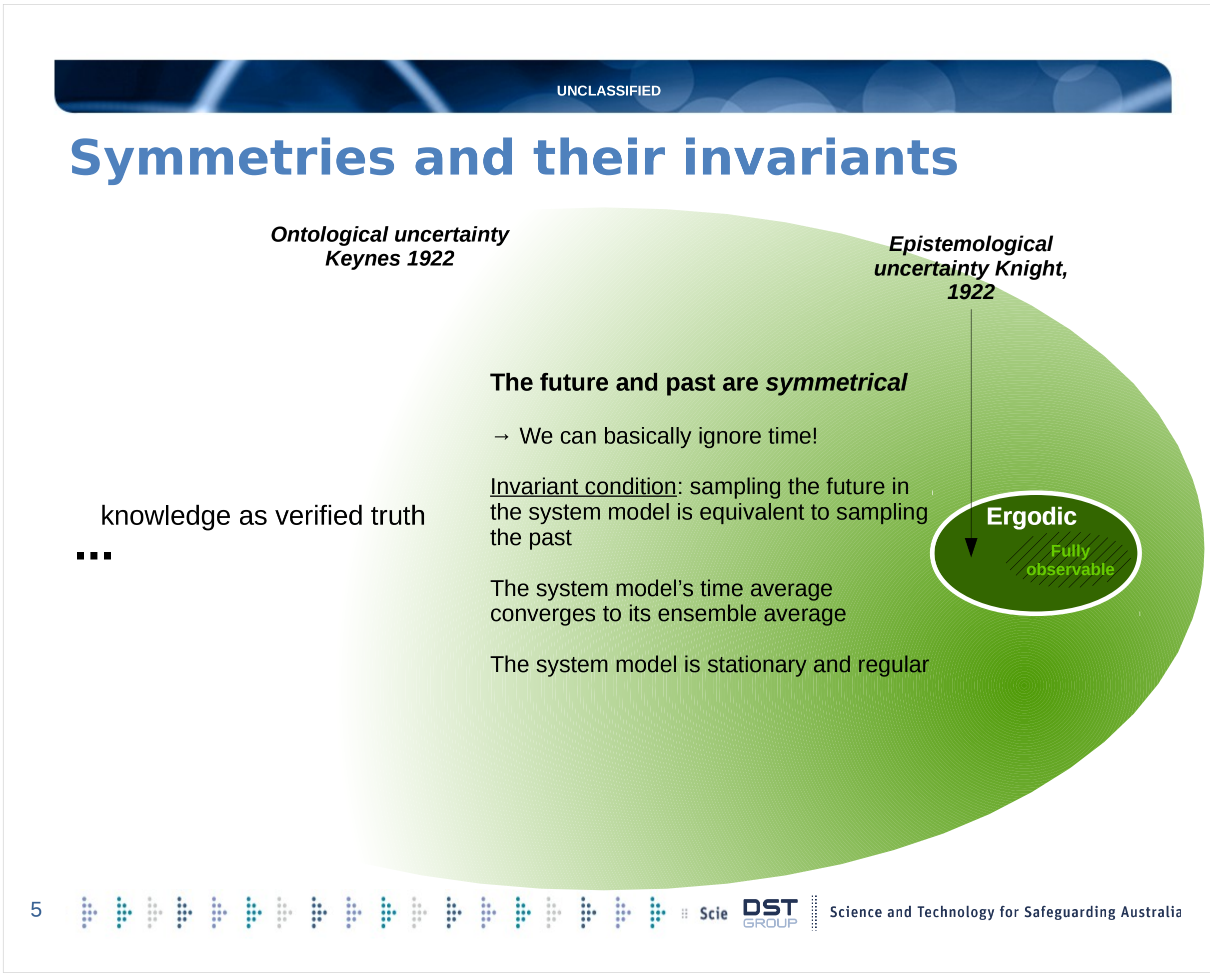

Returning to populate our picture of possible kinds of systems, I will talk here about how the failed conceptualisation of science as being about verified truth and thus with an unflinching obsession with prediction (and hence the need for assuming predictability) shows up in a concrete sense across our fields of study in the form of ergodic assumptions (though they are not always explicitly identified, such is the hold that this view of science has had on our thinking). Notice here too the loss of context (and I want people to hold this thought in their minds as we proceed): this view of science shows up in as prediction. Notice the lack of unqualification by context. It means predictability of anything we might just consider, and, in particular, it means the strongest kind of predictability of all: the predictability of future states or future outcomes. A good question at this point to start to think about would be to ask "prediction, exactly, of *what*, other than future outcomes, suffices to inform decision-making?

The main point here is to take another look at ergodic assumptions and to see them as an underlying *symmetry* assumption: the future is (statistically) the same as the past. Consequently, ergodic systems are effectively static and we can effectively ignore time. As with symmetry conditions that play a fundamental role throughout physics, this symmetry condition shows up as some kind of invariance condition. Note that in physics, symmetries typically manifest as some kind of *conserved quantity*; however, more generally, we see that conserved quantities are just particular cases of invariant conditions. Not every invariance condition is a conserved quantity, but every conserved quantity is an invariance condition. In the case of ergodicity, the invariance condition is that the statistical properties of the system are convenient and unchanging; an ergodic system is one that is effectively static and thus we can make predictions about future outcomes in the form of stationary expectations. The problem of working in such a system reduces to the problem of simply discovering what that expectation is, and we can do it by merely collecting up data to improve our expectation estimate until further improvements don't matter anymore.

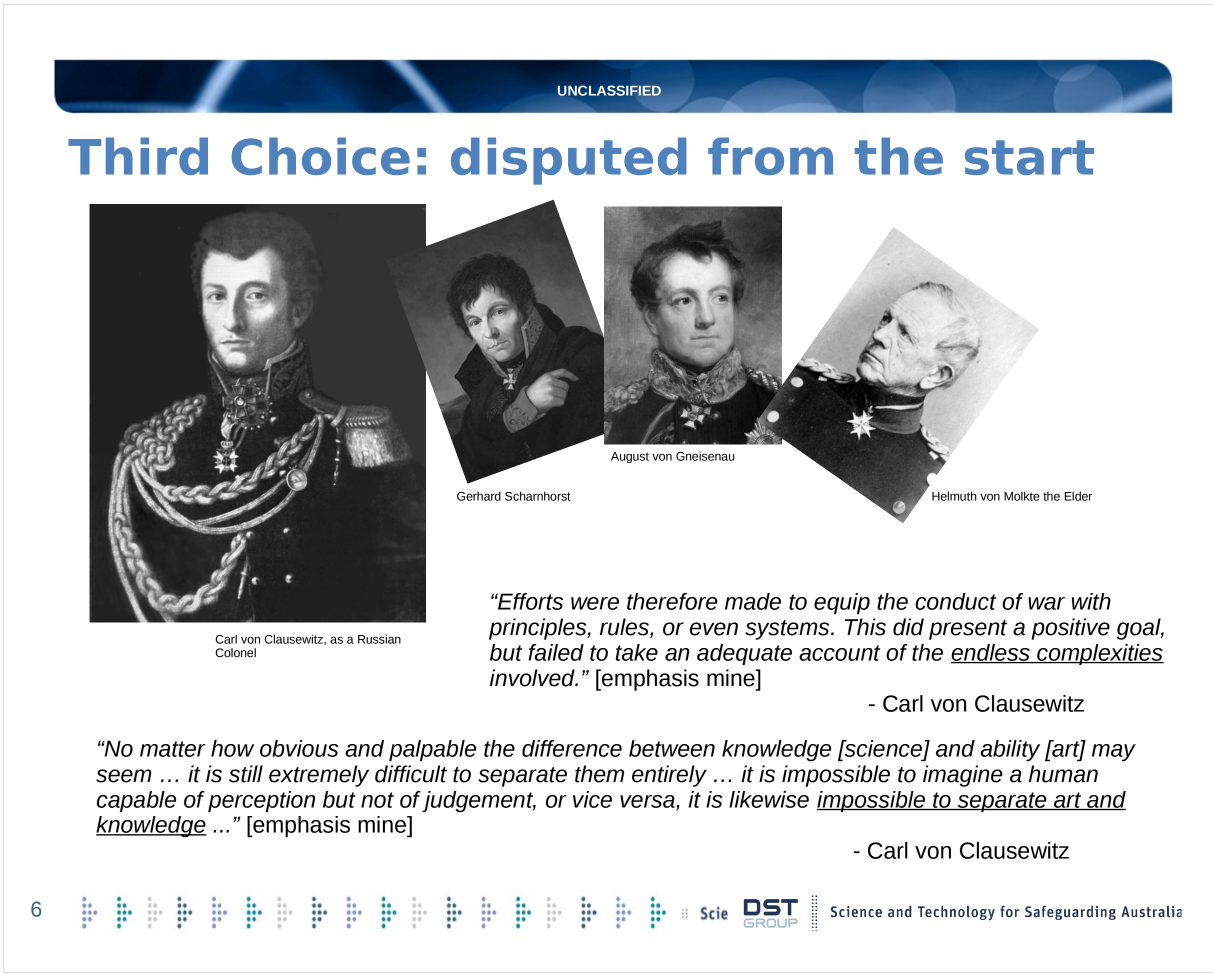

Jomini's contemporary, who he clearly saw as his major rival, the Prussian Carl von Clausewitz, vigorously disputed the Third Choice from the outset. Clausewitz was regarded as a superb staff officer yet unsuited to command because of his openness to progressive ideas and humanitarian leanings meant he was considered less than politically reliable and because of his pronounced inclination to subject ideas to intense intellectual scrutiny. (I used a picture here of him in Russian uniform deliberately – he objected on principle to the forced alliance of Prussia with Napoleon and fought with Russia against him). Aside from resenting Jomini's self-promotion, he openly rejected the notion that there ever be a universal theory of war and battle along the lines Jomini was pursuing. Like Jomini, Clausewitz turned to earlier Enlightenment period thinkers for insight into the nature of rationality, but appears to have listened to a very different set of voices. Prominent in *On War* is the influence of Immanuel Kant's earlier rejection of the views that knowledge is arrived at by accumulating observational facts and that knowledge can be established as justified true belief, by any means, including appeal to principles of pure reason.

A word is in order on the perceived difficulty of Clausewitz. His notions are necessarily abstract and focus particularly strongly on what we might now describe as 'strategic analysis' – and hence the particular relevance to the Force Design problem – that attempts to theorise about the interplays between, and consequences arising from, the mess of military and political objectives on all sides of conflict. Furthermore, *On War* was compiled posthumously from a set of manuscripts, the dates of which are all uncertain, that were written at various times between about 1816 and 1830; it is generally considered that Book I was revised towards the end of his life and thus reflects his most mature thinking. His early death from Cholera prevented him from finishing, and his greatest work reflects his lifelong process of hypothesis, observation and revision. This frustrates readers who go to him seeking locked-down answers.

Perhaps ironically, Jomini's main objection to Clausewitz is an excellent start in really understanding Clausewitz: *On War* repudiates generalisation by accumulation of particular cases, empirical verification rooted in abstract truths, simplification by reduction to parts, and prediction of future states from information about the past, and thus, in the course of offering a theory of war and battle, denies the very postulates that would appear necessary to the attainment of such a theory. At least, these postulates appeared necessary to Jomini, and later to thinkers such as Fuller, Hart, Mahan and numerous others as well. Jomini's observation, though astute, is a manifestation of his underpinning assumptions about science rather than being native to *On War*.

The resolution of Jomini's observation: the dearth of locked conclusions in Clausewitz dovetails with his emphasis on uncertainty to reflect that his approach primarily *methodological* in orientation: *On War* concerns the development of workable solutions for different unique problem contexts involving tortuous complexity, rather than attempting to provide a recount of a single locked theory for all occasions. With his much more sophisticated understanding of the nature of science and his deep appreciation for what we now regard as complexity and uncertainty, Clausewitz overtly rejects the separation between art and science. To the extent that science is about establishing verified knowledge by proof, war can be *neither* art nor science. To Clausewitz, the problem here lies in a misconception about the nature of scientific rationality.

Where Jomini and subsequent theorists basically abandoned problems of complexity and uncertainty to non-science – problems, that is, that have since Rittel and Webber (in *Policy Sciences* 1973, I think) have been described as 'Wicked Problems' in Operations Analysis circles – Clausewitz instead provides an account of methodology for their adequate resolution.

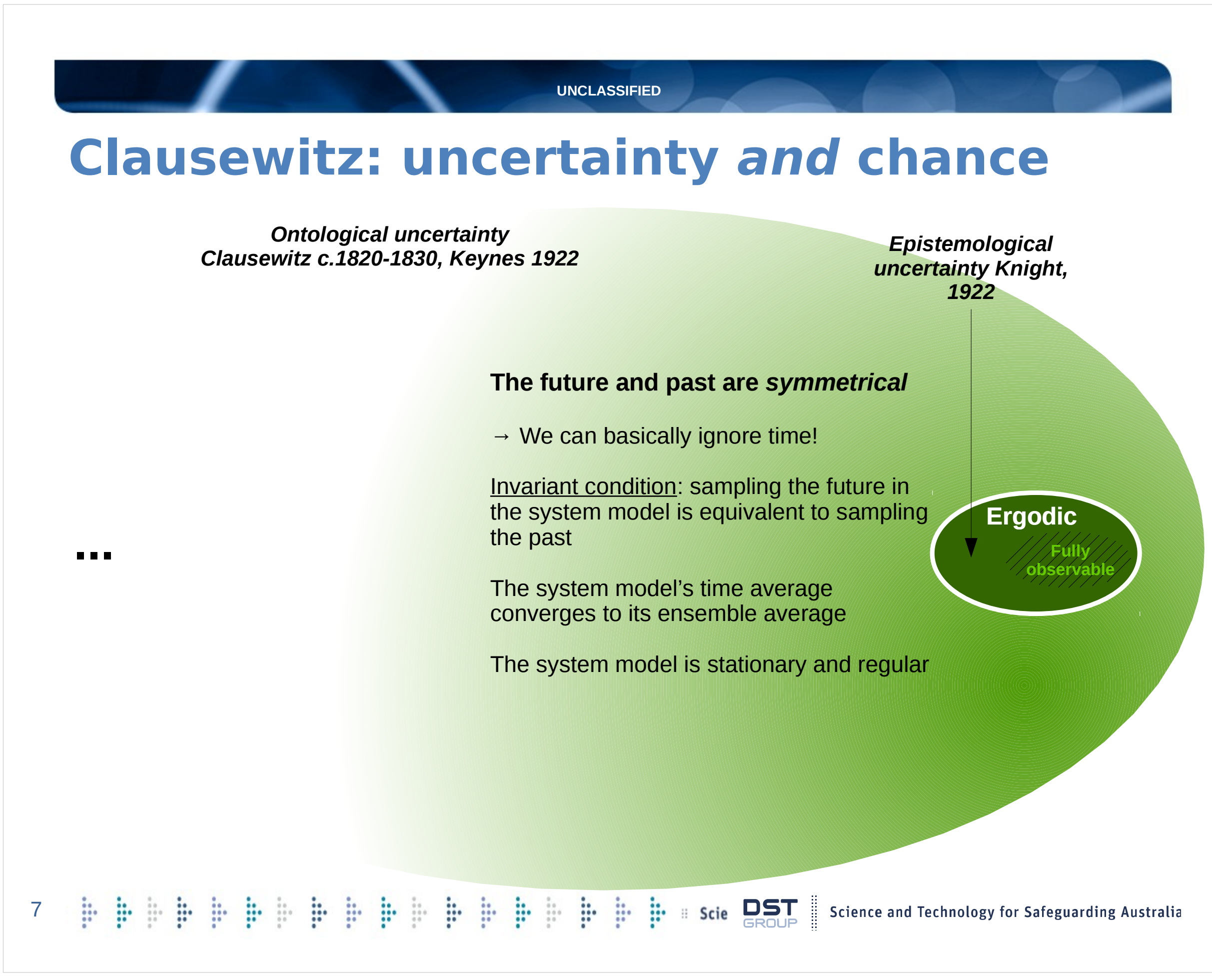

Just to add something to the growing picture of our degenerate little ergodic world within the much bigger world of all possible systems we might study and use in developing models that help us to resolve real decision-making problems: the notion of uncertainty in war and battle put forward by Clausewitz is precisely the same notion of uncertainty as that forwarded by Keynes, albeit in a different setting, and albeit about a century earlier. This notion of uncertainty was both motivated by and, in turn, animates, his vehement objections to both the positivist Third Choice and to the consequential split between the art of war and the science of war.

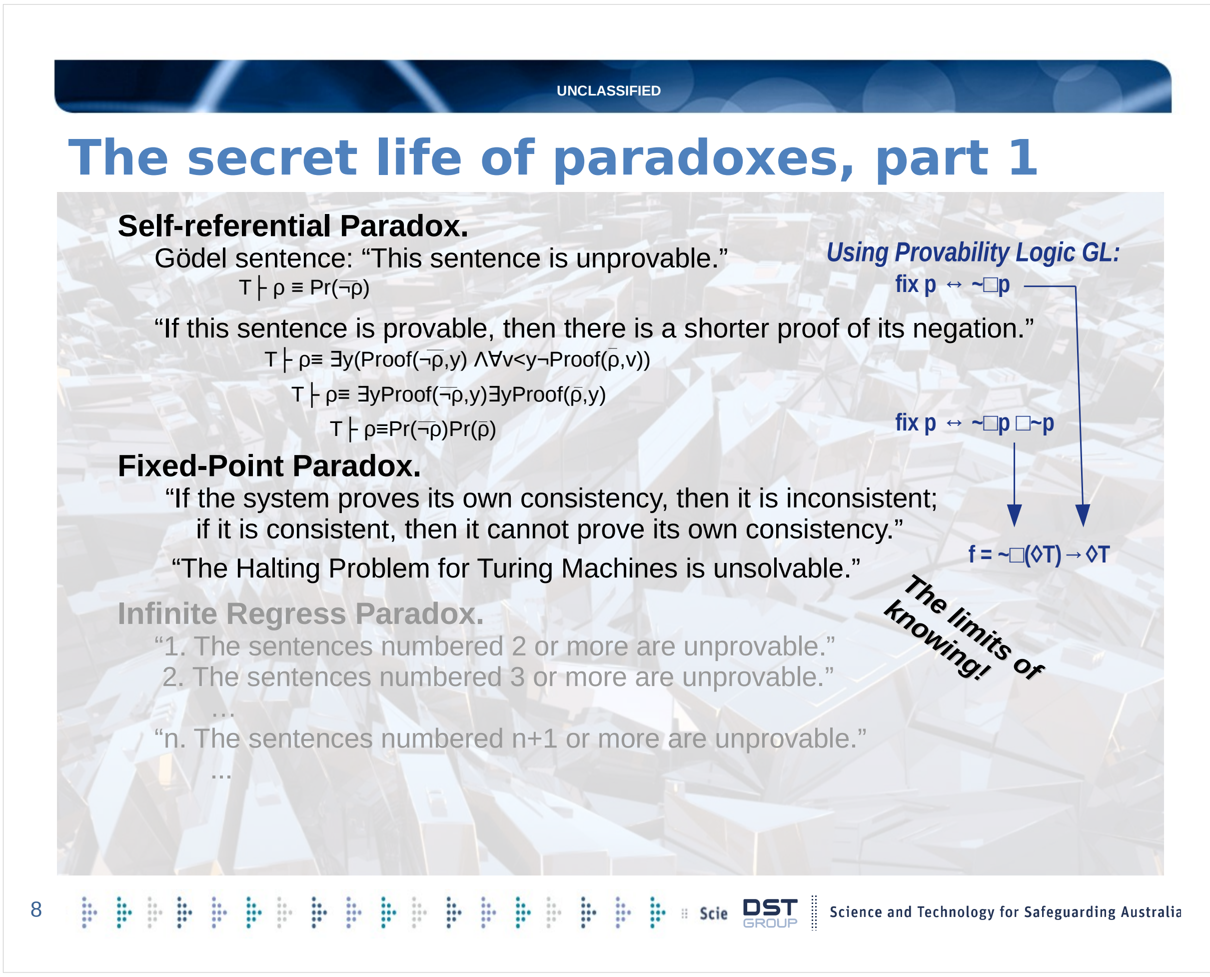

Note firstly that paradoxes are not contradictions, which are both true and false simultaneously; rather, a paradox is neither true nor false in the context of its statement. This attachment to context is both crucial and widely missed, so I will make a particular point to constantly emphasise it; understanding that paradoxes operate in a context rather than just free-float is crucial to understanding what uncertainty is all about.

The three forms in which paradoxes manifest (I'm not sure if there are others but they would be really just more complicated expressions of these forms anyway, so completeness should hold). Rather than do the Incompleteness Theorems in their original Peano Arithmetic setting, we can do the Second in other logics instead, which greatly simplifies it; we can't do the First directly though because we have no way to talk about numbers in the proof logic systems (the First says that you can state the self-referential Gödel sentence in number theory, though it was later supplanted with the Rosser sentence, also shown). What I'm going to do is work from the First and obtain the Second from it (the Second is really the deeper anyway). I've illustrated Gödel and Rosser sentences in a general logical form, and instantiated both the Gödel sentence and the last formulation of the Rosser sentence in Proof Logic GL, which is a simple modal logic consisting of the modal logic K and Löb's axiom: propositional logic, the K axiom, □(p→q)→(□p→□q), and Löb's axiom □(□p→ p)→□p. Inference rules are modus ponens, |- p→q and |- p then |- q, and distribution, |- q then |- □q. The T axiom □p→p does not hold! We could go to the more powerful proof logic R or other Rosser logics but I built my tableau prover for playing with incompleteness in Haskell for GL, which handles provability in Peano Arithmetic (which is the setting of the Gödel-Rosser Incompleteness Theorems).

In the general form, the provability predicate Pr(x) of a non-trivial theory T defines the set, say Ξ(T), of all provable sentences T |- φ⇔N |= Pr(φ). The provability predicate Pr(x) has the form ∃yProof(x,y).

What this all means is that we can state Gödel and Rosser statements, which are actually rather simple paradoxical statements in the scheme of things (you should see the mind-bending stuff the GL-prover can generate given more complex self-references), in **any** nontrivial theory! A corollary to this is that number theory, or Peano Arithmetic more properly, is actually a pretty weak theory in the scheme of consistent mathematical theories.

The first type of paradox is the explicitly self-referential paradox, which would be probably what springs to mind when I say "paradox". What if the sentence is false? Then it must be true. If it is true? Then it must be false. In a rare moment of creative inspiration, I named the fixed-point operator "fix", so the GL-prover statement fix p ↔ ~□p means "find a fixed-point solution to the (clearly self-referential) Gödel statement p ↔ ~□p".

But paradoxes need not manifest themselves in nature so obviously! The second form is what I will call the fixed-point paradox, which is a paradox that is not obviously self-referential at all, but takes the form of a statement about a limit to knowing; it is always generated by an underlying paradox of self-reference in the way I've illustrated. The fixed-point solution from either the Gödel sentence or the Rosser sentence is the famous Second Incompleteness Theorem! (We can't do the first in GL because we have no way to talk about numbers, but if I forget to explain the First Incompleteness Theorem before I get to here then I should remember at this point to describe how it works). A couple of other important fixed-point paradoxes are listed (though "important" is perhaps the understatement of the decade).

We can pump in paradoxical statements to a suitable theorem prover and get fixed-point paradoxes all day, generating arbitrary generalisations of the Incompleteness Theorems.

The fixed-point paradoxes are statements about the limits of knowing.

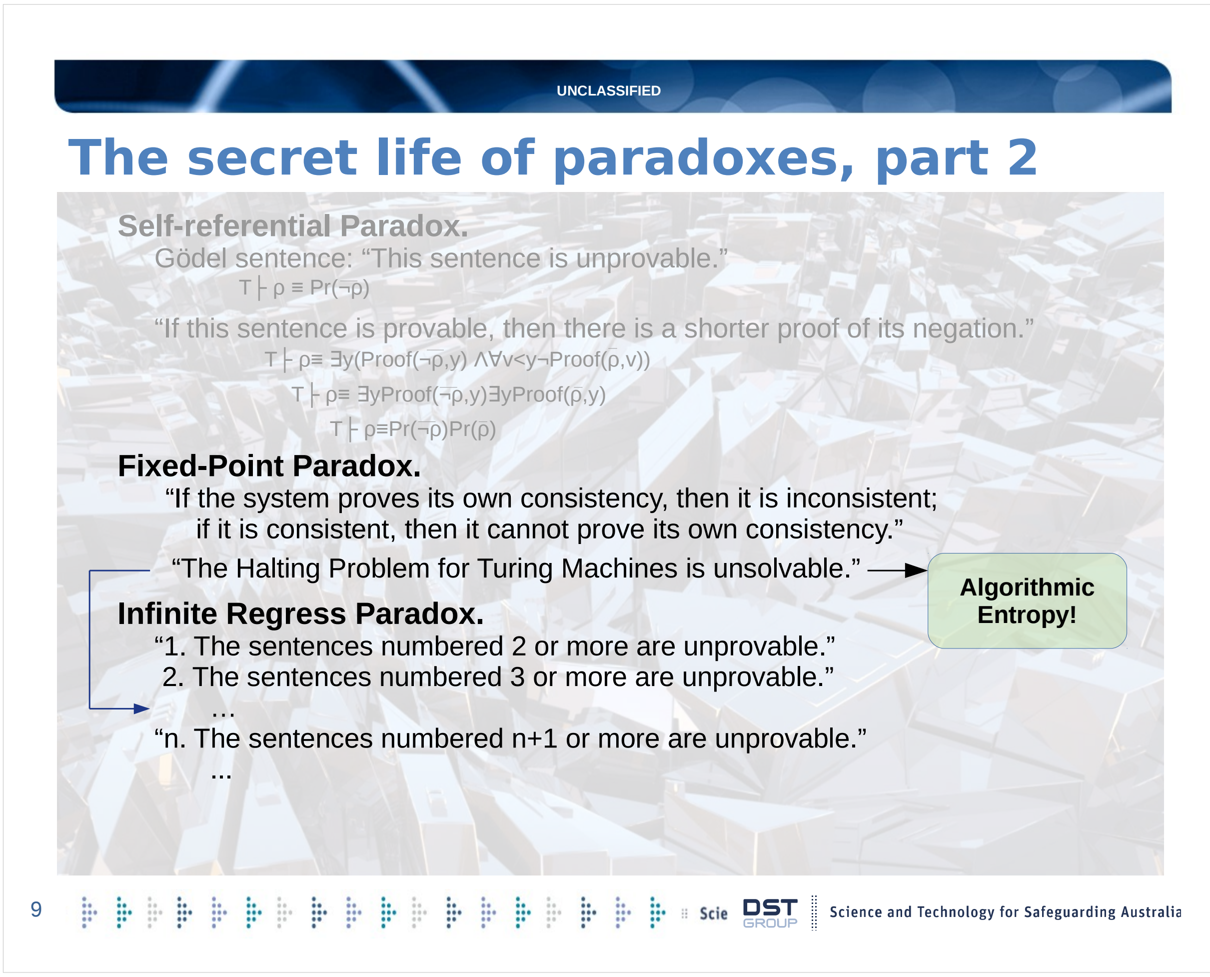

I will remember here to summarise from the previoius slide by stressing the importance of context: the mathematics reveals an important limitation in the general discussion of paradoxes in other literatures (philosophy and economics, for instance) where they are often presented as kind of free-floating things, but they are not free-floating at all. They exist in context, so what is a paradox in one context need not be so in another. For instance, the consistency of PA can be proved in Zermello-Fraenkel set theory. Yet this is no escape from the phenomena of uncertainty, for the bigger context contains more opportunity for self-reference, in some sense it has more incompleteness. The consistency of Zermello-Fraenkel set theory is not provable in Zermello-Fraenkel set theory. What it does mean is that the structure of paradoxical phenomena is intricate and rich, and spans different frameworks. Never boring.

If I have time and the inclination, I could talk here about causality and determinism. Contrary to common assumption, determinism does not give predictability, nor does non-derminism yield unpredictability. Furthermore, causality is, like the underlying paradoxes that drive uncertainty, context-dependent, not free-floating. Universal causality is not even a thesis about causality at all, but rather a thoroughly unreasonable thesis about predictability!

Now focussing on the third type of paradox: this is the irreducible infinite regress. This arises often in proofs – this is exactly the overall form of the Cantor diagonalisation proof of the unsolvability of the halting problem for Turing Machines, the Beta reduction problem in Lambda Calculus, and the uncountability of sets. It basically amounts to unpacking the fixed-point paradox by going around and around the hidden selff-referential loop, to reveal its self-referential nature.

The other point on this slide is a hook for the following slide: the unsolvability of the halting problem shows up in a very deep and fundamental set of theorems that essentially extend the incompleteness theorems with the concept of **algorithmic complexity** or **algorithmic entropy**.

Notice here that although we started with paradoxical statement without introducing time – we have demonstrated that the common somewhat implicit assumption that time is necessary to have paradoxical phenomena is not strictly true – we are already seeing here how a logical time emerges pretty quickly starting with paradoxes. This is something of an indication of the deeper connection between Gödel’s Incompletess Theorems and the Turing Halting Problem: the connection isn’t just the limits of knowing but also lies in the emergence of time ordering in the sequential (and logically parallel) application of computational operators. My personal view on the contemporary musing about the nature of time in physics as apparently not being prrimitive yet is nonetheless an extremely sticky emergent property is that time itself might be a manifestation of the inevitable and ubiquitous presence of underlying paradoxes (rather than it the other way around)! While time is not a necessary condition to having a paradox, in practice, it’s fair enough to say that it sure helps, even that there is a deep association underneath.

Entropy, reloaded

Entropies are invariants under class symmetries (isomorphisms)
Different ways of describing essentially the same system

Algorithmic entropy is about information content:
Invariant under effective procedure equivalence
K(ω) is the size of the shortest procedure yielding ω
K(ω|φ) is the size of the shortest procedure yielding ω from φ

Algorithmic entropy is about compression/decompression

We can use entropy to measure effectively the amount of information from *outside* the system we would need to solve our problem

* There are different entropy formulations dependent on context
** Entropy in economics: ‘exposure’

Now expanding on the algorithmic entropy view of the world and what it means. Again, we have symmetries at the heart of what we are doing. First a little bit about entropy: it is really not one single thing, but many things, depending on the framework (the context-dependence issue pops its head up again).

In our formal algorithmic framework, which we have chosen because of its power as a kind of fundamental or universal platform under the Church-Turing thesis, entropy is about the information required to do something. So we can talk about the information required to solve some problem, or the information required to solve a problem given a solution to another problem.

We can use algorithmic entropy to essentially measure the amount of information we require from outside the system in order to solve our problems of decision-making and control within it, for instance.

Note that economists effectively talk about what I think I would formalise as entropy when they talk about ‘exposure’. So rather than talk about what will or won’t happen in the future, we can talk about exposure to those possible events, not with respect to the events themselves but with respect to *how they might affect us.* This is essentially entropy, but again, not entropy of future events so much as entropy *as we experience the consequences of those future events.*

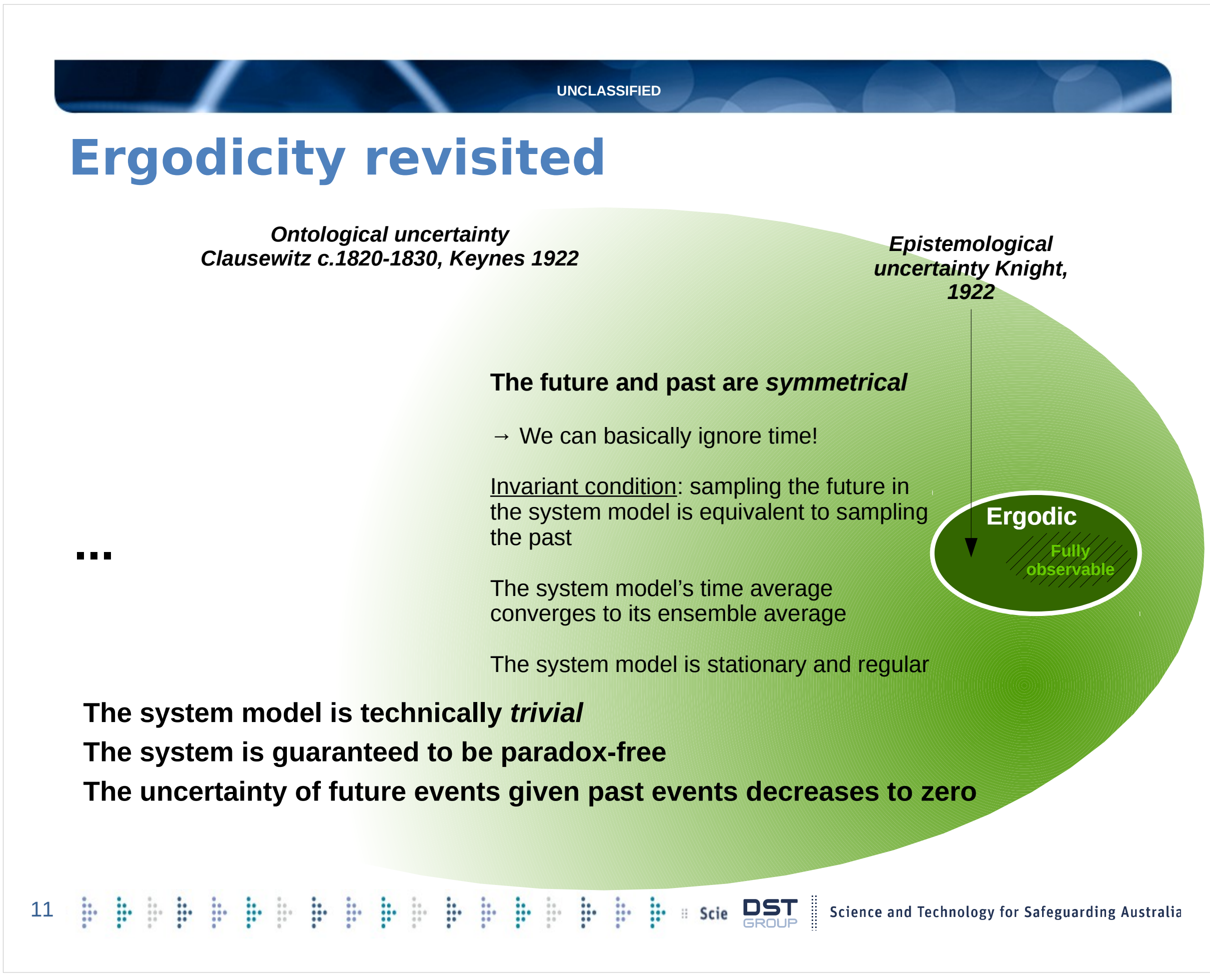

Back looking at ergodicity armed now with what we know about what fundamental uncertainty is all about. Ergodic systems are merely what we get when we assume the ridiculously strong symmetry condition that the future is (statistically) the same as the past. This amounts to asserting that the system contains no paradoxes, which we obtain by removing the possibility of self-reference (though noting that not all self-reference is paradoxical either, so we get an early sniff that ergodicity is overkill even with repsect to fairly naive goals). Restated now with what we know about entropy, we see that ergodicity means that the amount of information that we would need to fully solve our problem of determining future states of the system from outside the system (remembering the importance of context that we learned from our examinaton of incompleteness) decreases to zero as we collect data about the past.

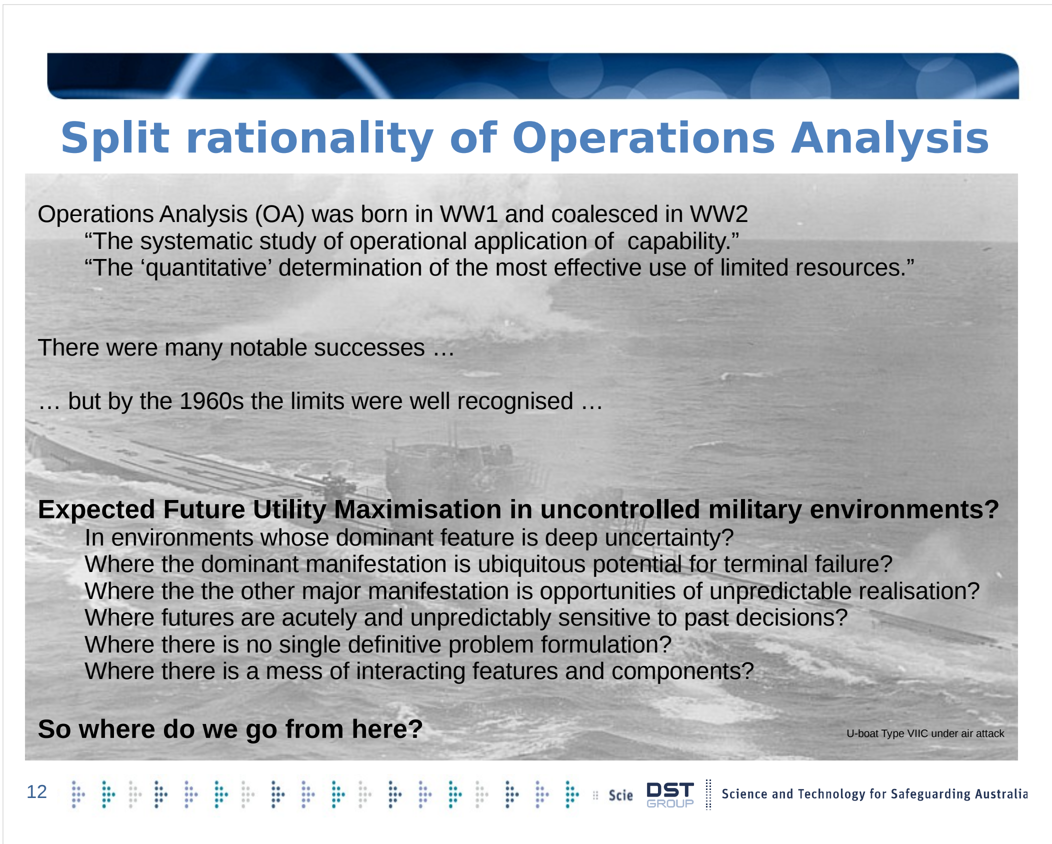

The Third Choice, embodying the ontological view that there is one objective truth that exists independently of the observer and that the goal of the researcher is to find that single truth, was institutionalised into what became the modern military professions and modern defence organisations. With it came the consequent split between art [ability] and science [rational knowledge], with art being everything that science, by this conception of science, seemed not to be able to handle. Theoretical debate continued, both inside military thinking and outside it more generally, about just what constitutes the correct method of science whereby truth might be so established – so we have multiple common schools of foundationalist thought including inductivism, reductionism, empiricism, so-called “scientific” rationalism, historicism and so-called “scientific” determinism vying for the title of “correct scientific method”. Yet these different accounts of method generally appeared smoothly blended together in practice; people often didn't look too hard for the lurking contradictions, in the interests of just getting on with solving the problem at hand. The institutionalising of these ideas was greatly hastened by the fact that they held sway for so long – and arguably still hold sway today – in military education. The stubbornness of these outlooks may lie in the fact that they are simple and intuitively appealing, particularly in terms of the premise of certitude.

Modern Operations Analysis was born in this environment, in the early 1940s, during World War 2, when science was applied systematically to problems of application of operational resources and not just, as had been the case previously, to problems of making new technologies for better equipment. Yet from its birth, the new field was constrained, highly, by the previously established and now long-institutionalised thinking about the nature of rationality. In particular, it was predicated on the view that science means verified correct models of quantitatively measurable things. The general method looks something like this (according to Rittle & Webber, who come back into the story later): “Understand the problem”, “Gather the information”, “Analyse the information”, “synthesise the information” or “wait for the creative leap” or “get the solution” or similar, which reflects the linearity of the foundationalist methodological explanations.

Yet by the 1960s the limitations of this conception were, predictably, raising their ugly heads, just as they had done in military thinking, and for the same reasons. Not everything that we might like to subject to the wonders of Operations Analysis appeared to be amenable to its methods. Not everything could be pinned down to fixed measures. Not everything could be solved with a linear method. Unwilling to admit defeat and throw it all back to the world of military art – the burgeoning sciences were showing that there was, indeed, science to be had here – so to resolve the conundrum, Operations Analysis simply mirrored the same spit between art and science. However, now it was relabelled as divide between “hard” quantitative Operations Analysis, where lies, in principle, indisputable answers to definitive problem formulations, and everything else that didn't fit into the regime of the separable and measurable was thrown over the fence into the new spooky demesne of “soft” Operations Analysis, where lurk the dark qualitative spirits of the social sciences.

Within the division of military and defence matters into analytical science and unfathomable art, the regime of geniuses, then, we have a another subsequent division of science that regurgitates the same conclusion with a new consequent dynamic tension, for the same base reason. Science now hard with a gooey, sticky coating. Science also split: quantitative, definitive “hard” and qualitative, gooey “soft” Operations Analysis branches.

This situation does not suffice.

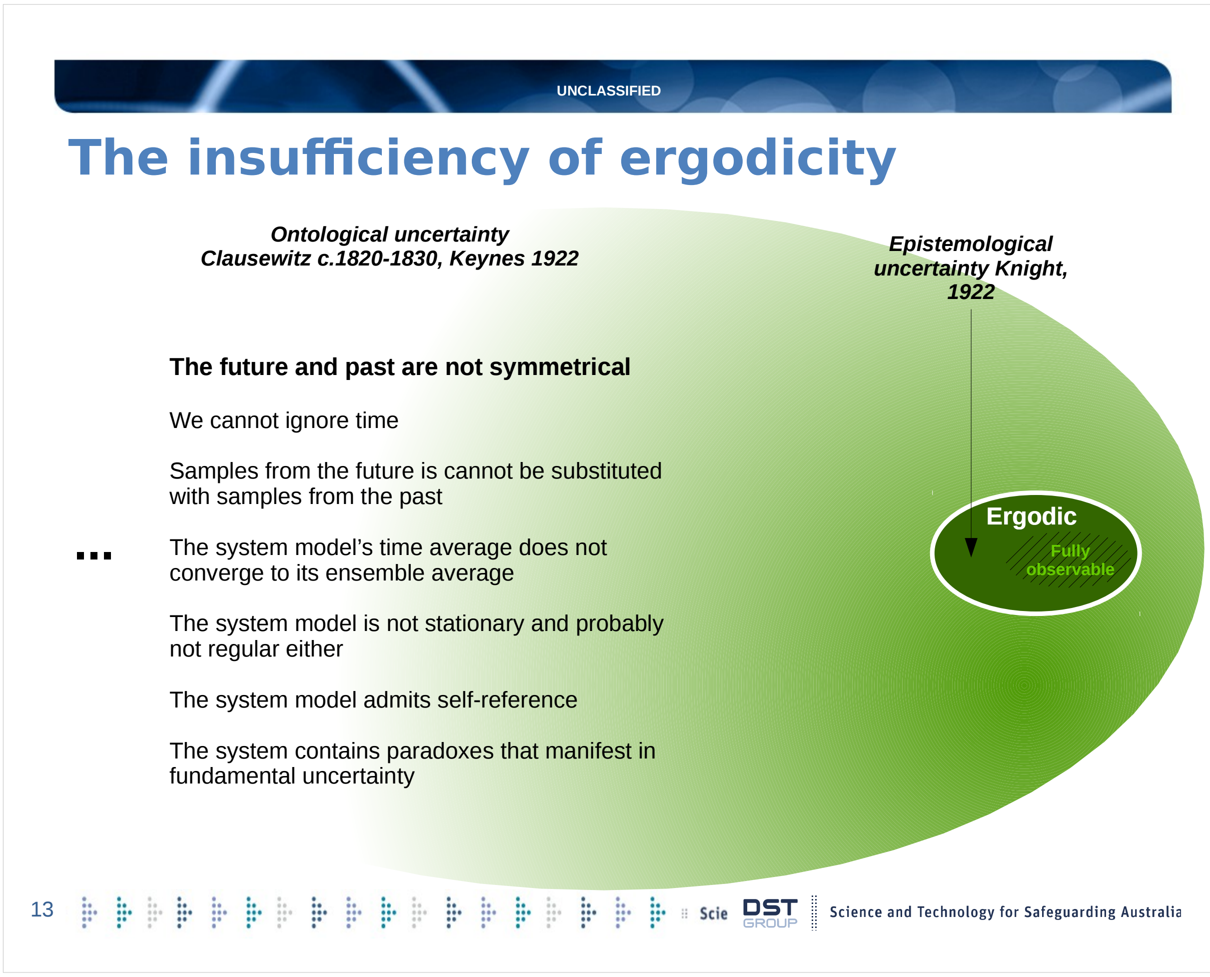

Back to our picture, with a summary of the major ideas so far, cast now as a rejection of the <u>sufficiency</u> of ergodicity. Rejection of its <u>necessity</u> will follow soon; the rejection of necessity is less widely appreciated and where our research is focussed. Rejection of necessity *is* the research direction we are implementing that addresses what we do about rejecting sufficiency.

Modelling Complex Warfighting

Motivation: Available current approaches are inadequate for the conditions of extreme uncertainty, complexity and uncertainty and terminal failure potential of our problem domains.

Purpose: to revolutionise Defence & military Operations Analysis, around complexity and uncertainty.

Topic: determining and evaluating options for the effective allocation of limited Defence resources under Force Design and Employment, where the consequences of the decisions are realised in futures that the decision-makers cannot know nor fully control at the time of the decision.

Approach: multiple transdisciplanary teams of DST with a diverse group of academic parters who are co-investing in building national capability in Operations Analysis research and application for Defence & National Security.

This slide is where I will talk about the ambitions of our Modelling Complex Warfighting (MCW) Strategic Research Investment (SRI). I've constructed the talk in this way rather than framing everything around MCW at the start (or for that matter any other SRI) to highlight that the programme doesn't come first with the research direction following, but rather the research programme and how it is set up and operates follows as a *response* to the base realisation that complexity and uncertainty matter enormously and that we aren't currently well equipped with approaches for dealing with these features. So rather than "here is our programme and this is what it contains", I'm talking about it as the logical consequence of the science and particularly of the mathematics of uncertainty.

The thing I will highlight here the most is the uncompromisingly ambitious nature of both the research and the organisational goals, rather than to go through the structure or the Strageic Responses' separate topics; the key feature I will highlight is the highly trans-disciplanary nature of the programme, as an instantiation of the dissolution of the split between hard and soft, and the underlying split between science and art.

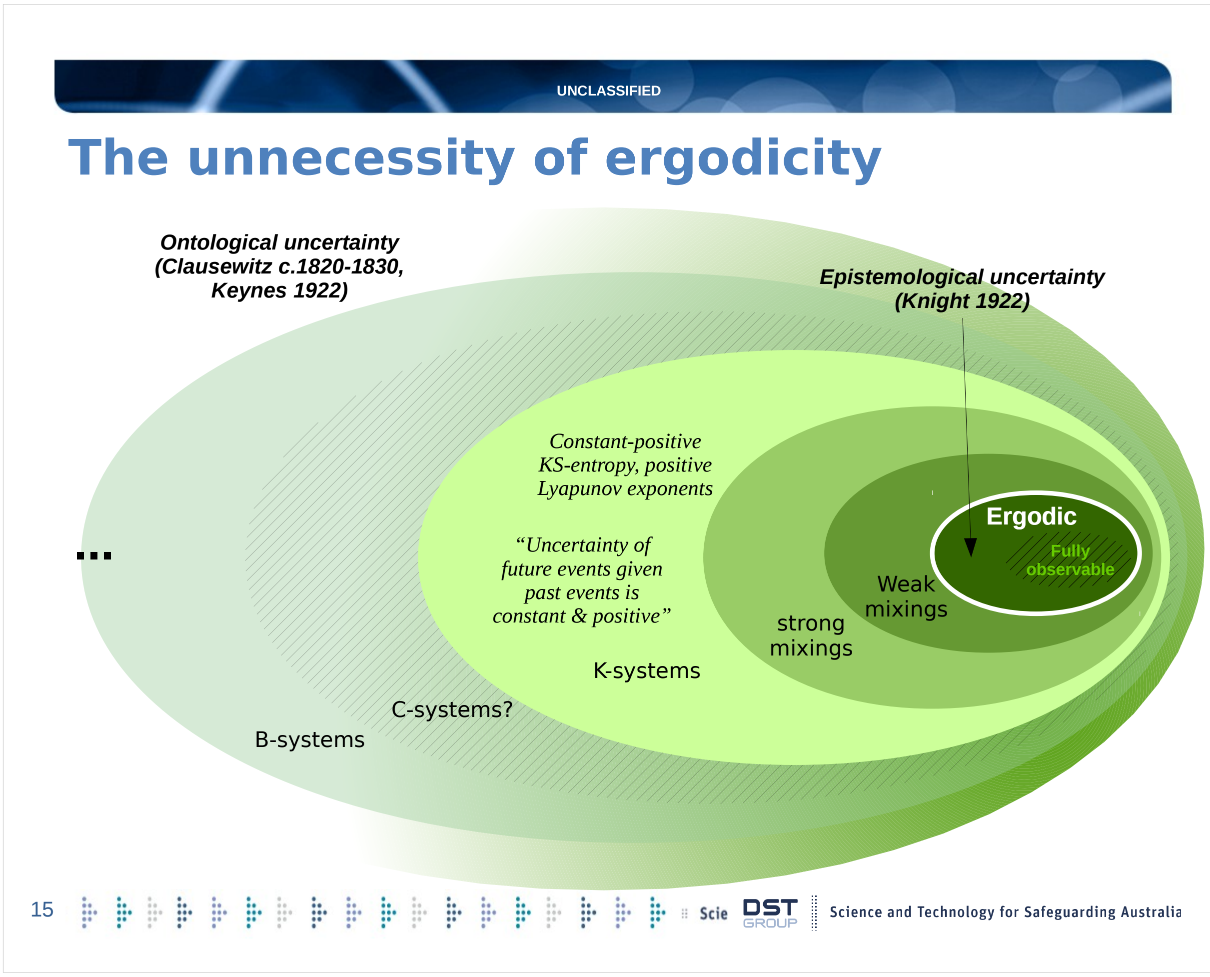

Remembering the central place that symmetries have in our conception of the problem: there are plenty of other symmetry conditions that are strictly weaker than ergodicity from which we can choose. Now taking a dynamical systems perspective, we get (here using Kolmogorov-Sinai entropy) a hierarchy of invariants to play with. We get the computational or axiomatic view of a dynamical system by breaking up the state space into cells of arbitray size; think of a kind of scalable grid, and then marking an X in each cell where the system crosses in phase space, in order. Provided our gridding system obeys certain mild conditions to ensure it is sensible, we turn the dynamical system into a computational system where we have our information entropy properties and computational properties – we can ask whether there is a compression and how good it is for reproducing the system's behaviour, so we can talk about its complexity in a fundamental way.

The central point here is that abandoning ergodicity isn't actually that bad; to the contrary, it's both interesting theoretically and of enormous practical benefit. So let's do it ...

Beyond ergodicity

"Time is a mechanism for stopping everything from happening at once"

- Introducing time and allowing non-stationarity
- Identifying paradoxes directly (often self-reference involving time)
- **Identifying alternate symmetries and their corresponding invariants**

  - the formulation of entropy follows from isomorphism class
- Formulating around entropy of events rather than the events themselves
- Formulating around entropy of costs rather than costs themselves
- Presence of terminal failure states is a primary feature
- Opportunities are just affordable failures with viable potential pay-off
- Feasible scenario spaces (path dependencies) rather than tuned scenarios

I will touch on as many of the points as I can in the available time to illustrate how the scientific adventure in the mathematics and economics of deep uncertainty manifests in our research programme.

The background is an example economic model: it is essentially a graphical representation of a path integral – a sum over histories – for a single asset with set starting conditions. It illustrates the generation of uncertainty to give non-stationary mathematical models to play with, both mathematically and experimentally.

Mathematical potentialities

- Measures of relative utility with respect to different command functions of force design options
- Determine crucial sensitivities to failure in force options (rather than maximising efficiency or expected utility)
- Feasible scenario spaces and path dependencies
- Generating novel events in simulations
- Entropic data sets to determine dynamic phase shifts

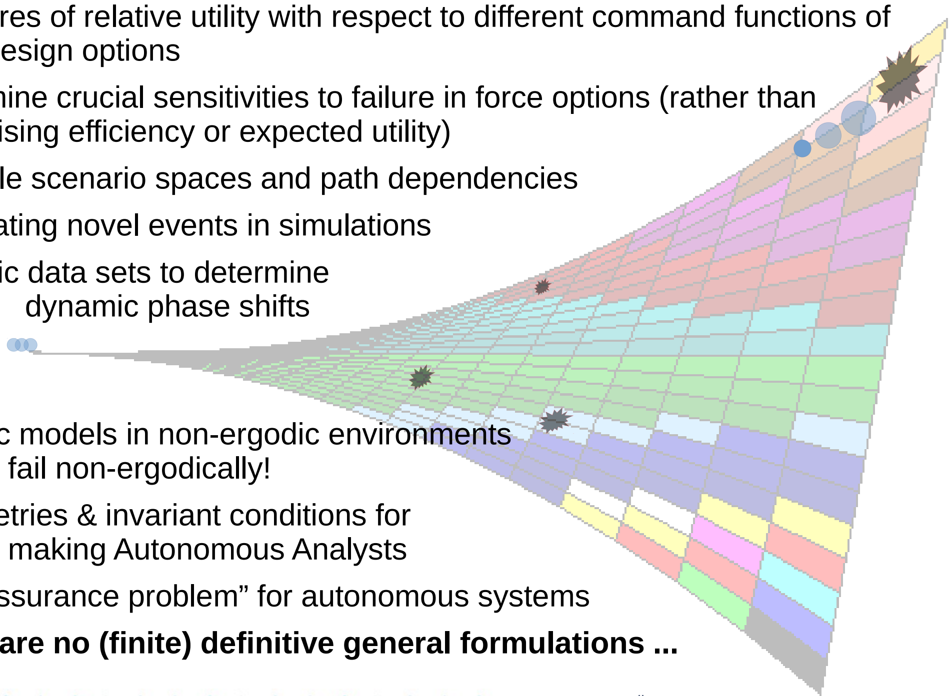

- Ergodic models in non-ergodic environments fail non-ergodically!
- Symmetries & invariant conditions for making Autonomous Analysts
- The “assurance problem” for autonomous systems
- **There are no (finite) definitive general formulations ...**

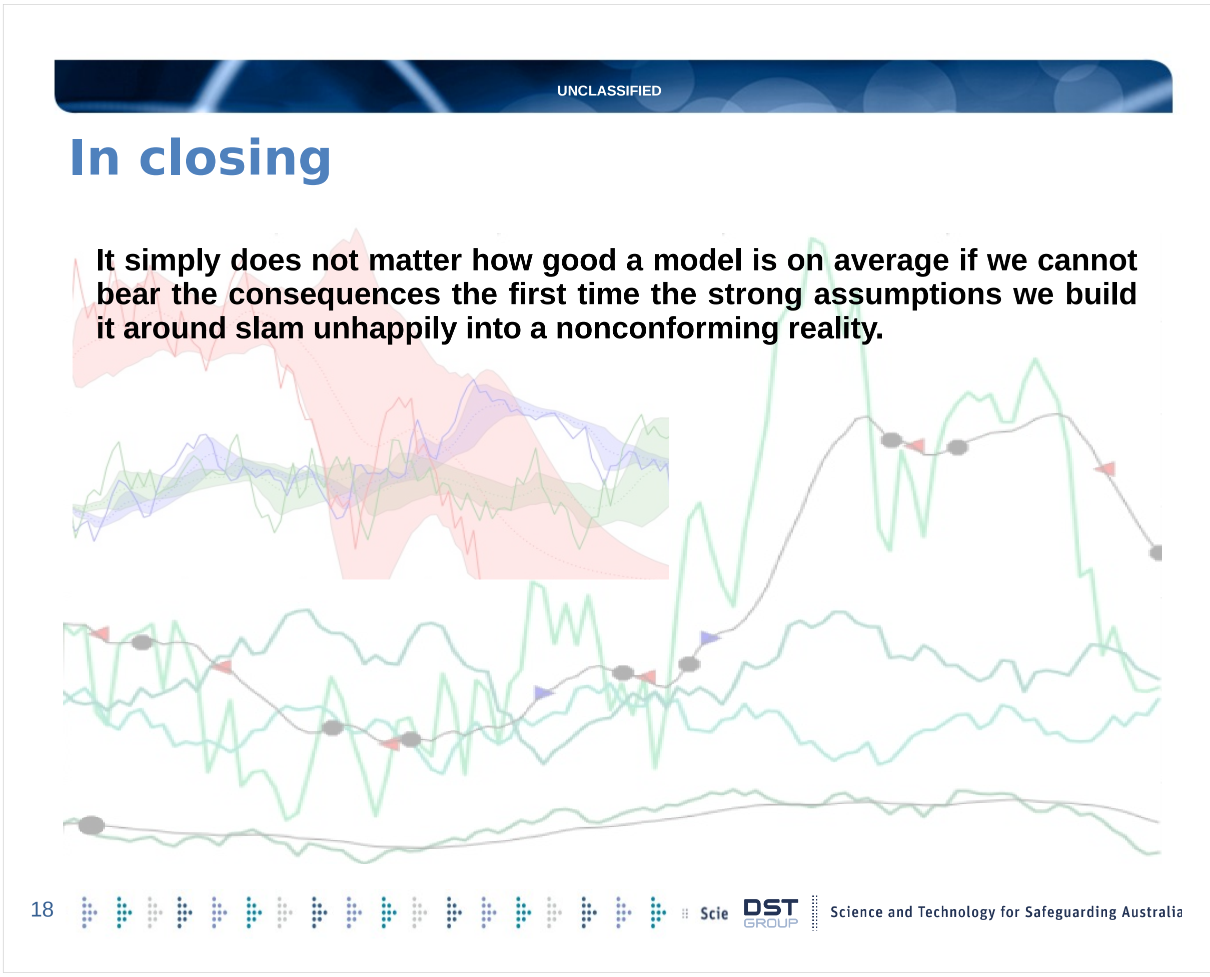

When it's all said and done, the conduct of science is not about a kind of convergence, by accumulating facts, towards some pre-ordained absolute truths, but rather it is an adventure into uncertainty. People often react to uncertainty as if it's bad, yet uncertainty throws up opportunity just as much as it delivers up threat. Rather, we should be afraid or at least suspicious of certainty: certainty is tyranny.

The background pictures show some of my exercises in showing how this is not just possible, it isn't even all that difficult: the top left is an online entropy estimation in a simulated multi-commodity market (just three assets shown); the bottom is the actions of a trading agent that consists of PQ micro-agents on each asset with entropy estimation driving the allocation of resources across the micro-agents, operating successfully in a highly non-stationary simulated market environment.

It follows that the worst crime in science is not to be wrong – indeed, in the ultimate scheme of things in empirical and applied science we cannot ever really be right – but rather the worst crime is to be boring. Complexity and uncertainty are not new things in the world; scientific progress is not the emergent property of data accumulation. Rather, knowledge is the consequence of theory development that then drives data collection for careful critical test; the limited ability we have to address complexity and uncertainty is not a lack of data but a lack of adequate theory by which to look for evidence and comprehend its meaning.

In this spirit, we are not really interested in making small incremental advances with better solutions to the problems we basically already understand. We are interested in scientific revolution, which means advancing the state of knowledge with better problem choices; big, messy, mind-bendingly complex, hairy, and relevant problems that nobody has ever tacked before. Relevance is not defined by immediate applicability within problem choices we already know about; relevance is about identifying the hard problem topics in our application domains that we have not been able to handle previously.

For us, across topics including Operations Analysis and Artificial Intelligence, the revolution is centred on expanding beyond conventional problem choices framed by narrow conventional symmetry assumptions, and dealing instead with problems involving complexity and uncertainty – problems where the usual symmetries break, problems where there is the possibility of self-reference, problems where paradoxes roam freely, problems where decisions have to be made such that the consequences will play out in futures the decision-maker cannot control nor predict at the time of the decision, problems where there is ever-present possibilities terminal faliure and opportunity. Problems where life is lived. In a delicious irony, having a very applied focus demands that we invest heavily in being deeply theoretical.

I hope that, in the course of talking about our outrageously ambitious research goals, I have managed to not be boring.