Managing Adversity Risks for Non-anthropogenic Systems

Ian Bryant
University of Warwick
Coventry CV4 7AL
UK

i.bryant@warwick.ac.uk

Tim Watson
University of Warwick
Coventry CV4 7AL
UK

tw@warwick.ac.uk

Carsten Maple
University of Warwick
Coventry CV4 7AL
UK

cm@warwick.ac.uk

ABSTRACT

It is important when considering a Non-anthropogenic System (NAS) to ensure that not only are all sources of Adversity (to both Safety and Security) to a mission captured and appropriately managed, but also that the management approach explicitly considers the mission view of Risks, the audience expectation of Scope and perceptions of likely Risks from what a technology in the early phases of adoption,

KEYWORDS

Adversity, Hazard, Risk, Security, Threat

1.0 INTRODUCTION

The Assurance of an Autonomous Unmanned Systems Mission is typically handled by two differing communities of interest, with one community being interested in Safety factors, and another in addressing the handling of Threats. This is both potentially wasteful of effort, and introduces the probability that both items will fail to be identified due to being perceived as lying outside the community's scope, and that emergent risks can arise from differential treatment approaches.

An alternative approach is to manage all sources of Adversity Risk under a single modelling technique, accepting that not only should this consider the Mission view of Risks, but also the audience expectations and perceptions of Risks, which may differ.

This paper is the latest in a series from the RASAE (Replicable And Scalable Adversity Enumeration) project, and provides both a summary of the preceding work, and the insights arising from this phase of the project, which were predicated upon the technologies and audiences for Non-anthropogenic Systems (NAS).

2.0 RISK FACTORS

In order for any entity to properly manage its risks – i.e. to use appropriate countermeasures to reduce exposure to an acceptable level – these risks first need to be described and/or enumerated, with the five main factors to be captured being:

- The set of Assets to be protected
- The set of Adversities that are faced
- The set of Compromise Paths that are exposed
- The Risk Appetite of the entity

- The set of Controls that are applied

3.0 ADVERSITY

3.1 Adversity Concept

The concept of adversity was proposed [1] to address the perceived, yet artificial "stovepipe" distinction between the views of the security and safety communities, in which the security community seeks to address threats (directed, deliberate, hostile acts) and the safety community seek to address hazards (undirected events).

This means that the security world assumes a deterministic threat model which typically ignore hazards, and is largely predicated upon characterisation of known types (if not necessarily details) of threat actor, which therefore has difficulties handling the full known-unknown-unknowable (KUU) model [2]. On the other hand, the safety community typically uses stochastic models to address hazards, and usually ignores threats.

Although there is a logical distinction between the two (a threat is normally viewed as being the result of a human actor who has Intent (Motivation) and Capability), an abstraction can be usefully taken to normalise them as the superset – called adversities – with associated probabilities of occurrence.

3.2 Modelling Adversity

In the RASAE approach [2], the Adversity Model (AM) provides a domain-neutral representation, which allows all sources of adversity to be captured, collated and simplified to produce a set of combined Adversity Classes (AC) that form the overall Adversity Set (AS).

This model means that early stages of analysis will cause a proliferation of factors to be considered, but the later stages then introduce simplification. This is illustrated in Figure 1, abstracted from a full AM. This abstract has been consciously selected to represent hazards and threats that, although largely and seemingly "non-technical", can nonetheless have significant impact on cyber or cyber–physical systems.

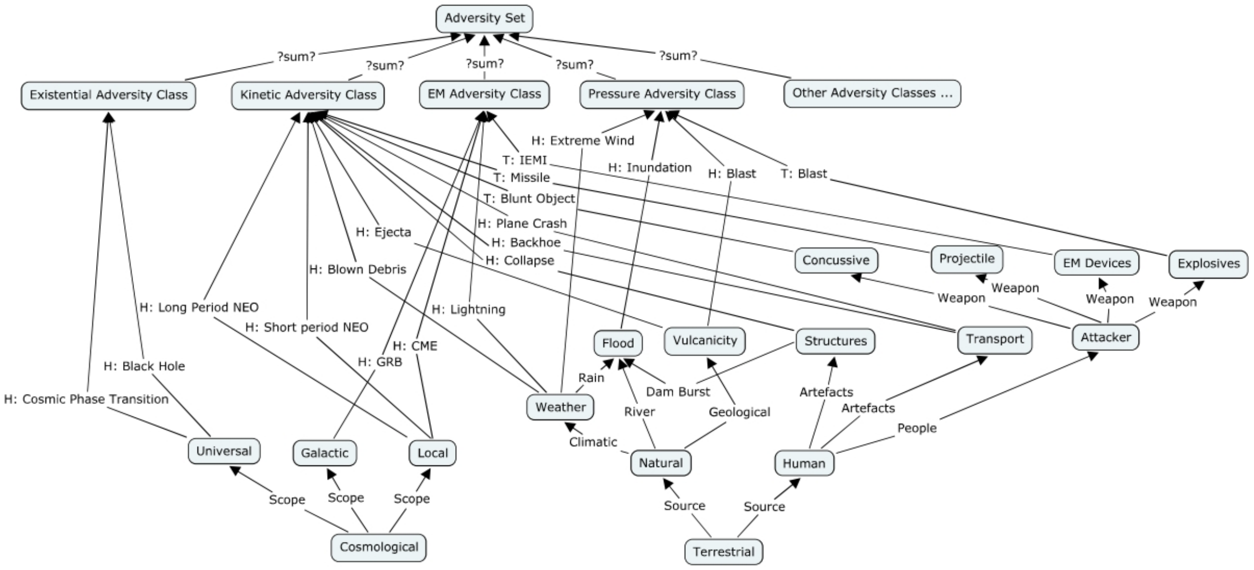

Figure 1: Adversity Model

The output of the modelling approach is a set of ACs, which are used as the baseline for high level analysis.

The generic set of ACs from the RASAE efforts is presented in Table 1.

Table 1: Generic Cyber Adversity Classes

Adversity Classes (AC)		Impact Class (IC)
AC.EX	Existential	IC.Cyber.Collateral
AC.KI	Kinetic Impact	IC.Cyber.Collateral
AC.PR	Pressure	IC.Cyber.Collateral
AC.IN	Inundation	IC.Cyber.Collateral
AC.TH	Thermal	IC.Cyber.Collateral
AC.CH	Chemical	IC.Cyber.Collateral
AC.BI	Biological	IC.Cyber.Collateral
AC.RD	Radiological	IC.Cyber.Collateral
AC.VA	Vanishment	IC.Cyber.Indirect
AC.TR	Trespass	IC.Cyber.Collateral
AC.EM	Electromagnetic	IC.Cyber.Indirect
AC.LD	Logical Disruption	IC.Cyber.Direct
AC.DD	Data Disruption (includes Destruction)	IC.Cyber.Direct
AC.DL	Data Leakage	IC.Cyber.Direct
AC.IM	Impede	IC.Cyber.Indirect
AC.FA	Failure	IC.Cyber.Indirect

Included within the AC definitions in Table 1 is the nature of impact on the subject cyber / cyber-physical systems was considered, with three impact classes being identified:

- Direct – where the adversity acts directly upon the logic-bearing function or data/information within the subject cyber / cyber-physical system
- Indirect - where the adversity acts directly upon the cyber / cyber-physical system, but not directly on the logic-bearing function or data/information within the subject cyber / cyber-physical system
- Collateral – where the adversity has an impact that impinges otherwise upon the cyber / cyber-physical system or its logic-bearing function or data/information

4.0 RISK EXPOSURE

4.1 Deleterious Result Scale (DRS)

RASAE uses Maslow's Hierarchy of Needs [4] to form a set of Deleterious Outcomes (DO) as follows:

- Regulatory Noncompliance
- Legal Offence
- Disrupt Relationships
- Disrupted Operations
- Reputational Damage
- Personal Distress

- Economic Damage
- Physical Damage
- Personal Injury
- Loss of Life

In order to allow comparison, a common scale is required, for which the most commonly encountered in risk terms is a monetary value. But many of the factors in the DO list do not have a direct monetary value, so "Value Of Statistical Life" (VSL: also known as Value of Preventing a Fatality) and Value of Preventing Injury (VPI) are used as a way to align between monetary and non-monetary scales, with the scaling based on a recent meta-analysis by OECD [5], which established a reasonable measure of central tendency for VSL.

The DRS alignment between monetary and non-monetary scales has therefore been pegged at the next highest rounded appropriate value in international currency units, of 1 Statistical Life = 2,000,000 XDR[1].

Such a value neatly illustrates the potentially large values that a DRS could assume, so to make the numbers more intelligible, a logarithmic, absolute scale is used, with 1 Statistical Life = DR7.0[2]. The logarithm-based approach of the DRS has the collateral benefit of facilitating the instinctive filtration of less relevant risks due to the intrinsic order of magnitude steps in the characteristic of the DR[3].

Audience testing of the DR approach with both producers and consumers of Cyber-Physical Systems (CPS) suggests that DR offers a useful way to aid mutual understanding in a Business to Business (B2B) context, analogous to the use of Micromorts [6] and Microlives [7] have been used for personal risks.

4.2 Risk Enumeration

DR is an expression of likely impact, and as such matters are fundamentally uncertainties, it is important to remember that DR, although based on an absolute scale, is a replicable value but not a precise value.

But it is not only magnitude of impact that is the subject of uncertainty – the likelihood of occurrence is also a variable value.

If estimates of both likelihood/frequency of occurrence and magnitude of impact are really probability distributions, then a way of conveying this information is required. For each of the two axes a 5-number-prediction (5NP), and its visual representation the Box Plot provides a means to summarise the expectation (rather than observation) of likelihood/frequency/magnitude across a variety of types of probability distributions, consisting predicted extreme lower value / first quartile / median / third quartile / extreme upper value.

This leads to an adaption [8] of the 5NP and Box Blot as "10NP" – a matrix of (2 * 5NP) and Conquad Plot (derived from the latin conquadro: square), as illustrated at Figure 2.

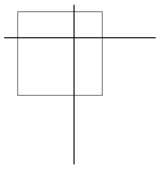

Figure 2 – Conquad Plot

[1] ISO4217 Special Drawing Rights

[2] Adjusted as (Log_{10} + 1.0) to allow 0 = 0.0

[3] And also facilitates discarding excess detail by ignoring the mantissa

4.3 Risk Expectation

Most Adversity assessments are clouded by practitioners who assume that the magnitude of adversity can be described in a generic manner, whereas in reality most adversities will vary with Entity, Locale, Archetype[4] and Time[5].

Different audiences inevitably have differing timescales in mind when discussing a risk. Taking the case of vehicles, for instance, we find that consumers risk focus is that for the time in which they expect to have use of the vehicle, whereas responsible producers should consider the whole lifespan of the vehicle.

To address the timescale focus challenge, the Annualised Expectation of Risk (AER) is preferable [8]. Like Expected Value (EV) from techniques such as Decision Trees, this is preferable to largely meaningless idea of "Risk" when unquantified, and in audience testing is understood to be an abstraction that will seldom if ever be the Actual Value. Furthermore its temporal span – Annularity – aids accommodation of LIHP and HILP concerns.

Moving to the scope of risk, most approaches such as the widely adopted Information Systems Management System (ISMS) methodology [9] are scoped to an Organisation, and as such would be preferably explicitly labelled as Entity Risk (R_E).

If a holistic view of risk expectation is the goal, the generalised statement [8] is provided in **Equation 1**:

$$AER_T = (AER_{AS.E}) + \sum_{1}^{n}(AER_{AS.P}) + \sum_{1}^{n}(AER_{AS.C})$$

where:

AER_T – the Total AER either for a single entity and all its externalities, or construct like AER_J (Joint) or AER_N (National)

$AER_{AS.E}$ – the AER from the internal (direct) Adversity Set

$AER_{AS.P}$ – the AER from the Partner(s) (indirect) Adversity Sets

$AER_{AS.C}$ – the AER from Collateral Adversity Sets

A challenge with production of the these aggregated risk sets remains that the mathematics of combining 10NPs, based at best incomplete actuarial data, and often on little more than expert opinion, is non trivial, as not only will accumulation and association have differing effects, but also Treatment Induced Risks (TIR) may arise as confounding factors [8].

5.0 ADDRESSING RISK

5.1 Feasibility and Flexibility

The classical approach to management of risk is to break the options down into categories, of which one of the best known is "T4": Tolerate, Terminate, Transfer or Treat [10].

Most of the focus of risk management is on treatment, but within the concept of treatment there is often less flexibility than is assume, as illustrated in Figure 3 [11], which maps concepts from the world of safety critical systems into a generic spectrum of risks.

[4] A function of the various types of persons engaged in the entity's operation, for instance frequent travellers having a differential risk to those who are static

[5] A function of time, for instance "Y2K" or a major sporting event in which the entity is engaged or is proximate to

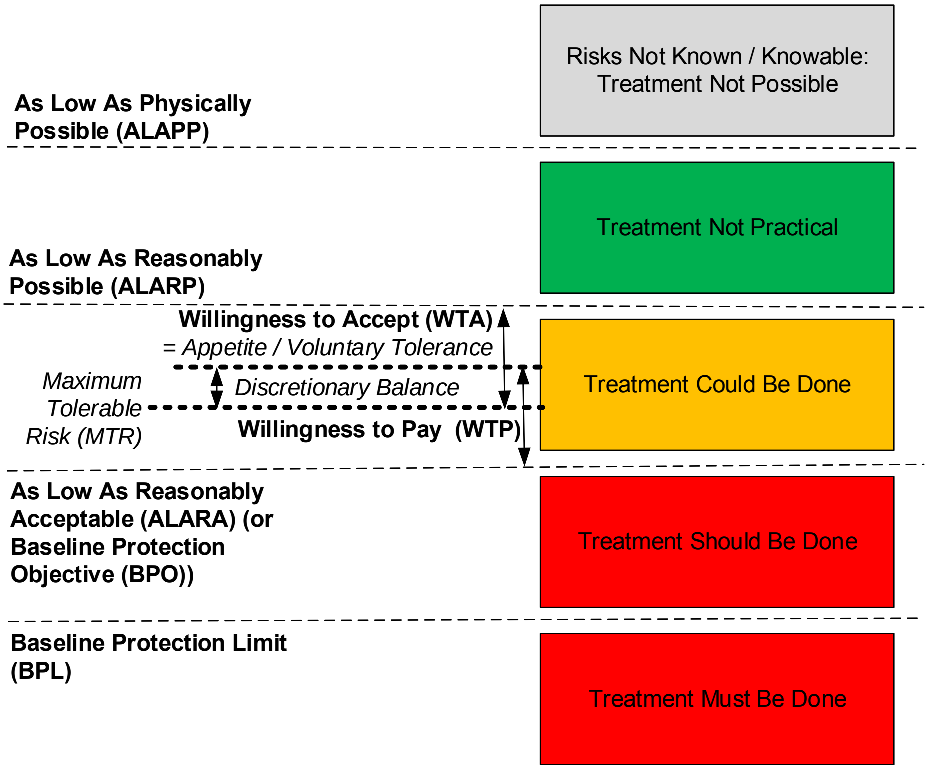
Risks Not Known / Knowable: Treatment Not Possible
As Low As Physically Possible (ALAPP)
Treatment Not Practical
As Low As Reasonably Possible (ALARP)
Willingness to Accept (WTA)
= Appetite / Voluntary Tolerance
Maximum Tolerable Risk (MTR)
Discretionary Balance
Willingness to Pay (WTP)
Treatment Could Be Done
As Low As Reasonably Acceptable (ALARA) (or Baseline Protection Objective (BPO))
Treatment Should Be Done
Baseline Protection Limit (BPL)
Treatment Must Be Done

Figure 3: Treatment Model

(Source: UK National IA Forum 2009)

This shows that for most practical systems, the degree of discretion in risk treatment is typically quite low, being only the balance of tolerance between Willingness To Pay (WTP) and Willingness To Accept (WTA).

5.2 Outliers

There is an overwhelming tendency to express and treat risks and their associated adversities as a single "magic number", which typically tends assume a simplistic, centre-tendency view of likelihood and uncertainty, or, in the limited number of cases where a probability distribution is assumed, this will typically tend to be modelled as a Gaussian [12, 13]

However, this view from both the producer (by implication) and consumer (by inference) is both naïve and unhelpful, especially as many cyber risks will have dramatically different underlying distributions:

- The large volume of Low Impact, High Probability (LIHP) adversities faced by many IOCT systems (e.g. network probes and malware infections) that for any individual instance would inherently count as an infinitesimal risk, yet the very existence of firewalls and anti-virus software (AVS) is a testament to the consensus that the aggregated risk is worth treating
- The underlying likelihood – which will typically be either unknown, or potentially even Unknowable – of new High Impact, Low Probability (HILP) adversities, a modern cyber-domain instantiation of the very essence of the Black Swan problem [14]

Work with both producer and consumer audiences has shown that there is little recognition of the specifics of LIHP and HILP risks, or even the spectral nature of risk tolerability.

To address this shortfall in understanding, an adaption of the classical Risk Heatmap is proposed [15] as shown in Figure 3, which extend the hierarchy of treatments defined in Figure 2 to include the LIHP and HILP factors.

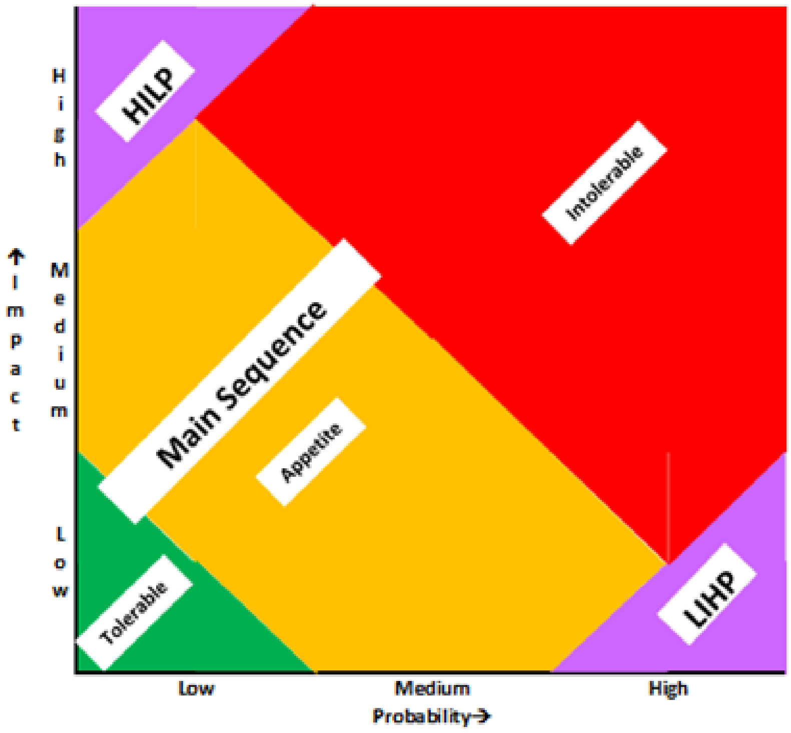

Figure 3: Bounded Risk Heatmap

Audience testing with cyber-physical systems users has shown this visualisation provides a useful way to discuss the LIHP and HILP issues that would otherwise be overlooked.

5.3 Appetite

It is a fact universally acknowledged that few people have identical appetites for risks, and that these appetites are contextually driven.

For instance, although an individual may be a Thrill Seeker – and thus willing to accept risk - in their private life, at work due to organisational pressures their stance may be diametrically opposite.

A number of codifications of risk appetites have be advanced, with one of the best known in the UK being the 5 layer categorisation defined by H M Government [16], as summarised in Table 2 overleaf, which assumes that risk can be broken down into a discrete distribution.

A challenge with this approach is it assumes a monolithic attitude towards all type of risk, which is probably unrealistic. This lack of realism is easily illustrated by reference to practice in the insurance markets, where customers are typically happy to accept an “excess”, which means that they get no benefit from the insurers for minor occurences where the downside risk to them customer is manageable, but with customer liability capped for larger occurences which would be intolerable to the individual, with the bulk of the risk falling to insurers.

Table 2: Generic Risk Appetites
(Source: H M Treasury)

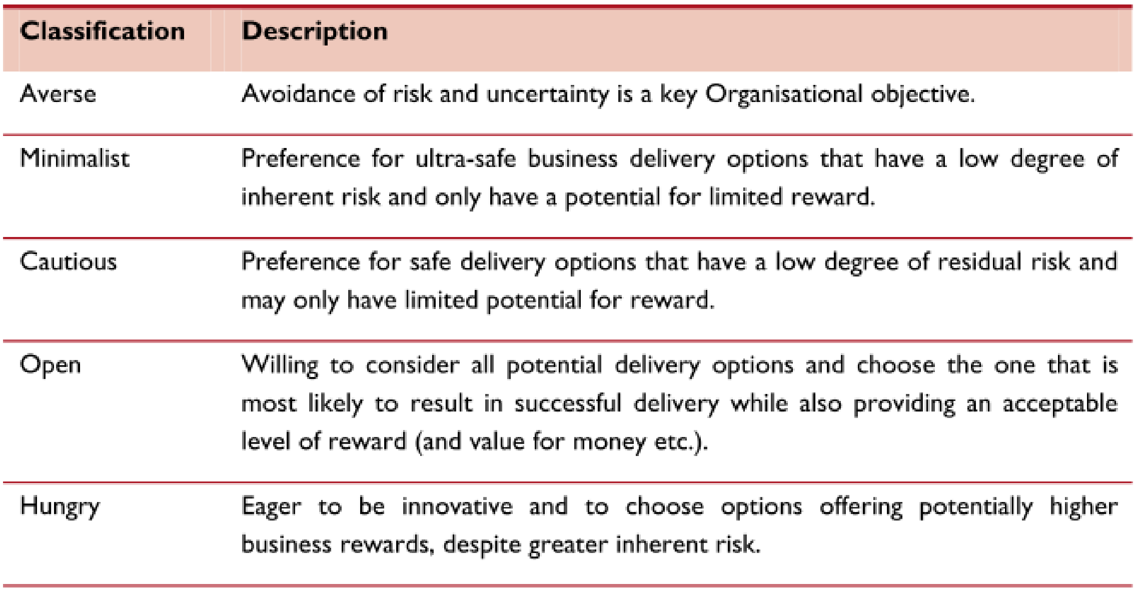

Classification	Description
Averse	Avoidance of risk and uncertainty is a key Organisational objective.
Minimalist	Preference for ultra-safe business delivery options that have a low degree of inherent risk and only have a potential for limited reward.
Cautious	Preference for safe delivery options that have a low degree of residual risk and may only have limited potential for reward.
Open	Willing to consider all potential delivery options and choose the one that is most likely to result in successful delivery while also providing an acceptable level of reward (and value for money etc.).
Hungry	Eager to be innovative and to choose options offering potentially higher business rewards, despite greater inherent risk.

A better expression of risk appetite is therefore that afford by Expected Utility Theory [17, 18], which defines risk appetites as being a continuous distribution, in three effective categories:

- Risk Aversion where the shape of the Utility Function is concave
- Risk Neutrality where the shape of the Utility Function is linear
- Risk Preference where the shape of the Utility Function is convex

If a continuous distribution is assumed, it would therefore be reasonable to represent a Risk Appetite, as well as individual and accumulated Risks, as a Conquad plot.

5.4 Tradeoffs

The three-way categorisation based on Expected Utility Theory maps very neatly to the three types of requirement typically found in Standards, including those used to control risks, of Mandatory (typically starting “shall …”), Preferred (typically starting “should …”), and Discretionary (typically starting “may …”).

From this categorisation, a generic approach to dealing with the need to tradeoff trustworthiness requirements is proposed, for instance inside a Statement Of Applicability (StOA), as shown in Table 3.

Table 3: Tradeoff Approach

Requirement Type	Tradeoff Justification
Mandatory (“Shall”)	A full set of details of rationale and consequences, to be formally presented, reviewed and endorsed by Senior Responsible Owner (SRO) or equivalent role, and maintained in system documentation.
Preferred (“Should”)	A summary of rationale and consequences, to be endorsed by Project Manager (PM) or System Operating Authority (SyOA), and maintained in system documentation.
Discretionary (“May”)	An explicit annotation to included and maintained be within system documentation.

6.0 NON-ANTHROPOGENIC SYSTEMS

6.1 Non-Anthropogenic Vehicles

The term Non-anthropogenic Vehicles (NAV) is proposed as a gender-neutral term to describe the whole set of vehicles that do not primarily rely on a human presence to discharge their movement functions, and a therefore a specialist subset of Cyber-Physical Systems (CPS).

These vehicles can potentially operate in one of six domains:

- Space (NAV-S)
- Air (NAV-A)
- Maritime (NAV-M)
- Land (NAV-L)
- Underwater (NAV-U)
- Hypogean[6] (NAV-H)

Although NAV is currently regarded as an emergent technology, the reality is that both remotely controlled and genuinely autonomous vehicles have existed in multiple domains for many years; indeed, the early history of rocketry is entirely one of NAV-A or NAV-S.

6.2 Non-Anthropogenic Payloads

Non-anthropogenic Payloads (NAP) are sensor or actuator systems, which can be sited on both NAV, or on Anthropogenic Vehicles (AV).

6.3 Non-Anthropogenic Systems

Both NAV and NAP – collectively Non-anthropogenic Systems (NAS) – can be characterised by the Level of Autonomy.

Research on Cyber-Physical Systems (CPS) in general, and with industrial partners in on-road motor vehicles in particular, now proposes extending the motor industry Levels of Autonomy [19] to a form usable for both types of NAS (NAV and NAP), as summarised in Table 3.

Table 3: NAS Levels of Autonomy

Level	**Name**	**Narrative definition**	
Human monitors the environment			
0	No Automation	The full-time performance by the human driver of all aspects of the control task, even when "enhanced by warning or intervention systems"	
1	Single Function Assistance	The mode-specific execution by an assistance system of single function	using information about the environment and with the expectation that the human performs all remaining aspects of the

[6] To avoid an acronym clash, which would be the case for Underground and Subterranean

2	Partial Automation	The mode-specific execution of more than one assistance systems	task
Automated system monitors the environment			
3	Conditional Automation	The mode-specific performance by an automated system of all aspects of the task	with the expectation that a human will respond appropriately to a request to intervene
4	High Automation		even if a human does not respond appropriately to a request to intervene
5	Full Automation		under all environmental conditions that can be managed by a human

7.0 MANAGING ADVERSITY FOR NAS

7.1 Scope of Risks

Much of the work on risks to NAS is predicated on the increasing number of access points for nefarious actors to corrupt collection integrity, inject false data, or modify data, with the intent to deceive or deny the mission.

But this Threat-based view represents only a subset of the ways in which a NAS mission can be endangered, and that the risk management of autonomous cyber-physical systems should take a holistic approach the capture and modelling of all sources of Adversity that can cause disruption to a NAS mission.

For NAS, it is important that the scope of the Adversity assessment is sufficiently wide, to include:

- The Non-Digital Information used by a NAS that originates from outside the NAS, for instance geospatial inputs

- The Digital Information Stored, Processed, or Forwarded by the NAS – the autonomy feature themselves

- The Non-Digital effects from the Digital Information Stored, Processed, or Forwarded by the NAS, for instance the analogue signals that cause movement of control actuators

7.2 NAS Adversity Enumeration

Adversity Set (AS) modelling is intend to allow the accumulation of a number of different Adversity Factors (AF) that cause damage or disruption on a cyber or cyber–physical system, such as a NAS, to simplifying review and treatment.

The first level of composition of an AS is into an Adversity Class (AC), and some worked examples of ACs relevant to NAS are:

- **AC.EM (Electromagnetic)**: combining the risks from both hazards (ranging from extreme Space Weather (SpW), though thunderstorms, to Unintentional Electromagnetic Interference (UEMI) from the likes of broadcast TV transmitters) and threats (with categories of

Intentional Electromagnetic Interference (IEMI) ranging from High-altitude electromagnetic pulse (HEMP) [20] to simple COTS cellular signal “jammers”)

- **AC.KI (Kinetic)**: combining the risks from both hazards (ranging from impact of meteoroids, though volcanic ejecta and wind-blown debris, to vehicle accidents) and threats (ranging from impact of a missile to impact of a sledgehammer)

In both cases the risk treatment measures are either coincident and/or overlapping, and therefore the AS view is less likely to result in nugatory effort and complexity that those arising from stovepiped views from, for instance, the “safety” and “security” communities.

7.3 Acceptability of Risks

In addition to the practitioner-focused work on quantifying and enumerating risk, a companion effort is ethnographic study of the alignment between such a model-based approach in contrast to the “audience” credibility and acceptability of such Adversity Risk judgements.

It is a reality of the modern world that public perception – and any consequential potential detrimental Reputational Impact – can cause adjustment or even abandonment of otherwise technically sound initiatives.

The initial findings of ethnographic work has thus far identified:

- That the audience expectation of Scope differs ssignificantly from custom and practice for Information/Cyber Security Risk, with the former expecting that both indirect (2nd party) and collateral (wider environment) impacts be taken into account, yet the typical organisational practice – as exemplified by ISO/IEC 27xxx Statements of Applicability (StOA) – being typically focused solely on the target organisation. This aligns with existing RASAE recommendations to seek AER_T (the Total AER for a single entity and all its externalities) rather than just $AER_{AS.E}$ (the AER from the internal (direct) Adversity Set)

- That the objective and subjective view of Risks frequently diverge, with former being modelled based on probable Deleterious Result (DR), yet the latter adding Perceived Susceptibility (PS – the non-expert misjudgement of a DR), and False Perception (FP – the non-expert misjudgement of a non-existent risk)

- That all forms of Automation are perceived as immature technologies, and, as such, the higher the Level of Autonomy, the wider the divergence between the pseudo-quantitative assessment of Adversity Risks and the perceived Risks

When it comes to considering the acceptability of behaviour of NAS, it would be remiss to not mention a challenge that falls outside of the DR / PS / FP taxonomy: that of problems associated with conferring a decision making function on automota, where no possible decision can lead to a non-deleterious outcome.

The classic example of such a challenge is the widely known Trolley Problem in moral psychology, which was framed over 100 years ago and relates to the “Who to kill?” conundrum for which there is no consensus answer.

The most widely known first extension of this to behaviour of automata was in the 1942 “Rules of Robotics” [22], for which it has been acknowledged there can be situations where no possible decision can lead to a non-deleterious outcome

One possible solution if all the alternatives are equally bad is to choose at random, but this latter approach would be likely to be regarded unfavourably by external audiences.

7.3 Acceptance of Innovation

Noting that the divergence between objective and subjective view of Risks can be correlated with the degree of maturity of a technology, it is interesting to reflect on previous work, such as the Diffusion of Innovations Model [23], and the Hype Cycle [24], which was combined as a Gartner-Rogers Model [25] shown in Figure 4.

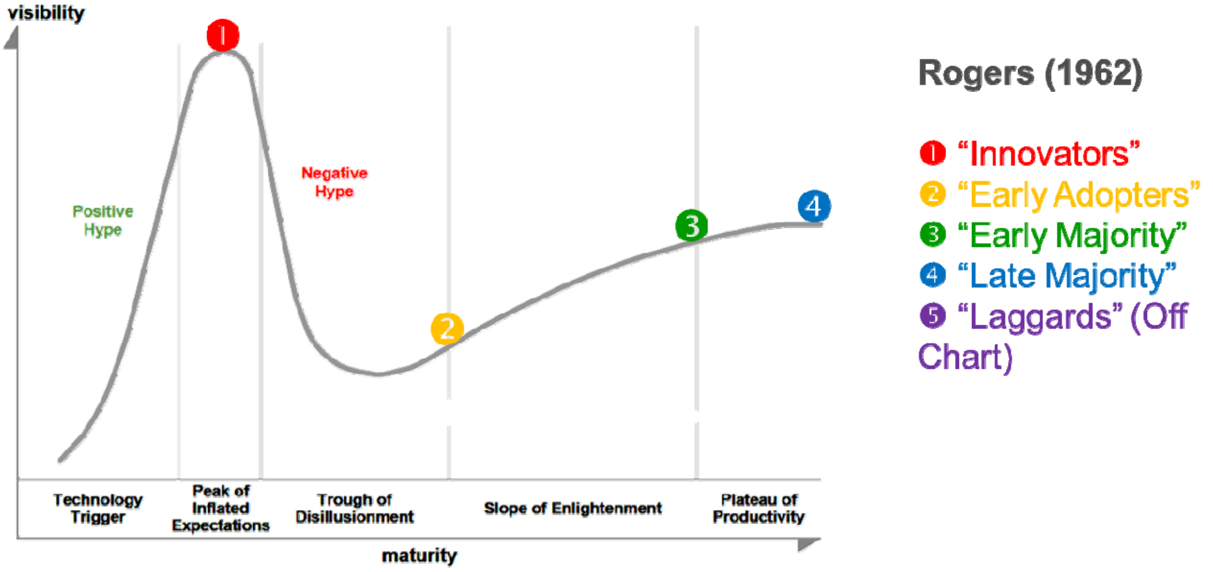

Source: Gartner (June 2005)

Figure 4: Gartner-Rogers Model

(Source: Bryant 2006)

A consequence of this is that in the case of NAS (and in particular NAP), the relative immaturity of the technologies increases the chances that perception of Risks (in particular from Safety factors) will inherently differ from the pseudo-quantitative assessment of Adversity Risks.

8.0 UNRESOLVED ISSUES

A number of areas from this phase of the RASAE project have been identified as requiring further investigation.

8.1 Risk Utility and Indifference

The work on NAS has clearly identified that audience acceptability of risks of such innovation does not fall neatly within the concept of there being a discrete model of Risk Appetite, but rather needs the more nuanced view such as that offered by understanding the shape of the utility function, as potentially summarised by a 10NP / Conquad plot.

Furthermore, the concept of indifference curves, when taken into consideration with a risk appetite expressed as a utility function, may offer a replicable and robust way in which to judge the optimal bundle of protective controls to be applied to manage a risk.

It is noted that valuable effort has already been made in this area within the insurance community, and an option for further work is to investigate how this can be transferred into both the NAS community in specific, and the wider CPS and cyber / information systems communities in general.

8.2 Messes and Wicked Risks

In addition to the problems of varying underlying distributions, the combined and/or blended nature of many real-world high-impact events [26] means that many cyber risks will be “Messes” [27] or “Wicked” [28].

Such forms of Messes and Wicked risks are not always amenable to a systematic treatment, which leads to the needs to consider approaches such as the Soft Systems Methodology [29] as a way of addressing changing, ill-defined problem situations.

8.3 Knowledge Transfer

The RASAE techniques have been used for small sample sets of CPS context, both in terms of the practitioners and audiences, but the ability to transfer the knowledge, and to scale to widespread use, remains a focus for future effort. This may need to include the need for automation support.

9.0 CONCLUSIONS

This phase of the RASAE project has established:

- The need to consider Risk Appetites as being a continuous rather than discrete distribution
- The recommendation for standardised ways to handle the need for Tradeoffs
- A taxonomy of sub-types of NAS (including NAV and NAP)
- A generic set of NAS Levels of Autonomy, as an extension to that widely used for motor vehicles
- The utility of the Bounded Risk Heatmap in helping audiences to understand Outlier Risks
- That in addition to objective probable Deleterious Result (DR), audience behaviours require consideration of Perceived Susceptibility (PS – the non-expert misjudgement of a DR), and False Perception (FP – the non-expert misjudgement of a non-existent risk)
- The utility of the RASAE Adversity Model (AM) approach (of Factors (AF); Classes (AC) and Sets (AC)) to simplify understanding of a complex set of Threats and Hazards when dealing with CPS systems such as NAS, but has reinforced the challenges in the underlying mathematics in respect of combining 10NPs, and dealing with Treatment Induced Risk (TIR)
- The utility of the RASAE recommendation to seek AER_T (Total AER for a single entity and all its externalities)

10.0 DEFINITIONS AND ABBREVIATIONS

10.1 Definitions

For the purposes of this document, the following terms and definitions apply:

- Adversity : the superset of Hazards and Threats
- Stakeholder : a person or organisation who will have an interest in or be a user of the RASAE approach

10.2 Abbreviations

For the purposes of this document, the following abbreviations apply:

AC	Adversity Class
AER	Annualised Expectation of Risk
AF	Adversity Factor
ALAPP	As Low As Physically Possible
ALARA	As Low As Reasonably Acceptable
ALARP	As Low As Reasonable Possible
AM	Adversity Model
AS	Adversity Set
B2B	Business To Business
BPL	Baseline Protection Limit
BPO	Baseline Protection Objective
BWD	Bryant-Watson Diagram
COTS	Commercial Off The Shelf
CPS	Cyber-Physical System
DRS	Deleterious Result Scale
HEMP	High-altitude electromagnetic pulse
HILP	High Impact, Low Probability
IEMI	Intentional Electromagnetic Interference
LIHP	Low Impact, High Probability
MTR	Maximum Tolerable Risk
NAP	Non-anthropogenic Payload
NAS	Non-anthropogenic System
NAV	Non-anthropogenic Vehicle
PM	Project Manager
RASAE	Replicable And Scalable Adversity Enumeration
SpW	Space Weather
SRO	Senior Responsible Owner
StOA	Statement of Applicability
SyOA	System Operating Authority
UEMI	Unintentional Electromagnetic Interference
VPI	Value of Protecting from Injury
VSL	Value of Statistical Life
WTA	Willingness To Accept
WTP	Willingness To Pay

REFERENCES

[1] “A Pareto Approach to Software Dependability”, I R C Bryant, NATO Research and Technology Organisation (RTO) Information Systems Technology (IST) Symposium (IST-111/RSY-026) on Information Assurance and Cyber Defence, September 2012

[2] “The Known, the Unknown and the Unknowable”, R Gomory, Scientific American, June 1995

[3] “Replicable and Scalable Adversity Enumeration (RASAE)”, I R C Bryant and T P Watson, NATO Research and Technology Organisation (RTO) Information Systems Technology (IST) Symposium (IST-122/RSY-030) on Cyber Security Science and Engineering, October 2014

[4] "A Theory of Human Motivation", A Maslow , Psychological Review, 1943

[5] OECD “The Value Of Statistical Life: A Meta-Analysis”, Ecole Nationale de la Statistique et de l'Administration, January 2012

[6] "On making life and death decisions - Societal Risk Assessment: How Safe Is Safe Enough?”, R A Howard, General Motors Research Laboratories, 1980

[7] “Using microlives’ to communicate the effects of lifetime habits”, D Spiegelhalter, BMJ, 345, December 2012

[8] “A Cross-Disciplinary Approach to Modelling and Expressing Adversity”, I R C Bryant C Maple and T P Watson, European Conference on Cyber Warfare and Security, July 2016

[9] ISO/IEC 27001:2013 “Information technology -- Security techniques -- Information security management systems – Requirements”

[10] “Orange Book: Management of risk - Principles and Concepts”, H M Treasury, October 2004

[11] “Risk Appetite, Balance, and Tolerance”, UK National IA Forum, March 2009

[12] “The Flaw of Averages”, S L Savage, Harvard Business Review, November 2002

[13] “The Flaw of Averages: Why We Underestimate Risk in the Face of Uncertainty”, S L Savage, Wiley, July 2009

[14] “The Black Swan: The Impact of the Highly Improbable”, N N Taleb Random House, July 2007

[15] “Formalising the Risk Heatmap”, I R C Bryant and T P Watson, University of Warwick, 2016

[16] “Thinking about risk - Managing your risk appetite: A practitioner's guide”, H M Treasury, November 2006

[17] "Exposition of a New Theory on the Measurement of Risk", D Bernoulli, 1738

[18] "Theory of Games and Economic Behavior", J von Neumann and O Morgenstern, Princeton, 1944

[19] SAE J3016 "Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles", September 2016

[20] IEC 61000-4-23 et seq “High-altitude electromagnetic pulse (HEMP)”

[21] “A Study of the Influence of Custom on the Moral Judgment”, Frank Chapman Sharp, Bulletin of the University of Wisconsin No. 236, June 1908

[22] "Runaround", Isaac Asimov, Astounding Science Fiction, March 1942

[23] Rogers, Everett M. (1962). “Diffusion of innovations”, E M Rogers, Free Press of Glencoe, 1st Edition, 1962

[24] “Hype Cycle for Emerging Technologies”, Gartner, 1st Edition, 1995

[25] “Warning the Non-technical Audience”, I Bryant, Annual Computer Security Applications Conference, December 2006

[26] "Normal Accidents: Living with High Risk Technologies", C Perrow, Princeton University Press, 1984

[27] "Redesigning The Future", R L Ackoff, Wiley, 1974

[28] "Issues as Elements of Information Systems", H W J Rittel and W Kunz, Heidelberg, 1970

[29] "Towards a systems -based methodology for real-world problem solving", P Checkland, J.Sys.Eng. 3, 87-116, 1972