NASA/TM– 20220006812

Directed Acyclic Graph Guidance Documentation

Human System Risk Board

Erik L. Antonsen, MD, PhD
Baylor College of Medicine, Houston, TX
NASA, Johnson Space Center, Houston, TX

Robert J. Reynolds, MS MPH PhD
KBR Wyle Services, LLC

Ahmed Abukmail, PhD
University of Houston Clear Lake

Avalon Monti, MPH
KBR Wyle Services, LLC

Kristina Marotta
Georgia Institute of Technology

Jacqueline Charvat, PhD
KBR Wyle Services, LLC

Daniel Buckland, MD, PhD
NASA, Johnson Space Center

Devan Petersen, MPH
KBR Wyle Services, LLC

Wilma Anton, PhD
KBR Wyle Services, LLC

Erin S. Connell
Leidos Innovations

Charlotte Brown
University of Michigan

Mark Shelhamer, PhD
Johns Hopkins University

Mary Van Baalen, PhD
NASA, Johnson Space Center

May 2022

NASA/TM– 20220006812

Directed Acyclic Graph Guidance Documentation
Human System Risk Board

Erik L. Antonsen, MD, PhD
Baylor College of Medicine, Houston, TX
NASA, Johnson Space Center, Houston, TX

Robert J. Reynolds, MS MPH PhD
KBR Wyle Services, LLC

Avalon Monti, MPH
KBR Wyle Services, LLC

Jacqueline Charvat, PhD
KBR Wyle Services, LLC

Devan Petersen, MPH
KBR Wyle Services, LLC

Erin S. Connell
Leidos Innovations

Mark Shelhamer, PhD
Johns Hopkins University

Mary Van Baalen, PhD
NASA, Johnson Space Center

Ahmed Abukmail, PhD
University of Houston Clear Lake

Kristina Marotta
Georgia Institute of Technology

Daniel Buckland, MD, PhD
NASA, Johnson Space Center

Wilma Anton, PhD
KBR Wyle Services, LLC

Charlotte Brown
University of Michigan

National Aeronautics and
Space Administration

Johnson Space Center
Houston, Texas

May 2022

Acknowledgments

This report is available in electronic form at

http://

ABSTRACT

For over a decade, the National Aeronautics and Space Administration (NASA) has tracked and configuration-managed approximately 30 risks to astronaut health and performance that occur before, during and after spaceflight. The Human System Risk Board (HSRB), a Health and Medical Technical Authority (HMTA) Board at NASA Johnson Space Center, is the entity responsible for identifying, assessing, analyzing, and monitoring the official understanding of the risk or risk posture for each of the Human System Risks and determining – based on evaluation of the available evidence – when that risk posture changes. The ultimate purpose of tracking and researching these risks is to find ways to reduce the risk that astronaut crews face during spaceflight. Historically, research, development and operations relevant to one risk have been conducted in isolation from other risks; these individual risk 'silos' enabled initial characterization of each specific risk. In spaceflight however, the impact of exposure to risk for astronaut crews is cumulative, and not independent of exposures or other risks, as all the adverse effects of the spaceflight environment begin at launch, continue throughout the duration of the mission and in some cases across the lifetime of the crews. In January of 2020, the HSRB at NASA embarked on a pilot project designed to assess the potential value of causal diagramming as a tool to facilitate understanding these cumulative and interdependent effects as applied within Human System Risk management. This process uses directed acyclic graphs as a means of formalizing a shared mental model of the causal flow of risk among Risk Board stakeholders. Initially this model was to improve communication among those stakeholders, but the potential value exceeds communication alone. Formalization of the process for creating these causal diagrams will enable the creation of a composite risk network that is vetted by members of the NASA community and configuration managed. The causal diagrams are formulated as directed acyclic graphs (DAGs) to function as a type of knowledge graph for reference for the board and its stakeholders. This document outlines the pilot process, the standardized approaches, and guidance for risk custodian teams when creating and updating DAGs as a part of the NASA Human System Risk Management process.

Table of Contents

Figures

Tables

Introduction:

Background

The Human System Risk Board (HSRB) at National Aeronautics and Space Administration (NASA) is tasked with identifying, assessing, tracking and reporting on the Human System Risks faced by astronaut crews due to spaceflight. The background and approaches used are detailed in the JSC 66705 Human System Risk Management Plan which is publicly available for download at https://ntrs.nasa.gov/. This NASA Technical Memorandum is intended to provide guidance for the creation and maintenance of a set of causal diagrams in the form of Directed Acyclic Graphs (DAGs) for the current set of Human System Risks that are managed by the HSRB. Summarized background of the HSRB approach is provided here, but for complete insight please refer to the JSC 66705 HSRB RMP.

NASA currently tracks approximately 30 health and performance risks that spaceflight crews face in human spaceflight missions, all of which derive in some way from five main Hazards which are unchangeable aspects of spaceflight harmful to humans [1]–[3]. These Human System Risks are evaluated against eight Design Reference Mission (DRM) categories that are mission templates for assessing risk posture applicable to a range of human space exploration missions relevant to NASA goals [2]–[4]. The DRM categories include short duration missions in Low Earth Orbit (LEO), missions to lunar space or the lunar surface, and excursions to Mars that can last up to three years. While approaches to understanding and mitigating these risks continue to evolve, the fact remains that there is significant uncertainty surrounding how much risk astronaut crews will face as exploration mission durations increase and extend beyond LEO.

In addition to the operational and clinical issues encountered during any given mission, NASA is also responsible for research priorities and clinical care for astronauts post-mission, and to some extent, post-career. Research on spaceflight Human System Risks has traditionally been organized into specialized topic areas or risk “silos.” This was initially advantageous for understanding the underlying nature of each risk by subject matter experts (SMEs), but risk in the human system is multifactorial and cumulative. This has led programs to increase cross-risk research efforts, which in turn has created the need for better methods of characterizing and representing the complex interactions between risks [5].

Why are we using Directed Acyclic Graphs for Human Spaceflight Risks?

Directed Acyclic Graphs (DAGs) are network maps which have unidirectional arrows (directed) and do not allow feedback loops (acyclic). The properties of DAGs make them a format well-suited to drawing causal diagrams. In the context of the HSRB, DAGs are used to represent the chain of events that lead from spaceflight exposures to negative mission-level outcomes. This enables two immediate uses as well as sets the stage for further evolution of the causal networks as tools of inference.

DAGs for communication of complex human spaceflight risks

The HSRB uses DAGs to serve as knowledge graphs displaying the set of factors that originate and propagate risk. These factors originate from the five Hazards of human spaceflight, and, through multiple chains of causation, terminate in Mission Level Outcomes. The Mission Level Outcomes are health and performance outcomes that are relevant to NASA as an agency and the individual astronauts assigned to missions. Four of the five Hazards—including *altered gravity*, *space radiation, isolation and confinement*, and the *hostile closed environment* — are ever-present features of the spaceflight environment. These Hazards are the ultimate cause of the process of degrading astronaut bodies and capabilities, which begins as soon as astronauts are launched into space. The fifth Hazard, *distance from Earth*, limits the provision of in-mission support capabilities and resources, while simultaneously increasing the need for them. Limits on these capabilities and resources stem from constraints on mass, volume, power, and data bandwidth allocations available to the vehicle's systems/habitats used by astronauts; the further a mission takes astronauts from Earth, the greater these constraints and thus the less support capability they will have. The need for capabilities and resources is increased because the further a mission goes from Earth, the longer astronauts are exposed to degradation by the spaceflight environment.

DAGs aid in prioritizing research and development

To mitigate the complex risks of human spaceflight, the stakeholders of the HSRB must evaluate the set of Human System Risks to prioritize where the investment of limited research, surveillance, and technology development resources should be targeted. Historically, these have been prioritized based on likelihood and consequence (LxC) scores (and associated 'Red, Yellow, Green' stoplight colors) assigned by the Board on an individual risk basis. However, this approach fails to consider the complex synergies between the risks and the effects that interactions among multiple risks have for the cumulative risk that astronauts face in-mission. For example, though a minor medical issue may be of small consequence on its own or in the immediate moment, the interconnectedness of risks means that it could create or amplify risks in other body systems and/or at a later time. Because of this potential for “downstream” effects, the Board created prioritization principles that are outlined in Section 3.8 of the JSC 66705 RMP [3]. Through the analysis of the structure of DAGs, nodes can be identified as the “important” factors in the causal network. In this context, “importance” can be thought of as those factors that have many causes, those that have many effects, those that bridge or join risks together, or those that exist “in the middle of the action”. The practical implication of scoring highly on these measures is that such factors are the ones that have the broadest influence (across the entire network) on the likelihood of one or more risks. The Human System Risk DAGs were created to enable a systematic and repeatable approach to indexing and comparing the drivers of risk in a complex and interconnected system.

How can DAGs be used by the stakeholder community?

Communication using Narrative DAGs

The set of DAGs that were originally created by and are continually maintained by the risk custodial teams for use at the HSRB are termed Narrative DAGs. They are primarily intended to serve as communication tools documenting the communal understanding of the sources and flow of the individual risks. Through this process of continual curation, the DAGs are “crowdsourced” products, as they represent knowledge and expertise from NASA scientists in the Human Health and Performance Directorate, as well as their best understanding of the evidence from the scientific literature. These DAGs are used by the Board to provide a shared mental model of the important causal pathways for each risk across the wide variety of SMEs, scientists, clinicians, operators, and program managers.

Network analysis using Detailed DAGs

Narrative DAGs may also be expanded to become Detailed DAGs, which include more granular depictions of the causal systems they represent. Detailed DAGs have additional nodes and arrows to show otherwise-omitted intermediate steps in causal pathways or individual nodes nested inside category nodes. When implemented with a common set of node names, Detailed DAGs enable the composition of a single risk network that intrinsically carries shared nodes and links between the risks within its structure. It is this structure that provides the opportunity to apply network analysis techniques to interrogate the complex structure and relationships. As such, the DAGs carry the potential to provide a quantitative, systematic, and repeatable approach to prioritizing those research, surveillance, and development investments that are most likely to provide the best return on investment as we look ahead to future NASA missions.

The rest of this document outlines the guidance for creating and interpreting DAGs used by the HSRB. These are intended to be consistent with the HSRB Charter and JSC 66705 RMP for the Board [3].

Methods

Directed Acyclic Graphs (DAGs)

A graph data structure is composed of a set of *vertices* (nodes), and a set of *edges* (links) G = <V,E>. Each edge represents a relationship between two nodes, making the nodes *adjacent*. There can be two types of relationships between nodes, and thus there are two types of links between adjacent nodes: *directed* and *undirected*. A graph typically has either directed or undirected links, making it either a *directed graph* or an *undirected graph*. For example, if an edge exists between two nodes *A* and *B* in a graph and the edge is undirected (A–B, no arrow) then *A* and *B* are said to be adjacent to one another; *A* is adjacent to *B* and *B* is adjacent to *A*. However, if a directed link exists between *A* and *B* with the arrow pointing only from *A* to *B* (A->B) then it is said that *A* is adjacent to *B* but *B* is not adjacent to *A;* another directed link from *B* back to *A (A<-B)* would be needed to make *B* adjacent to *A.* A directed graph can potentially contain a *cycle*, meaning that, from a specific node, there exists a path that would eventually return to that node. A directed graph that has no cycles is known as *acyclic.* Thus, a graph with directed links and no cycles (A->B) is a DAG.

DAGs therefore can be a type of network diagram which, through specific conventions in their construction, represent causality in a visual format [6], [7]. Specifically, each directed arrow connecting one node to another on a DAG indicates a claim of causality. In general, a reverse relationship can be shown by a two-sided arrow, but these would violate the acyclic requirement implemented here. Figure 1 shows an example of a directed acyclic graph. This graph tells a story about causality.

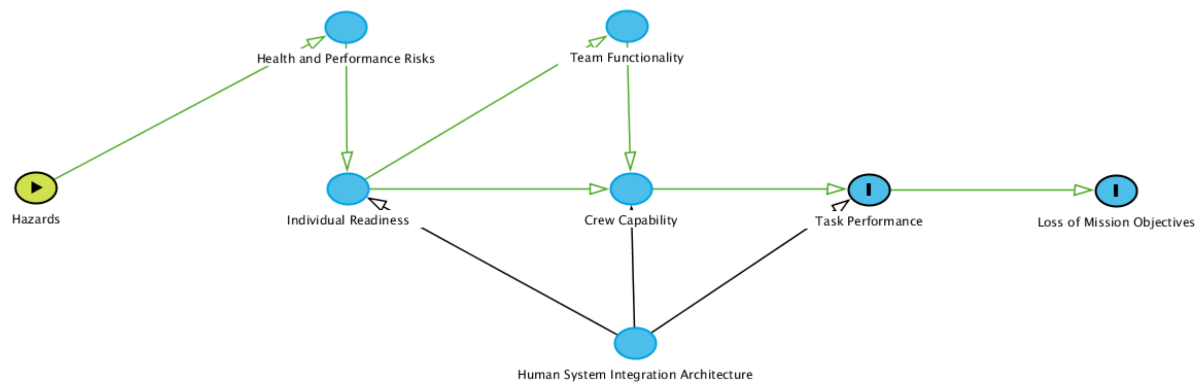

Figure 1: Example of a Directed Acyclic Graph. This is a simplified illustration of how and the individual, the crew, and the system contribute to the likelihood of successful task performance in a mission. Individual readiness is affected by many of the health and performance oriented risks followed by the HSRB, but the readiness of any individual crew is complemented by the team and the system that the crew works within. Failures of task performance may lead to loss of mission objectives if severe.

Starting at the left side of the graph, the green *Hazards* node is shown representing the Hazards of human spaceflight. The causal flow between the cause (*Hazards*) and effect (*Health and Performance Risks*) is established by the directional arrow between them. Most of the Human System Risks exert a deconditioning effect on an individual crewmember in the spaceflight mission. This causal effect is shown by drawing an arrow from the *Health and Performance Risks* node to the *Individual Readiness* node. The *Individual Readiness* of one crew member does not always determine mission success or failure. That individual contributes to *Team Functionality* and to the *Crew Capability* overall. If one crewmember is incapacitated, it may be possible that another member of the team can take over the responsibility for specific mission tasks. Whether this happens depends on many factors that are rolled into the terms *Team Functionality* and *Crew Capability*. For example, if another crewmember was trained in that specific task and is competent, then the task may be completed successfully. But if no other crew member was trained, the likelihood of failure of task performance increases, because the *Crew Capability* was inadequate without a backup. Whether or not a single task is successful usually does not determine if there was a *Loss of Mission Objectives*, but in some cases it may. Of note, successful *Task Performance* is not dependent on the crew alone, but also on the system they function within. The *Human System Integration Architecture* (HSIA) node denotes the umbrella effects of all of the design decisions and compromises within which the crew must perform their tasks. If there are inadequate spares to repair a broken part, then a repair task may not be performed regardless of how ready and able the crew is to perform the task. The *HSIA* node affects *Individual Readiness*, *Crew Capability*, and ultimately influences *Task Performance* both separately and combined with those factors.

This DAG in Figure 1 tells a basic story about the relationships between factors that contribute to successful *Task Performance* and by extension, those factors that most directly contribute to the possibility of *Loss of Mission Objectives*. *Health and Performance Risks* denote how the *Hazards* affect an individual and their *Individual Readiness*; *Team Functionality* and the individual contributions of crewmembers determine *Crew Capability*; the *Human System Integration Architecture* node defines how well the system enables crew to perform their jobs. Each of these steps is a causal step, and each of them represents a hypothesis that can be supported or disproven by assessing the evidence relevant to the causal claim indicated by the arrows. This example is not an HSRB approved DAG for a specific risk, it is simply the visual representation of a story about the factors that contribute to performance in spaceflight and the outcomes of concern to the agency. This story is presented here as an example to help the reader become familiar with the interpretation of DAGs. Because this story is now written down in graphical format, it can allow mathematical analysis of the structure of the relationships and potentially the strength of influence of those relationships if there is sufficient evidence to support the assignment of quantitative values to the nodes and edges. Importantly, each of the arrows that link one node to the next is a falsifiable hypothesis. This means that while this shows the best representation of our current understanding, it can always be challenged with evidence and improved. If the evidence base evolves to suggest that a causal connection does not exist where we have drawn one, then that connection can be removed. The DAGs that have been created for each of the Human System Risks have been broadly evaluated and commented on by the larger HSRB, stakeholders and community, and are configuration-managed stories similar to this one. These will be discussed in more detail below.

Causality is defined here to mean that any change or variation in the dependent variable (effect) is influenced by variation in the independent variable (cause). In simpler terms, we can imagine a binary factor that we suspect is causal: the probability of realizing a specific outcome is different when the factor is present than when it is absent. Table 1 shows an example of causal and non-causal relationships between two binary variables, A and B.

Table 1: Illustration of potential causal relationships between Factor A and Outcome B

	P(B=1 | A) when…	
Value of A	A not a cause of B	A is a cause of B
Absent	0.6	0.6
Present	0.6	0.8

As the first column in Table 1 shows, when A is not a cause of B, the probability that outcome B occurs, P(B=1 | A), is 0.6 whether factor A is present or not. However, if A is a cause of B, then the probability of event B occurring is greater in the presence of factor A (0.8) than it is in absence of factor A (0.6). Note that in Table 1 Factor A does not need to increase the probability of Event B to be a causal factor; it is causal if it effects a *change* in probability, either positive or negative. Also, a causal factor does not need to be the sole cause of an outcome but may be one of several contributors.

A more concrete spaceflight example could be the idea that exposure to altered gravity (i.e., gravity less than that on Earth, 1g) causes skeletal unloading, which leads to changed bone remodeling and decreased bone density in the body. This in turn can lead to skeletal fragility and increased likelihood of bone fractures. In practice this means that the probability of observing a given amount of bone remodeling is different depending on the gravity field to which astronauts are subjected, e.g., lunar (1/6g) gravity, Mars (3/8g) gravity, or microgravity (0g). Likewise, in this example, the degree of skeletal fragility is determined by the bone architecture after bone remodeling, and the probability of a bone fracture occurring is dependent in part on skeletal fragility. This example spans a Hazard of spaceflight (altered gravity exposure) to cellular level changes (bone remodeling) to functional changes (skeletal fragility) to an outcome (bone fracture). That outcome can lead to further outcomes such as affecting individual readiness of a crew member that, if severe enough, can affect Mission Level Outcomes. The next section discusses the derivation and standardization of DAGs for the unique problem of Human System Risks in spaceflight.

DAG Derivation for Human System Risks

The pilot study to create DAGs for each of the Human System Risks is a derivation challenge. It requires subject matter expertise across each of those risks to logically map out what factors contribute to the flow of risk. It requires common starting and ending points for all risks that are relevant to the NASA human spaceflight missions. It also requires presenting the flow of risk at a level of communication that is appropriate for non-experts to understand, but which can also be expanded in detail to be useful to the SME. The process therefore suggests two types of DAGs.

- Narrative DAGs are those intended to convey high-level and more aggregate concepts linking the key components of the causal flow to downstream effects in the risk domain. These are used to facilitate communication and shared mental models at a board or stakeholder meeting.

- Detailed DAGs allow a more in-depth exploration of the risk but are still intended to support the fidelity of the Narrative DAGs' high level 'story'.

Both Narrative and Detailed DAGs adhere to strict guidelines for standardized representation of relationships, and prescriptive terminology to ensure compatibility with other risk DAGs for the purposes of network analytical evaluation.

As mentioned previously, one primary feature of a DAG is that no feedback loops are permitted. In addition to keeping with the requirement that causes must precede effects, the "no feedback loops" requirement is included for two other reasons. First, as it is our experience that inclusion of too much detail often derails effective communication about risk, it is intended to ensure that the diagrams show a simplified flow of space flight Hazards to Mission Level Outcomes that enables understanding at a high level. Second, the requirement encourages the creators of these diagrams to articulate the most important steps that are likely to lead to undesirable Mission Level Outcomes while avoiding unnecessary complexity. In the case where the graph creators note that a feedback loop is present, they must choose the most likely predecessor node based on relevance to the spaceflight experience.

Creating the DAG Field

Exposure Set - Hazards

At NASA, the Human System Risks have historically been conceptualized as deriving from five Hazards present in the spaceflight environment [1]–[3], [8]. These Hazards can be thought of as unchangeable aspects of the spaceflight environment that are encountered when someone is launched into space and therefore are the starting point for causal diagramming of spaceflight-related risk issues for the HSRB. Whether an astronaut is launched to suborbital space for a short amount of time, or to Mars for a years-long mission, exposure to these Hazards (Figure 2) defines how and why physiologic changes happen to humans in spaceflight. As such they provide the common starting point for mapping how spaceflight risk evolves after entering the spaceflight environment. These Hazards are often interpreted in light of physiologic changes that occur in humans as a result of the exposure. However, it is not solely physiologic changes that lead to risk in human spaceflight. The interaction between human crew– which may be degraded due to the spaceflight environment – and the vehicle and mission systems that the crew must operate – are also dependent on these Hazards.

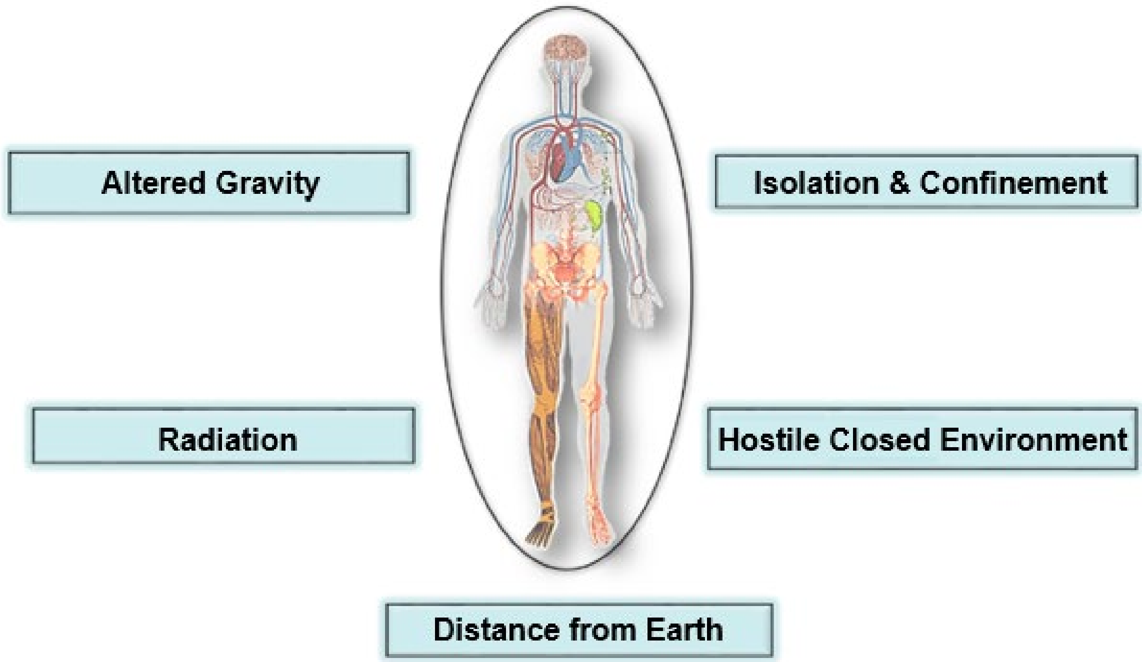

Figure 2: The five Hazards of human spaceflight - features of spaceflight from which Human System Risk derives.

Outcome Set – Mission Level Outcomes

Agreement on the ultimate effects of causal pathways in DAGs is a critical driver for communication among stakeholders of Human System Risks. Thus, at the terminus of the causal pathways are a set of Mission Level Outcomes, which are pre-specified outcomes common to all risk visualization. These are outcomes that rise to clinical or operational significance for the agency.

Common Start and End Points

Each Human System Risk DAG is intended to show the causal flow of risk from Hazards to Mission Level Outcomes. As such, the structure of each DAG starts with at least one Hazard and ends with at least one of the pre-defined Mission Level Outcomes (Figure 3). In between are the nodes and edges of the causal flow diagrams that are relevant to the Risk under consideration. These are called 'contributing factors' in the HSRB terminology, and include countermeasures, medical conditions, and other Human System Risks.

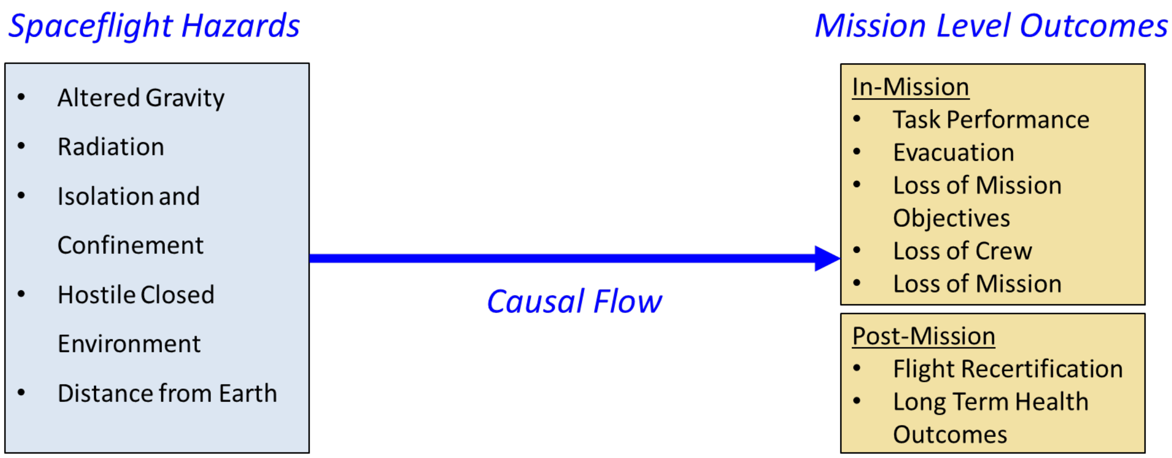

Figure 3: Common starting and ending points for the visualization of risk is the first step to a community-wide agreement on causal flow that is best supported by the available evidence.

Basic Requirements for Human System Risk DAGs

A contributing factor should be included on a DAG if it meaningfully affects risk posture, as determined by either: 1. Available high-quality evidence; or 2. SME concern. This is subject to additional guidance described below. Since the inclusion criteria for both nodes and edges to be eligible for a particular DAG are based on an assessment of evidence, it is important to review how the HSRB assesses evidence. The HSRB uses a Level of Evidence (LOE) assignment methodology that applies to risk posture as well as DAG inclusion and connections.

The LOE assigned by the board is based on a subset of the causal guidelines from the A. Bradford Hill Criteria, shown in Table 2. Deeper explanation for these and their applicability to the DAGs is provided in Appendix G of the JSC 66705 HSRB RMP.

Table 2: Sir A. Bradford Hill's causal guidelines employed by the HSRB for level of evidence assessment [3].

Criterion	Definition	Notes
Temporality	The effect must occur after the cause (and if there is an expected delay between the cause and expected effect, then the effect must occur after that delay).	This is necessary for all posited causal effects, even speculative ones.
Analogy	The use of analogies or similarities between the observed association and any other associations.	Analogues can be in exposure, population, or both.
Mechanism	If there is a plausible theoretical mechanism that can explain how the causal effect works, then the posited causal connection is more likely to be true.	
Reproducibility	Consistent findings observed by different persons in different places with different samples strengthens the likelihood of an observed effect being causal.	
Specificity	Causation is likely if there is a very specific population at a specific site and disease with no other likely explanation. The more specific an association between a factor and an effect is, the bigger the probability of a causal relationship.	This is the classic Person/Place/Time of epidemiology.
Coherence	Coherence between epidemiological and laboratory findings that validate the mechanistic assumptions increases the likelihood of an effect.	This is translational science.

For the DAGs, the application of the causal guidelines is needed to set a minimum criterion for what qualifies for inclusion in the DAG in terms of nodes and edges. NASA risk custodians are expected to bring their knowledge of spaceflight data and relevant terrestrial literature forward to assign a LOE for DAG creation and review. Generally, there are multiple sources of data or published evidence that are considered, and the causal guidelines are used to apply a LOE score to the full set of evidence that supports a specific node or edge. The assignment of the LOE score is visualized in Table 3 below.

Table 3: Excerpt from the JSC-66705 document showing the scoring criteria for Levels of Evidence Assignment. Note that any one of Mechanism, Reproducibility or Specificity can qualify for the Weak Level of Evidence.

LOE	Temporality	Analogy	Mechanism	Reproducibility	Specificity	Coherence
4 - Speculative	✓	✓				
3 - Weak	✓	✓	✓			
2 - Moderate	✓	✓	✓	✓	✓	
1 - Strong	✓	✓	✓	✓	✓	✓

The LOE score is used both in terms of filtering nodal inclusion (whether an event or factor actually occurs in the risk system) as well as to describe the properties of edges (whether or not a particular node has a causal influence on any other nodes); these procedures are described in more detail below. However, it is important to note that, as a necessary precursor to the evaluation of the body of evidence for a particular node or link, the ‘quality’ of the available data or publications must be evaluated by the risk custodians on a piece-by-piece basis before data or publications should be admitted to the evidence base. Evaluation of the quality of evidence demands a critical review and assessment of the data or publication relevant to the Human System Risk being considered. In plain terms, just because a peer-reviewed paper asserts a claim does not mean the board should simply accept its conclusions. If the publication that asserts specific claims shows poor methodology, inappropriate scoping of experiment to conclusions drawn, or does not appropriately discuss limitations or data that could lead to other conclusions, then the quality of that evidence is automatically considered low and in the LOE approach it should not be used to support a level assignment higher than ‘Weak’. To put it simply, the quality of evidence looks at whether a single publication or piece of data is of sufficient ‘quality’ to support the conclusions reached. The LOE is an assessment of the full body of evidence available to support or refute the inclusion of a node or edge in the causal diagram. Additional guidelines for evaluating quality of evidence are provided in the JSC 66705 HSRB RMP in Section 3.2.7 and in Appendix F, specifically the sections titled Animal and Cellular Models and Clinical Research Studies.

What qualifies a node for inclusion on a DAG?

A node is eligible to be considered for inclusion in a particular DAG if minimum nodal evidence criteria are met.

The minimal nodal evidence criteria for relevance to the risk under consideration based on available evidence includes:

1. The available evidence meets the minimum criteria for the 'Speculative' LOE (Temporality and Analog).
2. The proposed node has at least 1 connection of 'Speculative' or stronger LOE, to at least one other node within the Risk's DAG.

Note that this is the minimum level and does not guarantee that a node should be included. Stakeholders should discuss the importance of a proposed node in terms of its impact on the risk and in terms of the appropriate level of detail for the DAG to enable communication. Higher quality evidence that shows the relevance of the node to the risk causal flow (i.e., network structure) increases the value of inclusion in the DAG. Expectation of a significant magnitude of impact to the risk of Mission Level Outcomes (i.e., the node's contribution to outcome likelihood) also increases the value of the node's inclusion in the DAG. It should be noted that the contribution to likelihood may come from a large proximal contribution to an outcome itself, or small contributions to many more distant factors that eventually sum to a larger "downstream" contribution to the outcome.

What are the properties of a node?

Inclusion of a node in a DAG means that the node:

- Meets plausibility criteria, as defined above;
- Is one of the pre-determined set of Hazards or Mission Level Outcomes;
- Is considered by the HSRB or risk custodian team to be a factor worth noting in the causal flow of a specific Human System Risk towards Mission Level Outcomes.

What qualifies an edge for inclusion on a DAG?

An *edge* is a connection between nodes that signals a causal link exists between those two nodes. A given edge is eligible for inclusion if minimal causal criteria are met. The minimal causal criteria for an edge are that the available evidence supporting a causal relationship between the two connected nodes in the direction of the arrow meets the minimum criteria for the 'Speculative' LOE (Temporality and Analogy) (Table 3).

What are the properties of an edge?

If an edge is included in a DAG, then that edge represents a causal claim. This is a hypothesis that a causal link between the two nodes exists according to the A. Bradford Hill Criteria (Table 2). There are two basic properties relevant to the hypothesized causal link:

1. The likelihood that the cause-effect relationship claimed does indeed exist can be assessed by an LOE assignment. The LOE must be at the 'Speculative' level at a minimum for inclusion. The edge can be assigned an LOE score that denotes the strength of evidence associated with the causal claim. Higher LOE scores denote an increased likelihood that the proposed causal connection exists.
2. The strength of influence that the origination node has on the termination node can be determined if appropriate evidence is available. This is reflected by path coefficients (which are derived from covariation between the cause and the effect variables) or as entries in joint probability tables that express the changing probability of realizing various values of the outcome variable depending on the values of the causal variable.

Note that the criteria above are the minimum entry criteria and not a guarantee that the board will agree to the inclusion of a node or edge. In simpler terms, Temporality means "the proposed effect comes after the cause," and Analogy means "we've seen things like this elsewhere so it could happen here." Taken together Temporality and Analogy can be summarized as "It's not impossible, but we don't yet have the evidence to be sure that it's true." As the minimum entry criteria, these are only used to justify 'Speculative' level concerns that SMEs have for representation in the graph. Proposed nodes that meet these criteria do not have to be accepted into the DAG for a given risk. Just because something isn't impossible does not mean that including it brings value to the discussion about risk. Too much speculative detail can serve to derail effective communication. If a proposed node is supported by the published literature, consideration is given to the quality of the publication and the science expressed in it (guidelines for assessing this are provided in the JSC 66705 HSRB RMP). Support for proposed nodes may also come from unpublished data in the spaceflight domain. In these cases, consideration must be given to SME interpretation of those data and their relevance to the risk being diagrammed.

For a Narrative DAG there must be concern that the magnitude of a proposed effect is sufficient to create a measurable downstream effect on one or more adjacent nodes. These requirements help limit the size and scope of the DAG to key nodes and relationships of interest. Imposing this criterion helps to ensure that key factors believed to be important to a given risk can be visually related to the outcomes of interest and are relevant to NASA risk concerns.

While basic requirements for a DAG are necessary, they are not sufficient to enable like-to-like comparison across risks or to reliably illustrate how risks interact. Once the proposed set of contributing factors for a given risk are articulated, and the first draft of causal connections drawn between them, there are additional follow-on steps to ensure that the risk DAG is interpretable and can interface with the other risk DAGs. Standardizing the representation of Human System Risk DAGs is discussed in the next section.

Standardized DAG Representation

Between the Hazards and the Mission Level Outcomes lie the nodes and edges of the causal diagram that are intended to illustrate key relationships in risk propagation toward Mission Level Outcomes. The last section noted that key nodes would be included and provided definitions for Human System Risk items that qualify for node and edge inclusion. This section discusses how we name and depict these nodes and the edges between them in a standardized fashion. These are structured to align with the needs of the HSRB and its stakeholders.

Basic Drawing Guidance

Nodes are represented by circles. Edges are represented by arrows in the DAG and show causal relationships. Arrows are drawn starting from causes and extending into effects. Nodes that have one or more arrows coming out of them, with no arrows coming into them, are called exogenous nodes (meaning they are not caused by any nodes in the network). Nodes that have one or more arrows coming into them are called endogenous nodes (meaning they have causes within the network).

Contributing Factors

Contributing factors are items including any Hazard or operational design and human system variable that influences the outcome(s) of concern for a risk impacting human health and performance in spaceflight. It can be seen as a system variable whose state can contribute to mission success or failure and considered alterable through the implementation of risk mitigations. In the case of Human System Risks we refer to the risk mitigations as countermeasures. These are discussed further below.

By these definitions, any node that precedes the Mission Level Outcomes can be defined as a contributing factor. Within a specific risk DAG, we rename exogenous nodes as external factors to provide continuity with existing Board terminology. We also rename endogenous nodes as integral factors. Integral factors help complete the causal paths between spaceflight Hazards and Mission Level Outcomes and are essentially intermediate factors on the path from a particular exposure to an outcome of interest. For the purpose of Human System Risk DAGs, the ultimate exposures are the set of spaceflight Hazards and the ultimate outcomes are the set of Mission Level Outcomes. We say ‘ultimate’ here because outcomes and exposures are defined contextually between sets of nodes: any node that is upstream of another node can be called an exposure and any node that is downstream can be called an outcome. Therefore, generally the use of the term exposure or outcome will refer to upstream or downstream relative placement, and when referring to the starting points and ending points of a given risk DAG we will use the terms Hazard and Mission Level Outcomes.

Graph Theory vs. HSRB

The terms endogenous and exogenous nodes are a standard notation in graph theory. For the purposes of Human System Risk DAGs we have adopted a slightly different terminology to be more familiar with the standard risk terminology used by NASA. The HSRB has a long history of using the terms contributing factors and countermeasures and a relationship is defined here between the DAGs and the existing terminology. The definitions for contributing factor and countermeasure are provided in the JSC 66705 HSRB RMP in Appendix D [3]. Figure 4 below visually depicts graph theory terminology (A.) and the applied Human System Risk terminology (B.).

A. graph theory terminology using *Exposure*, *Endogenous* and *Exogenous Nodes* and *Outcomes*. B. shows the NASA Human System Risk notation used for the same DAG including *Hazards*, *External Factors*, *Integral Factors*, and *Mission Level Outcomes*.

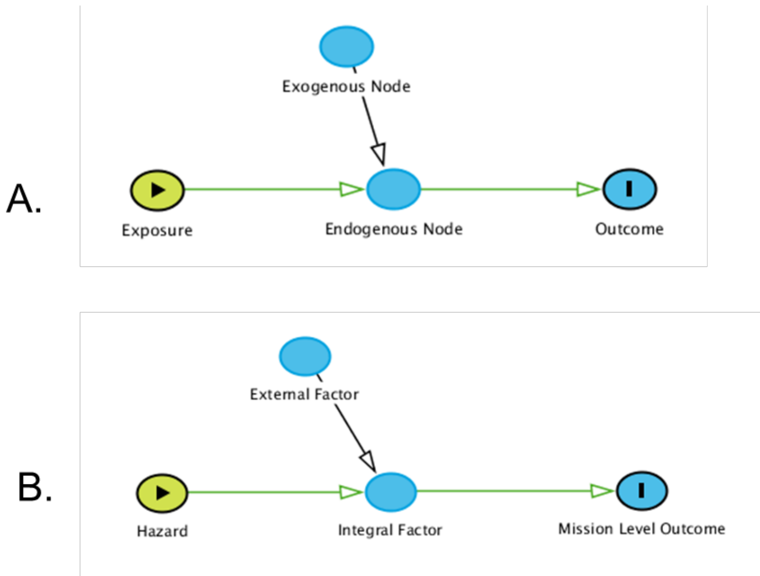

Figure 4: Example DAG which shows A.) an Exposure, an Endogenous Node, an Exogenous Node, and an Outcome and B.) which shows a Hazard, Integral Factor, External Factor, and Mission Level Outcome. The Notation in B. is used for Human System Risk DAGs.

Of note, the green oval is used to represent any node that is chosen to be an exposure for the purpose of immediate communication. The blue oval with a vertical line is used to represent any node chosen to be an outcome. Green lines for arrows show integral paths from exposure to outcome. Black lines for arrows show paths from external factors. These are artifacts of the Dagitty program (dagitty.net) that is used to create the initial versions of the DAGs and may change in the future as software packages evolve. Figure 5 shows the legend for the current display options used for the DAGs.

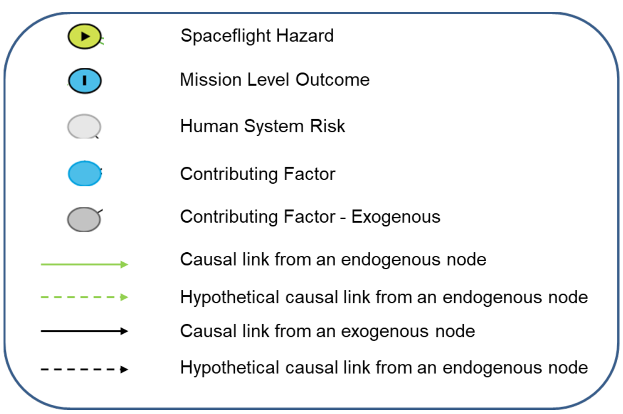

Figure 5: Legend showing the meaning of the node and edge depictions used with the Dagitty software [9].

Countermeasures

Countermeasure is a common NASA term used to describe any item or design feature intended to mitigate a specific aspect of a given risk. For example, anti-nausea medication is a countermeasure for space adaptation motion sickness. While these terms are useful from the Risk Board perspective, they can cause confusion from the perspective of graph theory. Therefore, we will discuss these as nodes within a specific graph. By definition, all countermeasures are also contributing factors. However, countermeasures receive a data label that distinguishes them as countermeasures and a sub-label that notes which type of countermeasure. A countermeasure is endogenous (an integral factor) when it is included in the spaceflight vehicle. For example, an ultrasound machine that focuses beams to push a kidney stone requires pre-planning and design so that it has power and fits within the mass budget of the vehicle. Countermeasures are exogenous (an external factor) when they are performed on the ground (either before or after flight). For example, pre-flight exercise regimens and post-flight health surveillance do not impact the vehicle design.

At the HSRB, countermeasures are broken into three categories – monitoring, prevention, and intervention [3].

Prevention Countermeasures

Prevention countermeasures include those countermeasures that prevent a deleterious impact on health or performance. For example, exercise prevents muscle atrophy and bone loss in an altered gravity environment [10]–[12]. It also plays a preventive role in behavioral health by improving mood [5]. When drawing prevention countermeasures in a Human System Risk DAG we assume that an exposure happened, and a prevention is performed to decrease the likelihood that the next factor downstream will occur. This is illustrated in Figure 6A using general terminology. Figure 6B shows an example case using *Landing Loads* and HSRB terminology. *Landing Loads* are an exposure during landing and may lead to *Traumatic Injury*. In this case, *Occupant Protection Measures* are prevention countermeasures – intended to prevent *Traumatic Injury* from occurring. The arrow from the prevention countermeasure is drawn to the node that should have a decreased likelihood of occurring if that countermeasure is implemented.

Intervention Countermeasures

Intervention countermeasures are those that are implemented in response to a recognized problem that has occurred for human health and performance. In the case of medical care this is often called 'treatment,' whereas in the case of vehicle systems failures this is referred to as 'repair'. Both are interventions performed to mitigate risk from a change to the human system state. Figure 6B also shows an *Intervention Countermeasure*, in this case *Medical Treatment*. Once the *Traumatic Injury* from *Landing Loads* has occurred, *Medical Treatment* is used to decrease the *Outcome* of that injury. In this case, that is a decrement in crew capability either because of a functional impairment (loss of the use of an injured extremity) or because of pain. The arrow from the intervention countermeasure is not drawn to *Traumatic Injury* – that would signal an attempt at prevention. Once an injury has occurred, the arrow is drawn to the next outcome downstream. In terms of probability, preventive countermeasures decrease the probability of the *Traumatic Injury*, intervention countermeasures decrease the probability of the consequences of *Traumatic Injury*, assuming it has already occurred.

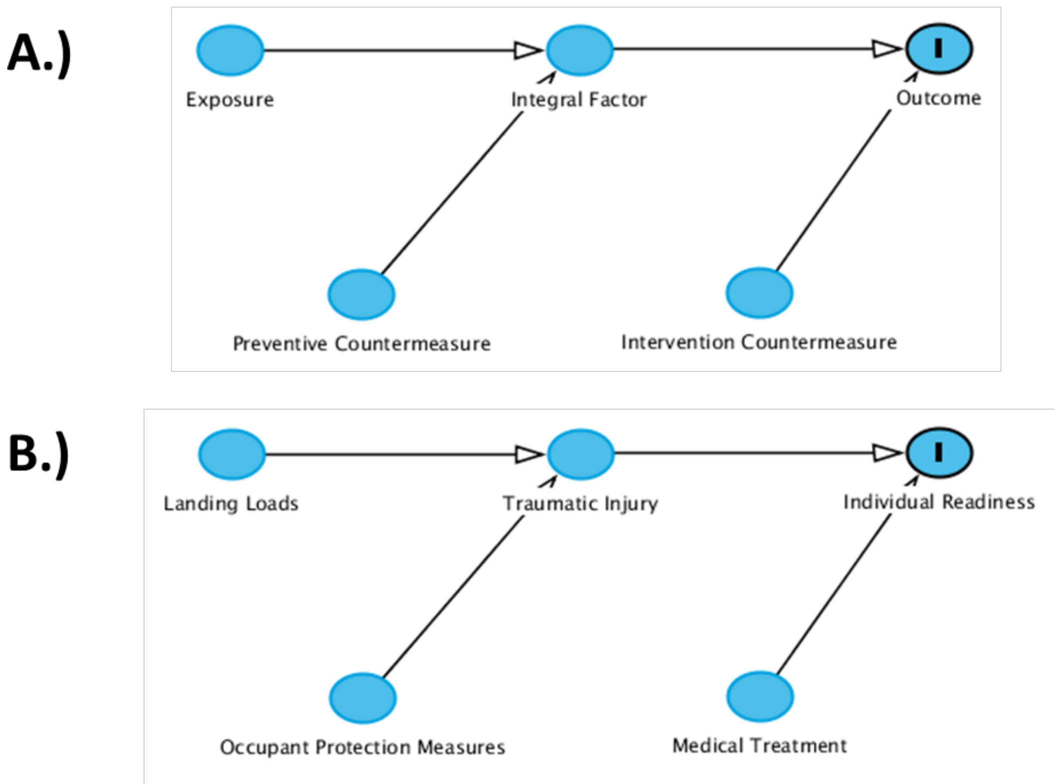

Figure 6: Illustration showing A. how to diagram prevention and intervention countermeasures and B. an example case using traumatic injury from landing loads as an example.

Monitoring Countermeasures

Monitoring is included as a countermeasure due to the recognition that many decisions to implement preventive or intervention actions are based on monitoring of the state of the human system or the crew-vehicle system. Monitoring sensors and data management must be designed into the system through effective human system integration processes [13], [14]. If they are not considered as part of vehicle systems or Crew Health and Performance (CHP) system design effectively, then the human becomes the monitor. This means that the human will detect potential environmental injury only when symptoms develop. In many DRM categories, this suggests that deleterious changes in the human system state are not able to be detected early enough to prevent those injuries at some level. These situations must be explained from a risk perspective to mission planners so that they can make risk-informed decisions about what to include in vehicle systems. The inclusion of monitoring hardware and software in the vehicle depends on decisions made early in the systems engineering life-cycle [1]. Therefore, monitoring capabilities are included by both the HSRB and in DAG creation as a necessary predecessor to various intervention countermeasures.

Monitoring is a unique case of countermeasure in human spaceflight. Often people overlook it because it does not meet the intuitive sense of a countermeasure. However, as a CHP System is designed for a vehicle, significant challenges are faced in the systems engineering processes that ultimately produce the vehicle. A CHP System can only mitigate risk up to the level it was designed. This means that if a countermeasure was not included in the design, then the capability is not present in the vehicle and risk is not mitigated. Monitoring capability and diagnostic capability require sensors, data systems, and software that must be considered in the vehicle and system design phases. Monitoring capability relies on sensors to track parameters. Parameters can be human-oriented like heartrate and blood pressure, or they can be environment oriented like air quality, carbon dioxide levels, water quality, acceleration loads, etc. In all these cases, obtaining quantitative information about the state of exposures or the state of the human system requires insight into the causal role they play in risk reduction.

Figure 7 shows guidance for presentation of monitoring or similar capabilities in the HSRB DAGs. In general, any specific monitoring capability included in the diagram enables detection of something that has occurred. Figure 7A shows this for a general case to illustrate the point and Figure 7B shows this for a specific example from spaceflight. In Figure 7A, the *Integral Factor* is the occurrence of interest. Both the *Monitoring Capability* node and the *Integral Factor* node that warrants monitoring point to the *Detection of Integral Factor* node because both modify the probability that the *Integral Factor* is detected. Once the *Integral Factor* is detected, then a *Treatment/Intervention* can be brought to bear to affect the likelihood that the *Mission Level Outcome* will occur. Importantly, removing the monitoring capability decreases the likelihood of detection and decreases the likelihood of treatment/intervention. The simplified example below now applies specific HSRB language.

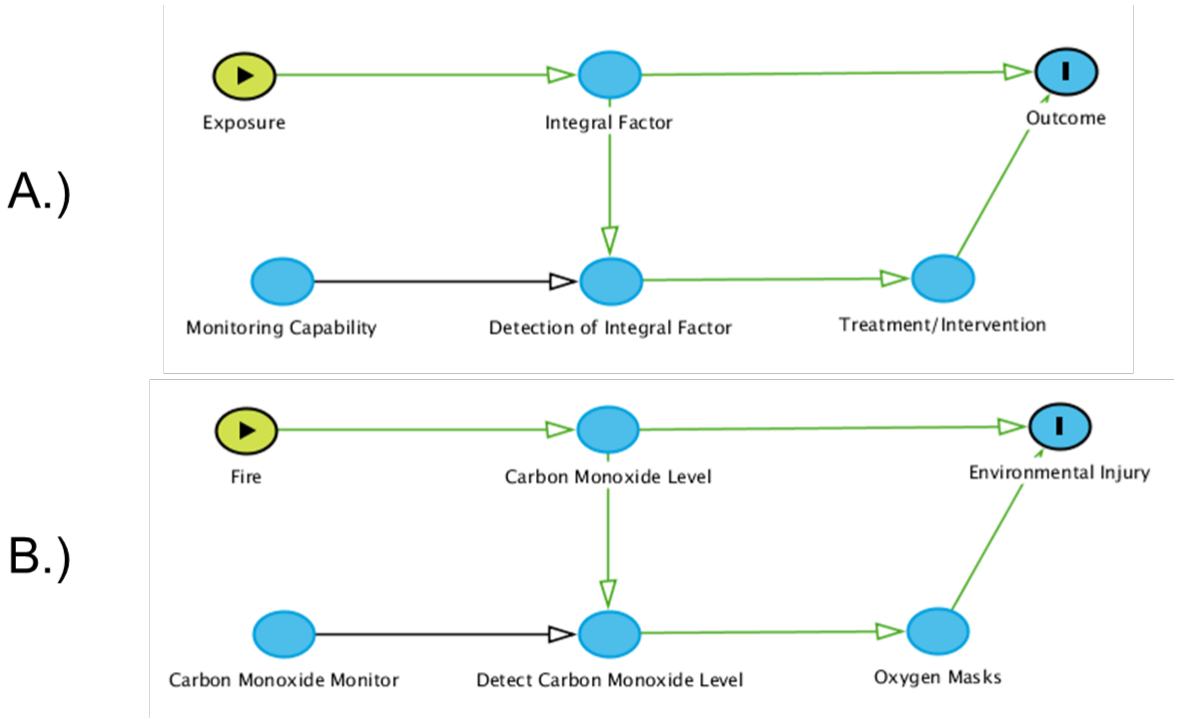

Figure 7: Approaches to visualizing carbon monoxide exposure as a risk. A.) shows the general approach to graphing monitoring capability and B.) shows a spaceflight specific example for carbon monoxide exposure.

Figure 7B shows a simplified example. Consider carbon monoxide (CO) exposure for the astronauts which can occur if there is a fire. While fire can lead to other issues such as smoke and burns, this example focuses solely on CO exposure because of the value in demonstrating how to graph monitoring capability.

Composition Example

To show how these can be included and visualized in a systematic way, consider the following example. If a crewmember has insomnia, it increases the probability of a sleep disturbance, which in turn increases the probability of fatigue, which affects crew capability. Figure 8 shows a simplified notional example of this causal flow associated with two countermeasures. *Zolpidem* is a sleep medication used to prevent *Sleep Disturbances* (by treating *Insomnia)* and has been used in the space program [15].

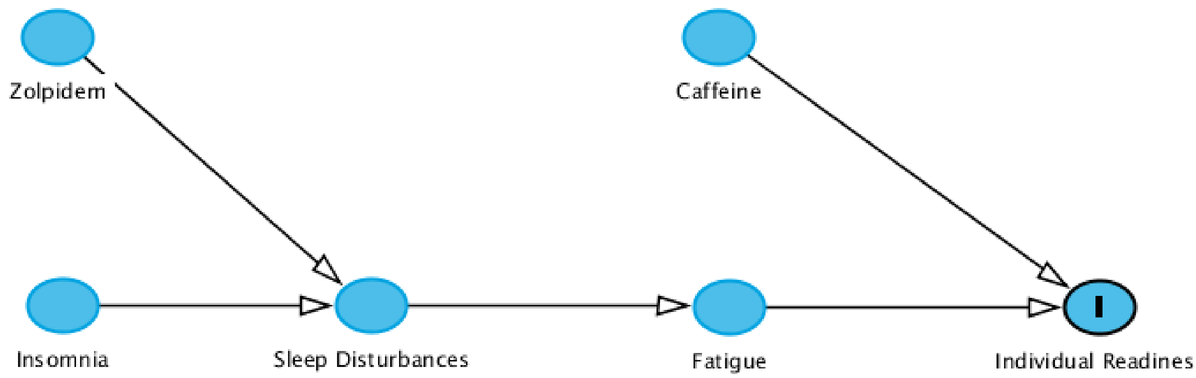

Figure 8: Example notional DAG for sleep deficits showing depiction of treatment/intervention countermeasures.

Here the arrow is drawn to a *Sleep Disturbances* node because the treatment for *Insomnia* is also intended to decrease the probability of the *Sleep Disturbances*, i.e., prevention. In this sense, a treatment for one thing may be considered a prevention for the downstream nodes. The medication may or may not be effective at preventing *Sleep Disturbances*, i.e., there is some probability of success that is less than 1.0. If it does not completely work, the astronaut has some probability of experiencing *Fatigue* as a consequence. How much *Fatigue* the astronaut feels may influence the decision to take *Caffeine* as an intervention (treatment) countermeasure. *Caffeine* as an intervention seeks to 'make up' for some of the impact of poor sleep by imparting additional energy to the crewmember. As anyone who has been addicted to coffee knows, too much caffeine can affect the ability to focus and concentrate. Zolpidem dosing is also important as too much can impart sleep inertia and affect how well crew can perform cognitively [16]. Therefore, both *Zolpidem* and *Caffeine* can become contributing factors for changes in *Individual Readiness* elsewhere in the larger picture of risk. While this is a theoretical example intended only to convey standard diagramming approaches and does not represent the entire Sleep Risk DAG, it illustrates the importance of linking together the specific Risk DAGs into a larger risk network. Many of the countermeasures used to decrease risk in one domain can lead to elevations of risk in another domain. The rest of this section shows how the sleep deficit and the CO example (Figure 7B) tie together two separate risks at the level of *Individual Readiness* and is intended to illustrate the power of combining separate Risk DAGs together to approach the risk hierarchy and risk dependency challenges outlined in the introduction.

Consider first the example shown in Figure 7B. If a *Fire* occurs in a closed environment such as a spacecraft, *Carbon Monoxide Levels* in the atmosphere may increase. This leads to headaches, a common symptom of CO toxicity and one of many possible *Environmental*

Injuries. In this case, the node *Carbon Monoxide Detector* is shown with an arrow pointing into a node marked *Detect Carbon Monoxide Level*. The node *Carbon Monoxide Level* also points into the node for *Detect Carbon Monoxide Level*. Logically, the ability to detect an occurrence has two prerequisites – the occurrence itself and the monitoring capability intended to measure that occurrence. Here the occurrence is a change in the *Carbon Monoxide Level* in the atmosphere of the vehicle. The *Carbon Monoxide Detector* includes sensing equipment that can measure the *Carbon Monoxide Level* and display or log those data. To some readers this may seem obvious, but the decision to use available mass, power, and volume for sensing equipment for CO levels is a system trade that is considered by systems engineers in the design phase of the spacecraft. There is a very real possibility that a risk-informed decision may be made to use that allotment of mass, power, and volume for something else that is deemed more important. The potential value of including sensing equipment is demonstrated not only by the *Detect Carbon Monoxide Level* node shown, but also by the follow-on decision that it informs. In this case *Oxygen Masks* may be donned by the crew at a pre-determined CO level. These are an intervention or treatment countermeasure and are shown with an arrow pointing into *Environmental Injury*. *Environmental Injury* includes medical conditions and symptoms such as headache and altered mental status that can occur with CO toxicity. (Note that these medical conditions and symptoms are 'nested' under the *Environmental Injury* node. This is discussed further below.) *Environmental Injury* from CO toxicity contributes to the likelihood that *Individual Readiness* for an exposed crewmember will be affected. Note that in the figures above, the outcome shown is an intermediate outcome, i.e. one that is not a Mission Level Outcome. These are used when we are 'zooming in' to look at a subset area of a larger risk DAG. In the case of the sleep issues shown in Figure 8, or CO exposure in Figure 7B, intermediate outcomes such as mild *Sleep Disturbances* or minor *Environmental Injury* may not be of high concern to program managers or flight directors because they may not significantly impact *Individual Readiness*. When *Fatigue* or *Environmental Injury* become significant enough to affect *Individual Readiness*, and possibly affect Mission Level Outcomes, then they begin to attract the attention of flight directors and program managers. Figure 9 shows the combination of these two theoretical examples into a single DAG showing how both pathways can affect *Individual Readiness*.

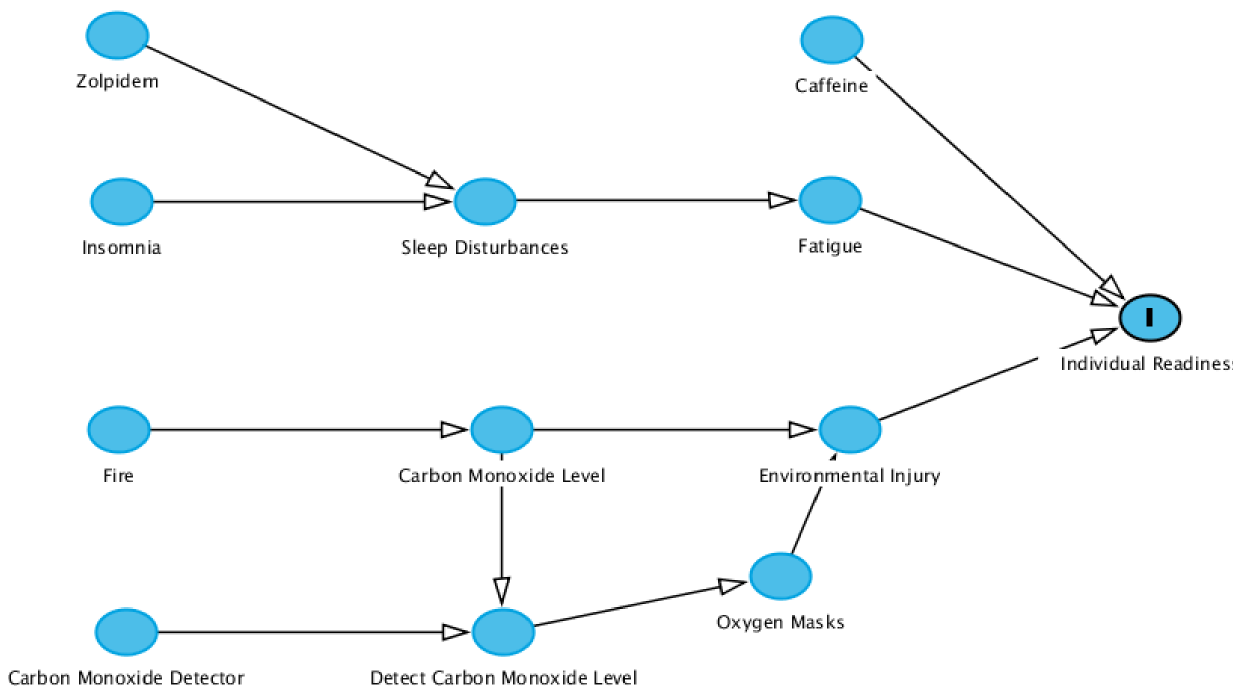

Figure 9: Combining examples above, it is apparent that there are multiple paths to affecting Individual Readiness.

When combined and structured in this way, DAGs provide the basis for a Bayesian network where joint probabilities of these types of effects can be combined to quantitatively approach downstream effects on outcomes of interest. The DAG in Figure 9 shows an oversimplified representation that highlights some features of Detailed DAGs. First, it introduces *Individual Readiness* as a local outcome. This completes a causal chain that can be quantified from the probability of having a *Fire*, the probability of experiencing a high *Carbon Monoxide Level* (given there has been a fire), and the probability of experiencing an *Environmental Injury* given a particular *Carbon Monoxide Level*. All of these factors contribute to *Individual Readiness* by way of this causal chain. Likewise, *Individual Readiness* is also affected directly by the probability of *Fatigue* and/or *Caffeine* use and further upstream by the chain of conditional probabilities involving *Sleep Disturbances*, *Zolpidem* use, and *Insomnia*. Both of these paths could be occurring at the same time in a given mission. This example is drawn from real-life example of the fire that occurred on the MIR space station [17].

Not shown here in these simplified examples are other well-known connections that are captured either in the larger DAG or in other Risk DAGs. For example, there are other causes of environmental injuries such as headaches caused by carbon dioxide (CO_2) [18] that can be difficult to distinguish from CO without appropriate environmental monitoring equipment. CO_2 requires different atmospheric monitoring equipment to determine inhaled levels either in the vehicle or spacesuit. A host of other causes and effects could be drawn into and around each of these nodes shown in these examples that would affect *Individual Readiness*. This demonstrates the point of creating DAGs – **to enable the addition or subtraction of relevant pieces of the story as evidence becomes available to support them or as communications needs arise**.

As potential new nodes and edges are identified, it is important to consider if the magnitude of those impacts is likely to become problematic at the mission level. If not, they can be safely ignored and left out of the DAG. However, if the answer to that question is 'we don't know', then the exercise has identified a knowledge gap that should be considered for further characterization. As DAGs become more complicated and difficult to analyze, the techniques of network analysis can be brought to bear to assess the influence of nodes or groups of nodes on the larger risk network.

Harmonized and Neutral Terminology

There are two guidelines for standardizing DAGs that are important to note at this point:

- Two different teams of experts may use different terminology to refer to the same concept. Here we force them both to use a common term so that we can effectively tie their DAGs together at a point that is relevant to both. Imposing this restriction on terminology is referred to as harmonization of terminology and is key to ensuring that software used to capture and analyze DAGs is functional.
- In the example case above, the node name *Carbon Monoxide Level* is used as opposed to 'Carbon Monoxide Increase'. *Carbon Monoxide Level* conceptualizes CO as a continuous variable, capable of taking on any value greater than or equal to zero, whereas 'Carbon Monoxide Increase' is a binary variable, scored as 0 when the CO level is less than a chosen threshold value and scored as 1 when the CO exceeds this threshold. The continuous form is chosen to maintain neutrality about possible effects from the fire. CO could conceivably go up, stay the same (clean burning fire), or decrease (i.e., through atmospheric removal). For the purposes of HSRB modeled DAGs, this neutrality of terminology is required across all risks whenever possible.

Individual DAGs may be combined into a larger DAG because some nodes appear in multiple risk DAGs. These nodes provide points of connection between the DAGs. However, for software to achieve this composition, the common nodes in all risk group DAGs must be named using harmonized terminology, as described above. An examination of terms used before the implementation of DAGs among risk groups to describe the Hazard “altered gravity” revealed 10 risk groups using six different variants of this term (e.g. microgravity, gravitational changes, weightlessness, etc.). The different terminology used by different groups is a source of confusion when discussing among the Human Health and Performance community at NASA and would make combining risk networks nonsensical if not resolved.

Harmonizing each of the terms used to “altered gravity” allows us to visually show how all of the risk groups are identifying the same factor in their risk networks and creates a connection between these risks via this common cause. The choice of altered gravity also allows for the continuous representation of gravity above or below Earth normal, 1g. This harmonization and neutrality of language are intentional processes as Human System Risk DAGs are built to ensure that different groups of experts use the same name for the same concepts.

As DAGs are built or modified, the HSRB ensures the use of a proscribed set of terminology to prevent terminology drift. Creators of the DAGs (i.e., risk custodial teams) can propose new terminology if the concept they seek to include is not already represented within the harmonized list. If a risk custodian team proposes creation of a new node term, the following steps are required:

1. The new node term and rationale must be provided to the HSRB office for review. This is typically done as part of a formal risk update and accompanied by evidence that supports the inclusion of the new node.
2. The HSRB office will review the request and the rationale as well as the current list of approved terminology. If the new node terminology is significantly different from all existing terms and is still at useful level for the Board’s shared mental model, the term will be discussed as part of the risk update at the Board.
3. If approved, the configuration managed set of terminology will be updated to include the new node and its definition. If the new node name is reasonably close to an existing node name, then the HSRB office will request a discussion on whether the creation of a new term is appropriate or use of the existing term should be considered.

In the end, the HSRB chair has the final determination regarding the creation of new nodes outside the accepted terminology list.

Nesting of Nodes

Narrative DAGs are specifically designed to convey high-level information. This can result in oversimplification of some of the details beneath high-level nodes. Nodal nesting is used as a labeling tag that highlights when a single node encompasses two or more sub-nodes. The term

‘Category Node’ is used to identify the highest-level node used in DAG depiction. Sub-nodes are assigned the data label of the Category node and any other sub-nodes that are hierarchically above them. Figure 10 shows an illustration of nodal nesting using the Category node titled ‘*Individual Factors’*. In this view there are two immediate sub-nodes, ‘*Modifiable Factors*’ and ‘*Non-Modifiable Factors’*. Beneath each of these is another layer of sub-nodes shown in list format that are not the complete set but are good examples.

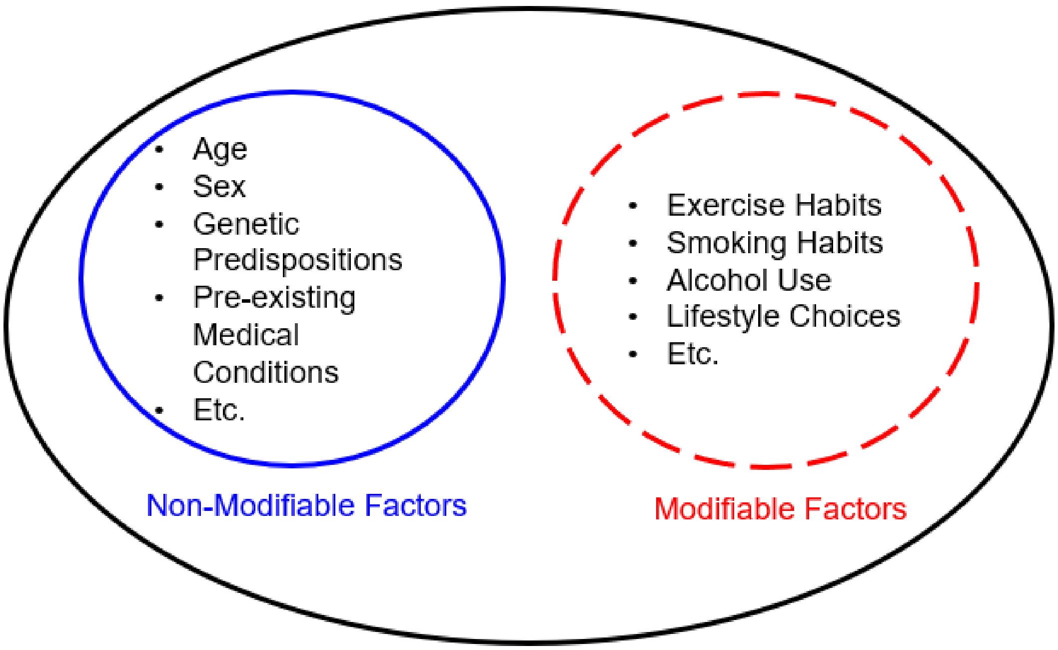

Individual Factors Category Node

Figure 10: The Category node ‘Individual Factors’ houses two direct sub-nodes – Modifiable Factors and Non-Modifiable Factors, which in turn house several sub-nodes at the level below shown here in bulleted list format.

This illustration shows why the term ‘nesting’ is used as the sub-nodes appear to be nested within the higher-level nodes. In many cases, showing the Category node in a DAG can help to simplify the story visually. These are generally used for Narrative DAGs to enable meaningful communication among stakeholders. For Detailed DAGs, the Category node can be expanded and replaced with the sub-nodes that are relevant to the specific risk as appropriate. Each arrow that enters and exits a Category node must also connect to at least one nested sub-node. Figure 11 shows an example of a subset of the Bone Fracture Risk DAG information centered on *Bone Remodeling*. (In this case only a subset of the larger DAG is shown to illustrate the nesting concept.) The high-level DAG (A.) shows the high-level approach which includes a *Bone Remodeling* node as an integral node that leads to Bone Density and Bone Structure in spaceflight. A more detailed breakdown of *Bone Remodeling* into *Bone Formation* and *Bone*

Resorption is shown in the Detailed DAG (B.). These two nodes are captured as sub-nodes of *Bone Remodeling* in labeling to enable the level of visualization that is most appropriate to communication needs at a given time.

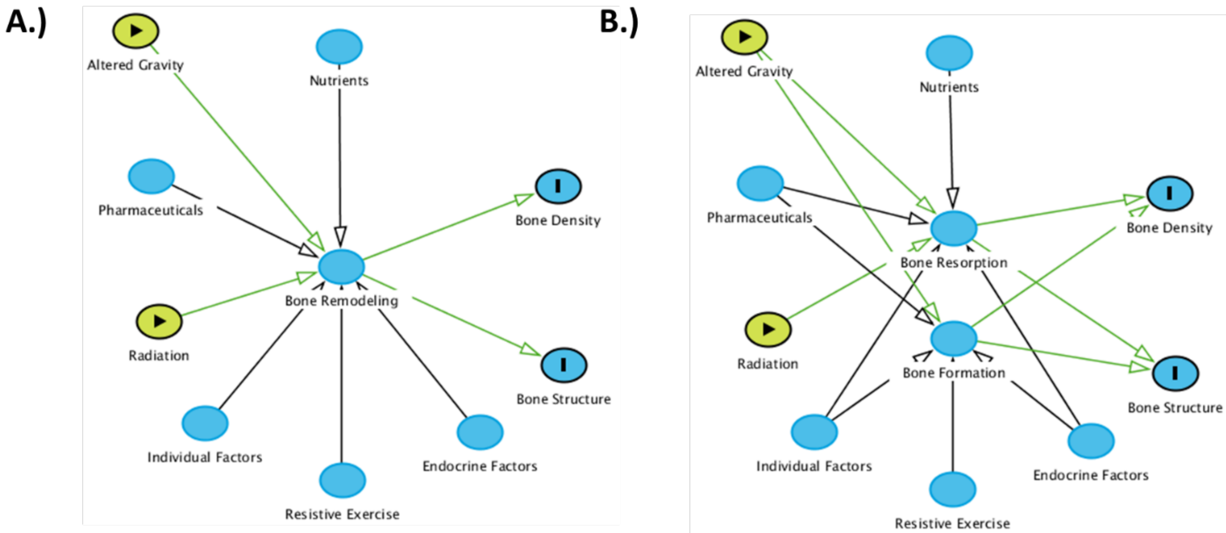

Figure 11: Nodal Nesting example – A. High level sub-DAG showing multiple contributing factors and countermeasures affecting the Bone Remodeling node. B. Shows a more detailed DAG that visually breaks out Bone Resorption and Bone Formation as sub-nodes of Bone Remodeling.

Note that in Figure 11, the *Resistive Exercise* node only points to the *Bone Formation* node in the detailed level graph. This is a key feature of nesting that must be recognized. At the highest levels, differences in connections at sub-nodes are not visually depicted. However, when a more detailed visualization is needed, it is critical to ensure that the arrows are drawn to the appropriate places and not just assumed that all sub-nodes have the same connections as the category node they are nested under. In this case, resistive exercise is known to stimulate osteoblast cells to produce bone but do not appear to have an effect on osteoclast cells responsible for resorbing bone. Ensuring faithful representation of the real-world understanding based on the best available evidence is critical to ensuring that the DAGs are useful.

Time and Feedback Loops

A common criticism of the acyclic requirement is that it glosses over feedback loops that not only exist in the real world but are important to scientific understanding of the causal components of risk. The response to this is two-fold. First, tracing a path through a DAG represents causal factors influencing effects *over time*. For a path to lead back to its source would violate the coherence of the time sequence inherent in the DAG. This is true for even the simplest feedback loop: two nodes with directed arrows pointing at each other. For this representation to be accurate, each node would be the cause of each other at *the same time*. In reality, feedback loops are never simultaneous; they occur in a sequence, even if the timespan of that sequence is quite small. Secondarily, if a cycle is truly needed for purposes of scientific understanding, it can be represented on the Detailed DAG using time-indexed variables to visually represent the concepts involved as nodes that occur repeatedly over time. An example of the time-indexed representation of feedback loops is presented in Figure 12.

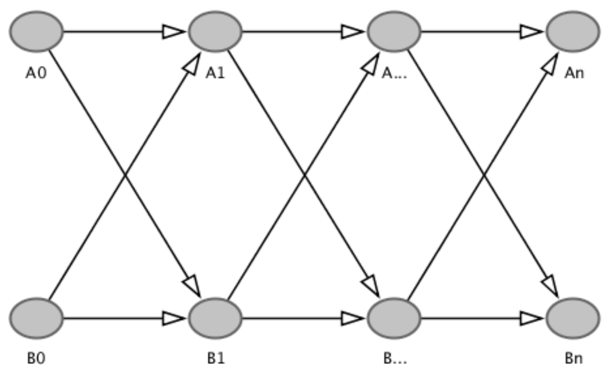

Figure 12: Example of a time-indexing approach to DAG construction that can be implemented in Detailed level DAGs.

Here we see a series of variables representing two distinct factors, A and B. Both have an initial value at time zero, denoted by A_0 and B_0. The feedback loop between these factors is represented at successive time points {t=1, …, n} by the arrows that point from A_t to B_{t+1} and B_t to A_{t+1} at each time index. The placeholders $A_{...}$ and $B_{...}$ represent the fact that we may show as many “time slices” as we wish, but the pattern (i.e., the feedback loop) remains the same across time. Thus, while deciding the proper sequence for communication purposes in a Narrative DAG may be challenging, feedback loops can be represented with high fidelity in Detailed DAGs. Because this level of detail can easily become confusing, these types of time loops are excluded from the high-level DAGs in favor of representations that facilitate easier communication.

Yet another method of representing time-dependent effects on DAGs is to add a ‘duration of effect’ node separate from any intermediate factor nodes to which it relates. For example, in the Extravehicular Activity (EVA) Risk a node named *EVA Duration* is included as an intermediate factor that affects the *Workload* and *Cognitive Function* nodes. This simply represents that there is a known time dependence on those nodes but does not go to the level of detail of creating a time-indexed approach to those relationships. Other nodes such as *Effective Mission Duration* and *Effective Exposure Duration* are used in appropriate places in other risks when a known or strongly suspected causal relationship with time is drawn.

Discussion of Key Nodes

This section is intended to provide insight for several key nodes that appear in multiple risk DAGs. These concepts are unique to Human System Risk management and are intended to show why specific terminology was chosen to apply across all risks.

Risk Names – Within a given Human System Risk DAG the links from other recognized Human System Risks are shown as light grey nodes. These nodes only identify a connection at the Risk level and do not carry information about what specific connections exist in detail. These are shown in Narrative DAGs to improve high level insight into known or suspected risk-risk connections.

Individual Readiness – *Individual Readiness* is the knowledge, skills and abilities of a given individual as well as the functional capacity of that individual. It denotes the complete readiness of any specific crew member to perform mission tasks assigned to them.

Crew Capability - *Crew Capability* is the representation of the readiness of the entire crew to perform tasks. The edge that connects *Crew Capability* to *Task Performance* shows the intersection of two key points. 1. the level of crew deterioration from the spaceflight Hazards over the effective exposure duration and 2. how the system was designed. This is the challenge of Human System Integration: the human must be considered in the functional state they will exist during the mission, and the design of the vehicles/suits/systems either accommodates that or fails to accommodate it. Successful *Task Performance* depends on these multiple inputs.

Figure 1 at the beginning of this document shows an oversimplified view of a key interaction in risk to astronaut performance and is reviewed here for relevance to the *Individual Readiness* and *Crew Capability* node. Most of the Human System Risks degrade *Individual Readiness*. Continued exposure to the Hazards of the spaceflight environment result in physiologic changes that, if left unchecked, functionally impair an astronaut in some way. That functional impairment results in a decreased *Individual Readiness* for any specific crew member affected, but the probability of successful *Task Performance* is not only dependent on individual crewmembers. The *Team Functionality* of the larger crew helps determine some of the resilience of the total crew and also contributes to *Crew Capability*. However, functional impairment of individuals or the team is not the only way that *Task Performance* can be threatened. The effects of system design limitations are captured in the *Human System Integration Architecture Risk* node. Inadequate system design can result in failures of *Task Performance* even when crew are 100% functional. For example, if there are insufficient spare parts to repair the spacecraft, then the repair tasks can fail. If the vehicle systems are designed in such a way that crew cannot access what they need for maintenance/repair issues, then the maintenance/repair tasks will fail.

Human System Integration is the process by which these types of considerations are taken into account throughout the systems engineering life cycle of a vehicle.

Individual Factors - The *Individual Factors* node is intended to comprise the influence that the individual crewmembers will have on biologic variability affecting risk. The *Individual Factors* node as a category node, houses subsets of individual factors. The first layer below *Individual Factors* is broken into two nodes - *Modifiable Factors* and *Non-Modifiable Factors*. *Non-Modifiable Factors* include traditionally unchangeable factors including age, sex, genetic predispositions, etc. *Modifiable Factors* include factors that can traditionally be changed by the individual including smoking status and healthy lifestyle decisions (healthy eating, exercise, etc.). Both of these sub-nodes (*Modifiable Factors* and *Non-Modifiable Factors*) can be further broken down into multiple nodes of those factors listed, depending on risk communication needs. Although *Individual Factors* are affected by *Astronaut Selection* decisions, there are many factors that come along with an individual that are not part of the selection process. Biologic variability in the response of the crew to the spaceflight environment is heavily dependent on the individual factors included in the crew.

Decision Nodes - There are two decision nodes pre-identified within the risk set. These are in the Medical and HSIA Risks. These represent decisions to either treat or to intervene that are at times are undertaken either completely or in part by experts in mission control under real-time guidance. As distance from Earth increases, real-time guidance options diminish until the only option for support from mission control is a store-and-forward format for messages and data. Real-time communication with mission control reduces cognitive load on the crew and reduces the need for on-board knowledge, skills and abilities regarding maintenance and repair of vehicle systems, as well medical knowledge for the care of fellow crewmembers. High levels of expertise are required to detect, diagnose and intervene appropriately in anomalies in highly complex systems. The risk of providing an inappropriate treatment/intervention can range from neutral to life-threatening. As it is impossible to predict which specific medical problems or vehicle system malfunctions will occur before a mission begins, it is critical to provide the crew access to the expertise needed to address a wide range of possible problems that can occur during a mission. As distance from Earth increases and real-time access to the expertise of mission control is lost, the expertise contained within the crew or the vehicle systems must be sufficient to deal with any emergency issues that cannot be stabilized or resolved within the time it takes for messages to go between the vehicle and Mission Control back on Earth. The ability to make a decision that does not cause harm is a critical part of the causal flow of risk in spaceflight. For these reasons, decision nodes are specifically included in the Medical and HSIA Risk DAGs

- **Treatment Decision** (Medical Risk) - The decision to provide or withhold a medical treatment to a crew member is generally made in real-time with flight surgeons and biomedical engineers in mission control. This is also true of behavioral health conditions with psychiatrists/psychologists at mission control. In the case of clinically significant behavioral health conditions, these are grouped under the Medical Risk. Understanding of the available options for treatment, the possible side effects, expected outcomes, monitoring needs, prognosis, identification of when treatment is failing, and more are considerations in making a decision to treat. This is because it is not only possible, but likely to cause more harm than good by making ill-informed decisions to treat.
- **Intervention Decision** (HSIA Risk) - The decisions surrounding identifying anomalies with the vehicle and systems on-board are generally informed by multiple experts at mission control in all LEO operational scenarios. Determination of whether alarms are 'false', due to sensor issues, or real are almost exclusively handled by analysis in mission control. The experts in mission control provide direction and guidance to crew members in flight to remedy identified anomalies and act to prevent ill-advised crew interventions from causing damage to the vehicle in the course of diagnosis and resolution.

Effective Exposure Duration - When considering the probability of a poor Mission Level Outcome, many of the risks are dependent on time of exposure to a Hazard (or other factor). Physiologic challenges to the body worsen as time in the spaceflight environment increases. While some of these can be wholly or partially mitigated with countermeasures, the extent to which those countermeasures are successful are often also dependent on time. Multiple risks are explicitly dependent on time. While it is possible to introduce time as a series of nodes that are denoted by a time variable, this would complicate visual representation, so that it would not serve communication purposes of the DAGs. Where specific nodes have a known time-dependence we include a node titled *Effective Exposure Duration* as an exogenous factor that contributes the change in probability at that point in the risk causal chain. This is used to specifically identify which nodes are likely to be heavily time dependent.

Crew Health and Performance (CHP) System – One of the vehicle systems that is the chief concern of the HSRB and the Human Health and Performance community. The CHP System includes the hardware, software, and data and communications support needed to maintain crew health and ensure that crew performance is sufficient to meet Mission Objectives.

As a term CHP System is intended to be mission-agnostic. The CHP System that exists on the International Space Station (ISS) is called the Crew Health Care System (CHeCS). This system is where mass, volume, and power allocations are provided at the system level and therefore the countermeasures identified in the DAGs must compete for inclusion with other system level needs.

Vehicle Design – This node represents the design of the vehicle that results from requirements set early in the systems engineering life cycle. It is also the result of the human systems integration processes that the agency implements throughout that lifecycle. This node imparts both the advantages and disadvantages of the decisions made years before a mission that result in mass and volume allocations to a mission's risk posture. Once a vehicle is launched, no further design changes can be made. Realistically, design freezes happen years before launch. In the simplest terms, if it wasn't designed into the vehicle, then the crew does not get the expected risk reduction benefits.

Suit Design – This is similar to vehicle design but framed for spacesuit design specifically. Several risks include spacesuit-specific dependencies including the EVA, Decompression Sickness (DCS), Hypoxia and other Risks. Like the vehicle, the design decisions for space suits are made long before a mission and will set the risk posture for the mission. Also like the vehicle, if it wasn't included in the design of the space suit, then the crew will not have expected risk reduction benefits from it.

Crowdsourcing and Configuration Management

There is an inevitable question that arises from any attempt to 'map' out the factors that are important to a Human System Risk: Who decides what is important? This is not a trivial question as those factors come from a variety of disparate fields including medicine, life sciences, pharmacology, food sciences, behavioral health, exercise physiology, engineering, human factors, and many more. This broad set of potential domains that contribute to Human System Risk posture requires that a broad community evaluate and provide input to DAG creation. This is critical to ensuring that the initial Narrative DAG reflects the needs of stakeholders involved in assessing and mitigating risk across all domains. The HSRB provides a forum where the larger Human Health and Performance community involved in risk assessment provides input and reviews each of the risk assessments on a regular basis [3]. Since January 2020 this includes a review of an official DAG for each risk. A risk custodian team is assigned for each Human System Risk including SMEs from operations, research, and epidemiology at NASA. This team, along with representatives from the risk management team and other SMEs as needed, confer to create the basic DAG. The HSRB includes 10 standing members representing expertise and organizations responsible for supporting human health and performance at NASA and typical meetings include between 100-200 members of the Human Health and Performance community at NASA. DAGs are walked through and discussed as part of risk updates in a continuing risk management process. The HSRB configuration manages each Risk including a formal comment period and comment resolution process that is tracked [3]. Any disagreements are discussed at, and resolved by, the Board or the Board Chair based on evidence standards used by the Board [3]. In this way, the Narrative DAGs receive crowd-sourced feedback that informs and ensures that these causal diagrams adhere to the most current evidence-based knowledge available regarding human health and performance in spaceflight. Detailed DAGs are produced by the SMEs that follow the structure of the Narrative DAG but may include more detail within nested nodes than are shown in the Narrative DAGs. These are not reviewed by the board as the level of detail generated is often beyond the expertise of the general board.

Full DAG Example

It is instructive to consider a full DAG that results from application of the guidance above. We will consider the Acoustics/Hearing Loss Risk DAG here because it is a relatively small DAG that can be explained in the context of the above guidance. Using Dagitty software [9] to generate the DAGs results in node and edge representations that require reference to the legend shown in Figure 5. Figure 13 shows the example of the Acoustics Risk DAG as baselined by the HSRB at NASA. On the left of the graph are three of the spaceflight Hazards – Altered Gravity, *Hostile Closed Environment*, and *Distance from Earth*. These are denoted by a green node with an arrow which indicates that these are the exposures of interest. On the right side of the graph are the *Mission Level Outcomes – Task Performance*, Loss of Mission Objectives, and *Long Term Health Conditions* denoted by the blue nodes with vertical lines. Green lines here denote the causal pathways from exposures to outcomes (i.e., *Hazards* to *Mission Level Outcomes*), while black lines denote pathways from external factors to integral nodes in the network.

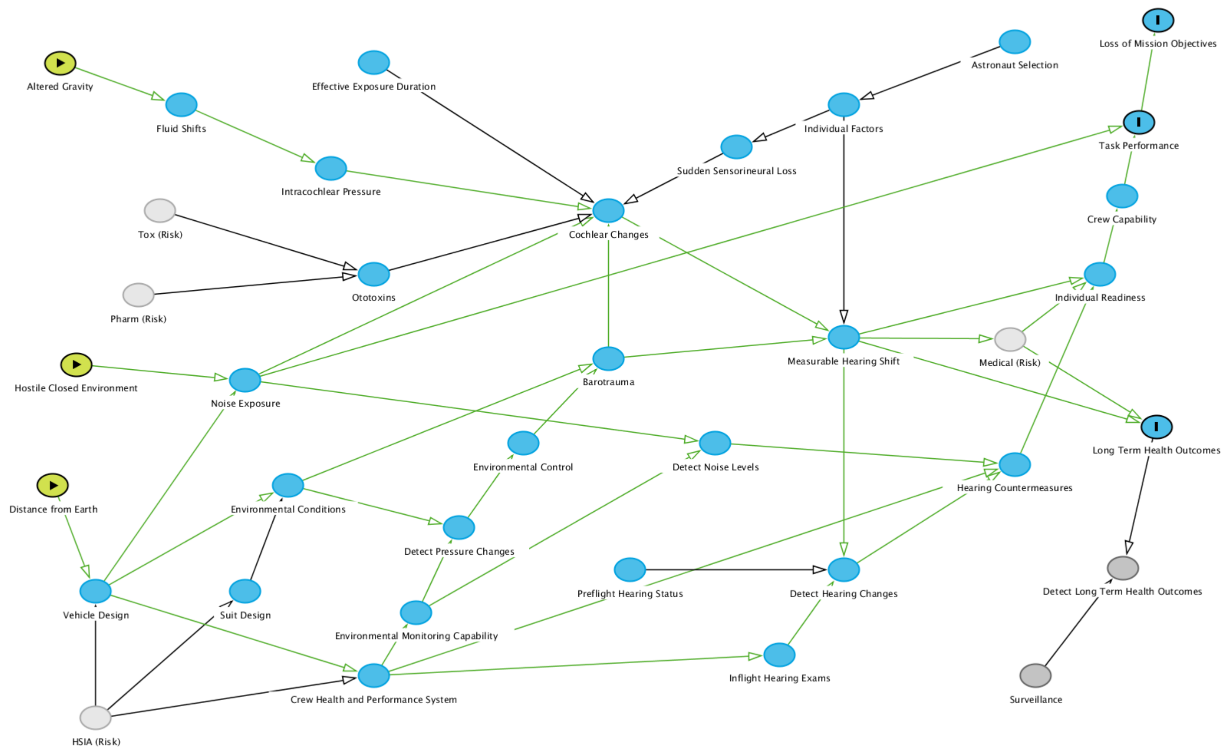

Figure 13: The DAG for the Acoustics Risk. Hazards are denoted as exposures while Mission Level Outcomes are denoted as outcomes. CHP – Crew Health and Performance, Tox – Toxic Exposure Risk, HSIA – Human System Integration Architecture Risk.

The causal flow of this particular risk to Mission Level Outcomes follows several Integral pathways and can be described in terms of several key nodes (Figure 14). For this walkthrough we will refer to the names of nodes in *italics*. The closed environment of a spacecraft results in a noise environment as well as a possible environment for *Ototoxins*. The noise in a spacecraft like the ISS is typically produced from operating equipment – mostly fans that circulate the atmosphere and pumps that circulate thermal control fluids [19], [20]. Therefore, *Vehicle Design* and the *Hostile Closed Environment* (acoustic in this case) within a spacecraft both point to the *Noise Exposure* that crew experience. The *Noise Exposure* node includes noise intensity level, duration of exposure, and frequency spectrum of exposure. *Noise Exposure* can be severe enough to affect the cochlea of the ear, inducing *Cochlear Changes* that lead to *Measurable Hearing Shifts*. Any time there is a *Measurable Hearing Shift*, there is a possibility that *Individual Readiness* is affected by that hearing shift. Any decrement to *Individual Readiness* among the crew may result in *Crew Capability* being affected (i.e. if one crewmember has more difficulty communicating, hearing alarms, etc.). In the case of the Acoustics Risk, audio communication among the crew and between crew and mission control are critical for many tasks. Impacts to hearing affect the probability that any task that depends on clear communication is negatively affected. In these DAGs, *Individual Readiness* is a high-level term that denotes the possible ranges of functional impairment (from 0% to 100%). A fully capable crewmember has no functional impairments, likewise, a fully capable crew has no impairments. If *Crew Capability* is negatively affected, then there is a chance that they may not be able to perform tasks that are required of them. This is the reason the *Crew Capability* node has an arrow pointing to the *Task Performance* node.

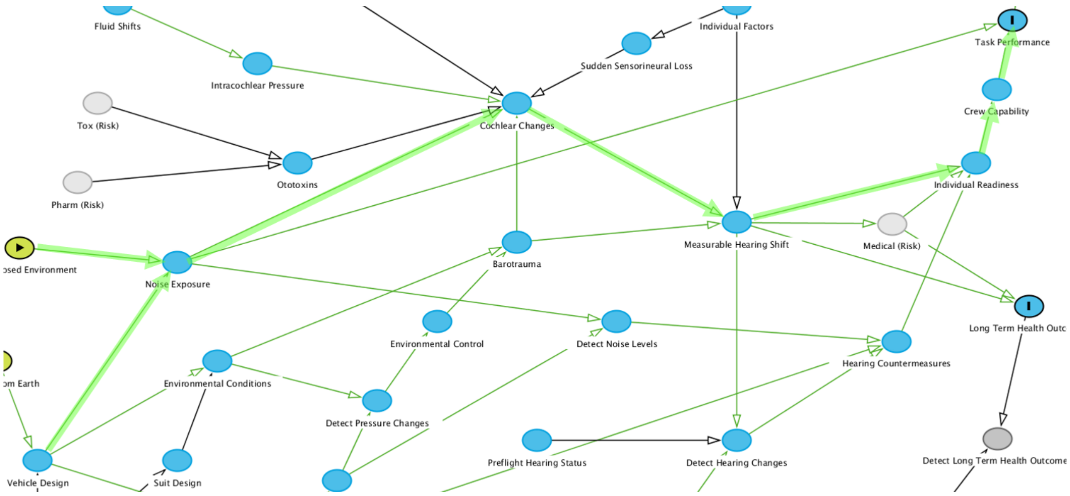

Figure 14: Causal path from the space flight Hazard "Hostile Closed Environment" to "Task Performance"

It is also possible that the *Noise Exposure* in the vehicle is such that it drowns out important sounds and communication even when there are no hearing issues experienced by the crew. This is shown by an arrow that directly connects *Noise Exposure* to *Task Performance (*Figure 15*)*. This is not a theoretical concern, as *Individual Readiness* and *Crew Capability* reductions have been reported in the news when ISS crew members noted that the noise environment on the ISS was such that they could not hear alarms or understand their fellow crewmembers [21]. Note that *Measurable Hearing Shifts* have occurred in astronauts in space though there have been no permanent hearing shifts reported in the literature [22], [23].

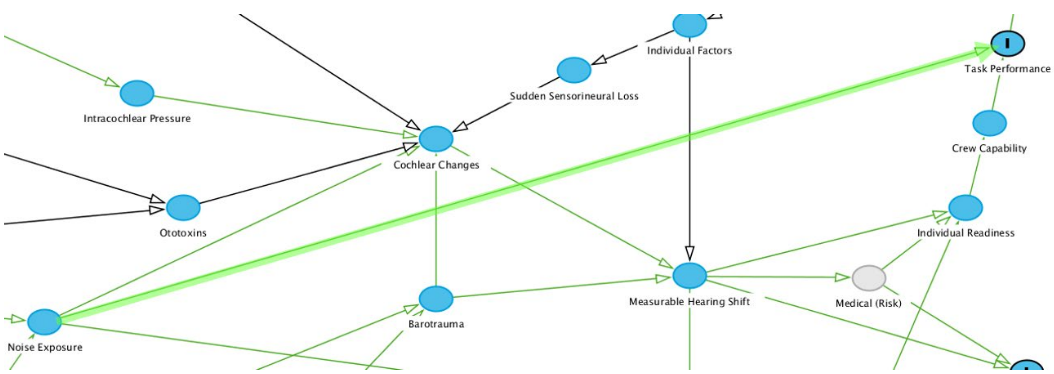

Figure 15: Causal path from Noise Exposure directly to Task Performance

There are several other pathways that can affect *Cochlear Changes* in the ear. Among these is the effect of *Altered Gravity* on the body, which leads to *Fluid Shifts* and may cause *Intracochlear Pressure Changes (*Figure 16*)*. The magnitude of this contribution to *Cochlear Changes*, if present, is currently not known but is likely dependent on the *Effective Exposure Duration*, i.e. how long are the crewmembers in the *Altered Gravity* environment? This pathway represents a falsifiable hypothesis that is based on the concern of SMEs and could potentially represent a research pathway to provide evidence either supporting or refuting this pathway.

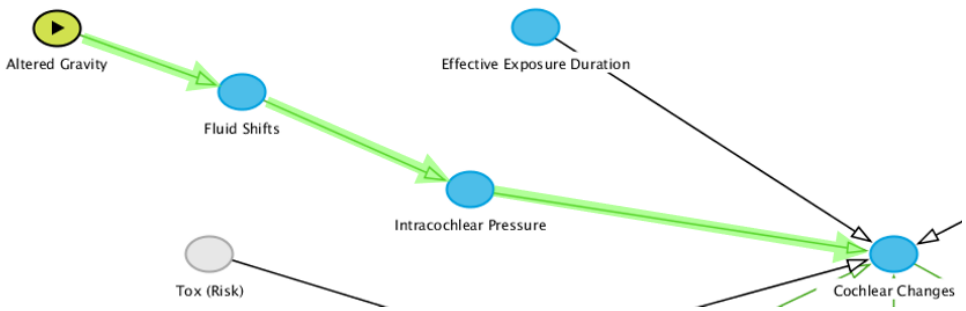

Figure 16: Causal path from the space flight Hazard of Altered Gravity to Cochlear Changes

Another possible path to *Cochlear Changes* in the ear comes from the *Tox (Risk)* and the *Pharm (Risk)* through *Ototoxins (*Figure 17). These are chemicals or medications that can result in damage to the cochlea and affect hearing [24], [25]. Note that this is the first interaction with another Human System Risk we have described in this DAG. For the purposes of communication, it is easier to represent other Human System Risks as a single node (grey in this case) in the DAG. They can be exogenous or endogenous depending on the connections. As with nesting discussed in earlier sections, the grey (Risk) nodes include those sub-nodes within the adjacent risk that have connection to the nodes in the risk under consideration. In the case of the Toxicology Risk, the release of toxic substances into the spacecraft or the ingestion of toxins could lead to ototoxic exposures. The toxicology group monitors the ISS for chemical ototoxins and there are over 20 medications in the medical kit that are known to have potentially ototoxic side effects. Because there are no known medically documented cases of ototoxicity in spaceflight this pathway represents speculation about possible Cochlear Changes based on possible causal pathways. Future flight surgeons and mission planners should be aware of this possibility.

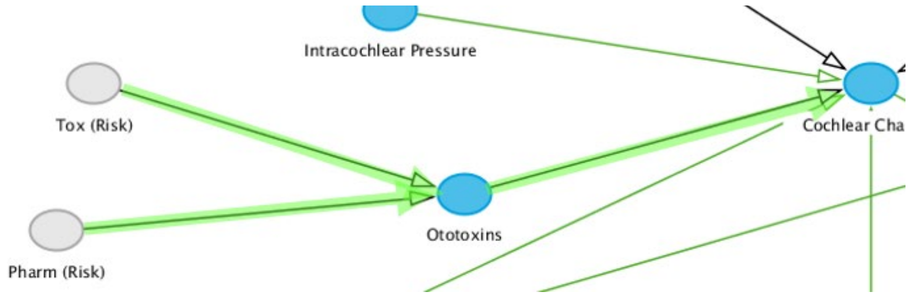

Figure 17: Causal path from the other Human System Risks to Cochlear Changes

Finally, *Sudden Sensorineural Hearing Loss* is a largely idiopathic disorder that affects around 27 of every 100,000 people in the US and can also lead to *Cochlear Changes (*Figure 18*)* [26]. Cases have been found in the active astronaut population. In cases where this is considered as a possible source of *Measurable Hearing Shifts* or hearing loss, medical treatments are available for consideration by flight surgeons [26]. Once a hearing shift is detected, the option to treat is represented by the grey node denoting the *Medical (Risk)*. In the cases where treatment fails there remains the possibility that an astronaut's hearing changes could result in a *Long Term Health Outcome* after their flight and career. NASA has not documented any cases of permanent hearing shifts in astronauts in the literature to date.

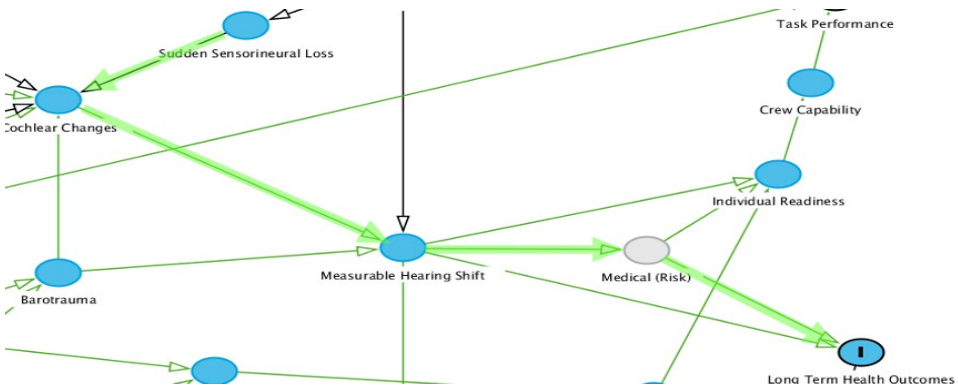

Figure 18: Causal path from Sudden Sensorineural Hearing Loss through Long Term Health Outcomes

There is one other pathway to *Cochlear Changes* and *Measurable Hearing Shifts* that we have not yet discussed, and that is the pathway through *Barotrauma (*Figure 19*)*. *Vehicle Design* and *Suit Design* affect the *Environmental Conditions* that any crewmember experiences during the mission. This includes atmospheric conditions such as atmospheric pressure. When pressure changes occur, there is chance that a crewmember can experience either middle ear barotrauma or inner ear barotrauma due to those changes. Crews experience atmospheric pressure changes when performing EVAs as the ISS environment is kept at 14.7 psia and the suit pressure is much lower. Medical evaluations to assess for fluid behind the eardrums are performed prior to EVAs in an attempt to prevent barotrauma from occurring. Inner ear barotrauma can result in *Cochlear Changes* that can lead to *Measurable Hearing Shifts*. Middle ear barotrauma generally does not affect the cochlea, but through damage to the tympanum or other parts of the middle ear it can lead directly to *Measurable Hearing Shifts*. Here both inner ear barotrauma and middle ear barotrauma are represented by the node titled *Barotrauma*.

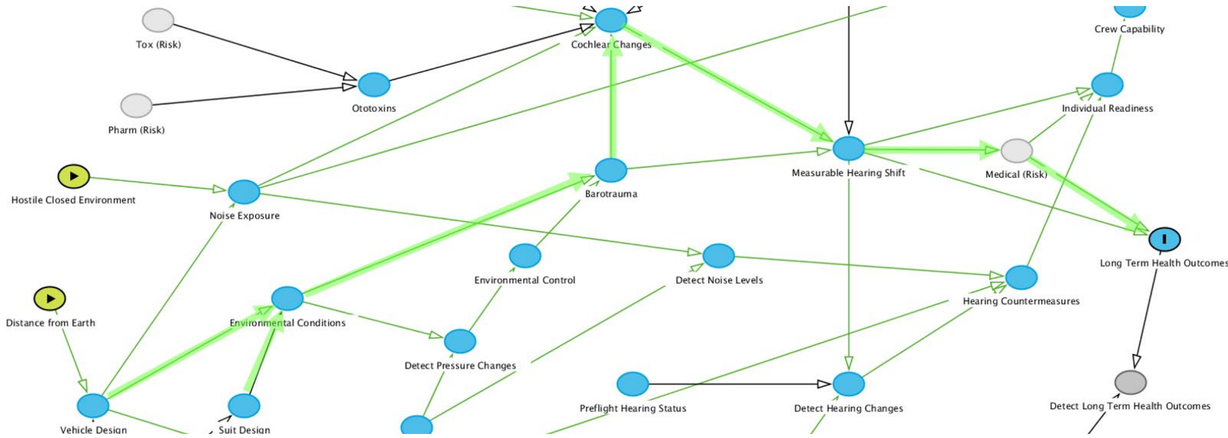

Figure 19: Causal path through Barotrauma to Long Term Health Outcomes

The final set of nodes on the bottom of the DAG image show causal paths involved in monitoring and detection of both noise and hearing (Figure 20). The leftmost nodes are *Pre-flight Hearing Status* and the hazard *Distance from Earth*. *Preflight Hearing Status* denotes the results of audiometric testing done on astronauts prior to flight. This establishes a hearing baseline against which hearing shifts can be measured. However, it is not guaranteed that either hearing shifts or the noise environment will be measured in spaceflight. Specialized equipment is required that must be included in the mission planning stages. The monitors for both of these are parts of the *Crew Health and Performance (CHP) System* which receives mass, volume, and power allocations from the spaceflight vehicle. Hence the node for *Vehicle Design*. *Vehicle Design* is limited in mass, power, volume, and bandwidth by *Distance from Earth*. Standards, Requirements, Human System Integration Processes, and Crew Vehicle Integration that all affect *Vehicle Design* and the *CHP System* are all sub-nodes of the Human System Integration Architecture or *HSIA (Risk)* that is shown as the grey node in the bottom left corner of Figure 14.

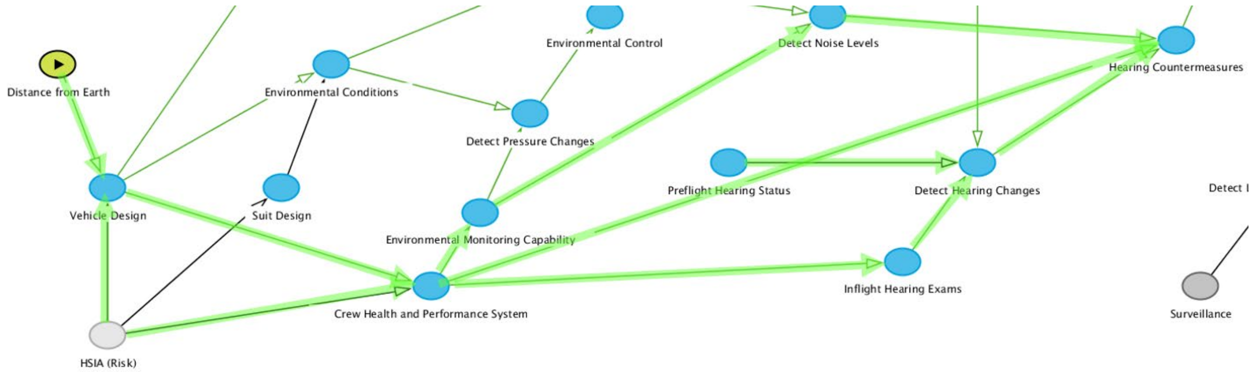

Figure 20: Causal paths involved in monitoring and detection of both noise and hearing

As in the monitoring example in Figure 7, the *Environmental Monitoring Capability* node and the *Inflight Hearing Exams* node point to detection nodes. A noise monitor is a subnode under Environmental Monitoring Capability that enables us to *Detect Noise Levels*. within the vehicle. Likewise, a pressure monitor is also a subnode that enables us to *Detect Pressure Changes*. In the case of *Detect Hearing Changes*, the pre-flight baseline for hearing must be known in order to determine if there are changes from pre-flight, otherwise detection of changes will only be comparable with other in-flight exams. For this reason both point to the detection node. Detection of either unacceptable Noise Levels or Hearing Changes can prompt initiation of *Hearing Countermeasures*. In this case hearing protection such as earplugs or headsets, or adjustment of the noise levels are all possible interventions that may be initiated by early detection of these issues. The arrow from *CHP System* to *Hearing Countermeasures* represents the requirement for those countermeasures to be included in the system planning and design or they will not be available when needed.

While this detailed discussion of the DAG provides insight in written form, for general presentation and brevity we create Narrative slides that accompany each DAG. These are intended to be bullet point formatted and very brief descriptions of the key points of the DAG. While additional detail is helpful to the reader, it can be distracting to high level decision-makers who require a brief high-level discussion of the risk. For the purposes of communication to high level decision makers such as boards, technical authorities, or program managers, the Narrative Slides are intended to provide a high-level walkthrough that can enable them to ask questions if they would like deeper information. In the case of the Acoustics Risk DAG, the Narrative Slides shown below (Figure 21) accompanied the DAG and were configuration managed along with it at the time of acceptance, January 13, 2022.

Hearing Loss Risk DAG: Narrative

- From a health perspective this DAG centers around **Cochlear Changes** which are changes inside the inner ear that can lead to issues with hearing. These culminate in effects on **Individual Readiness** and **Crew Capability**. This can be influenced by changes in
- **Noise Exposure** which includes **Noise Intensity Level**, **Noise Exposure Duration**, and **Noise Spectrum**.
- **Ototoxins** in the environment or in medications
- **Sudden Sensorineural Hearing Loss** which is dependent on **Individual Factors** and has been recorded in some astronauts.
- **Intrachochlear Pressure** caused by **Fluid Shifts** in **Altered Gravity** environments. In this case the **Effective Exposure Duration** accounts for the cumulative effect that the exposure will have for different Design Reference Missions.
- **Barotrauma** that can result from changes in pressure represented here by **Environmental Conditions**. This can result in Inner Ear Barotrauma that affects **Cochlear Changes** or Middle Ear Barotrauma that affects **Measurable Hearing Shifts** without affecting the cochlea. This is affected by **Suit Design.**
- From a performance perspective, **Noise Exposure** leads directly to **Task Performance** showing that the noise environment can affect performance by impacting effective communications without degrading astronaut health.
- **Vehicle Design** and the **Crew Health and Performance System** enable **Noise Monitoring** and **In-Flight Hearing Exams** if these are designed into the system. When designed into the system, they enable **Detect Noise Levels** and **Detect Hearing Changes**. **Inflight Hearing Exams** must be coupled with **Pre-Flight Hearing Status** to enable detection of changes. Detection of either inappropriate **Noise Levels** or actual hearing changes can prompt crews to use **Hearing Countermeasures** such as hearing protection, which must also be designed into the **Crew Health and Performance System** to enable risk mitigation.
- From the **Barotrauma** perspective, **Environmental Monitoring Capability** enables us to **Detect Pressure Changes**. Standards require that crew have **Environmental Control** over the rate of depressurization that can minimize the likelihood of experiencing **Barotrauma**.
- **Measurable Hearing Shifts** and **Hearing Countermeasures** both affect **Individual Readiness** and **Crew Capability**. In some cases, **Measurable Hearing Shifts** can lead to medical problems like **Hearing Loss** both In-Mission as well as **Long Term Health Conditions**.

Figure 21: Narrative Slide describing the causal flow of the Acoustics/Hearing Loss Risk DAG above. Approved in the Human System Risk Board January 13, 2022.

Applications of DAGs

- Once the causal flow in a DAG is agreed upon, metadata can be assigned to either the nodes or the edges as needed to help visualize important features of the risk. Nodes that exist in a given Human System Risk DAG form the set of nodes with the label of the risk name. Due to harmonization guidance, some of those nodes are likely to exist in the DAGs for other Human System Risks. Each node carries a set of labels for which Risks it belongs to as well as other potential labels. Nodal labels/visualizations can include Identification of:
- Hazard Set
- Mission Level Outcomes Set
- Human System Risk(s) Set
- Countermeasures Set
  - Preventive Countermeasures Subset
  - Intervention Countermeasures Subset
  - Monitoring Countermeasures Subset

- Nesting Labels - i.e. sub-nodes under other nodes carry a label that shows they are included in the set of sub-nodes under every node above them.

Other metadata can be assigned to Nodes and Edges. For example, the Level of Evidence (LOE) that is available to support any specific claims of causality within a DAG can be illustrated by varying the appearance of the connections between the relevant nodes. Figure 15 shows the initial DAG for the SANS risk that was accepted at the HSRB on April 23, 2020. (Note that this DAG was accepted while the guidelines for standardized DAG creation were still being developed and therefore does not follow all of the standardization guidance outlined in this document, and there have been subsequent official updates to this DAG not shown here.) This version provides an excellent example on the potential use of metadata in the DAGs. The DAG conveys a story about how changes to eye structures caused by the spaceflight environment lead to functional changes (visual impairments) which in turn can lead to Mission Level Outcomes.

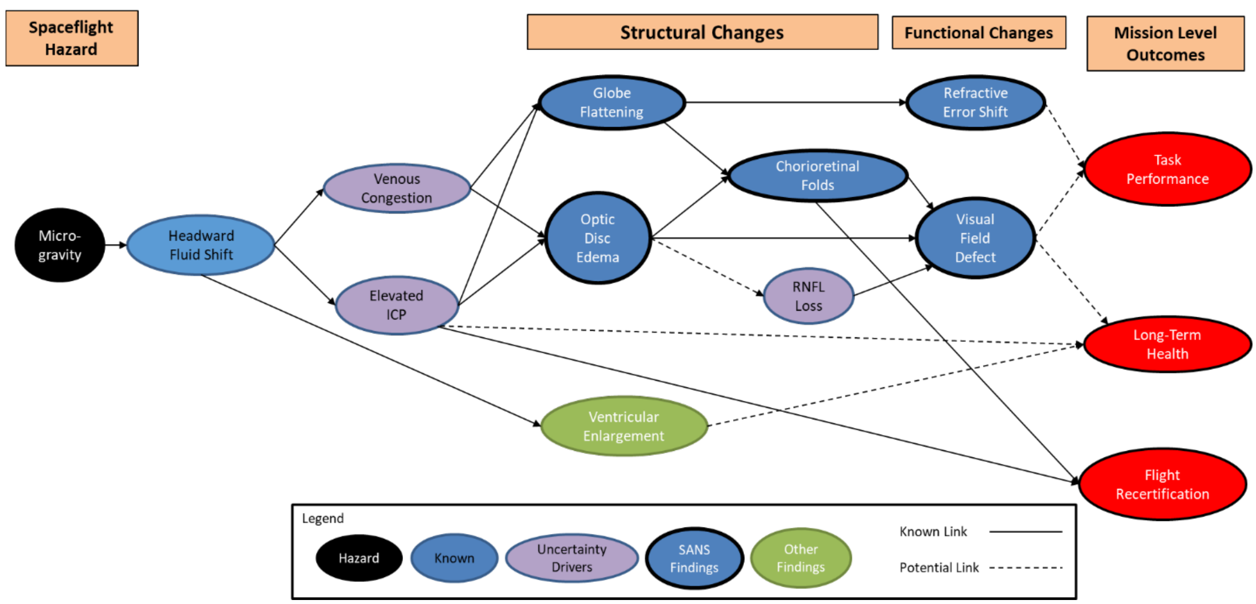

Figure 22: The initial DAG for SANS accepted by the HSRB. Metadata was added to nodes through color and border differentiation with a legend to illustrate different features of specific nodes. Metadata was added to edges to illustrate whether the causal connection was known (supported by reasonable evidence) or suspected (supported by 'Weak' or no evidence). RNFL – Retinal Nerve Fiber Layer.

The most important take-away here is the variety of node colors, node border thickness, and edge arrow features that are used to impart additional information about this risk. For example, integral nodes are broken into sub-categories of *Known* nodes, *SANS Findings*, and *Uncertainty Drivers*. This helps the reader understand more easily where there is 'Strong' evidence and where we are uncertain about contributions of some of these nodes. This diagram suggests that the structural changes observed in astronauts' eyes are caused by both Venous Congestion and Elevated Intracranial Pressure (ICP) [27]–[29]. However, the relative contributions of these two nodes are not a part of this DAG. If sufficient data are available, the contributions of each could be assigned as either a path coefficient – in a regression-style analysis – or as part of a joint probability distribution expressed over the structure of the DAG. This convention allows the Risk Custodian Teams to differentiate nodes that are well-documented from those that are currently areas of uncertainty. It also allowed differentiation of other findings that are potentially related to SANS and often brought up in the context of the discussion of the risk. Ventricular changes have been documented in post-flight MRIs of astronauts as well as other brain changes [30]–[32]. Evidence has yet to be presented that shows a link from those changes to measurable clinical effects in astronauts, post-flight. For this reason, a dotted line is applied to the edge. This highlights gaps in knowledge. Dotted lines show areas where a 'Weak' or 'Speculative' Level of Evidence supports the assertion that a causal connection exists between those nodes, and a solid line shows areas where a 'Moderate' or 'Strong' Level of Evidence exists between those nodes. This references Table 3 and illustrates how LOE can be applied as metadata to each arrow to clearly show where our available evidence base leaves SMEs unsure of the connections. For example, both ICP Elevation and Ventricular Enlargement have dotted line connections to Long Term Health Outcomes. ICP Elevation has some known links in terrestrial medicine to long term health issues, but it is unclear if these apply in spaceflight. And it remains a knowledge gap as to whether ICP is truly elevated in spaceflight as it has never been measured during spaceflight.

This example shows why it is important to enable different visualization capabilities based on the metadata that is applicable to each risk and to the larger risk network. The ability to visualize LOE through dashes, dotted lines, or different line thicknesses is a software feature that is planned for future capabilities. Not shown in this basic DAG are the relevant countermeasures, one of which is corrective lenses which enables adequate visual acuity for unaffected task performance so far in human spaceflight. This example highlights the potential strengths of metadata assignment to nodes and edges based on features of interest such as Relevant Standards or Level of Evidence. In future work guidelines for metadata assignment and visualization are planned as well as software updates that can enable further relationships to be captured within the DAG data structures.

NASA Pilot Study Discussion

It is critical to note that the acceptance of these DAGs at the HSRB is not the final word on whether the story is accurate. Rather, it is a starting point from which changes to the currently accepted story can be brought forward with sufficient evidence to inform additions or subtractions of nodes, or changes in the connections between nodes. It also allows discussion and inclusion of speculative areas of concern as potential falsifiable hypotheses in the larger network of risk. In a sense it is internal crowd-sourcing of expertise: it enables the community to keep track of how the evolving evidence from research, occupational surveillance, clinical care, and other sources of data impact our understanding of each of the Human System Risks. NASA engages in a continual risk management process for exactly this reason – as new evidence is generated in the nascent field of human health and performance in spaceflight, the accepted story should be reconsidered and modified based on the strength of that evidence. Low level evidence has an important role to play. It can identify places in the DAG where new nodes should be considered as a possible knowledge or capability gap for the human system in spaceflight. Identifying the relationships between proposed new nodes that have pathways to Mission Level Outcomes can help identify which areas of research investment are likely to have valuable return on investment for risk reduction in spaceflight. This provides research scientists with a way to frame their concerns with program managers in the context of risk management through a shared mental model. When they can point to a node or a connection between nodes as the location in the known risk map where their research will help address a knowledge or capability deficit, it automatically contextualizes the relationship between specific research and Mission Level Outcomes.

While this approach has strengths, there are challenges and limitations to consider as well. First, the health and human performance community at NASA are not the only contributors to characterizing Human System Risk in spaceflight. Other international space agencies, academic researchers, and the emerging commercial space enterprise hold sources of spaceflight data that should be considered. It is our hope that in publishing this approach (and placing the resulting DAGs in the Public Domain) that those entities will be empowered to bring forward their insights and help to refine and advance the story of risk begun here. NASA intends to provide methods for improved public interaction in the future through their websites. On the other hand, it should also be noted that once individual DAGs are created and documented, they could impart anchoring bias on the scientific community. As such it is important to continually challenge the story that each DAG presents as new evidence emerges and to add and subtract from the story when sufficient evidence is gathered to justify these updates. Naming conventions are unlikely to satisfy all stakeholders of risk at detailed levels, so there must be a means of settling debate on how exactly to name and portray nodes, categories, and relationships between nodes. This should be guided by the strength of evidence brought forward for consideration, but in cases of dissent the authority currently resides with the HSRB at NASA to determine what the ‘official’ DAG for a risk will be until the next update to the risk occurs as part of the continuous risk management process; such changes in node naming and parameterization are changes that can and should be documented as risks are updated. The appropriate level of detail for represented nodes can also be a source of disagreement. Some might argue that specific medical conditions should all be explicitly shown in individual risk DAGs. While this is possible for a Detailed DAG that is intended to be analyzed through computer-based data analytics, it is overwhelming from the perspective of a board or program manager who has to understand the high-level sources of risk. The appropriate level of detail for Narrative or Detailed DAGs depends on the purpose for which it is being employed and is likely to remain more of an art than a science in the near future.

The creation of configuration-managed DAGs not only creates a communal organizational memory of how our understanding of spaceflight risks has evolved, but it also enables the creation of a risk network that can help inform how the risks interact with one another. It is intuitively obvious that the 30 risks are interdependent, and NASA has sought methodologies in the past for understanding and documenting the inter-relationships among risks and the potential cumulative effects they might pose on Mission Level Outcomes [8], [33]. By diagramming the causal flow from common Hazards to common Mission Level Outcomes, a set of individual DAGs enables the creation of a risk network that can be structurally and computationally analyzed to gain insights that have not been available in the siloed approach to risk taken to date. While the HSRB at NASA has approved initial versions of individual DAGs for all 30 configuration-managed risks currently tracked (January 13, 2022), it is the board’s intention to continue curating the existing DAGs, and to develop new ones for any future risks yet to be defined. The standardized approach to representation and lexicon in DAG creation is intended to facilitate the creation of an integrated risk network and use it as the basis for data-driven decisions regarding risk characterization and mitigation as the evidence base for human spaceflight evolves.

References

[1] R. S. Williams, C. R. Doarn, and M. Oxford, Eds., *Engineering, life sciences, and health/medicine synergy in aerospace human systems integration: the Rosetta Stone Project*. Washington, DC: National Aeronautics and Space Administration, NASA Headquarters, 2018.

[2] E. Romero and D. Francisco, "The NASA Human System Risk mitigation process for space exploration," *Acta Astronautica*, vol. 175, pp. 606–615, Oct. 2020, doi: 10.1016/j.actaastro.2020.04.046.

[3] E. Antonsen, "Human System Risk Management Plan, JSC-66705 Rev. A." Johnson Space Center, NASA, Oct. 02, 2020. Accessed: Feb. 21, 2021. [Online]. Available: https://ntrs.nasa.gov/citations/20205008887

[4] J. Kahn, N. Conrad, P. Demitry, B. Dunbar, B. Harris, and D. Hoel, *Health Standards for Long Duration and Exploration Spaceflight: Ethics Principles, Responsibilities, and Decision Framework*. The National Academies Press, 2014.

[5] L. B. Landon *et al.*, "The Behavioral Biology of Teams: Multidisciplinary Contributions to Social Dynamics in Isolated, Confined, and Extreme Environments," *Front. Psychol.*, vol. 10, p. 2571, Nov. 2019, doi: 10.3389/fpsyg.2019.02571.

[6] S. Greenland, J. Pearl, and J. M. Robins, "Causal Diagrams for Epidemiologic Research," *Epidemiology*, vol. 10, no. 1, pp. 37–48, Jan. 1999.

[7] D. Koller and N. Friedman, *Probabilistic Graphical Models*. Cambridge, MA: The MIT Press, 2009. [Online]. Available: https://djsaunde.github.io/read/books/pdfs/probabilistic%20graphical%20models.pdf

[8] J. R. Davis, J. A. Fogarty, and E. E. Richard, "Human health and performance risk management—an approach for exploration missions," *Acta Astronautica*, vol. 63, no. 7–10, pp. 988–995, Oct. 2008, doi: 10.1016/j.actaastro.2008.02.004.

[9] J. Textor, B. van der Zander, M. S. Gilthorpe, M. Liśkiewicz, and G. T. H. Ellison, "Robust causal inference using directed acyclic graphs: the R package 'dagitty,'" *Int. J. Epidemiol.*, p. dyw341, Jan. 2017, doi: 10.1093/ije/dyw341.

[10] S. M. Smith, M. A. Heer, L. C. Shackelford, J. D. Sibonga, L. Ploutz-Snyder, and S. R. Zwart, "Benefits for bone from resistance exercise and nutrition in long-duration spaceflight: Evidence from biochemistry and densitometry," *J Bone Miner Res*, vol. 27, no. 9, pp. 1896–1906, Sep. 2012, doi: 10.1002/jbmr.1647.

[11] J. D. Sibonga, E. R. Spector, S. L. Johnston, and W. J. Tarver, "Evaluating Bone Loss in ISS Astronauts," *Aerospace Medicine and Human Performance*, vol. 86, no. 12, pp. 38–44, Dec. 2015, doi: 10.3357/AMHP.EC06.2015.

[12] L. L. Ploutz-Snyder *et al.*, "Integrated resistance and aerobic exercise protects fitness during bed rest," *Medicine and science in sports and exercise*, vol. 46, no. 2, pp. 358–68, Feb. 2014, doi: 10.1249/MSS.0b013e3182a62f85.

[13] J. Rochlis, P. Campbell, M. Miller, and E. G. Witt, "Human Systems Integration (HSI) Practitioners Guide, NASA/SP–2015-3709." NASA Johnson Space Center, Nov. 2015. [Online]. Available: https://ntrs.nasa.gov/citations/20150022283

[14] M. A. Seibert, D. S. S. Lim, M. J. Miller, D. Santiago-Materese, and M. T. Downs, "Developing Future Deep-Space Telecommunication Architectures: A Historical Look at the Benefits of Analog Research on the Development of Solar System Internetworking for Future Human Spaceflight," *Astrobiology*, vol. 19, no. 3, pp. 462–477, Mar. 2019, doi: 10.1089/ast.2018.1915.

[15] L. K. Barger *et al.*, "Prevalence of sleep deficiency and use of hypnotic drugs in astronauts before, during, and after spaceflight: AN observational study," *The Lancet Neurology*, vol. 13, no. 9, pp. 904–912, 2014, doi: 10.1016/S1474-4422(14)70122-X.

[16] D. F. Dinges, M. Basner, A. J. Ecker, P. Baskin, and S. L. Johnston, “Effects of zolpidem and zaleplon on cognitive performance after emergent morning awakenings at Tmax: a randomized placebo-controlled trial,” *Sleep*, vol. 42, no. 3, Mar. 2019, doi: 10.1093/sleep/zsy258.
[17] B. Burroughs, *Dragonfly: NASA and the Crisis Aboard MIR*, 1st Edition. NY, NY: Harper Collins, 1998.
[18] J. Law *et al.*, “Relationship between carbon dioxide levels and reported headaches on the international space station,” *Journal of Occupational and Environmental Medicine*, vol. 56, no. 5, pp. 477–483, 2014, doi: 10.1097/JOM.0000000000000158.
[19] C. Allen and S. Denham, “International Space Station Acoustics - A Status Report,” presented at the 41st International Conference on Environmental Systems, Portland, Oregon, Jul. 2011. doi: 10.2514/6.2011-5128.
[20] J. G. Limardo, C. S. Allen, and R. W. Danielson, “Status: Crewmember Noise Exposures on the International Space Station,” p. 13.
[21] M. Dunn, “U.S. Astronauts Sound Off on Space Noise,” *Los Angeles Times*, Dec. 02, 2001. Accessed: Jul. 04, 2021. [Online]. Available: https://www.latimes.com/archives/la-xpm-2001-dec-02-mn-10521-story.html
[22] J. C. Buckey, F. E. Musiek, R. Kline-Schoder, J. C. Clark, S. Hart, and J. Havelka, “Hearing loss in space,” *Aviat Space Environ Med*, vol. 72, no. 12, pp. 1121–1124, Dec. 2001.
[23] A. Nakashima, J. Limardo, A. Boone, and R. W. Danielson, “Influence of impulse noise on noise dosimetry measurements on the International Space Station,” *International Journal of Audiology*, vol. 59, no. sup1, pp. S40–S47, Jan. 2020, doi: 10.1080/14992027.2019.1698067.
[24] P. Ganesan, J. Schmiedge, V. Manchaiah, S. Swapna, S. Dhandayutham, and P. P. Kothandaraman, “Ototoxicity: A Challenge in Diagnosis and Treatment,” *J Audiol Otol*, vol. 22, no. 2, pp. 59–68, Feb. 2018, doi: 10.7874/jao.2017.00360.
[25] K. C. M. Campbell and C. G. Le Prell, “Drug-Induced Ototoxicity: Diagnosis and Monitoring,” *Drug Saf*, vol. 41, no. 5, pp. 451–464, May 2018, doi: 10.1007/s40264-017-0629-8.
[26] T. H. Alexander and J. P. Harris, “Incidence of Sudden Sensorineural Hearing Loss,” *Otology & Neurotology*, vol. 34, no. 9, pp. 1586–1589, Dec. 2013, doi: 10.1097/MAO.0000000000000222.
[27] A. G. Lee, T. H. Mader, C. R. Gibson, T. J. Brunstetter, and W. J. Tarver, “Space flight-associated neuro-ocular syndrome (SANS),” *Eye*, vol. 32, no. 7, pp. 1164–1167, Jul. 2018, doi: 10.1038/s41433-018-0070-y.
[28] T. H. Mader, C. R. Gibson, N. R. Miller, P. S. Subramanian, N. B. Patel, and A. G. Lee, “An overview of spaceflight-associated neuro-ocular syndrome (SANS),” *Neurology India*, vol. 67, no. 8, p. 206, May 2019, doi: 10.4103/0028-3886.259126.
[29] G. R. Clément *et al.*, “Challenges to the central nervous system during human spaceflight missions to Mars,” *Journal of Neurophysiology*, vol. 123, no. 5, pp. 2037–2063, May 2020, doi: 10.1152/jn.00476.2019.
[30] D. R. Roberts *et al.*, “Effects of Spaceflight on Astronaut Brain Structure as Indicated on MRI,” *N Engl J Med*, vol. 377, no. 18, pp. 1746–1753, Nov. 2017, doi: 10.1056/NEJMoa1705129.
[31] D. R. Roberts and L. G. Petersen, “Studies of Hydrocephalus Associated With Long-term Spaceflight May Provide New Insights Into Cerebrospinal Fluid Flow Dynamics Here on Earth,” *JAMA Neurol*, vol. 76, no. 4, p. 391, Apr. 2019, doi: 10.1001/jamaneurol.2018.4891.

[32] J. P. Bagian, “How Safe Is Safe Enough for Space and Health Care?: Communication and Acceptance of Risk in the Real World,” *JAMA Neurol*, vol. 76, no. 4, p. 399, Apr. 2019, doi: 10.1001/jamaneurol.2018.4405.
[33] J. Mindock and D. Klaus, “Application of a Spaceflight Contributing Factor Map for Definition and Assessment of Spacecraft Design Requirements,” presented at the 42nd International Conference on Environmental Systems, San Diego, California, Jul. 2012. doi: 10.2514/6.2012-3420.
[34] R. L. Boyer, M. Bigler, and J. H. Rogers, “Probability of Loss of Crew Achievability Studies for NASA’s Exploration Systems Development,” Honolulu, HI, Jun. 2014, p. 6. [Online]. Available: https://ntrs.nasa.gov/api/citations/20140007417/downloads/20140007417.pdf

Appendix A

Quick guide rules for drawing Human System Risk DAGs

1. Each Human System Risk DAG must have at least one Hazard and at least one Mission Level Outcome.
2. Each arrow between nodes represents a falsifiable hypothesis that a causal connection exists. If there is disagreement about an arrow, the team should bring forth the evidence that supports a change in the diagram. If there is conflicting evidence in the evidence base, the arrow should remain but should be indicated as a 'Weak' level of evidence (dotted line).
3. Draw countermeasures to the node that they affect. This means that a preventive countermeasure should be drawn to the node that it prevents - i.e. a preventive medication node points to the illness that it prevents. An intervention countermeasure node should be drawn to the nodes downstream of the node that you are treating or intervening upon - i.e. a medication to treat an illness is not drawn to that illness but to the Individual Readiness node that is downstream of the illness.
4. When drawing links to other Human System Risks in a given risk DAG, do not show outgoing risks, only show incoming risks. This is for visual clarity and to keep each risk DAG focused on the risk in question.
5. When questioning whether to include a higher level of detail in a DAG, ask yourself if that level of detail would be helpful to explain the risk to a decisionmaker who is not an expert in your field, such as a NASA program manager. If it might confuse them or take you down an irrelevant rabbit hole during discussion, then it is likely too much detail.

Appendix B

Pre-Specified Definitions

It is important to understand the definitions and terminology used by NASA. These are reviewed in this section and are formally defined for the agency in the Human System Risk Management Plan JSC-66705 [3].

Hazards are unchangeable aspects of spaceflight that are harmful to humans. The set of Hazards = {Altered Gravity, Radiation, Isolation and Confinement, Hostile Closed Environment, Distance from Earth}.

Mission Level Outcomes are those health and performance outcomes that matter at an agency level as defined by the HMTA. A brief description of each of the Mission Level Outcomes of importance to the NASA HMTA are as follows:

- Task Performance – impacts to crewmembers' ability to accomplish the tasks they are to perform manifest as risk to in-mission timelines and resources. In the worst case these deficits can lead to loss of mission objectives. To be eligible for consideration for inclusion in a DAG, decrements in Task Performance must be both plausible and measurable.
- Evacuation – injury or illness that rises to a sufficiently concerning level may result in consideration of evacuation of the crew from the mission to preserve 'life and limb'. Changing return times to Earth for different DRMs affects the resources required for successful evacuation. In Mars missions, evacuation is not available due to orbital mechanics, so any issues that rise to this level will either self-resolve or lead to death or permanent impairment.
- Loss of Mission Objectives – Mission Objectives include the agency purpose for sending astronauts on a given mission. Inability to accomplish these represents the loss of a significant reason for the mission and is high risk for the agency.
- Loss of Crew Life – Loss of an individual crew life is a possibility in the human health and performance domain due to injury or illness and represents a Mission Level risk Outcome.
- Loss of Crew – Loss of the entire crew, as opposed to a single individual is typically calculated at the mission safety level separate from health and performance risk calculations [34]. However, there may be cases where Loss of Crew could happen for health and medical reasons.

- Loss of Mission – Loss of Mission can result from loss of sufficient mission objectives or loss of crew and is dependent on agency assessment of goals. An example of this is the Apollo 13 mission, where the crew experienced Loss of Mission when they were unable to land on the moon, but they did not experience Loss of Crew, as they safely returned to Earth. In contrast, the loss of the Space Shuttle Challenger is an example of both Loss of Crew and Loss of Mission.
- Flight Recertification of Astronauts – NASA investments in astronaut training and skill sets are critical to mission success. When astronauts experience medical issues incurred from flight exposures, they may be unable medically to recertify for flight.
- Long Term Health (LTH) Outcomes – Spaceflight exposures that lead to post-mission medical conditions affect the long-term health and quality of life of astronauts. The Chief Health and Medical Officer at NASA also carries some responsibility for this risk. A common example is the risk of developing cancer from radiation exposures.

The set of Mission Level Outcomes = {Task Performance, Evacuation, Loss of Mission Objectives, Loss of Crew Life, Loss of Crew, Loss of Mission, Flight Recertification, Long Term Health Outcomes}.

Other Key Terms include Design Reference Missions (DRM) categories, contributing factors and countermeasures. These are commonly used in the human spaceflight community to describe what missions we are talking about and what assumptions we make (DRMs), where risk comes from (contributing factors) and what we do to try to mitigate it (countermeasures). Note that some of the countermeasures we use to reduce risk in one area can cause increased risk in other areas. Think about the side effects of medications for example – a medicine that helps reduce space motion sickness can also cause drowsiness at a time when a crew member is expected to perform a complex operation. Recognizing this, in the context of DAGs, all countermeasures are also categorized as contributing factors.

- Design Reference Mission categories - NASA mission categories, derived from a subset of risk drivers, loosely defined by destination, operating environment, and expected duration. These broad categories are scoped to allow the flexibility to provide risk characterizations and assessments that will be applicable to a range of human space exploration missions including those yet to be defined. There are currently four DRMs which are divided into long and short durations.

- Contributing Factor – an operational, design, or human-system variable (including spaceflight hazards) that can influence the likelihood and/or consequence of Human System Risks. For example, (degree of) crew autonomy is a contributing factor to the Risk of team performance and behavioral decrements; (amount of) in-flight exercise capability is a contributing factor to Risk of reduced muscle size and strength.
- Countermeasure – any action, hardware/software or capability provided pre-, in-, or post-mission that serves to reduce risk within the Risk Impact Categories. There are three types of countermeasures as applied to Human System Risks managed by the HSRB:
    - Monitoring Countermeasure – a countermeasure implemented during the course of a mission used either operationally or for occupational surveillance to provide actionable information to crew or clinicians on prevention effectiveness, and when to implement risk reduction interventions. For example, *Environmental Monitoring Capability* and *Inflight Hearing Exams* are monitoring countermeasure for the Acoustics Risk. *Environmental Monitoring Capability* here includes noise monitoring and atmospheric pressure monitoring.
    - Prevention Countermeasure – a countermeasure implemented pre-flight and during flight that decreases the influence of contributing factors and hazards on the Risk or on the scenario that enables the Risk to manifest. For example, *Environmental Control* is a prevention countermeasure for the Acoustics Risk. *Environmental Control* here includes control over noise levels and atmospheric pressure.
    - Intervention Countermeasure – a countermeasure applied after the risk scenario occurs intended to reduce the severity of the consequence. For example, *Hearing Countermeasures* is an intervention countermeasure for the Acoustics Risk. In cases where the noise exposure experienced by the crew becomes excessive, the crew can intervene by applying ear plugs. *Environmental Control* can also be an Intervention Countermeasure, in cases where the noise environment becomes too loud, the intervention may be to intervene to reduce the noise.

Appendix C

Acronyms and Abbreviations

Acronym	Meaning
CHeCS	Crew Health Care System
CHP	Crew Health and Performance
CO	Carbon Monoxide
CO_2	Carbon Dioxide
DAG	Directed Acyclic Graph
DCS Risk	Risk of Decompression Sickness
DRM	Design Reference Mission
EVA	Extravehicular Activity
EVA Risk	Risk of Injury and Compromised Performance Due to EVA Operations
g	gravity
HMTA	Health and Medical Technical Authority
HSIA	Human Systems Integration Architecture
HSIA Risk	Risk of Adverse Outcomes Due to Inadequate Human Systems Integration Architecture
HSRB	Human System Risk Board
Hypoxia Risk	Risk of Reduced Crew Health and Performance Due to Hypoxia
ICP	intra-cranial pressure
ISS	International Space Station
LEO	Low Earth Orbit
LOE	Levels of Evidence
LxC	Likelihood and Consequence
NASA	National Aeronautics and Space Administration
RMP	JSC-66705 Human System Risk Management Plan; Rev. A (Oct 2020)
SME	Subject Matter Expert