

A HOLISTIC, CROSS-GOVERNMENT ALL HAZARDS RISK ASSESSMENT

Simona Verga, PhD
Defence Research and Development Canada,
Centre for Security Science (DRDC CSS)
Ottawa, ON, Canada
email: simona.verga@drdc-rddc.gc.ca

ABSTRACT

This paper describes an approach to build a holistic, cross-government all-hazards risk assessment process, which aims to capture threats of all “stripes” and understand the extent to which they can become a risk to the safety and security of the Canadian population and society. The approach makes a conscious effort to consider the Canadian risk picture within the global risk environment. The federal All-Hazards Risk Assessment (AHRA) initiative aims to develop a mechanism for a comparative assessment and rating of risk events derived from all hazards (regardless of the source, whether malicious or non-malicious), in order to support emergency management planning in the federal domain. A standardized methodology is pursued in order to leverage the expertise of individual federal departments, share this expertise and knowledge across a community of practice, and generate a whole-of-government view of risks. Methods that attempt to build shared understanding of risks across organizations may be useful to multi-agency problems beyond the Canadian national context.

INTRODUCTION

Building a holistic, cross-government all-hazards risk assessment which aims to capture threats of all “stripes” and assess the associated public safety and security risk is a challenging task. Some of the challenges include the constantly evolving nature of man-made threats, the shifting patterns in threats derived from natural hazards, and increasingly complex ways even simple threats can lead to societal disruption, because they act on an increasingly interconnected and complex society.

While the approach described herein considers the Canadian risk picture, it has become more and more difficult to dissociate threats at home from the international threat environment. World Economic Forum’s recently published Global Risks 2012 [1] is very revealing in terms of how current and emerging risks transcend national boundaries and challenge the traditional emergency-driven risk management. The challenge is to come up with novel approaches that incorporate the global context, and make a conscious effort to link national risk concerns with global issues. This is a necessary step towards improving collaborative efforts, to the benefit of more resilient communities within our nations, and aiming to build a resilient global community.

The safety and security landscape is changing constantly, in Canada as elsewhere, either slowly eroded by less violent but constantly mounting pressures (the emergence of new diseases and the possibility of new pandemics; the exacerbation of weather events by global warming; unforeseen effects of emerging technologies; the asymmetric distribution of wealth and resources leading to increased demographic pressures; the increased complexity of our infrastructures stressing control mechanisms closer and closer

to the breaking point; and, not in the least, the greater and greater expectation of protection by the population from its government), or “shaken” violently by “quake-like” events, such as the terrorist attacks of September 11, 2001. The defence landscape has changed significantly as well, following the end of the Cold War, in response to the shifting geopolitical landscape and the continuous blurring of the boundary between defence and security concerns.

These examples underscore both the need for a robust, all-hazards approach to public safety and security, and the challenges associated with undertaking such an approach.

TOWARDS A CANADIAN HOLISTIC RISK PICTURE

In Canada, all levels of Government – federal, provincial/territorial and municipal – share the responsibility to protect Canadians and Canadian society. Within each jurisdiction, the governments’ public safety and security functions are shared among many departments and agencies. Hence, preparedness at the national level depends on synchronized efforts among many partners. For practical reasons, any attempt to formalize an overarching risk model needs to respect existing structures for ownership and responsibility in managing risks. As an example, in Canada, public safety and security functions are shared horizontally, within one jurisdiction, among several departments and agencies, and they cross jurisdictional boundaries based on severity of consequences of the event. This adds to the complexity of planning and managing even single emergencies that escalate across jurisdictions and organizational boundaries. The added challenge of an all-hazards planning comes from the lack of a coherent picture on the relative severity of the risks associated with various threats and hazards.

Organizational background

The Centre for Security Science (CSS) has made it its mission to generate and support scientific activities aimed at improving public security across the whole of government, and making a contribution internationally. CSS is organized as an extensive network of national and international Science and Technology (S&T) partners and public security communities, which includes both “producers” and “users” of S&T products and services, risk analysis being one of them. Within CSS, a Risk Portfolio has been established in response to growing interest in the risk field from across government and defence. The vision is to develop a risk resource centre to support the community with threat, vulnerability and risk assessments, gap analysis, foresight and future security visioning and other related activities and products. The author, a member of the CSS Operations Research (OR) team, has conducted research work on risk assessment methodologies and models, work that constitutes an important contribution to this enhanced risk “toolbox” housed within the CSS Risk Portfolio.

One of the initiatives where the Centre saw an opportunity to contribute was the development of an All Hazards Risk Assessment (AHRA) methodology, in close partnership with Public Safety Canada, the federal department with jurisdiction over coordinating federal public safety and security functions. The purpose of the AHRA is to enable federal institutions to perform risk assessments consistently, and to formalize a structure for combining departmental risk assessments to create a whole-of-government risk picture to support emergency management planning in federal institutions.

Given its mandate and the project’s importance, CSS supported the AHRA initiative, as one that has the potential of benefiting the entire safety and security community. Although this initiative is focussed for

now on the federal level of government, the Centre's position as a "network hub" provides it with a unique opportunity to bank on the variety of expertise available within the spectrum of partner organizations, and to advocate for models and tools that ensure interoperability with efforts at the local, regional, provincial/territorial, as well as the national and international level. Public Safety Canada is leading the collaborative AHRA project and provides the policy cover, while CSS provides technical support to developing the supporting methodology, which includes support to scenario development, risk analysis and evaluation, as well as effective presentation of the overall results. The author contributed significantly to the development of metrics and methods for risk analysis (likelihood and impact assessment) and risk evaluation, and generated the graphical depiction of the overall results for presentation to decision makers.

The Federal All-Hazards Risk Assessment

The Federal All-Hazards Risk Assessment (AHRA) aims to develop a mechanism for a comparative assessment and rating of risk events derived from all hazards (regardless of the source, whether malicious or non-malicious) that are significant enough to warrant federal interest, in order to support emergency management planning in the federal domain. The all hazards risk picture generated with the AHRA methodology initiative has the potential to inform decision-making at all levels.

Amongst the intended outcomes of the AHRA process, an important one is to generate a shared understanding of risks, their likelihood and potential consequences. Each federal institution has its own strategic and operational objectives, with each being exposed to its own unique set of risks, and each having its own information and resource limitations. The AHRA addresses the interconnected nature of Canada's risk environment by providing a means to produce a collective judgment of risks that may be of concern to more than one federal institutions. A standardized methodology is pursued in order to leverage the expertise of individual departments, share this expertise and knowledge across a community of practice, and generate a whole-of-government view of risks. The relative ordering of risk events based on their ratings and the process for assessing them will be used as a starting point for emergency management planning at the federal level, and to inform future actions and initiatives. This shared assessment of risks necessarily takes a high level view; however departments may continue to conduct more focused assessments of risks which fall within their immediate problem space, according to their mandate and responsibilities or interests, and as legislated in Canada by the Emergency Management Act and other relevant legislation and policies.

MULTI-ORGANIZATIONAL CHALLENGES TO PUBLIC SAFETY AND SECURITY RISK ASSESSMENT AND MANAGEMENT

Risk is an intellectual construct, contingent on the belief that human intervention can influence the outcome of future events, as long as an effort is made to anticipate these events. Oftentimes risk has a negative connotation and is associated with the possibility of future harm or loss [2].

In simple terms, risks are about events that, when triggered, cause problems. Most commonly, risks are discussed in the context of effects of uncertainty on an enterprise's objectives. Because risks refer to potential problems in the future, often there is a great deal of uncertainty with regard to how and to what degree such events may be realized. If the enterprise's interests are potentially affected, processes are set

up to manage the uncertainty and strategies are developed to minimize the effect of such future events on desired outcomes. *Risk management* is the structured approach to set up such processes and develop such strategies [3]. Although the details of the risk management approach vary widely across risk domains/organizations, the following steps identified in ISO/FDIS 31000, “Risk Management – Principles and Guidelines” [4] are commonly followed:

1. Risk identification: The process of finding, recognizing, and describing risks.
2. Risk analysis: The process of understanding the nature and level of risk (in terms of its impacts and likelihood).
3. Risk evaluation: The process of comparing the results of risk analysis with risk criteria to determine whether the risk and/or its magnitude is acceptable or tolerable.
4. Risk treatment planning: The process of identifying and recommending risk control (or risk treatment) options.

The outputs of these four steps should provide decision-makers with an improved understanding of the relevant risks (likelihood, impacts) that could affect the enterprise’s objectives, indicators of the effectiveness of risk treatment measures already in place, the potential effectiveness of additional risk treatment measures as well as an appreciation of the inherent uncertainties in all key aspects of the risk assessment process. Generally, risk management is a part of decision making and hence an integral part of organizational planning processes.

One difficulty with existing risk management models is that they are usually designed for single organizations. They can be difficult to extend to a collective of organizations, such as the multitude of public, private and voluntary organizations that have a role in public safety and security. In the context of public safety and security, one must consider carefully what constitutes the “enterprise”, and what is the role of risk assessment in informing collective planning processes to manage risks to society that arise from all hazards. Effective risk management of societal risks must, therefore, also address the multi-organizational planning context of public safety and security. The challenge this poses to the overall risk management paradigm is two-fold:

1. The risk domain is distributed amongst multiple organizations; as such, the assessment requires “assembling” expertise and knowledge across the risk practitioner community.
2. Decision making and planning processes that include risk control and risk treatment options, which complete the overall activity of risk management, are also distributed across various organizations, with their own levels of risk managers and decision makers.

Given the complexities surrounding the multi-organizational structure of the public safety and security domain, it is important to delimitate the scope and role of the AHRA process within overall activities of risk management.

The AHRA process and methodology are focused primarily on the assessment component of the overall risk management paradigm. The hazard risk domain is covered by the AHRA methodology. However, the strategic risk domain (e.g., political risks, reputational risks) and the operational risk domain (e.g., day-to-

day issues confronting the institution) are not, although these aspects may be considered and factored in assigning impact ratings. Figure 1 highlights the role of the AHRA process within the overall process of emergency management, which happens government-wide and involves intra- and inter-departmental activities.

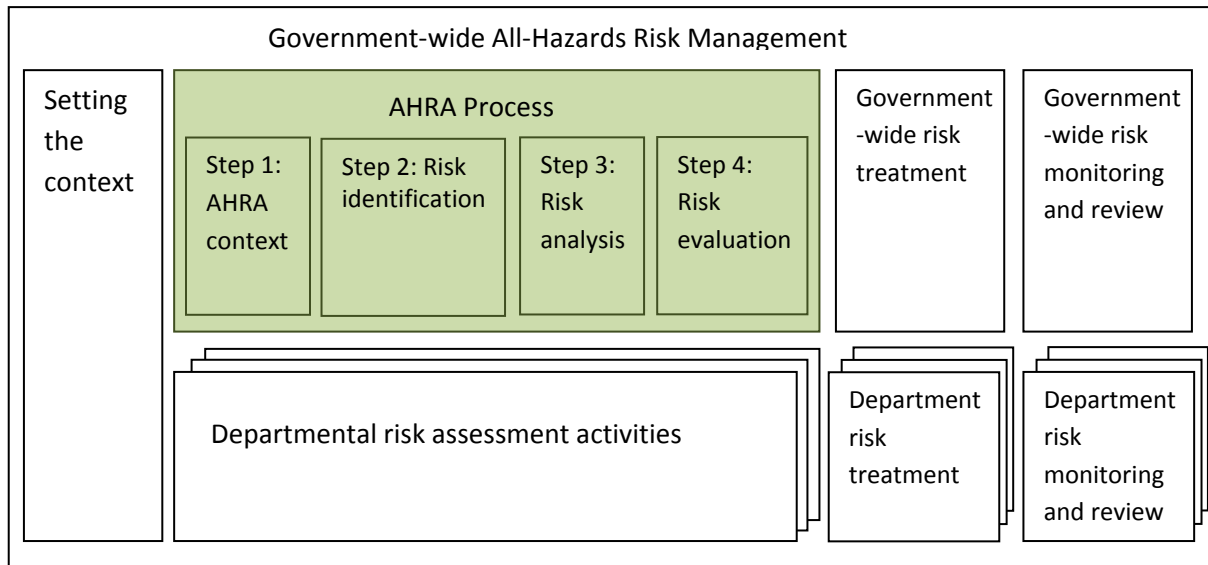


Figure 1: Government-wide All-Hazards Risk Management

SCOPING AND STRUCTURING THE ALL-HAZARDS RISK DOMAIN

Risks are events or circumstances that, if they materialize, could negatively affect the achievement of objectives. In the context of the AHRA, the objective is to ensure that Canadians are protected from the gamut of threats and hazards that can become a risk to their safety and security; to that end, the goal of AHRA is to produce a whole of government all-hazards risk picture that can potentially inform strategies to minimize the effects on society.

In order to make the entire process more manageable, a risk taxonomy was developed [5], which breaks down all-hazards risk by types of threats and hazards of significant potential impact within the field of view of the Federal Government. The most current version of the AHRA taxonomy is shown in Figure 2 below.

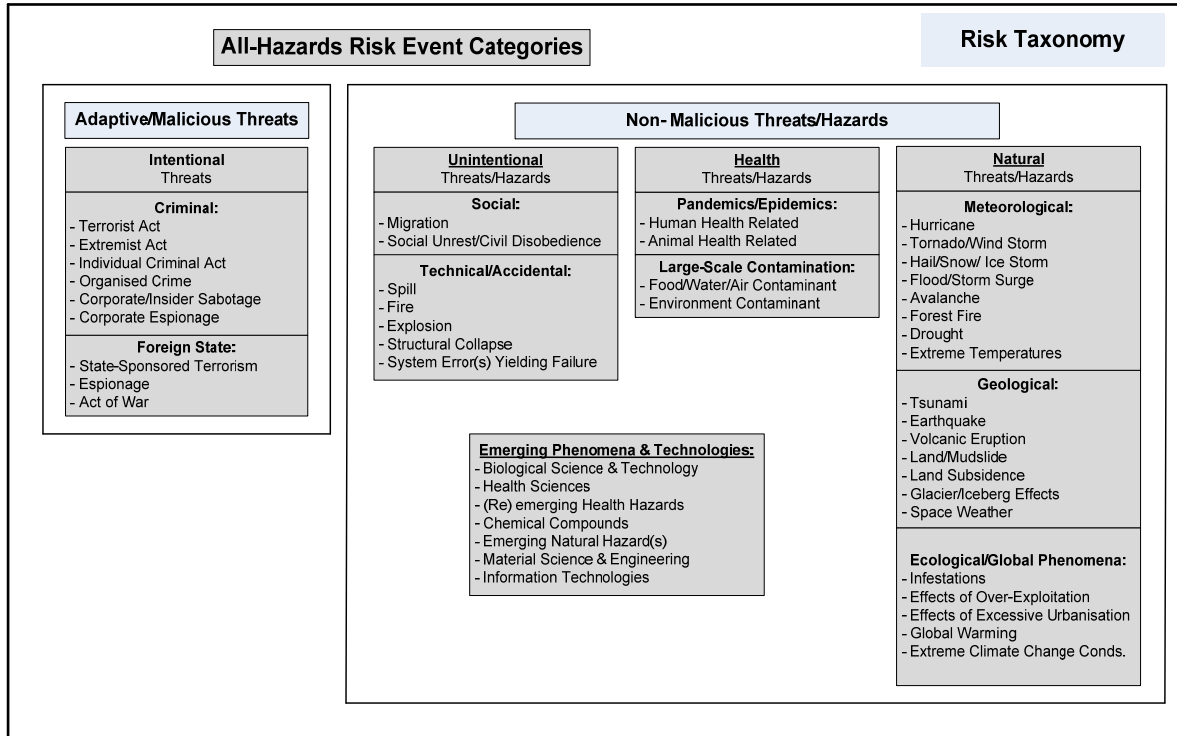


Figure 2: All-Hazards Risk Taxonomy

An effective taxonomy has three key attributes: it provides a classification scheme; it is semantic; and it can be used as a map to navigate the domain [6]. The proposed scheme has been designed with those key attributes in mind. The All Hazards Risk Taxonomy defines the risk domain in which there are two major classes; the Malicious/Adaptive and the Non-Malicious/Non-Adaptive. Malicious threats are intentional and originate from threat actors like terrorists, organized crime actors, or foreign states. Non-malicious threats include unintentional man-made, health, and natural disasters, and cross-cutting risks such as those related to emerging technologies or climate change. To aid the tackling of so vast a domain, the AHRA taxonomy partitions threats/hazards into logical categories, and provides a “blueprint” for further analytical and methodological development. The AHRA taxonomy was developed together with a risk lexicon [7], which can help the collaborative effort by establishing a common terminology and ensure mutual understanding. Both initiatives benefited from input from the project partners, representing many federal organizations with national security, public safety and emergency management objectives. The AHRA taxonomy helps structure the vast problem space, while capturing its breath and complexity.

THE AHRA CONTEXT AND RISK IDENTIFICATION

Invariably, risk assessment starts with risk identification. An important prerequisite to risk identification is establishing the context. This implies selecting the domain of interest, establishing the identity and objectives of stakeholders, the scope of the process, and the basis upon which risks will be evaluated (i.e., assumptions, constraints).

The scope of the approach described in this paper is “all-hazards”, which means that all types of hazards, as captured in the AHRA taxonomy, are considered as sources of risks, although not all may be retained

for further assessment and evaluation. For the federal AHRA project, the focus is on emergencies that are likely to require federal involvement.

Also, in order to employ a common time frame for analysis, the AHRA considers events that are possible within the next five years. This does not exclude rare events, which might have low (but not zero) probability of happening in the next five years. An example would be the disruption of satellite communications by solar storms; although rare, this could occur at any time within the next five years. However, within the current scope of the approach, potential events that are not considered in the AHRA are those that cannot realistically be expected to occur within the next five years. This may be because the conditions that can generate the risk may require more time to emerge as threats (examples include: new technologies, whose effects on human health and environment have not been tested; an asteroid impact, assuming that it is known that there are no asteroids in sufficient proximity in this time frame; etc.).

The federal AHRA project has annual cycles. Each year, federal organizations are asked to bring forward their top threats and hazards. An Interdepartmental Risk Assessment Working Group is mandated to select those threats and hazards that will be retained for further assessment. For example, the federal AHRA project has gone through two cycles (the second cycle is under way): during the first, pilot cycle (2010-2011), 16 threats and hazards were retained, while for the current cycle (2011-2012), 12 threats and hazards were retained, as shown in Table 1 below.

The AHRA is a scenario-based approach. The advantage of such an approach is that scenarios can provide context and specificity to identified sources of risk. The challenge is ensuring that the risk space is appropriately characterized. In developing the set of scenarios one should aim for an “importance sampling” of the risk domain. On the one hand, this means developing a set of scenarios large enough to ensure that drivers for capabilities are adequately captured. On the other hand, the set of scenarios should be reduced to a number that is realistic, in terms of required stakeholder and expert engagement, and in terms of what can be achieved during the assessment process. In other words, one should strive to cover the entire all-hazards risk space, albeit sparsely, to enable a parametric analysis of the whole domain [8], rather than concentrate a disproportionately large number of scenarios on a few narrow risk intervals.

Given the goal of the AHRA project of emphasizing the interconnected nature of Canada’s risk environment, each of the proposed scenarios must present a level of complexity that places it beyond the sole responsibility of one federal organization. Each scenario is a collaborative effort, with multiple institutions involved in its development. The desired output of the scenario development exercise is a realistic set of incident scenarios that reflect serious, plausible threats and hazards that federal organizations with an emergency management mandate must be prepared to address. However, because of the challenges associated with setting up a complex process such as the AHRA, as well as those associated with ensuring the availability of the interdisciplinary expertise required during the assessment process, the number of scenarios developed and assessed was kept at 10 per year, at least for the first two AHRA cycles, as illustrated in Table 2. Generally, new scenarios are sought each year, to build a comprehensive set of scenarios over time.

List of Threats & Hazards (Unrated)	
<p>2010-2011:</p> <ul style="list-style-type: none"> • Terrorism • Foreign interference/espionage (cyber) • Breach of information • Explosions • Oil incident • Accidental and/or intentional chemical event • Aircraft disasters • Marine disasters • Rail disasters • Pandemic (e.g. pandemic influenza) • Emerging respiratory infectious disease outbreak • Foodborne infectious disease outbreak • Zoonotic infectious disease outbreak • Non-zoonotic animal disease • Floods • Hurricanes 	<p>2011-2012:</p> <ul style="list-style-type: none"> • Cyber attack • Terrorism • Influx of Illegal Migrants • Radiological or Nuclear Accident • Marine pollutants/Accidental Chemical Event • Earthquakes • Floods • Hurricanes • Pandemic • Foodborne Outbreak • Extreme Weather • Extreme Space Weather

Table 1: Threats and hazards identified during the first two federal AHRA cycles

For some hazards, a single narrative scenario may be found insufficient to characterize the continuum of associated risk. A more complete assessment would consider the continuum in the magnitude and likelihood of triggering events and variation in the consequences associated with different contexts (locations, weather conditions, economic cycles, day versus night events, etc.). At the same time, the efforts required of a full risk assessment are not practical for the purposes of the AHRA process, particularly in the first few cycles. The possibility of developing composite scenarios with variants has been considered, in order to find the right balance between the reduced number of scenarios that are considered practical in terms of being able to address in a collaborative project, and the more complete view of risk that may ultimately be required to properly inform priority-setting and emergency planning. A composite scenario is in fact a set of related scenarios of varying magnitudes, representing potential realizations on the continuum of risk generated by the same particular hazard, and chosen in such a way to best characterize that continuum for emergency planning purposes.

List of event scenarios	
<p>2010-2011:</p> <p>Health:</p> <ul style="list-style-type: none"> • Emerging zoonotic respiratory pathogen outbreak • Listeriosis outbreak • Foot-and-mouth disease outbreak <p>Natural:</p> <ul style="list-style-type: none"> • Hurricanes • Flood <p>Unintentional:</p> <ul style="list-style-type: none"> • Marine oil spill <p>Malicious:</p> <ul style="list-style-type: none"> • Aircraft disaster (cargo) • Aircraft disaster (passenger) • Chemical release incident • Cyber incident 	<p>2011-2012:</p> <p>Health:</p> <ul style="list-style-type: none"> • Pandemic human disease • Foodborne outbreak <p>Natural:</p> <ul style="list-style-type: none"> • Earthquake • Hurricane • Ice storm <p>Unintentional:</p> <ul style="list-style-type: none"> • Nuclear accident/ technical failure • Marine pollutants – accidental chemical event <p>Unintentional Social/ Intentional Criminal</p> <ul style="list-style-type: none"> • Influx of Illegal migrants <p>Malicious:</p> <ul style="list-style-type: none"> • Cyber attack • Terrorism event

Table 2: Scenarios developed during the first two federal AHRA cycles

RISK ANALYSIS

Once a risk has been identified, it must be analyzed, which means understanding the nature and level of risk, usually in terms of likelihood and impacts or consequences. “Likelihood” refers to the risk event’s chance of occurrence, and “consequences” measure the severity and extent of the risk event’s effects. A challenge facing an “all-hazards” methodology is defining and finding good measures for “Likelihood” and “Consequences” consistently across risks of very different nature.

The AHRA methodology proposes to investigate the consequences of a risk event, as described by a scenario, along a number of dimensions, chosen to capture the spectrum of risk interests of the federal government. There are six impact categories, each designed to cover a significant segment to federal emergency management institutions. Public Safety Canada consolidated a list of impact dimensions, together with relevant criteria to be considered within each impact category, by consultation with the

federal emergency community. The author's role was to develop metrics for risk analysis (likelihood and impact assessment), and to ensure that the approach was as robust and consistent as possible.

Overall, the scoring approach is an "order of magnitude" approach. The scoring method uses a logarithmic (half-log-10) scoring scale, and attempts to apply it consistently across the impact categories, at least to a practical extent. A logarithmic scale is appropriate to allow for the extent of variability and uncertainty associated with a large problem space. It is also appropriate for the purposes of maintaining an underlying mathematical linkage between the simplified scoring approach used for the collaborative work and a full quantitative risk, which might exist and be used by expert analysts in more focused risk areas across the many departments. In other words, the approach is fundamentally (mathematically) compatible with more formal quantitative risk assessment approaches, but does not presume that they exist. Maintaining a mathematically rigorous underlying structure for the simplified scoring approach allows for the current "working" version to be enhanced over subsequent phases, in order to increase the level of compatibility with more formal methods, without need to lose compatibility with previous generations of work.

RISK ANALYSIS - IMPACTS

The six impact categories in the AHRA are: People, Economy, Environment, Territorial Security, Canada's Reputation and Influence, and Societal and Psycho-Social effects. In theory, the six categories were chosen to cover "orthogonal" dimensions in the consequence space; in practice, there were many cases where the effects cannot be clearly split in independent components, which created a challenge in terms of avoiding double-counting of effects. For example, the "environment" impact category is intended to capture unique effects on natural environmental assets; in practice, environmental effects are often tangled up with costs associated with the loss or degradation of those environmental assets, which can be counted as an economic impact. In such cases, the "orthogonality" requirement was enforced during the assessment process, by asking participants to decide which category provides the best fit, and only include those effects once.

People

The first impact category is the "People" category, which captures the impact following a given risk event in terms of fatality, injury, and disease. To account for non-fatal health effects, alongside fatal ones, a composite measure of human health impact was employed, by estimating the proportion of indirect deaths caused by the scenario, which may be thought of as fatality-equivalents. They can be the result of physical injuries, chronic illness, mental illness and being displaced (or lacking basic necessities of life). Fatality-equivalents may also be inferred by defining the relative severity of the injury, illness or displacement and duration of these effects, using metrics such as the Disability-Adjusted Life Years (DALY) lost, which provide quantitative means for taking into account the burden imposed by premature mortality and morbidity on populations [9]. The severity scale ranges between 0 (perfect health) and 1 (death). The duration of the injury is represented in years. For the AHRA project, relative severity is defined as mild, moderate and severe, while the duration includes the options of short term, long term and permanent. A short term injury/illness is defined as being less than four weeks, while a long term injury greater than four weeks. Permanent injuries/illnesses are present for the remainder of the person's life which is assumed to be 40 years for the average adult.

The DALY measure is also used to measure the impact of fatalities. The severity of a fatality is always equal to 1, while an estimate of the number of years of life expectancy at the time of death is the appropriate measure of duration. An adult fatality, on average, can be assigned a score of approximately 40 (for 40 years of lost life). This provides a way to calculate the number of fatality equivalents from DALYs.

Once the total number of fatality equivalents has been tallied up for a given risk event, the following table shows the conversion into an impact rating on a scale from 0 to 5.

Impact Score	Fatality-Equivalents
No Impact	No Impact
0	1
0.5	3
1	10
1.5	30
2	100
2.5	300
3	1,000
3.5	3,000
4	10,000
4.5	30,000
5	100,000

Table 3: Rating table for the “People” impact category

Economy

The second impact category is the “Economy” category. This category captures the dollar value following damage(s) or loss to economically productive assets and disruptions to the normal functioning of the Canadian economic system, following a risk event. This economic loss is broken down into:

- Direct loss, which accounts for the immediate economic damage generated by the disaster, measured in repair or replacement costs for physical assets (buildings, infrastructures, equipment etc.)
- Indirect loss (flow losses), measured in costs associated with the disruption of flows of goods and services, due to damages or disruption to productive assets and economic infrastructure, relative to the duration of the disruption.

When identifying contributions to the economic loss, the assessment should ensure that no double-counting takes place, and certain built-in mechanisms and behavioural changes (e.g. consumer-demand shift, substitution of inputs and/or reallocation of resources, etc.) should be considered, which can mitigate to a certain extent the losses. Taking all identified “positive-” and “negative-costs” into account, all costs are added, and the rating for this category is based on the final dollar figure, as illustrated in Table 4 below.

Impact Score	Total Economic Loss
No Impact	No impact
0	\$10M
0.5	\$30M
1	\$100M
1.5	\$300M
2	\$1B
2.5	\$3B
3	\$10B
3.5	\$30B
4	\$100B
4.5	\$300B
5	\$1,000B

Table 4: Rating table for the “Economy” impact category

Macroeconomic studies provide a complementary way to assess the repercussions of direct and indirect economic losses. For instance, estimates of macroeconomic effects would take into account that some indirect effects could be exacerbated or mitigated in the aggregate by changes in prices or flexibility in the production process (e.g. through reallocations in spending/production across sectors or through the mobilization of production factors if production is not at full capacity). Estimates of high-order impacts require the use of more sophisticated economic models.

It must be noted that for the “people” as well as the “economy” category, in order to minimize round-off errors, the actual numbers in units chosen for the category (fatality equivalents, and dollars, respectively) are elicited during the assessment process; the ratings derived from those numbers are only used during the later stages in the process, when all the assessments are “assembled” into an overall measure of risk for each risk event considered, as described by the corresponding scenario.

The remaining impact categories, however, are less quantifiable; the next few paragraphs describe briefly the qualitative approaches to assessing the “Environment” category, “Territorial security”, “Society and psycho-social” effects, and “Canada’s reputation and influence”. For each of these categories, a 0 to 5 rating scale was developed based on sets of qualitative criteria to be considered within each category, as identified by the federal community (although the same order of magnitude approach was followed in the more quantifiable aspects of impact for some categories). A more detailed description of the rating methods for these categories can be found in [10].

Environment

The “Environment” category rating scale focuses on environmental damage caused by a risk event. In the context of the AHRA, environmental damage refers to non-economical aspects associated with the loss of

environmental assets or environmental quality. This category will exclude assessing economic aspects created by such loss, as they are better captured under the “Economy”, in order to avoid double counting. The Environment impact category relates to the preservation of specific components of the environment pertaining to air, water and soil ecosystems, including fauna and flora. The rating of this category considers four elements that characterize the size and severity of environmental damage from a risk event or an emergency: the magnitude of an environmental response required (local, regional, multi-jurisdictional, general, specialized, etc.); the geographical extent of the damage; the magnitude of damage based on adverse effects to different components of the environment; and the duration of the damage including the level of recovery efforts.

Territorial Security

The “Territorial security” category is intended to capture the impact of a risk event on the Government’s ability to maintain safety and security functions within all of its territory. This dimension captures conditions in which there is a loss or disruption in the Government’s ability to secure its territory or its borders, and to secure the safety of citizens. Challenges can come from abroad (e.g. terrorist attacks, challenges to Arctic sovereignty) or from natural disasters (e.g. hurricanes, earthquakes, infectious diseases).

For this category, the rating is determined by the area affected, with factors including the duration of disruption and population density.

Canada’s Reputation and Influence

“Canada’s reputation and influence” category captures shifts in views towards Canada by foreign governments, international actors and populations following a risk event in Canada or involving Canadians abroad. The rating is based on qualitative descriptions of a non-exhaustive list of situations that can demonstrate effects on Canada’s international position. Examples include: damage or loss of control over Canada’s embassies, suspension of international agreements, protests against Canada, imposition of travel restrictions to Canada, deterioration of bilateral political relations, etc.

Society and Psycho-Social

The “Society and psycho-social” category measures the extent of disruption to normal societal function following a risk event leading to sustained adverse behaviour change in the population. Societal and psycho-social effects might be rooted in people’s understanding and perception of the incident as well as their sense of control over the outcome, which may lead to changes in their individual pattern of behaviour over the short or long term, and may even lead to social actions, such as protests, civil disturbances or vandalism.

The rating of this category is based on a qualitative assessment that focuses on two criteria: public outrage and public anxiety. The descriptors for each of these criteria consider the number of people affected, the nature and severity of disruption, and the possibility of short to long-term psycho-social effects in the population.

RISK ANALYSIS - LIKELIHOOD

Whether it refers to a natural hazard or a malicious threat, likelihood assessment attempts to estimate the chance of an event occurring. The assessment method, however, is by necessity very different between the two major classes in the AHRA taxonomy:

- For non-malicious threats/hazards: Quantitative approach by which experts draw on historical data to determine the probability of a risk event as described in the scenario.
- For malicious threats: Through elicitation, experts provide qualitative judgment by considering overall capability (technical feasibility and enabling capabilities) and intent of the malicious actor(s) carrying out the threat.

Likelihood – non-malicious events

To provide estimates for the likelihood of non-malicious risk events, experts draw on historical data of comparable Canadian and international cases, as well as modeling and analysis. The likelihood estimate associated with the risk event as described by the scenario can be expressed by using the table below:

Likelihood Score	Estimated frequency, once every X years, where X is:
0	100,000
0.5	30,000
1	10,000
1.5	3,000
2	1,000
2.5	300
3	100
3.5	30
4	10
4.5	3
5	1

Table 5: Rating table for likelihood of non-malicious scenarios

As for the “people” as well as the “economy” category, in order to minimize round-off errors, the actual numerical estimates, expressed either as in the table above (i.e., one event every X number of years) or, alternatively, as annual probabilities or as the probability of the event happening in the next five years, are elicited during the assessment process; the ratings derived from those numbers are only used at the end, when estimating the overall risk for each risk event considered, as described by the scenario.

Likelihood –malicious events

For malicious scenarios, although records of past attacks are very useful, they cannot be the sole basis for estimating the likelihood of a possible future attack. What makes this category unique is the fact that malicious actors learn, both from past experiences and from advances in technology, as do security organizations trying to prevent them from carrying out attacks, and both sides continuously adapt and adjust their strategies and methods.

Estimating the likelihood of malicious scenarios is considerably different than for non-malicious ones, as these estimates must take into account the determined and adaptive nature of an intelligent adversary. Such an adversary will make a choice to carry out an attack based, on the one hand, on the statement they want to make, in accordance with the individual's or the organization's ideology. To capture this dimension, the current approach relies on the intelligence community to provide expert judgment on an individual's or organization's intent to carry out an attack, such as the one described in the scenario. On the other hand, the adversary's choice of an attack is also based on considerations of whether mounting an attack is technically feasible, as well as whether the adversary has adequate organizational and support means to carry it out. Again, the current approach relies on judgment from domain experts to assess various components of the technical feasibility of a malicious attack scenario, and on the intelligence community to provide expert judgment on whether an individual or organization has sufficient capability to carry it out. The combined assessments of feasibility, capability, and intent are used to generate an overall assessment or composite judgment of likelihood. Figure 3 shows the components and steps involved in producing an estimate for the likelihood of a malicious scenario.

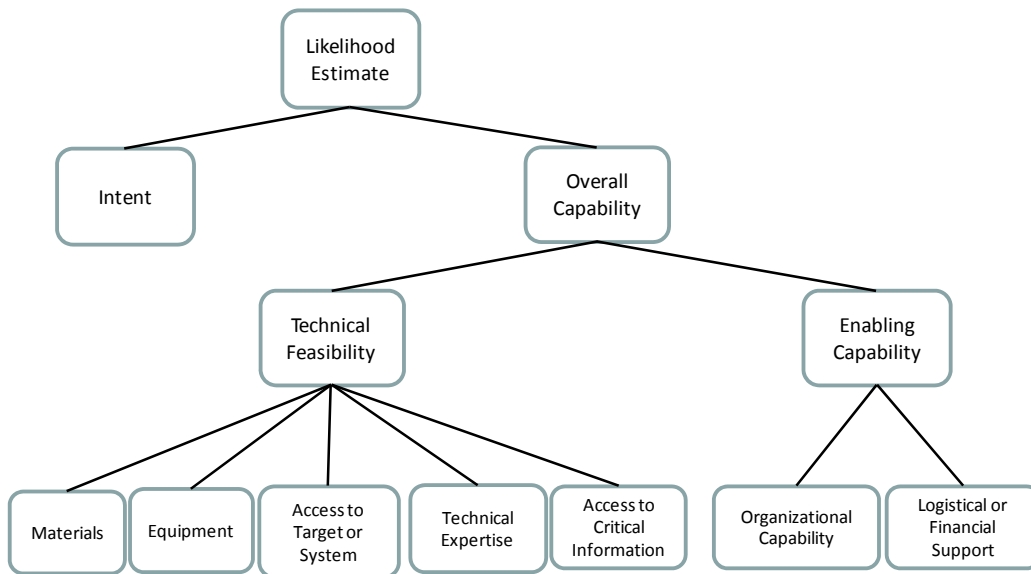


Figure 3: Malicious Likelihood – Components

The overall likelihood score is based on the principle of the “weakest link”, meaning that the final rating is determined by selecting the lowest component rating, across all components. A successful adversarial attack cannot occur if one of the elements is absent, lacking or unobtainable; in other words, an attack is assessed as unlikely if the level for one element of the overall capability of the malicious actor(s) is below a necessary level to being materialized, or if there is a lack of intent to carry it out.

The final estimate is expressed on a rating scale of 0 to 5, similar to the rating scale of non-malicious scenarios. However, the two estimates are very different in nature, and there are no grounds for the assumption that an identical rating indicates an equally likely malicious scenario compared to a non-malicious one. A more detailed description of the rating method for likelihood of malicious scenarios can be found in [10].

RISK EVALUATION – COMPARING RISK SCENARIOS

The goal of the risk evaluation stage in the AHRA process is to bring diverse risks into the same high-level view. It generally comprises the following steps:

- a. Determination of the risk magnitude (i.e. likelihood and consequences) for the risk.
- b. Aggregation and consolidation of risk assessment results into a whole-of-government AHRA.
- c. Production of selected AHRA communication products and graphical representations of results.

The determination of the overall magnitude of the risk associated with each scenario includes “aggregating” the impact assessments into a single measure of consequences, which, combined with the likelihood assessment, provided an estimate of risk for each scenario. Presented collectively, once the set of selected scenarios has reached “critical mass”, in terms of proper representation of the risk space, these risk estimates will generate a picture of “all hazards” risk to the federal government, and can be used to inform federal emergency management planning.

As discussed in previous sections, for those impact categories more amenable to a quantitative assessment, an “order of magnitude” rating approach was employed, using a logarithmic scale with half scores; for example, for the “People” category, for 1 fatality, the risk score is $s_{I_p} = 0$, since $1=10^0$, while 300 fatalities will produce a rating of $s_{I_p} = 2.5$, since $300=10^{2.5}$.

Rolling up the impact categories into a single composite score requires strong societal value judgements. Establishing equivalency of ratings in different categories is necessarily subjective. For example, the rating scale for the “Economy” assumes a monetary equivalent of 10 million dollars for one fatality, or, in other words, a risk event causing an economic loss of about 10 million dollars is about the same magnitude as a risk event causing a fatality. Although the use of such conversion factors is a sensitive and controversial issue, they are commonly used by government agencies to optimize government decisions and make cost-effective changes to policies. In most cases, the monetary value is based on what the society is willing to pay in order to save a human life; a common measure used to determine such value is the Value of Statistical Life (VSL). The research literature on the topic of monetary valuation of human life is vast, with a great variability in the values obtained from different studies, particularly in the United States [11], [12]; the variability in methods and values may be the reason why federal agencies use different monetary valuations [13].

According to one study published in 2009 [12], which looked at 40 studies from 9 countries, the average value of VSLs stands \$9.5 M and the median at \$6.6 M, expressed in units of \$US 2000; among the 40 studies, there were seven done in Canada, with an average VSL value of \$9.2 M and a median of \$4.0 M. In the US, one of the prolific academic contributors to the field (K. Viscusi, [14]) suggests that the value

of \$8.7 M (in \$US 2011) is appropriate, while in 2004, the US Office of Management and Budget instructed federal agencies to use values between US\$1 million and US\$10 million per life lost [13].

Considering such evidence from academic research and government practice alike, and considering currency exchange rates over recent years, an equivalency factor of C\$10 million per life lost was adopted for the AHRA.

Unfortunately, not much evidence was found in order to support equivalency of ratings with the remaining categories. During the rating workshops, the participants were prompted to examine whether a similar rating in a different category would indicate an event of similar magnitude. However, as it matures, AHRA would benefit from including dedicated “calibration workshops”, where expert judgement is sought in order to establish such conversion factors and calibrate ratings across all impact categories.

Assuming equivalency of ratings has been established (i.e., impact categories have been calibrated), one can calculate a “Consequence Score” S_C :

$$S_C = \log\left(\frac{1}{6} \sum_{i=1}^6 10^{S_i}\right) \quad (1)$$

Non-malicious likelihood tables are formulated as “one event every 10^{5-S_L} years”, where S_L is the associated likelihood score. The rating scheme was designed to produce the highest rating, 5, for the highest frequency considered, which is once a year or more frequent: e.g., $1/10^{(5-5)}$

As noted in a previous section, malicious likelihood scores, based on qualitative judgments, are also expressed on a 0 – 5 scale; however, the two scales are not necessarily equivalent. A separate “calibration” exercise was conducted to place the malicious scenarios on the same likelihood scale as the non-malicious ones; however, the results for the malicious scenarios are classified, and they are not included in the current paper.

The joint presentation of both the likelihood and consequence dimensions on a scatter plot provides a graphical means for presenting a high level view of diverse risks. This graphical depiction can be used to compare similar events (e.g., possible variants of a hurricane scenario), or very diverse events (e.g., pandemic influenza versus marine oil spill). Figure 4 presents such a scatter plot, where the isorisk contours are given by:

$$S_R = S_C + S_L \quad (2)$$

It must be noted that the uncertainty “bubbles” around the risk estimate for each scenario come from elicitation of uncertainty around each of the estimates provided by experts during the elicitation process (i.e., uncertainty around the likelihood estimate, as well as around estimates for each of the impact categories). To come up with an uncertainty interval around the “aggregated” consequence measure, the author simply looked at the impact estimate driving the consequence measure (due to the “order of magnitude” approach) and used its associated uncertainty; although more sophisticated calculations can be used, it was deemed sufficient for providing a rough measure of uncertainty around the estimates. The “bubbles” are simply PowerPoint ellipses with the uncertainty intervals for likelihood and consequences,

respectively, as minor and major axes. Having an illustration of uncertainty included on the graph, together with the point estimates for overall risk, was considered more important than improving the accuracy of the calculation for marginal improvements in uncertainty results. In the author's opinion, the important point was that the uncertainty be presented as part of the overall risk picture, rather than strive for a level of accuracy in uncertainty calculations which would not be supported by the low accuracy in elicited uncertainty estimates. For validation of results, the author collaborated with experts in the AHRA Interdepartmental Working Group on Risk Evaluation and Visualization.




Figure 4: High level view of AHRA results for non-malicious scenarios

Summary

In summary, this paper started out by highlighting the complex and interconnected safety and security risk landscape, both from a global as well as Canadian perspective. Then, the paper argued for an increased need for collective planning processes to manage risks to society derived from “all-hazards”, and discussed some of the challenges associated with such an endeavor. The biggest part of the paper was devoted to describing the Canadian Federal All-Hazards Risk Assessment, a collaborative effort within the Canadian Federal community aiming at developing a mechanism for a comparative assessment of risk events and to generate a picture of “all hazards” risk to the federal government, and to inform federal emergency management planning. The AHRA approach is scenario-based, where scenarios are used to provide context and specificity to identified sources of risk. A risk taxonomy is used to structure the complex and vast problem space, and to guide the scenario selection and development towards an appropriate sampling of the risk domain. Each scenario is assessed in terms of likelihood and consequences. Likelihood refers to the risk event's chance of occurrence; different methodologies are

used for its estimation, depending on whether the scenario involves a malicious or a non-malicious threat/hazard. Consequences of a risk event, as described by a scenario, are investigated along a number of dimensions. The scoring approach is an “order of magnitude” approach, and attempts to apply consistently across the impact categories. For those impact categories more amenable to a quantitative assessment, the scoring method uses a logarithmic scoring scale. Some of the impact categories, however, are less quantifiable; and for those categories, the rating was based on sets of qualitative criteria identified by the federal community.

The impact assessments, “aggregated” into a single measure of consequences and combined with the likelihood assessment, determine the magnitude of the overall risk associated with each scenario. Presented collectively, once the set of selected scenarios has reached “critical mass” in terms of proper representation of the risk space, these risk estimates generate a picture of “all hazards” risk to the federal government, and can be used to inform federal emergency management planning. Most importantly, the approach described in this paper attempts to build shared understanding of risks across organizations, and may be useful to multi-agency problems beyond the Canadian national context.

References

- [1] Global Risks 2012, World Economic Forum, http://www3.weforum.org/docs/WEF_GlobalRisks_Report_2012.pdf.
- [2] Rowe, W.D. (1977), *An anatomy of Risk*, Systems Engineering and Analysis Series (Chestnut, H., Ed.), John Wiley & Sons, 1977.
- [3] Haimes, Y.Y., *Risk Modeling, Assessment and Management*, 2nd ed., Systems Engineering and Management Series (Sage, A.P., Ed.), John Wiley & Sons, Hoboken, New Jersey, 2004.
- [4] International Organization for Standardization, “Risk management – Principles and guidelines on implementation”, Draft International Standard ISO/DIS 31000, 2008.
- [5] Verga, S., “All-hazards risk framework – an architecture model”, in S. Martorell, C. G. Soares and J. Barnett(Eds.), *Safety, Reliability and Risk Analysis: Theory, Methods and Applications*, vol. 1 (Proceedings of the European Safety and Reliability Conference, ESREL 2008, and 17th Society for Risk Analysis-Europe, Valencia, Spain), pp. 315-323, CRC Press/Balkema, Taylor and Francis Group, London, UK, 2008.
- [6] Lambe, P., *Organising Knowledge: Taxonomies, Knowledge and Organisational Effectiveness*, Oxford: Chandos Publishing, 2007.
- [7] Verga, S., “Intelligence Experts Group All-Hazards Risk Assessment Lexicon”, DRDC CSS Note, DRDC-Centre for Security Science-N-2007-001, 2007.
- [8] Davis, P.K., *Analytic Architecture for Capabilities-Based Planning, Mission-System Analysis, and Transformation*, RAND Monograph Report MR-1513-OSD, 2002.
- [9] Mathers, C.D., Vos, T, Lopez, A.D., Salomon, J, and Ezzati, M (ed.), *National Burden of Disease Studies: A Practical Guide*, Edition 2.0, Global Program on Evidence for Health Policy, Geneva: World Health Organization, 2001.

- [10] The All Hazards Risk Assessment Methodology Guidelines, 2011-2012, Public Safety Canada, Emergency Management Planning Unit, AHRA-ETR@ps-sp.gc.ca.
- [11] Viscusi, W. Kip and Aldy, Joseph E., “The Value of a Statistical Life: A Critical Review of Market Estimates Throughout the World”, *Journal of Risk and Uncertainty*, vol. 27, 2003, p. 5-76.
- [12] Bellavance, F., Dionne, G. and Lebeau, M., “The value of a statistical life: A meta-analysis with a mixed effects regressions model”, *Journal of Health Economics*, 28: 444-464, 2009.
- [13] Appelbaum, Binyamin, “As U.S. Agencies Put More Value on a Life, Businesses Fret”, *The New York Times*, New York, NY, USA, 16 February 2011.
- [14] Viscusi, W. Kip, “The Value of a statistical life,” Discussion Paper, n° 517, Harvard Law School, Cambridge, 2005.