Chapter 6 - Safety Goals

Introduction

In Chapter 1, we touched briefly on the topic of *safety goals.* A safety goal partially answers the question "how safe is safe enough?". In this mini-chapter, we will develop a worked example of how a safety goal might be derived, and then compare it to safety goals adopted by other organizations. We shall also point out some of the pitfalls of using a safety goal as the *only* safety criterion.

Basis of Numerical Safety Goal

In order to make meaningful decisions, a safety goal should be expressed in quantitative terms. A safety goal such as "Make the reactor as safe as possible" will mean different things to different people, and provide no guidance to the designer. A goal such as "The reactor must never have a severe accident" is probably physically impossible and sets up expectations that cannot be met.

Here is how one requirement tries to get around the public's desire to have absolute safety, and the technical fact that severe accidents are rare but possible. It is one of the requirements for the European Pressurized Reactor, taken from the French regulator's 2006 Annual Report[1]:

"Accidents liable to lead to significant early radioactive releases, in particular accidents involving high-pressure core meltdown, must for their part be 'practically eliminated'".

As is, this cannot be used by designers, who must seek further guidance to interpret what "practically eliminate" really means.

Here is a possible starting point for a safety goal which at least has numerical requirements:

"The annual risk of death to the most exposed member of the public due to accidents in a reactor should be small in comparison to his/her total risk of premature death."

Let's break this down. First, it tries to compare like for like. The consequence of an accident in a nuclear reactor is ultimately a (small) risk to an individual that his life might be shorter than if the reactor were not there, everything else being equal. The safety goal therefore compares the risk from the reactor to other risks to life expectancy. There is no point comparing radiation dose in Sieverts to chemical exposure in mg/ml since they do not share a common currency.

Second, it compares the risk from a nuclear reactor to *all* other risks of premature death. This is not at all obvious. One could set up the comparison to all other risks of electricity generation, or all energy generation, or all (average) industrial activity. In fact the first two would be much more logical, since they compare the risks of producing the same product (energy). The choice of all other risks as the benchmark reflects a social, not a scientific, basis: that nuclear power is "new" and "special" and therefore to reassure the public, its risk must be small compared to *all* other risks; and it allows one to reassure a person that nuclear power will have minimal negative impact on his/her life. One could have a safety goal which relates the risk of nuclear power to its *benefits*, say: the benefit/risk ratio should be no worse than other similar activities. Even this requires elaboration: does the risk of competing technologies include the effects of greenhouse gases and release of carcinogens to atmosphere? Does the comparison include only the production facility (the generating station) or does it include the whole fuel cycle, from mining to waste disposal? Note that there is no "correct" safety goal; you may well disagree with this one and propose a better one. However it is the basis for most of the safety goals declared to date.

Third, the safety goal (in this example) is for the most exposed *individual*. It does not consider social effects such as exposure to a large number of individuals, evacuation, land contamination, and effects on the environment such as on animals and plants. One can argue that the impact of Chernobyl due to evacuation of the surrounding population, banning of food, psychological trauma and disruption of many people's lives and livelihoods was far greater than the immediate harm to the public. However it can also be argued that if the plant had been designed and operated adequately to protect the most exposed members of the public, then society at large would also have been protected. Thus an assumption in our safety goal is that protection of the most exposed individual member of the public also provides sufficient protection for society at large and for the environment[a].

Fourth, the goal refers to members of the public, not workers in the plant. It is generally accepted (as discussed below) that people will 'trade' a higher risk to life and limb if it comes as part of their job - i.e., if there is a direct benefit as part of their acceptance of risk. There have been attempts to set (nuclear) safety goals for plant workers in nuclear power plants, but the risk to such workers is dominated by conventional industrial risk. Notwithstanding, statistics show across the board that the hazard to workers in the nuclear industry is much less than the industrial

---

[a]You might want to challenge this assertion. While it is generally true, certain chemicals can get concentrated in some animal food chains, posing a greater risk to some of the animals than to humans. This leads to another social question: should animals be protected to the same degree as humans? If so, does the protection apply to a species, a local group, or an individual animal? These are not idle philosophical questions since the answer chosen will form the basis of regulations on nuclear power and can influence the design, siting and operation.

average.

Fifth, the goal refers to the risk of nuclear power in isolation. If one did *not* have the nuclear power plant, its risk would be zero; but because one needs energy, it would be replaced by a coal plant, say, or by an energy shortage, both of which have their own risks which may be greater than the risks of the nuclear power plant. Examining nuclear power in isolation is actually fine, as long as the competing technologies do their risk calculations the same way, and that someone calculates the risk of 'doing without'. Just as there is a risk to having nuclear power, there is also a risk to *not* having it.

Derivation of Numerical Safety Goal

You will recall from Chapter 1 that the public risk from accidents in a nuclear power plant is predominantly from radiation, and that the effects of radiation on life expectancy can be quantified (even at low doses, as long as one assumes a dose/effect hypothesis such as the linear hypothesis). To reflect the two effects of radiation - acute (non-stochastic) and latent (stochastic), we can break down the safety goal proposed above into two sub-goals:

"The annual risk of *prompt* death to the most exposed member of the public due to accidents in a reactor should be small in comparison to his/her total annual risk of prompt death due to all accidents",

and

"The annual risk of *fatal cancer* to the most exposed member of the public due to accidents in a reactor should be small in comparison to his/her total annual risk of fatal cancer due to all causes."

In Canada in 1997, accidents were the fifth leading cause of death, over the whole population, at a rate of 27.6 deaths for every 100,000 people[2]. Thus the average person's risk of death from an accident is ~3 x 10^{-4} per year (note that the rate for males is almost double that for females).

We could then say that the risk from a nuclear power plant of premature death to this individual should be much less than 3 x 10^{-4} per year, say a factor of 100 (this may be too conservative, but we'll use it for illustration), or 3 x 10^{-6} per year. Since the only way of causing prompt fatalities to the public in a nuclear accident is via a core melt and failure of containment[b], this suggests that

[b]Recall however that many nuclear plants, like fossil plants, have chlorine on site to treat the cooling water, and its release could be hazardous outside the site boundary.

our safety goal should be:

"The likelihood of a large release from a nuclear power plant in an accident should be less than 3 per 10^6 reactor years".

We have now got a goal that can be used in design: it is relatively straightforward using PSA to calculate the likelihood of a large release (core melt plus failure of containment); to see (if the goal is exceeded) where the dominant contributors are; and to fix them if needed to meet the safety goal.

Note however that risk is not uniformly distributed; Table 6-1 shows risk (in the U.S.) as a function of occupation[3].

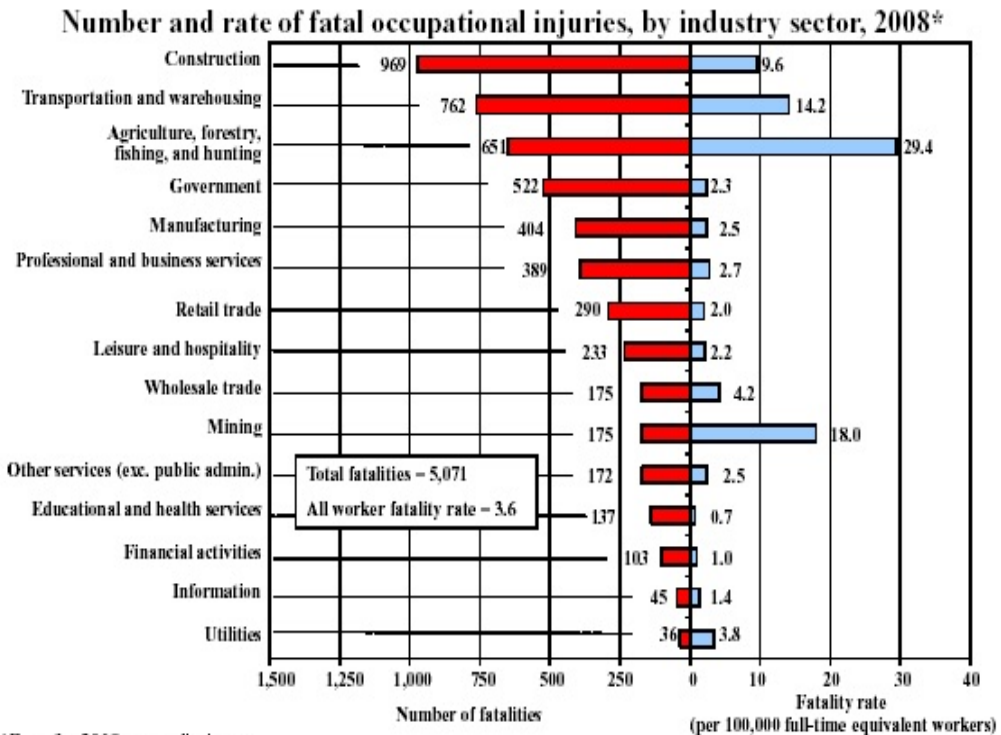

*Data for 2008 are preliminary.
NOTE: In 2008, CFOI implemented a new methodology, using hours worked for fatality rate calculations rather than employment. For additional information on the fatality rate methodology changes please see http://www.bls.gov/iif/oshnotice10.htm.
SOURCE: U.S. Bureau of Labor Statistics, U.S. Department of Labor, 2009.

Table 6-1 - NATIONAL CENSUS OF FATAL OCCUPATIONAL INJURIES IN 2008

For non-occupational activities, there is likewise a wide range of risk of death, although generally much lower than that for occupational activities, as in Table 6-2 below[4], from 2004:.

Table 6-2 - Cause of Death in Canada (Accident, non-Occupational)

Cause of Death	Mortality rate (/100,000-year)
Motor vehicle accidents	8.7
Falls	5.4
Poisoning	2.8
Homicide	1.7
Drowning	0.8
Fire	0.7

Note that lightning killed one person in Canada in 2005[5] - a risk of 0.003 / 100,000-year.

Cancer fatalities can be considered in a similar fashion. In the same year in Canada (2004) as the Figure above, malignant neoplasms were the second leading cause of death over the whole population, at a rate of 173 deaths per 100,000 people. Thus the average person's risk of dying from cancer is 1.73×10^{-3} per year (or about 13% over a 75-year lifetime). Recall from Chapter 1 that 100 person-Sv will produce about 5 fatal cancers in the exposed population, or a risk of 5×10^{-2} fatal cancers per Sv. The average person's annual risk of dying from cancer due to "natural" causes would be equalled by a dose (averaged over a large population) of 35 mSv per person per year. This is about 35 times natural background and almost twice the time-averaged annual occupational dose *limit*. It is also based on what is probably a flawed hypothesis - the linear dose-effect hypothesis. If we divide by 100 again, then the maximum time-averaged individual dose from accidents should be less than 0.35 mSv per year, averaged over a group of people, or about 35% of natural background radiation. The paradox with this safety goal is that it would make nuclear power safer than natural background radiation, so the goal may be somewhat too tight.

This is not as useful a safety goal as the previous one, since it does not tell us anything about the frequency distribution of accidents. However it can likewise be validated via summing *all* the events in a PSA which cause a release of radioactivity.

Other Safety Goals

In 1983, the Advisory Committee on Nuclear Safety (ACNS), an independent group of experts which advises the Canadian Nuclear Safety Commission, proposed risk-based numerical assessment criteria for nuclear electric generating stations. In their document ACNS-4[6], a probabilistic safety assessment of a nuclear reactor is required, and the consequences are to be grouped into six dose intervals (see Figure 6-1). The sum of the estimated frequencies of all accidents whose consequences fall within each dose interval must be less than a specified amount, as shown in the figure. One could infer that if one summed the maximum dose times the summed frequency in each interval of the histogram in Figure 6-1, then one could get an upper bound to the risk. However events below a frequency of 10^{-7} per year are not included, regardless of dose. If such events are assumed not to contribute significantly to risk, then ACNS-4 implies a maximum acceptable annual dose from accidents to the most exposed individual of 2.5 mSv per year, not too far off our safety goal model above.

In both the U.S. and the International Atomic Energy Agency (IAEA), safety goals have been set for both existing and future reactors. These goals state two criteria: the likelihood of a core melt, and the likelihood of a large release. For *existing* reactors, the goals are:

The frequency of a core melt (severe core damage) accident must be less than 10^{-4} per reactor-year

and

The frequency of a large release must be less than 10^{-5} per reactor-year.

For new reactors, the frequencies have been reduced by an order of magnitude each.

The second requirement above is equivalent to stating that the conditional containment failure probability following a severe core damage accident must be less than 0.1.

In some types of safety goals, the goal is expressed as a 'band' of either consequence or frequency. For example, in a given dose range, there may be a range of permissible summed frequencies. If the summed frequency is above the upper limit of the band, the result is unacceptable; if it is below the lower limit, it is acceptable; if it is in the band (between the two limits), then the designer must strive to lower the frequency or show why it is impractical to do so. Cost-benefit analysis is often used to demonstrate that further safety optimization is impractical.

The United Kingdom expresses its probabilistic safety goals[7] this way.

Maximum effective dose (mSv)	Total predicted frequency, per year	
	Basic Safety Limit	**Basic Safety Objective**
0.1 - 1	1	10^{-2}
1 - 10	10^{-1}	10^{-3}
10 - 100	10^{-2}	10^{-4}
100 - 1000	10^{-3}	10^{-5}
>1000	10^{-4}	10^{-6}

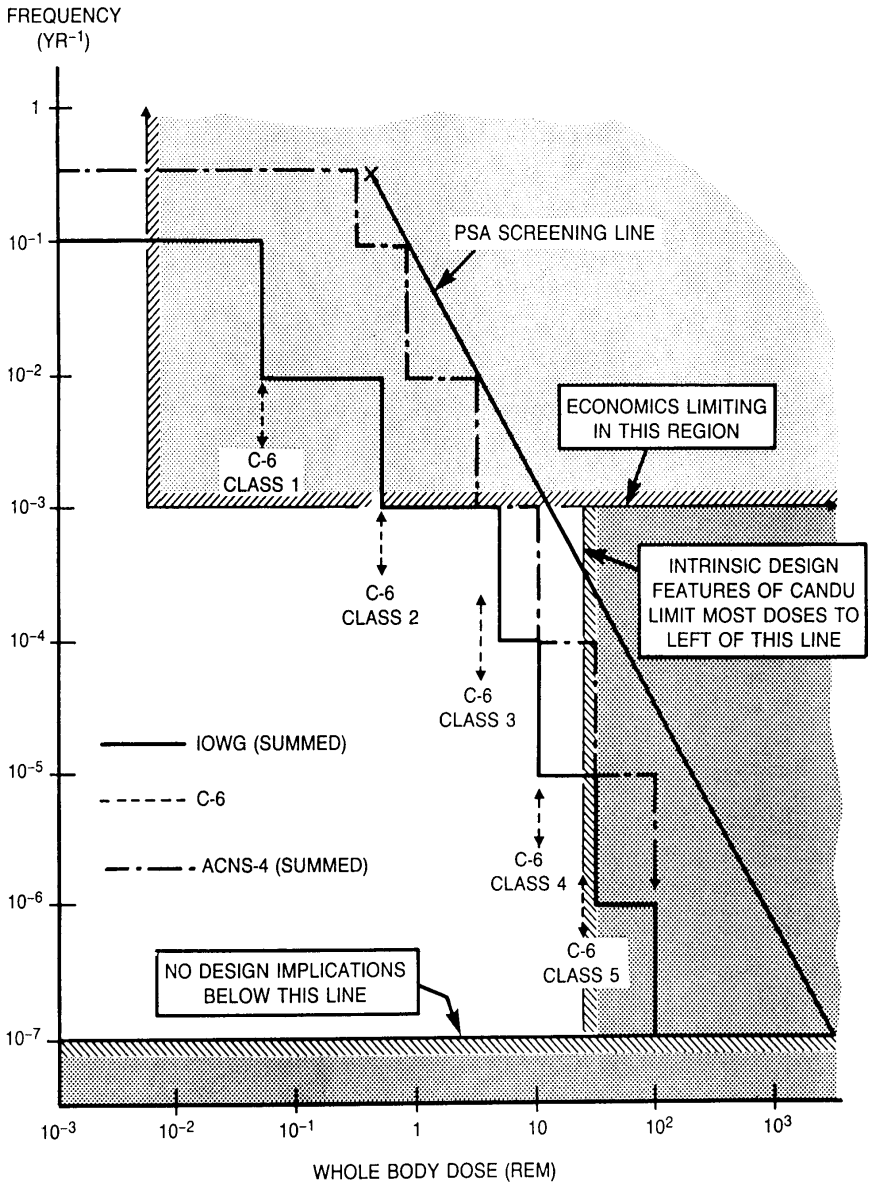

Figure 6-1 - Safety Goals in Canada

New Safety Goals in Canada

The CNSC has recently published safety goals[8,9] for new reactors built in Canada. The rationale is similar to what we have seen before. Specifically CNSC states:

"A limit is placed on the societal risks posed by nuclear power plant operation. For this purpose, the following two qualitative safety goals have been established:

1. Individual members of the public are provided a level of protection from the consequences of nuclear power plant operation such that there is no significant additional risk to the life and health of individuals; and

2. Societal risks to life and health from nuclear power plant operation are comparable to or less than the risks of generating electricity by viable competing technologies, and should not significantly add to other societal risks."

Further the CNSC define *three* safety goals:

1 **Core Damage Frequency:** The sum of frequencies of all event sequences that can lead to significant core degradation is less than 10^{-5} per reactor year

2 **Small Release Frequency:** The sum of frequencies of all event sequences that can lead to a release to the environment of more than 10^{15} becquerel of iodine-131 is less than 10^{-5} per reactor year. A greater release may require temporary evacuation of the local population.

3 **Large Release Frequency:** The sum of frequencies of all event sequences that can lead to a release to the environment of more than 10^{14} becquerel of cesium-137 is less than 10^{-6} per reactor year. A greater release may require long term relocation of the local population.

This is largely consistent with international practice for new reactors, discussed above, with one new concept - the small release frequency. CNSC has explained that it is intended to address those accident scenarios which may result in limited core damage, leading to small but significant releases. These accidents require emergency measures such as sheltering or short term evacuation of an area around the plant, and the small release frequency sets a limit on those.
.

Limitations of the Risk Approach

An approach to safety based *only* on a safety goal has a number of limitations:

- To determine compliance with a risk target, *all* significant events have to be identified and summed. It is difficult to do this summation early in the design, so that the designer is left without a yardstick until his design is mostly complete. Since the risk can be dominated by low-frequency severe accidents, involving multiple failures, cross-links and subtle dependencies, by the time design changes are identified, it may be too late to implement them.

- There is no *risk aversion* in a summed safety goal. Risk aversion proposes that risk should decrease with increasing event severity - based on increasing levels of uncertainty in predicting the consequences of such events, and perhaps on a perceived social aversion to disasters (even if the risk is small) - i.e., a weighting against large consequences regardless of frequency. Risk aversion is used by many regulators in setting safety requirements. An example is that aeroplane safety requirements far exceed car safety requirements, per passenger mile, due to people's fear of having many deaths at one time and place (aeroplane crash) rather than the same number spread out over many locations (car accidents), even if the risk and the number of deaths are the same. Note that putting in risk aversion is a social decision, not a scientific one; the example we gave of UK safety goals does *not* use risk aversion (how can you tell?).

- In principle there should be no frequency cutoff in proving that a safety goal is met. In practice analysis of events with frequency below about 10^{-8} per year is not very useful: in many cases the consequences are speculative, and frequencies that low are meaningless - e.g., comparable to one event in 100 times mankind's existence on the planet. There is however a reasonable upper limit on consequences, since one cannot get 'worse' than releasing all the mobile fission products outside containment. Such an approach was taken in the early days of reactor safety to bound the consequences of an accident. Highly pessimistic and unphysical assumptions resulted in predictions of about 50,000 prompt casualties. A more realistic upper limit is provided by Chernobyl, where the number of prompt casualties was 31, all on-site. Such debates do not however provide much useful guidance to the designer.

- Safety goals are meaningful only for events whose frequencies and consequences are reasonably calculable. In practice this includes most 'internal' events for which actual data exists or for which fault trees can be calculated; or where a reasonable extrapolation from a historical record can be made. However if the design has innovative features, with little operating experience, it may be difficult to support the reliability values and hard to

spot the cross-links. Passive safety systems pose a particular challenge in this regard since they can be difficult to test, and therefore it is hard to build up a reliability database.

- Not all (rare) events can be assigned a frequency and consequence with confidence, for example:
  - massive structural failure
  - massive failure of pressure vessels
  - very low-frequency, high-consequence external events such as earthquakes beyond historical record-keeping
  - sabotage, terrorism and war

  The approach to the first two is usually to design to accepted engineering codes and standards. Then from experience (mostly non-nuclear, since there is more of it), one can infer the likelihood of sudden failure of structures and components so designed. Failure of a LWR pressure vessel would be a catastrophic event, since it would lead to an immediate release of fission products and probably damage containment at the same time. Some calculations have been done to show that such massive failures are less frequent than 10^{-8} per year, but in the author's opinion, such calculations for a *single event* are meaningless and unsupportable (particularly since rare events can happen by sneak paths not anticipated - e.g. erosion of the pressure vessel wall in the Davis Besse plant). By the same token, historical records allow one to define the intensity of earthquakes down to about the thousand-year return frequency. More severe, rarer earthquakes are hard to characterize. One can do a "seismic margin" analysis to calculate the likelihood of survival of an earthquake somewhat more severe than the "Design Basis Earthquake" with the 1000-year return frequency; much beyond that, about all one can say is that the effects of damage to the nuclear plant would be small compared to the havoc wreaked by such an earthquake on the rest of society. Finally for events resulting from hostile human actions, the approach has generally been to design according to rule (e.g. the plant inherent defences plus the local security force should be able to delay an attack of *x* people armed with *y* type of weapons for *z* minutes); x, y, and z are indeed chosen based on reasonableness (likelihood) but the historical database of hostile acts against nuclear power plants is not good enough (fortunately) to support a true risk approach. In any case the defences being built into new plants for severe accidents are also helpful against malevolent acts.

For these reasons, those regulators that have safety goals use them *in addition to* whatever deterministic criteria they have developed. They do however provide a powerful rationality check

Exercises

1. Develop a set of high-level safety goals for a military use nuclear submarine. They can be, but do not have to be numerical. The most important part of your answer is to explain and justify it, not whether or not it matches someone else's 'official' goals. Consider any differences due to docked versus at-sea; and peacetime versus war.

2. Small reactors could be used in remote northern communities, for heating/electricity production. They would replace very expensive diesel generators, the fuel for which has to be flown in, whereas the reactor could be designed to be refuelled once in twenty years. Small reactors have also been used for powering unattended remote military installations. Propose safety goals for each case, with reasons.

3. Propose a way of translating the CNSC safety goals on page 9 into targets that a designer can use for pipe failure frequency, and reliability targets for shutdown, ECC, and containment. (Hint: how can one get a core melt? How can one get a large release?)

4. Develop a set of high-level safety goals for a nuclear-powered satellite. They can be, but do not have to be numerical. The most important part of your answer is to explain and justify it, not whether or not it matches someone else's 'official' goals.

5. Find out what safety goals have been used in the design of your reactor project and explain how they were obtained.

6. Calculate the risk to the public implied by the UK SAPs.

7. The International Atomic Energy Agency has a project (called INPRO) now underway, which is attempting to set requirements for reactors 50 years from now[10]. The top-level safety requirements were proposed by a group of experts to be as follows:

   "Installations of an Innovative Nuclear Energy System shall:

- Incorporate enhanced defence-in-depth as a part of their fundamental safety approach and ensure that the levels of protection in defence-in-depth shall be more independent from each other than in existing installations.
- Excel in safety and reliability by incorporating into their designs, when appropriate, increased emphasis on inherently safe characteristics and passive systems as a part of their fundamental safety approach.
- Ensure that the risk from radiation exposures to workers, the public and the environment

during construction/commissioning, operation, and decommissioning, shall be comparable to that of other industrial facilities used for similar purposes.

Further, the development of an Innovative Nuclear Energy System shall:

- Include associated RD&D work to bring the knowledge of plant characteristics and the capability of analytical methods used for design and safety assessment to at least the same confidence level as for existing plants."

Comment on these requirements from the points of view of:

- usefulness to designer
- meeting public safety requirements 50 years from now

References

1. "Annual Report 2006", Autorité de Sûreté Nucléaire, Section on EPR reactor project Safety

2. Statistics Canada, "Selected leading causes of death, by sex", 1997.

3. United States Department of Labor, "National Census of Fatal Occupational Injuries in 2008"

4. Statistics Canada, "Mortality, Summary List of Causes – 2004", Table 1-1: Deaths by selected grouped causes, sex and geography — Canada; Statistics Canada – Catalogue no. 84F0209.

5. Statistics Canada, Table 102-05401,2,3,4,5 - Deaths, by cause, Chapter XX: External causes of morbidity and mortality (V01 to Y89), age group and sex, Canada, annual (number), 2005.

6. "Recommended General Safety Requirements for Nuclear Power Plants", Atomic Energy Control Board Advisory Committee on Nuclear Safety, ACNS-4, June 1983.

7. United Kingdom Health and Safety Executive, "Safety Assessment Principles for Nuclear Plants", 2006.

8. "Design of New Nuclear Power Plants", CNSC Report RD-337, November 2008

9. I. Grant, T. Viglasky G. Rzentkowski and D. Miller, "Development of Licensing Basis for Future Power Reactors in Canada", presented at the Technical Meeting on "Experience Feedback on the Application of IAEA Safety Standards on Design of Nuclear Power Plants" November 7-11, 2005.

10. International Atomic Energy Agency, "Guidance for the Evaluation of Innovative Nuclear Reactors and Fuel Cycles", IAEA-TECDOC-1362, Vienna (2003).