



Short communication

Reassessing the safety of nuclear power[☆]Spencer Wheatley^a, Benjamin K. Sovacool^{b,c,*}, Didier Sornette^a^a ETH Zurich, Department of Management, Technology and Economics, Switzerland^b Department of Business and Technology, Aarhus University, Birk Centerpark 15, DK-7400 Herning, Denmark^c Science Policy Research Unit, University of Sussex, United Kingdom

ARTICLE INFO

Article history:

Available online 28 January 2016

Keywords:

Nuclear energy
Nuclear power
Nuclear safety
Nuclear accidents

ABSTRACT

We summarize the results of a recent statistical analysis of 216 nuclear energy accidents and incidents (events). The dataset is twice as large as the previous best available. We employ cost in US dollars as a severity measure to facilitate the comparison of different types and sizes of events, a method more complete and consistent than the industry-standard approach. Despite significant reforms following past disasters, we estimate that, with 388 reactors in operation, there is a 50% chance that a Fukushima event (or more costly) occurs every 60–150 years. We also find that the average cost of events per year is around the cost of the construction of a new plant. This dire outlook necessitates post-Fukushima reforms that will truly minimize extreme nuclear power risks. Nuclear power accidents are decreasing in frequency, but increasing in severity.

© 2015 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

It has been more than four years since an earthquake and tsunami caused an accident at the Fukushima Daiichi nuclear power plant in Japan resulting in repeated fires and three reported core meltdowns. At the latest count, the accident had caused \$166 billion in damages¹ [1] and at least 573 immediate deaths from the evacuation, along with hundreds of future deaths related to cancer anticipated to occur [2]. Somewhat sweeping industry reforms were called for, and public acceptance of the technology plummeted [3]. Supporters of nuclear power were quick to point out that a complete phase out would complicate efforts at mitigating greenhouse gas emissions from the electricity sector [4] and could lead to cumulative global losses in global gross domestic product [5].

The March 2011 Fukushima nuclear accident is a poignant reminder that disasters of enormous consequences can occur in the nuclear industry. But how often and with what severity? These two questions constitute the core of sound risk management, which requires identifying and quantifying such potential losses and their frequencies. For most natural and human-made catastrophes such as earthquakes, meteorites, avalanches, mountain collapses, forest

fires, hurricanes, epidemics, health care costs, war sizes, terrorist intensities, cyber risks, dam failures, industrial disasters, and so on, plentiful historical data has allowed scientists and engineers to determine the distributions of losses.

The admittedly favorable situation of a paucity of nuclear accidents, combined with scantily available public historical data, has prevented any such statistical analysis. Nuclear engineers have thus resorted to the classification of hypothetical accident scenarios deemed credible and of their potential consequences. The common industry approach to assessing nuclear accident risk depends on a technique known as probabilistic safety analysis (PSA), which assigns probabilities and damage values to particular failure scenarios. Nonetheless, such techniques are known to poorly predict events and to under-appreciate incidents that cascade into failures [6–11,12].

Similarly, the IAEA (International Atomic Energy Agency) provides the INES (International Nuclear Event Scale) to communicate the severity of nuclear accidents on a progressive discrete scale of 1 (anomaly) to 7 (major accident), meant to correspond to the amount of radiation released by order of magnitude. Yet its approach has been critiqued for offering relatively crude scores, for reporting only a fraction of known events, for not being transparent in its methodology, and for being more of a public relations tool (propaganda) than a meaningful metric [9,13]. For instance, there are about 12,000 events reported by French operators every year, of which 600–800 are classified annually as “significant for nuclear safety,” yet little to none of these show up on the INES database, and such unreported events occur at just 15% of the currently operating world nuclear fleet [14].

[☆] One of the authors of this paper is an editor for Energy Research & Social Science. They were not involved in managing the peer review process for this article.

* Corresponding author. Fax +45 3032 4303.

E-mail address: sovacool@vt.edu (B.K. Sovacool).

¹ Updated to US\$2013 and adjusted to monetize human fatalities. Originally reported as \$150 billion in \$2010 damages.

In this study, we summarize the results of a statistical analysis of a dataset of 216 events (incidents and accidents) occurring in nuclear energy systems [15], a dataset that is twice as large as any of the previous best ones available in the scientific literature [8,16], but we refrain from using the INES data directly. Instead, we use the estimated cost in USD (US dollars) as the common metric that allows one to compare often very different types of events across the nuclear fuel cycle. This dataset has more than three times the number of accidents compared with studies using solely the INES data, providing a much better basis for statistical analysis and inference, and a better comparative tool for reassessing the safety of nuclear power. Following Chernobyl, several authors proposed utilizing a monetary value of damage severity to make events comparable, and use a rate measure normalized by the number of reactor operating years to consider frequency [17–19]. This is what we have done here, but extending the range of analysis well beyond 1986 to include Fukushima and other nuclear events leading up until the end of 2014. The dataset has been published online², where the public is encouraged to review and recommend additions and modifications with the intention of continually expanding and improving the quality of the data.

2. Methods

There are many ways to quantify the risk of accidents in nuclear energy systems. The Farmer curve is one of the standard tools of nuclear risk assessment, with the risk defined as “probability $\times$ consequences” [20]. Typical Farmer plots display the annual frequency of fatalities or of property damage from human made sources of risk. Remarkably, the nuclear risks reported in Farmer plots are fundamentally different from all previously mentioned risks, in that the distributions for nuclear event losses are always thin-tailed and Gaussian-like, presenting a downward concave shape in the standard log–log representation.

The appearance of the Soviet Union’s Chernobyl accident in 1986 and of Japan’s Fukushima Daiichi nuclear power plant accident, after the tsunami on 11 March, 2011, seem at odds with the statistics implied by the Farmer curves. Actually, following the Chernobyl accident, Hsu [17] and Sengor [18,19] suggested a different approach, based on the reasoning that the number of fatalities is an incomplete, if not misleading, metric for measuring nuclear losses given the difficulties in assessing long term real mortality in addition to early morbidity and mortality. Indeed, this metric misses many other dimensions and also prevents quantitative comparisons. Hsu in particular made the point that the statistical analysis of earthquake risks, for instance, would have missed the fundamental Gutenberg–Richter magnitude–frequency law [21] if seismologists had focused on only the few large earthquakes. By considering a range of event sizes above which the data is known to be sufficiently complete, or at least representative, one can identify possible statistical regularities that are relevant to the largest events.

Here, we analyze the distribution of losses resulting from all possible types of nuclear events from 1952 to 2014. To be consistent with both the INES, as well as earlier peer-reviewed studies [8,9], we assessed events across the entire nuclear fuel cycle—that is, not only at nuclear reactors and power plants but also at uranium mills, fuel enrichment facilities, reprocessing stations, and nuclear waste repositories. In addition to maintaining consistency, this inclusion of non-reactor events is also necessary to trace the full impact of nuclear power technology on society as well as to account for the fact that many sites prone to accidents concentrate multi-

ple elements of the fuel cycle in one location.³ Searching historical archives, public utility commission filings, regulatory reports, and other sources explained in SM1, we created a unique dataset of 216 nuclear events, with 104 of these events having at least \$20 million in inflation-adjusted cost.⁴ In addition, whenever events had the same dependent cause, such as Fukushima, we treated them as a single occurrence. As it is important to evaluate the number of accidents relative to the number of reactors in operation, we have normalized our assessment to operational reactor data from the IAEA [22].

To be fair, a few caveats and limitations deserve mentioning. In this study, we focus only on damage and loss of life from nuclear accidents, and not other externalities such as lung cancer risks from coal mining or particulate pollution from petroleum-fuelled automobiles. Consequently, our study details the risks present from continuing to operate existing reactors, it does not assess the risks from not operating them (such as greater reliance on fossil fuels) [4]. Also, as is typically the case in data such as this, there is an event severity level below which events are less frequently reported, or even noticed—making our analysis conservative because of incomplete data. We base our analysis on the current reactor fleet, heavily tilted towards older light water reactors (often called “Generation II” technology), not state-of-the-art designs such as the European Pressurized Reactor or “paper” units at the conceptual stage such as small modular reactors, primarily because there is insufficient operating experience for their statistical analysis, but also since the adoption of these designs is uncertain. Our characterization of the current risk level, and its use for forecasting, presumes that 388 reactors remain in operation, and does not include any potential improvements in response to Fukushima. Any significant nuclear renaissance or massive build-out would alter our characterization, as would any massive phase-out. Lastly, we limit our assessment to nuclear generated electricity and its fuel cycle, and thus exclude risks posed by nuclear explosives and nuclear weapons, except for those facilities (such as reprocessing spent fuel) that are dual use.

3. Results and discussion

We quantify four identifiable dimensions of risk: (i) historical frequency of accidents, (ii) historical costs, (iii) the presence of so-called “dragon kings” and extreme events, and (iv) expected future costs.

In terms of frequency, panel (I) of Fig. 1 plots the number of events with at least \$20 million in damage (and standard errors) per reactor per year, calculated on 5 year windows spanning 1960 to 2014. The main message here is that the rate of events has dropped substantially since the 1960s, and may have stabilized since the late 1980s. In panel (II) of Fig. 1 the rate of events is calculated running away from the Chernobyl accident in both directions. From here it is clear there was a significant decline in event frequency after the Chernobyl accident, and the rate of events since that drop has been roughly stable, indicating that Chernobyl was a catalyst for change that decreased the rate of events, but not necessarily the size of each event. Rate estimates for 2014 remain in a conservative range of 0.0025–0.0035, or 1–1.4 events per year over the entire nuclear fleet. The methodology used here is described in SM2.

³ Sellafield in the United Kingdom, for instance, is home to commercial reactors, research reactors, waste repositories, and reprocessing facilities and Fukushima Daiichi in Japan was home to commercial reactors and waste repositories.

⁴ The analysis here is focused on events with at least \$20 million USD in damage. These events are more visible and thus the dataset is more likely to be complete above this threshold. Therefore statistics on this subset will be more reliable than when considering smaller events. Further, these large events are most relevant as they drive the total risk level. For instance, the ten most costly events contribute approximately 94% of total costs to date.

² See https://tasmania.ethz.ch/index.php/Nuclear_events.database.




Fig. 1. Frequency, damage, and severity of nuclear power accidents.

Panel (I) plots the rate (annual number of events with cost in excess of USD 20 million per reactor) calculated for five year periods, with standard error bars given.

Panel (II) plots the rate estimate running both backward and forward from the Chernobyl accident in 1987, bounded by standard errors.

Panel (III) plots the NAMS (Nuclear Accident Magnitude Scale) and damage/cost data according to their CCDF (Complementary Cumulative Distribution Function). The CCDF of the cost data is given for periods both before and after the accident at Three Mile Island (TMI) in 1979. The dashed line overlaid on the NAMS CCDF is an Exponential fit for the 15 largest points. The dashed line overlaid on the post TMI cost CCDF is a Pareto fit for the 90 largest points. The x marks indicate events that are outlying.

Panel (IV) plots the logarithm of damage/cost versus the INES (International Nuclear Event Scale) scores for each event. The median and quartiles of damage are given for each INES level with enough points. A linear regression fit to the scatterplot is overlaid.

In terms of historical severity, panel (III) of Fig. 1 plots both cost and the Nuclear Accident Magnitude Scale (NAMS) [9] according to a complementary cumulative distribution function (CCDF) described in SM3. As the figure demonstrates, the damage CCDFs corresponding to the periods of before and after the Three Mile Island (TMI) major accident of 1979 are different. It is most plausible that this change was a reaction to TMI, which involved both improving safety standards as well as reporting more events.

We also find that the heavy tailed Pareto distributions are insufficient to account for the extreme empirical tails in the sense that a few exceptional events are “outliers”, or better said, are dragon-kings, revealing the existence of transient amplification mechanisms. Such dragon-kings are found to “coexist with power laws in the distributions of event sizes under a broad range of conditions in a large variety of systems” [23]. As described in SM4, the presence of dragon-kings provides a diagnostic for the existence of causal factors behind accidents not apparent from the main Pareto model used for the distribution. The dragon-kings are shown with X marks in panel (III) of Fig. 1. The main point here is that post-TMI moderate severity events are suppressed but extreme events escalate to the extent that statistically significant dragon-kings emerge in both NAMS and damage, exhibiting a runaway disaster regime.

Next, bringing together models for rates and magnitudes, we quantify the current risk level for the existing nuclear fleet, which may be used as a status-quo characterization of the future risk level using the methods described in SM5. Presuming a low rate $\lambda = 0.002$, and without considering the effect of dragon-kings, the 0.99 quantile is \$54.3 billion, almost five times the estimated damage from Three Mile Island. Presuming the moderate rate $\lambda = 0.003$, with the dragon-king effect, this quantile is \$331.6 billion, which is almost double the estimated damage of Fukushima. In other words, there is a 1% probability each year that an accident occurs that leads to a loss of at least \$331.6 billion. Such large numbers do not appear to be taken into account in standard calculations on the economics of nuclear power [24]. Moreover, according to our analysis, with 388 reactors in operation, there is a 50% probability of a Fukushima-like event (or more costly) every 60–150 years, and a Three Mile Island event (or more costly) every 10–20 years.

Finally, panel (IV) of Fig. 1 compares our estimated costs with INES scores, indicating inconsistencies where events deviate from the exponential growth in cost qualified by the line in the logarithmic scale. The multitude of dots above or below the INES scale strongly suggest it fails to adequately capture the magnitude of events. For instance, Fukushima (the largest event) would need to have an INES score of 10.6 to be consistent. Further, there is

considerable uncertainty in the INES scores as evidenced by the overlapping costs.

4. Six conclusions and policy implications

Our study reveals six important conclusions about the risks of nuclear power. First, concerning event frequency, our analysis shows that the rate of civil nuclear accidents over time since 1952 decreased significantly from the 1970s, reaching what appears to be a stable level of around 0.003 events per plant per year. In this sense, nuclear power is getting safer, although this improvement could be offset by the construction and operation of many new facilities. We find concrete evidence of a history of learning from previous accidents within the industry, especially the significant reduction in event frequency after the Chernobyl accident in 1986, and a suppression of moderately large cost events after TMI.

Second, however, is that these past reforms, rather than minimizing risk, have apparently spawned the prevalence of dragon kings and accidents with major costs. Chernobyl and Fukushima are both such dragon kings, as they together represent 84 percent of the total damage in our dataset. The morphology of nuclear accident risk has altered from more frequent, less costly events to less frequent, more costly events.

Third, existing databases are woefully incomplete when it comes to the reporting of nuclear incidents and accidents. For instance, only half of the events in our database have INES scores, and thousands upon thousands of small events – but with the potential to cascade into larger ones – remain unreported. As the authors of [14] concluded, “many nuclear safety related events occur year after year, all over the world, in all types of nuclear plants and in all reactor designs and that there are very serious events that go either entirely unnoticed by the broader public or remain significantly under-evaluated when it comes to their potential risk.” A fully transparent, centralized source of reliable data on nuclear accidents is needed; one that enables planners, investors, and even nuclear regulators to better comprehend, and then weigh, nuclear risks. Such full disclosure will need to be balanced with the legitimate security concerns of the nuclear industry and the need to avoid promoting a culture of panic and hysteria.

Fourth, apart from being incomplete, industry standard tools such as the INES scale of the IAEA are inadequate and inconsistent at identifying and projecting nuclear accident risk, especially related to dragon kings. For the costs to be consistent with the INES scores, the Fukushima disaster would need to be between an INES level of 10 and 11, rather than the maximum level of 7. To use an analogy, the INES scale is like the antiquated Mercalli scale for earthquake magnitudes, which was replaced by the continuous physically-based Richter scale. Instead of INES, we recommend the use of continuous scales genuinely based on relevant physical variables (radiation emission as in NAMS) and/or economic metrics (dollar costs as proposed here) and that these scales be publicly disclosed for as many events as possible, including all of those in our database.

Fifth, we need to better understand “near misses,” “false negatives,” “minor mishaps,” and “residual risk” [14]. Our study has focused only on “extreme risk,” that is, accidents that precipitated at least \$20 million in damages, but an entire class of narrow escapes exist, unplanned or unanticipated events and warnings that never resulted in damage [25,26]. In the European Union, for example, legislation called the Seveso directive⁵ has emphasized, since 1982, the importance of near-misses for hazardous accidents on

land, especially in the oil and gas industry. A similar directive ought to be considered for the nuclear industry, and it requires a complete data set of both small and large events to properly quantify the frequency with which small events escalate into larger ones.

Sixth, future frequency and severity of accidents are perhaps unacceptably high. While the nuclear industry can be characterized by an impressive improvement in incident prevention and safety procedures, our thorough analysis of this new data shows that, when a nuclear event of at least \$20 million in damage occurs, the probability that it transforms into a catastrophe with damage larger than one billion dollars is almost ten percent. Under the status quo, we project at least one Fukushima-scale dragon king (or larger) accident with 50% probability every 60–150 years. And, more common but still expensive events of about \$20 million will occur with a frequency of about one per year—making accidents a relatively routine part of nuclear power's future.

In conclusion, although the frequency of events per reactor has become less common, the relative frequency with which events cascade into “dragon king” extremes is large enough that, when multiplied by severity, the aggregate risk to society is still very high. To effectively reduce this risk, the possibility of Chernobyl and Fukushima sized events needs to be better anticipated and then more effectively managed.

Acknowledgments

To solicit as much critical feedback as possible, we posted a working paper using a different methodology and dataset on this topic in April, 2015 available here <http://arxiv.org/abs/1504.02380>. The working paper does not replicate the data presented in this study, although it does host publicly our dataset so that readers and others can continually improve the robustness and completeness of its contents. We also thank seven anonymous reviewers for extremely helpful comments on earlier versions of this draft, as well as colleagues MV Ramana from Princeton University, Per Peterson from the University of California Berkeley, Mycle Schenider from the World Nuclear Industry Status Report, Mark Cooper and Peter Bradford from Vermont Law School, and Andy Stirling and Gordon MacKerron from the University of Sussex. Despite their input, the findings and conclusions in this study derive only from the authors.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at <http://dx.doi.org/10.1016/j.erss.2015.12.026>.

References

- [1] Reuters Special Report: Help wanted in Fukushima, October 25, (2013), available at <http://www.reuters.com/article/2013/10/25/us-fukushima-workers-specialreport-idUSBRE99004320131025>.
- [2] John E. Ten Hoeve, Mark Z. Jacobson, Worldwide health effects of the Fukushima Daiichi nuclear accident, *Energy Environ. Sci.* 5 (2012) 8758.
- [3] Lei Huang, Ying Zhou, Yuting Han, James K. Hammitt, Jun Bi, Yang Liu, Effect of the Fukushima nuclear accident on the risk perception of residents near a nuclear power plant in China, *PNAS* 110 (49) (2013) 19742–19747.
- [4] P.A. Kharecha, J.E. Hansen, Prevented mortality and greenhouse gas emissions from historical and projected nuclear power, *Environ. Sci. Technol.* 47 (2013) 4889–4895.
- [5] Nico Bauer, Robert J. Brecha, Gunnar Luderer, Economics of nuclear power and climate change mitigation policies, *PNAS* 109 (42) (2012) 16805–16810.
- [6] D. Lochbaum, Nuclear plants risk studies: failing the grade, *Union of Concerned Scientists Report* (2000).
- [7] Laurent Stricker, Nuclear safety chief calls for reform, *Nature* 274 (April 7343) (2011).
- [8] D. Sornette, T. Maillart, W. Kroger, Exploring the limits of safety analysis in complex technological systems, *Int. J. Dis. Risk Reduct.* 6 (2013) 59–66.
- [9] D. Smythe, An objective nuclear accident magnitude scale for quantification of severe and catastrophic events, *Phys. Today: Points View* (2011).
- [10] M. Ha-Duong, V. Journe, Calculating nuclear accident probabilities from empirical frequencies, *Environ. Syst. Decis.* 34 (2014), 2.

⁵ <http://ec.europa.eu/environment/seveso/>.

- [11] L. Escobar Rangel, F. Leveque, How Fukushima Dai-ichi core meltdown changed the probability of nuclear accidents? *Saf. Sci.* 64 (2014) 90–98.
- [12] Wolfgang Kroger, Didier Sornette, Reflections on Limitations of Current PSA Methodology, ANS PSA 2013 International Topical Meeting on Probabilistic Safety Assessment and Analysis, Columbia, South Carolina, USA, September 22–26, 2013, on CD-ROM, American Nuclear Society, LaGrange Park, IL (2013), invited article for the Probabilistic Safety Analysis 2013 (PSA2013) (accepted 05.07.13) (<http://www.psa2013.org>).
- [13] G. Brumfiel, Nuclear agency faces reform calls, *Nature* 472 (7344) (2011) 397–398.
- [14] Georgui Kastchiev et al., Residual Risk: An Account of Events in Nuclear Power Plants Since the Chernobyl Accident in 1986 (May, 2007).
- [15] S. Wheatley, B.K. Sovacool, D. Sornette, Of disasters and dragon kings: a statistical analysis of nuclear power incidents & accidents, *Risk Anal.* (2016), <http://dx.doi.org/10.1111/risa.12587>.
- [16] M. Hofert, M.V. Wuthrich, Statistical review of nuclear power accidents, *Asia-Pac. J. Risk Insur.* 7 (1) (2013) 1–18.
- [17] K. Hsu, Nuclear risk evaluation, *Nature* 328 (1987) 22.
- [18] A. Sengor, Evaluating nuclear accidents, *Nature* 335 (1987) 391.
- [19] A. Sengor, Predicting the threat of nuclear disasters, *Nature* 472 (2011) 169.
- [20] Bal Raj Sehgal, Light Water Reactor (LWR) Safety, *Nuclear Eng. Technol.* 38 (8) (2006) 697–732.
- [21] C.F. Richter, *Elementary Seismology*, W.H. Freeman, New York, 1958.
- [22] International Atomic Energy Agency Nuclear power reactors in the world 2010. Reference data series no. 2. 2010.
- [23] D. Sornette, Dragon-Kings, Black Swans and the Prediction of Crises, *Int. J. Terraspace Sci. Eng.* 2 (1) (2009) 1–18.
- [24] World Nuclear Association, The Economics of Nuclear Power. 16.02.15, available at <http://www.world-nuclear.org/info/Economic-Aspects/Economics-of-Nuclear-Power/>.
- [25] T.W. van der Schaaf (Ed.), *Near Miss Reporting as a Safety Tool*, Butterworth-Heinemann, 1991.
- [26] W. G. Bridges, Gains from Getting Near Misses Reported, Presentation at 8th Global Congress on Process Safety, Houston TX April 1–4, 2012 (Process Improvement Institute, Inc., 2012).