# What Is The Risk Of Nuclear War?

*Presentation given at the Vienna Conference on the Humanitarian Impact of Nuclear Weapons, 8 December 2014.*


## Seth Baum, Ph.D.

Executive Director, Global Catastrophic Risk Institute
http://gcrinstitute.org • http://sethbaum.com • seth@gcrinstitute.org • @SethBaum


**Presentation Slides**
Available at:
http://www.bmeia.gv.at/fileadmin/user_upload/Zentrale/Aussenpolitik/Abruestung/HINW14/Presentations/HINW14_S2_Presentation_Seth_Baum.pdf

**Presentation Video**
Available at:
https://www.youtube.com/watch?v=MTzbIE69Q4U
Presentation begins at 49:40 in the video.


**Abstract**
It is wrong to assume that just because nuclear war has never happened before, the probability of nuclear war is insignificant. To the contrary, risk analysis shows that deterrence can fail and nuclear war could occur. Furthermore, even if the probability of nuclear war is low, the consequences of it happening would be so severe that it makes for a significant risk. The risk also increases over time. Taken together, this underlines the urgency of actions to reduce the possibility, especially as humanity may never get a second chance if nuclear war occurs.


**Speaker Bio**
Seth Baum is Executive Director of the Global Catastrophic Risk Institute (gcrinstitute.org), a nonprofit think tank that Baum co-founded in 2011. Baum's research focuses on risk, ethics, and policy questions for major threats to human civilization including nuclear war, global warming, and emerging technologies. He is based in New York City.

**Correction added 26 February 2015:** The talk gives the wrong numbers for the probability of Russia-US inadvertent nuclear war. The talk says: "We found the probability to be somewhere from once per century to once per 10,000 years." However, the correct range is once per 15 years to once per 100,000 years. The correct range implies more uncertainty in the probability of Russia-US inadvertent nuclear war. The corrected range also implies a significantly higher "best guess" probability of around once per century, instead of around once per millennium as the talk suggested. In plain terms, Russia-US inadvertent nuclear war is more uncertain and more likely than the talk implied. I regret this error and apologize for giving the wrong impression.

The corrected numbers are obtained by taking the reciprocal of the annual probabilities 0.07 and 0.00001, which are the higher and lower numbers for the 90% confidence ranges presented in the underlying research. These annual probabilities are found in Barrett et al., 'Analyzing and reducing the risks of inadvertent nuclear war between the United States and Russia', *Science & Global Security*, vol. 21, no. 2 (2013), pages 106-133. The relevant text begins at "There is significant uncertainty in the model-estimated annual inadvertent nuclear war probability…" It appears in page 120 of the journal version of record or p.11-12 of the online preprint: http://sethbaum.com/ac/2013_NuclearWar.pdf.

**Annotated Presentation Text**

[slide 1] I thank the Government of Austria for inviting me. It is an honor to be here.

First, a disclaimer. The views presented here are my views alone, and not the views of the Global Catastrophic Risk Institute or anyone else.

[slide 2] When I was 17, I almost died in a car crash. Indeed, if I wasn't wearing a seatbelt, I probably wouldn't be here today. So there are many risks all around us. There is no shortage of things for us to worry about. Risk analysis can help us identify which risks are most important and what the best opportunities are to reduce risks. In simplest terms, risk is the possibility of something bad happening. It is measured as the probability of that bad thing happening multiplied by the severity of the consequences if it does.[1]

[slide 3] Professionally, I work on nuclear war risk, because this is one of the largest risks and has some of the best opportunities to reduce risk—certainly much more so than car crash risk.

That is quite a bold claim, saying that nuclear war risk is larger than car crash risk. Car crashes really are a large risk. Each year about one million people die from car crashes worldwide.[2] In contrast, the only people who have ever died from nuclear war are those killed by that bomb dropped on Nagasaki and the other one on Hiroshima. These were great tragedies, for sure. But that's about 200,000 people across the entire 70 year history of nuclear weapons.[3] That's quite small compared to one million people per year from car crashes.

But these numbers are misleading. During World War II, there only were those two nuclear weapons, plus one used for testing. Today, there are 16,000.[4] If a nuclear war happens now, millions of people could die from the initial explosions, and potentially billions could die from the global nuclear famine that would follow, as Michael Mills discussed earlier today.[5] The worst case would see the permanent destruction of global

---

[1] A classic paper on the definition of risk, which is still relevant today, is Stanley Kaplan and B. John Garrick, 'On the quantitative definition of risk', *Risk Analysis*, vol. 1, no. 1 (1981), pages 11-27.

[2] The World Health Organization puts the exact number of car crash deaths per year at 1.24 million. See *Global Status Report on Road Safety 2013*, http://www.who.int/violence_injury_prevention/road_safety_status/2013/en. The distinction between one million and 1.24 million is not important for this talk.

[3] See e.g. Atomic Archive, *The Atomic Bombings of Hiroshima and Nagasaki*, http://www.atomicarchive.com/Docs/MED/med_chp10.shtml.

[4] See Federation of American Scientists, Status of World Nuclear Forces, http://www.fas.org/programs/ssp/nukes/nuclearweapons/nukestatus.html.

[5] Michael Mills' talk is available at http://www.bmeia.gv.at/fileadmin/user_upload/Zentrale/Aussenpolitik/Abruestung/HINW14/Presentations/

human civilization or even human extinction, which as far as I can tell is entirely possible.[6] In this case, the impacts would not just be to those who die in the war. It would also be the countless members of future generations who would be denied the opportunity to ever live.[7] Car crashes simply cannot do this.

But even this does not make nuclear war the larger risk. If nuclear war could never occur, then it doesn't matter how severe the impacts are. And there are people who say that, thanks to nuclear deterrence (which, by the way, didn't exist during World War II), the probability of nuclear war is basically zero.[8] We should not believe this. Look, nuclear deterrence works. The problem is, it does not always work. And when the consequences could be so severe, that's a big problem.

There's a it of a paradox here. A large nuclear deterrence failure, resulting in a large nuclear war, could potentially kill us all, in which case we wouldn't be alive to observe it, except for that brief moment when we're dying. The fact that we're here having this conversation requires that no nuclear war has ever happened.[9]

But there's more to it than that. Historical experience shows that deterrence can fail. The Cuban missile crisis is perhaps the clearest example of this. We made it out of that due in no small part to pure, simple luck.[10]

[slide 4] And then there is inadvertent nuclear war. This is when one side misinterprets a false alarm as a real attack and launches nuclear weapons in what it believes is a counterattack, but is in fact the first strike.[11] Inadvertent nuclear war is important because it means nuclear war could happen even if deterrence works perfectly —and it doesn't—but even if it does, here the other side actually was deterred, yet the one side believed it was under attack anyway. Indeed, according to deterrence theory, if you believe you're under attack, then arguably you're supposed to launch in response. It's that commitment to launching in response that creates the deterrence effect in the first place.

My colleagues and I studied the probability of inadvertent nuclear war between the United States and Russia.[12] We used historical data like the four cases shown here and

---

HINW14_S1_Presentation_Michael_Mills.pdf. See also papers available on Mills' website, http://acd.ucar.edu/~mmills, and the website of Alan Robock, http://climate.envsci.rutgers.edu/nuclear.

[6] I discuss these issues in Timothy M. Maher Jr. and Seth D. Baum, 'Adaptation to and recovery from global catastrophe', *Sustainability*, vol. 5, no. 4 (April 2013), pages 1461-1479, http://sethbaum.com/ac/2013_AdaptationRecovery.html.

[7] A classic discussion of this point is found in Carl Sagan, 'Nuclear war and climatic catastrophe: Some policy implications', *Foreign Affairs*, vol. 62, no. 2 (Winter 1983), pages 257-292.

[8] For example, international relations scholar Kenneth Waltz claimed that 'The probability of major war among states having nuclear weapons approaches zero'. See Kenneth N. Waltz, 'Nuclear myths and political realities', *American Political Science Review*, vol. 84 (1990), pages 731-745, at page 740.

[9] This paradox is known as anthropic bias or the observation selection effect. See e.g. Milan M. Ćirković, Anders Sandberg, and Nick Bostrom, 'Anthropic shadow: Observation selection effects and human extinction risks', *Risk Analysis*, vol. 30, no. 10 (2010), pages 1495-1506.

[10] See e.g. Ward Wilson, *Five Myths About Nuclear Weapons*, Houghton Mifflin, 2013; Benoît Pelopidas, 'Remembering the Cuban missile crisis, with humility', European Leadership Network, 11 November 2014, http://www.europeanleadershipnetwork.org/remembering-the-cuban-missile-crisis-with-humility_2118.html.

[11] Anthony M. Barrett, Seth D. Baum, and Kelly R. Hostetler, 'Analyzing and reducing the risks of inadvertent nuclear war between the United States and Russia', *Science & Global Security*, vol. 21, no. 2 (2013), pages 106-133, http://sethbaum.com/ac/2013_NuclearWar.html.

[12] Ibid.

modeled the process of going from false alarms to the decision to launch nuclear weapons in response. We found the probability to be somewhere from once per century to once per 10,000 years.[13] This range is so large because this is a very uncertain risk. It's not like car crash risk, which has excellent data. Once per century to once per 10,000 years is still fairly rare, but it's not zero. And this is just for one type of nuclear war between one pair of countries. The total probability across all types of nuclear war and all countries will be higher.

[slide 5] For sake of discussion, let's take a number somewhere in the middle of that range for the total probability of nuclear war, say, one per 700 years. And suppose that 700 million people die in the war. Then on average, one million people die each year from nuclear war, which is the same rate that people die from car crashes.

These numbers—700 years and 700 million deaths—are entirely reasonable estimates for the probability and severity of nuclear war. But they are not the only reasonable estimates we could make. Perhaps these numbers are too pessimistic, in which case the risk is smaller. Perhaps. However, my professional opinion, having studied this, is that both numbers are too optimistic, that nuclear war risk is larger than this. That means each of us is on average more likely to die from nuclear war than from car crashes. I say on average because the risk will still vary from person to person.

And this is just the risk to the present generation. The risk to future generations from nuclear war is much larger, because only nuclear war threatens the existence of future generations. That is why nuclear war is the much larger risk.

Ultimately, what matters is not the risk itself, but the opportunity each of us has to reduce the risk. Here nuclear war risk wins easily. Now, I imagine there are some disenchanted diplomats and activists in the audience who have a hard time believing that nuclear war is an easy risk to reduce. But in comparison, it really is. Car crash risk comes from millions of drivers scattered all over the world. In comparison, only a handful of countries have nuclear weapons, and the weapons are controlled (more or less) by central governments. That means we can come together in international conferences like this and address a large portion of the total risk all at once.[14]

Indeed, when I look at this conference, and the broader initiative it's part of, and other activities across the international community, I'm honestly optimistic. With enough effort, this seems like a solvable problem.

I'll make one last point. One million deaths per year, 700 million deaths per 700 years—the risk increases with time. In other words, the longer we wait, the more likely nuclear war is to occur. And it's not like car crashes, where one person dies, and the rest of us learn from it, adjust traffic laws and so on. Nuclear war could kill all of us at once. We may get no second chance. We cannot learn from experience. We must get it right the first time, every time. Nuclear war is a matter of life and death for all of us. The longer

[13] Actually, the range of probabilities is even wider than this. About 90% of the probabilities we found fell within this range. I believe it is reasonable to assume that the probability is within this range. UPDATE, 26 FEBRUARY 2015: PLEASE SEE THE CORRECTION AT THE TOP OF THIS DOCUMENT.

[14] It was quite remarkable speaking these lines at the HINW14 conference. In my mind, as I'm saying them, I'm seeing the risk decrease by some nontrivial amount on account of the conference. Here lies another paradox of nuclear war risk analysis. As we analyze the risk, and in particular as we share the analysis with the international community, the risk changes, so the analysis is no longer correct. I don't have a good estimate for how much HINW14 reduced nuclear war risk, in part because it depends on how HINW14 influences future events. But the conference certainly seems worth the effort, and I am indeed honored to have contributed to it.

we wait, the more likely we are to fail. So why should we address nuclear war risk now? Because we are still alive.

[Slide 6] Thank you.