

Appendix: Additional nuclear risk research ideas

This is an appendix to my series of posts on nuclear risk research ideas. This appendix contains additional ideas that didn't make it into that post. I strongly recommend you read the [summary & introduction post](#) before reading this appendix.

The first section contains ideas that didn't make it into the main series of posts simply because I ran out of time to look into or explain them in as much detail or with as much clarity as the ideas in the series were described. I have no particular reason to believe the ideas in the first section are notably lower priority than those in the main series.

The second section contains ideas that I tentatively believe are lower priority than those in the main series and in the second section of this appendix. However (a) I could be wrong, and (b) a project could be lower priority for most readers but higher priority for a particular reader due to considerations of personal fit for a particular reader. I therefore decided to flag these ideas anyway and let people choose for themselves whether to read, skim, or ignore them.

The final section links to other collections of research project ideas.

Many parts of this appendix are essentially just quickly written rough notes.

Ideas that are simply less polished

These are in very approximately descending order of how high-priority I see them as, accounting for importance, tractability, and neglectedness.

Given conflict, how many warheads and what targets?

What is this idea? How could it be tackled?

I started researching this topic and have notes in [this doc](#). A project here could involve building on what I have so far.

Some relevant previous work

- [Ladish's \(2019\) "Does the US nuclear policy still target cities?"](#)

Other people to consider talking to

- A person at RAND I know of (could probably share their name on request)

What places might be targeted in a countervalue, counterforce, and/or tactical nuclear strike (given various pairs of countries)?

Rough notes:

- Has anyone collected together best guesses on this?
 - [FLI did](#), but the info was from the 1950s or something
- If no one has done this, then it could be good for someone to do it?
 - CERI fellow or equivalent?
 - Maybe suggest it to Nuclear Notebook?
 - Maybe suggest someone does it with an Open Phil or LTFB grant?
 - Would this be harmful in some way?
 - Destabilising?
 - But I think the info we'd get is probably no better than what militaries probably already have in classified docs?
 - Makes militaries, governments, etc., hate EA?
 - Seems unlikely?
- [Update as of 6.3.22: I think some people are now doing or are interested in some version of this in reaction to the Russian invasion of Ukraine.]

Risk analysis for country pairs/combo from [Which nuclear wars should worry us most?](#) other than US-Russia

- The highest priority would probably be China-US, China-Russia, China-India, and India-Pakistan
- Luisa Rodriguez, Will Aldred, and I have each done some work on this and have drafts/notes that could probably be shared with someone interested in picking the topic up
- These analyses could separately cover (a) the odds of each pair/combo of countries entering into nuclear conflict and (b) the expected consequences if that occurs

What would be the ideal size and composition of nuclear arsenals?

What is this idea? How could it be tackled?

I started looking into this topic and have notes in [this doc](#). A project here could involve starting out by following my skeleton plan.

Different ways nuclear conflict could start and play out

Such a project could include forecasts - especially of how many warheads would end up used if each way/path occurs, soliciting forecasts, and expert elicitation.

What is this idea? How could it be tackled?

I started looking into this topic and have notes in [this doc](#). A project here could involve starting out by following my skeleton plan.

What sort of person might be a good fit for this?

I expect any good generalist researcher could make progress on this cluster of questions. I expect someone to be a stronger fit the more they already know about various things relevant to nuclear risk (since this project would ideally address a diverse array of variables and pathways) and the more experience they have with forecasting and expert elicitation.

Causal diagram of the various possible paths from nuclear weapons to harms (or benefits) to the long-term future

See "[Causal diagrams of the paths to existential catastrophe](#)" (Aird, Shovelain & Kristoffersson, 2020) for examples of causal diagrams.

Why might this research be useful?

I and other longtermists often assume that it's far more important to prevent very unusually large-scale nuclear conflicts (e.g., multiple thousands or >10,000 warheads used, with a large portion against cities) or improve resilience to and recovery from those than to prevent other nuclear conflicts, e.g. a smaller India-Pakistan exchange.

But maybe the latter would have quite bad effects via *responses* to it, e.g. increasing the chance of other conflicts or the development of other WMDs or making geopolitical situations less conducive to cooperation or international governance on e.g. AI issues. This doesn't have to be very likely to be a relatively big deal, since relatively direct x-risk from even large-scale nuclear war seems decently unlikely anyway, and most individual contributors to x-risk seem decently unlikely.

But I don't know if this is at all plausible - maybe we should even actually be neutral about whether a small scale exchange improves or worsens the long-term future. So an initial stab at that could inform prioritisation between nuclear risk and other areas and prioritisation of intermediate goals and interventions within the nuclear risk area.

Further notes:

- Could draw on Stephen Clare's diagrams in [Modelling Great Power conflict as an existential risk factor](#), and maybe Kit Harris's work.
- Also overlaps with GCRI's diagram(s) in [A Model For The Impacts Of Nuclear War](#)

Might nuclear weapons play an important role in the aftermath of a GCR/collapse scenario?

Rough notes:

- I seem to recall Ord alluding to this in The Precipice. Though I think maybe he mostly dismisses it as not a big deal
- How easy would it be for random non-state actors to target, launch, and/or detonate nukes, if there was state collapse?
 - Maybe easy, if the main barrier normally is armed guards, physical barriers they could get past easily if they had time unobserved, etc.?
 - Maybe hard, if there are physical barriers that would stop them getting in and that they wouldn't be able get past just because of state collapse, or if launching, etc., is itself a very technically challenging process and they wouldn't be able to find instructions?

Risks of nuclear war triggering other catastrophes, scarring of humanity's values, or similar

Another framing: **“How much should we worry about nuclear risks other than direct extinction risks?”** E.g., risks of nuclear war triggering other catastrophes (such as dangerous development and/or deployment of other techs), further conflict, scarring of humanity's values / lock-in of bad political systems, or similar.

Something that could be part of this or a separate question is **“How would nuclear conflict affect the odds of future nuclear conflict?”**

Miscellaneous notes:

- I'm excited about work on roughly this
 - But it could instead be done for war more generally
 - See also Stephen Clare's diagrams in [Modelling Great Power conflict as an existential risk factor](#)
 - See also some related questions in [Politics, Policy, and Security from a Broad Longtermist Perspective: A Preliminary Research Agenda](#)

- Would this be very similar to just trying to answer the question in relation to GCRs/great power wars/collapse more generally (similar to Luisa Rodriguez's work [here](#) and the other work she was doing on this kind of thing)?
- Luisa said she thinks it could be useful to write up an example scenario to try to describe really concretely how one goes from the initial catastrophe to a later catastrophe, or to values change
 - So you can see what assumptions are being made
 - E.g., this group finds this bomb, then deploys it, motivated by X
 - She thinks lots of people have assumptions - like “obviously people would cooperate” or “obviously there'd be conflict.” But it'd be easier to generate intuitions if someone has looked at a particular scenario
 - So she thinks this could be a useful complement to her more scenario-neutral work
- [“It's my job to worry about any way nukes could get used” - 80,000 Hours](#)
 - “Robert Wiblin: So there’s a lot of different scenarios that you might worry about. You might worry about this dirty bomb case. I guess you’ve got like small or medium scale or nuclear exchanges, the kind of thing you might get with North Korea, or between India and Pakistan. And then you’ve got that just all out nuclear exchange between China and America or America and Russia.
 - And suppose I’m not entirely sure, which one is best to focus on. In the past, I thought it’s definitely the all out nuclear exchange between Russia and America, that kind of scenario because just the consequences would be so much worse. But then what someone pointed out, it might be much more likely for someone to use this kind of dirty bomb.
 - And even though it would only kill a fairly limited number of people directly, you can imagine just the social consequences being quite catastrophic that countries will become very unwilling to trade with one another because they would worry that a bomb would be imported in shipping container.”

Advice on whether and how to focus one’s career on nuclear risk reduction

What is this idea? How could it be tackled?

80,000 Hours wrote an “exploratory” problem profile on nuclear security in 2016 ([McIntyre, 2016](#)). But that was only a fairly quick investigation, and the profile has some issues. Furthermore, the nuclear security landscape has since changed: funders have exited (see [“How should EAs react to funders pulling out of the nuclear risk space?”](#)), EA now has more money and more people, and new research has been carried out in the meantime (e.g., [Rodriguez](#),

[2019a](#); [Rodriguez, 2019b](#); [Rodriguez, 2019c](#)). It could therefore be useful for someone in 2022 to figure out and disseminate some valuable career advice on the topic of nuclear risk.

Things this project could cover:

- Which careers to pursue?
- What career capital to build?
- Which countries matter most?
- Which sectors, departments, etc., matter most?

Possible output: Update and extend what 80k currently have published

- I made some rough notes on flaws and areas for improvement in 80k's current profile on nuclear security, which I could share on request
- It seems plausible they'd be open to someone helping them update their profile

Further notes:

- Should we try to encourage more people to go into things like diplomacy and arms control?¹
 - I have some brief notes on this which I could share on request
- There are important issues where the time isn't right to do government work on the issue. Examples are long-term AI governance, engineered pandemic policy, and maybe nanotechnology. What are the best fields to work in as "training grounds" for later working on these important issues, and how does nuclear risk stack up as a training ground?
 - I worry that we might get excited about nuclear risk as a training ground (see [Aldred, 2022](#)) without sufficiently considering alternative training grounds like cybersecurity, near-term/low-stakes AI issues, emerging tech policy, international relations / national security more broadly, and possibly climate change (e.g., because that involves difficult coordination problems and major externalities)

¹ Relevantly, Jeffrey Ladish writes in [One Hundred Opinions on Nuclear War](#): "73) The US, China, and Russia need to cooperate to prevent more countries from obtaining nuclear weapons! The US and the Soviet Union were able to do this despite their conflict. It's unfortunate that the current climate is so hostile to this. 74) Arms control is underrated, especially today. I want to see more bright young students study arms control. It's a dark time for diplomacy, time to turn that around!"

Subproject or particular thread: Should we be very keen to get EAs pursuing influential positions in non-Western nuclear-armed states in order to reduce risks from nuclear exchanges?

- Would pursuing that goal be net-negative via creating a perception in these states that this is an attempt by foreign agents to interfere with their national security policies?
 - That perception could make it harder to make further progress on nuclear risk reduction or other important issues
 - Are there some methods of pursuing that goal that create a minimal level of this risk?
 - Is it net-negative even to discuss this goal?

What portion of EA resources should be allocated to nuclear risk reduction or working out how much to prioritise that?

What is this idea? How could it be tackled?

I started looking into this topic and have notes in [this doc](#). A project here could involve starting out by following my skeleton plan.

20 things everyone (considering) working on nuclear weapons issues should know

What is this idea? How could it be tackled?

I started looking into this topic and have notes in [this doc](#). A project here could involve starting out by following my skeleton plan.

Aggregating, summarising, or drawing inferences from existing relevant forecasts

- My [Nuclear Risk Horizons Project](#) and [Nuclear Risk Tournament](#) collectively contain a huge number of questions, and Metaculus also has many other relevant questions
- My plan had always been to follow the tournament with one or more writeups summarising the key results and my main takeaways from them
 - But, very annoyingly, I ended up not having time!
 - I expect many decision-makers won't have the time, interest, or necessary background knowledge to go through the big list of forecasting questions themselves and work out what inferences to draw

- One could also perhaps produce an op-ed styled write-up on the Metaculus nuclear risk tournament
 - (linked [here](#) is the Metaculus journal, which does not yet include write-ups on tournaments - at least not to my knowledge - but which does give a flavour for what such a write-up could look like)

Commissioning forecasts and/or writing forecasting questions for forecasting platforms like Metaculus, Good Judgment, INFER

Public, etc.

- See my [Nuclear Risk Horizons Project](#) and the [Nuclear Risk Tournament](#) for relevant questions that already exist on Metaculus
- One project option would be commissioning forecasts from strong forecasters on some subset of the nuclear risk questions on Metaculus.
 - This could focus on the questions that seem most important and/or those that seem most neglected or where the forecasts seem most uncertain/scattered.
 - This could include paying the forecasters to write up (parts of) their reasoning
 - I know some people who've done or plan to do this, so could share more info or people to talk to about this on request
- Another project option would be helping write additional forecasting questions on Metaculus or other forecasting platforms.
 - I have a big list of rough forecasting question ideas that would require further operationalization before posting to the site
 - (And there could of course also be many ideas I haven't yet come up with!)
- See also [Impactful Forecasting Prize for forecast writeups on curated Metaculus questions](#)

How might nuclear risk change over time? Why? What can be done?

- Some things that might increase risk from nuclear weapons include:
 - Various technological developments
 - See also [Technological developments that could increase risks from nuclear weapons](#) and [my research project idea on that topic](#)
 - Increases in arsenal sizes
 - Increases in [yields](#) (of regular nuclear weapons)
 - Changes in targeting policies (e.g., toward an increase chance of large numbers of [countervalue](#) strikes)
 - Other changes in policy/diplomacy that increase the chance of nuclear war
 - Increases in geopolitical tensions or
 - Population increase, urbanisation, and wealth increases (which will tend to increase fuel density in targeted areas)

- Of course, risk could also *decrease*
 - Many of the things that could cause that are the same variables listed above or the opposites of them

Rough notes:

- [Carl Shulman has written:](#)

“I agree it's very unlikely that a nuclear war discharging current arsenals could directly cause human extinction. But the conditional probability of extinction given all-out nuclear war can go much higher if the problem gets worse. Some aspects of this:

-at the peak of the Cold War arsenals there were over 70,000 nuclear weapons, not 14,000

-this [Brookings estimate](#) puts spending building the US nuclear arsenal at several trillion current dollars, with lower marginal costs per weapon, e.g. [\\$20M per weapon and \\$50-100M all-in for for ICBMs](#)

-economic growth since then means the world could already afford far larger arsenals in a renewed arms race

-current US military expenditure is over \$700B annually, about 1/30th of GDP; at the peak of the Cold War in the 50s and 60s it was about 1/10th; Soviet expenditure was proportionally higher

-so with 1950s proportional military expenditures, half going to nukes, the US and China could each produce 20,000+ ICBMs, each of which could be fitted with [MIRVs](#) and several warheads, building up to millions of warheads over a decade or so; the numbers could be higher for cheaper delivery systems

-economies of scale and improvements in technology would likely bring down the per warhead cost

-if AI and robotics greatly increase economic growth the above numbers could be increased by orders of magnitude

-radiation effects could be intentionally greatly increased with alternative warhead composition

-all-out discharge of strategic nuclear arsenals is also much more likely to be accompanied by simultaneous deployment of other WMD, including pandemic bioweapons (which the Soviets pursued as a strategic weapon for such circumstances) and drone swarms (which might kill survivors in bunkers); the combined effects of future versions of all of these WMD at once may synergistically cause extinction”

- How might stockpile sizes change?
 - “Arms control may also have a significant impact on the nuclear modernization pathway to 2030. Deeper reductions would reduce requirements, though probably not significantly between now and 2030. An end to arms control might increase them. But [even] in the absence of an arms control framework, an arms

race driven by a desire for quantitative supremacy by any of the three major powers seems unlikely between now and 2030, given economic factors constraining all of them, among other factors.” [Fit for Purpose? The U.S. Strategic Posture in 2030 and Beyond](#)

- “strategic competition brings with it the possibility, even likelihood, of arms races. There are already signs of intensifying competition. Russia and China have competed to redress damage they believe the United States did to strategic stability during its “unipolar moment;” they have also begun to exceed the requirements of the status quo ante (the prior balance of strategic force). In response, the United States is now deciding how to compete. Arms races are not necessarily a bad thing. An arms race may be necessary to redress a new instability or to signal, as an alternative to war, the resolve to stand up to a challenger. But arms races also bring uncertainty, fear, and temptation. They also come with risks—including, in the contemporary case, a heightened risk of crisis stability associated with the apparent need to strike first and hard to gain a decisive advantage early in a mounting crisis. In a free-for-all among three powerful actors, everyone might lose. This begs the question: is there an alternative to a strategic free-for-all? The collapse of the legacy arms control regime casts a dark shadow over this question, as does the failure so far to start meaningful dialogues among the three about possible future forms of strategic restraint and common security. But this only serves to underscore the urgency of finding practical means to mitigate dangers” [Fit for Purpose? The U.S. Strategic Posture in 2030 and Beyond](#)
 - [Major Power Rivalry and Nuclear Risk Reduction](#) mentions “The demise of the Cold War-vintage arms control regime in the context of major power rivalry and the heightened risk of an arms race, potentially triangular.”
 - From “[Nuclear winter and human extinction: Q&A with Luke Oman](#)”:
 - Q2: If nuclear arsenals become much larger in the future, e.g. 100x as large, damage would presumably scale sublinearly (only so many cities to ignite). Could the detonation of millions of nuclear weapons make a material difference to your estimate?
 - A: Yes, it would make a difference but as you state I would definitely think it would scale sublinearly. The largest thing that I would think, more so than the number above a certain point, would be how much the Southern Hemisphere is involved. In the 2007 paper scenario it is assuming largely NH mid-high latitude injection so there is likely large difference in black carbon aerosol amounts in the respective hemispheres. This is one of the largest differences between the 150 Tg of BC scenario and that of Toba, which was a tropical eruption and presumably spread much more evenly over both hemispheres.
- Changes in geopolitical relations
 - Proliferation

- <https://cgsr.llnl.gov/content/assets/docs/Major-Power-Rivalry-and-Nuclear-Risk-Reduction.pdf> notes “Continued development of nuclear capabilities and doctrine in South Asia amidst continued conflict and confrontation.”
- See Metaculus question(s)
- The conclusion of [Fit for Purpose? The U.S. Strategic Posture in 2030 and Beyond](#) is relevant
 - “By 2030, the balance of central strategic forces among the United States, Russia, and China may have altered in various ways. But it is unlikely to have changed in a way that fundamentally calls into question the ability of the United States to retaliate in a devastating manner. The United States is likely to have a nuclear posture much like today’s force. It is highly unlikely to be larger or more diverse; it may be somewhat smaller and less diverse. Russia will have a fully modernized force that is unlikely to be substantially different from today’s force. China’s force may be double its current size and fully modernized and its precise role and purpose are likely to remain ambiguous. In 2030 there will be significant uncertainties about where all three will be headed with force size and function in the decade to 2040. The imbalance that may be consequential in 2030 is the imbalance we’re beginning to recognize in 2020: the imbalance in strategic thought about the requirements of effective deterrence in the kinds of conflicts we might face in an era of major power rivalry”
 - “The extended deterrence balances in the Transatlantic and Transpacific Alliances look much more problematic. NATO has made important progress since 2014 in adapting its overall deterrence and defense posture to new Russian challenges, but its 2020 posture is not robust against the emerging Russian threat to the U.S. power projection strategy. Without significant new inputs, the theater/strategic balance is likely to degrade from a NATO perspective. In East Asia, U.S. alliances have also made progress—but mostly vis-à-vis North Korea, not China. The conventional balance vis-à-vis China is worrisome and eroding. The United States and its allies must do more to compete to keep up, while making some hard choices about where specifically to compete and where not. The Pacific Deterrence Initiative, like the European Deterrence Initiative, makes a positive but limited contribution.”

Something about AI/TAI/AGI and nuclear weapons

- This area seems important
- But it would be partly covered by [my technological developments work-in-progress](#) or other work on that topic
- And I don’t really have a crisp question or clear scope in mind; this is just an area, rather than a fleshed out research project idea

- See also my brief discussion of the [“Intersections between nuclear risk and AI”](#) and the collection of relevant Metaculus questions therein

TPNW-related research

- Could work on what Luisa intended as part 2 to her [“Will the Treaty on the Prohibition of Nuclear Weapons affect nuclear deproliferation through legal channels?”](#)
 - Where part 2 would address TPNW’s effects (or not) on nuclear deproliferation through non-formal channels
 - Could also work on questions in the [further research directions](#) section from Luisa’s part 1
- My rough notes, from when I was planning on doing part 2 myself:
 - "Luisa published a post and had notes and draft sections for the followup post.
 - I would now guess doing this would take ~4 weeks, but apparently earlier I said 12, so I’ve gone with 6 as a compromise. Probably I was envisioning something grander, whereas now I’m envisioning something a bit more like just wrapping up Luisa’s thing in the quickest way that is still passable.
 - [But really] a key factor in whether I’d recommend supporting TPNW/ICAN is things like “ideal size and composition”, which this won’t cover. So this [...] won’t in my view fully address the underlying question of “Should I fund ICAN?”

Investigating various possible treaties other than the TPNW

- What are the salient options?
- How valuable would each be?
- How tractable would bringing each into effect be?
- How can that best be done?

Ideas that seem lower priority

I have further notes on most of these, which I could share on request.

Nuclear terrorism

- [Rodriguez \(2019\)](#) writes:
 - “While a nuclear detonation by a non-state actor (terrorist) looks plausibly quite harmful in expectation, it’d be very difficult to analyze, as there’s no single terrorism scenario to consider. I therefore leave a discussion of the potential harm caused by nuclear terrorism for future work.”

- [Major Power Rivalry and Nuclear Risk Reduction](#) says: “Of note, for the moment at least, there is a sharp decrease in concern about the risk of nuclear terrorism.”
- There are some relevant questions in my [Nuclear Risk Tournament](#) and my [Nuclear Risk Horizons Project](#)

Yields over time

- Specific questions:
 - For each nuclear armed state (or just the US and Russia), what has been the mean and median yield of warheads in their stockpile in various years?
 - What has been the total yield of their stockpiles in various years (maybe broken down by alert status vs deployed vs reserve vs retired)?
 - What has been the largest yield warhead in their stockpile in various years?
- Maybe ideally produce graphs like [OWID's](#) one for overall arsenal sizes
- This might already exist, but I wasn't able to find it quickly
- This seems very tractable and well-scoped, and a decent job might be doable in even just a day, very likely less than a week
- Not sure how useful this actually is. But I often found myself wishing I had access to this info when trying to write about how risk has changed over time and what sort of base rates we should have in mind for future changes
- Possible data sources:
 - Nuclear Notebook has info on when various countries' arsenal size peaked and when their total yield peaked. The yield always peaks earlier. I have Anki cards for this. Can't remember if they had more detailed data.
 - [Turco et al. writing in 1983](#) that “A review of the world's nuclear arsenals (20-24) shows that the primary strategic and theater weapons amount to ~12,000 megatons (MT) of yield carried by ~17,000 warheads.” This implies a mean yield of 700kt

How much funding and labor goes to nuclear risk reduction?

- Create an updated version of [this post](#) in light of changes since late 2021 (in particular, the [New Nuclear Security Grantmaking Programme at Longview Philanthropy](#) and the emergence of the FTX Future Fund²)

² I don't yet know if the Future Fund will fund much nuclear risk work, but it seems likely they'll fund at least some and plausible they'll fund a lot.

- It could perhaps makes sense to create a version of this post quickly once per year, by copying the post and then adjusting the estimates and perhaps the takeaways
- Further potential ways to improve the post:
 - Attempt to get Peace and Security Funders Group to improve or clarify their data (see [here](#)), attempt to find or create a better data source for levels of total non-governmental funding, and/or attempt to figure out how to adjust the estimates given in this post to better account for the issues with the PSFG data
 - Estimate non-EA non-governmental labor devoted to nuclear risk and governmental funding, labor, and attention devoted to nuclear risk
 - Use Guesstimate rather than point-estimate-based Fermis
 - Both to capture uncertainty and to make it easier for people to make their own tweaks to individual assumptions and see what results from that
 - Check various things more closely and search for info more comprehensively
 - E.g., contact EA funders to ask if they've made other donations that the post hadn't accounted for yet
 - Create and integrate more forecasts, as opposed to just explanations of current or historical rates
 - I already created some relevant forecasting questions, which are listed at [Nuclear risk funding & labor - Metaculus](#)
 - But I didn't integrate forecasts from all of those questions into this post
 - And other useful questions could also be created

How might nuclear energy stuff affect nuclear weapons risk?

- Some nuclear energy capabilities can help a state on the path towards developing its own nuclear weapons.
- On the other hand, more than 30 states already have nuclear energy capabilities ([IAEA, 2021](#)), and most of them neither have nuclear weapons nor seem likely to develop nuclear weapons anytime soon. Additionally, one of the NPT's three "pillars" (along with nonproliferation and disarmament) is "the right of all Parties to develop nuclear energy for peaceful purposes and to benefit from international cooperation in this area, in conformity with their nonproliferation obligations" ([Non-Proliferation Treaty, 2010](#)). Thus, frustrating that right could undermine the legitimacy of the NPT or give states an excuse to engage in proliferation efforts.

Additionally, I think the Future Fund's size and intended rapid pace of giving reduces the extent to which we should focus on cost-effectiveness and increases the extent to which we can pursue many things in parallel (e.g., both nuclear risk reduction and other existential risk reduction) as long as we have sufficiently good ideas and the relevant talent available. This reduces the importance of posts like this one.

- There's an EA-aligned academic - Simon Friedrich - who wrote [a paper](#) that has some discussion of this (in section 7.2)
- I commented on a draft of that paper saying that I might break this topic down into the following questions (though I should note that I haven't thought about this topic much yet):
 - 1. Does a country having a nuclear energy programme increase the chance that they develop nuclear weapons?
 - I might want to consider how this differs depending on what components of a nuclear energy programme the country has. E.g., do they produce their own HEU or import it? Do they have reprocessing facilities in their territory or do they just export the spent fuel for reprocessing elsewhere? (Maybe no such importing/exporting occurs - I don't know.) What specific types of reactors/reprocessing facilities do they have, if that's relevant?
 - 2. Is it bad for a new country to have nuclear weapons?
 - I might want to break this down into:
 - a) the effects of that on civil war, militarised interstate disputes in general, large-scale non-nuclear war, nuclear war, one-sided nuclear strikes, and stability of bad regimes (e.g., a country like NK continuing to be a bad place to live for longer because it's harder for other actors to influence it because it has nukes)
 - b) how bad each of those things are (this is only relevant if nuclear prolif increasing some but decreases others, such that tradeoffs are required)
 - I'd guess it's bad, but it doesn't seem totally obvious, and I believe the question is debated in the literature.
 - 3. How does more countries having nuclear energy programmes affect nuclear risk through things like effects on norms and diplomacy? E.g., maybe adding legitimacy to the NPT since that says countries have a right to this.
 - I think for this "country has a nuclear energy programme" is less relevant than what specific interventions are being proposed that would cause/prevent that. E.g., actively blocking countries from getting nuclear energy programmes might have a notable effect here, whereas whether philanthropists invest to help them probably matters less.
 - And here I think the effects on nuclear risk wouldn't just be about proliferation but also about disarmament, arms buildups, and maybe other things like confidence building measures and generally good relations.

(So "Proliferation of nuclear weapons" would perhaps not quite be right as the heading, but it seems close.)

Evaluation of [ALLFED](#)

- A post that's like an evaluation of ALLFED, a cost-effectiveness analysis of them or their proposed interventions, and/or a critique of their cost-effectiveness analyses/arguments
- Nuno Sempere has already carried out a [shallow evaluation of ALLFED](#) (2021), which could be used as a starting point

Other relevant collections of research project ideas

- Some Metaculus questions or groups of questions in the "[nuclear threats](#)" category, perhaps especially those in my [Nuclear Risk Tournament](#) and/or my [Nuclear Risk Horizons Project](#), could be worth doing research projects on
- My "[Politics, Policy, and Security from a Broad Longtermist Perspective: A Preliminary Research Agenda](#)"
- [Baum \(2015\)](#), Confronting the Threat of Nuclear Winter, Section 5.1
- Some people have identified research questions they feel it could be useful for someone to look into in light of the 2022 Russian invasion of Ukraine
- For a list of collections of research project ideas on *many* topics (not specifically nuclear risk), see my "[A central directory for open research questions](#)"

(This is of course not an exhaustive list.)