

AI Strategy, Policy, and Governance | Allan Dafoe | Beneficial AGI 2019 Conference

To learn more, check out the [Centre for the Governance of AI \(GovAI\)](#).

Discussion prompts - pick whatever sounds interesting!

1. First impressions: What did you think of this talk? Did you like it? Was anything confusing?
2. What are you skeptical of? Did this talk leave out any important information? What unanswered questions do you have?
3. What did you think of Putin's statement “whoever leads in AI will rule the world”?
4. Do you think the governance of AI will be extraordinarily difficult?
5. “Fermi thought that the conservative thing was to play down [his 10%] possibility that [a nuclear chain reaction] may happen,” while Szilard thought we should take precautions anyways. How should policymakers prepare for low-probability, high-impact risks?
6. Will the robots take our jobs? How seriously should we take concerns about advanced AI causing unemployment?
7. How might advanced AI be different from previous technological breakthroughs?

8. Do you agree that the long-term political challenges from AI will resemble the near-term political challenges from AI that we see today?
9. Should the US government promote faster or slower advancement of AI?
10. How might our world incentivize companies to advance AI as quickly as possible while neglecting safety considerations?
11. The Windfall Clause is a commitment that companies can make: if a company gets extremely large profits from advanced AI, it will donate part of this “windfall” towards benefiting humanity. Why might companies agree or disagree to make this commitment?
12. How might advanced AI increase the risk and damage of international conflicts between countries with powerful militaries?