

Running Program Sessions

Thank you for agreeing to run sessions for the program! It's a great way to have an impact while developing your own skills (and hopefully having fun).

As a facilitator we hope you'll make sure your cohort:

- Learns a lot about EA, specifically the key points we've identified
- Has fun and feels welcome

This document aims to help you do this. The first part contains some general advice on your role in the program and tips for facilitation. The second part contains a week-by-week guide to running each individual session.

[Your role](#)

[What does it mean to facilitate?](#)

[Before the Program Starts](#)

[Running Sessions Well](#)

[Before the Session](#)

[During the Session](#)

[Setting the group culture](#)

[Building a social atmosphere](#)

[Facilitation tips](#)

[After the Session](#)

[Ending the Program](#)

[Week-by-Week Guide](#)

[Week 1: The Effectiveness Mindset](#)

[Key Points](#)

[Tips for this session](#)

[Session](#)

[Icebreaker \(15-20 mins.\)](#)

[Give overview of program](#)

[Set discussion norms](#)

[Discussion questions](#)

[Week 2: Differences in Impact](#)

[Key Points](#)

[Tips for this session](#)

[Session](#)

[In-Session Exercise \(10 mins.\)](#)

[Discussion questions](#)

[Week 3: Expanding Our Compassion](#)

[Key Points](#)

[Tips for this sessions](#)

[Session](#)

[Icebreaker suggestion \(10 mins.\)](#)

[In-Session Exercise \(15 mins.\)](#)

[Discussion questions](#)

[Week 4: Longtermism](#)

[Key Points](#)

[Tips for this session](#)

[Session](#)

[Icebreaker \(10 min.\)](#)

[Discussion questions](#)

[Week 5: Existential Risk](#)

[Key Points](#)

[Tips for this session](#)

[Session](#)

[Discussion questions](#)

[Week 6: Emerging Technologies](#)

[Key Points](#)

[Session](#)

[Icebreaker \(10 mins.\)](#)

[Discussion questions](#)

[Week 7: Increasing the Accuracy of our Judgements](#)

[Key Points](#)

[Session](#)

[Icebreaker \(10 mins.\)](#)

[In-session exercise \(20 mins.\)](#)

[Discussion questions](#)

[Week 8: Putting it into Practice](#)

[Key Points](#)

[Tips for this session](#)

[Session](#)

[Survey](#)

[Discussion questions](#)

[Shoutouts \(15 mins.\)](#)

Your role

What does it mean to facilitate?

To facilitate means to make a discussion more useful and pleasant for the discussion participants. Your goal in the sessions is therefore not to *teach* the content of the week, but to encourage participants to discuss it between themselves and to provide useful input and structure to the conversation. Just by being there and starting the discussion on each topic, you're already facilitating and adding value to the conversation!

Facilitating sessions well is a skill that you will develop over time. It isn't easy, and there will likely always be room for improvement. For that reason, it's important to spend time every week reflecting on what went well and what might be improved next time – and then acting on those thoughts. Talking to other people about any challenges you've had or how best to run sessions can also be very helpful.

The rest of this document provides advice on how to facilitate in general and how to run each individual session. We encourage you to use these resources, but also to adapt to your particular group of participants and your personal style. Knowing the particular context, you will have better information than we can when writing a general guide like this, so **use your best judgement** to determine how you can provide the most value to the group. Feel free to use the suggested discussion questions and activities exactly as they are, and feel free to use them as guidance. The same goes for the facilitation advice in this section – some groups will benefit from highly structured discussions, others prefer free-flowing conversation. The advice in this document are suggestions for ways to make your sessions even better, not requirements to stress about keeping to.

Before the Program Starts

Make sure that you:

- Know when and where your sessions will be
- Know how to communicate with other people involved with the program
- Have read this guide
- Have an overview of the program curriculum
- Have thought about the goals of the program and your role in achieving them
- Know about anything else you need to do before the sessions

Running Sessions Well

Before the Session

The key thing to do to prepare for a session is to think about what you want to cover in the session, how you want to structure it and what directions the conversation could go in. Use the **week-by-week guide** provided below, as well as the readings from the curriculum, to help with this.

You should also read the core and recommended readings for the week, or re-familiarise yourself with them if you've read them before. Think about how each contributes to the key points listed in this guide.

Some other things that may be useful preparation:

- Look up any relevant definitions that you're uncertain about.
- Think back to how last week's session went and what you might try doing differently this time.
- Draw up a rough structure for the session with time estimates for each activity or topic. This is useful for keeping an overview during the session, although you don't need to stick to it exactly – deviate whenever something else feels more useful.

Here's an example of a 1.5 hour session (your plan can be more or less detailed):

- Intro chat, people arrive (5 min)
- Icebreaker (10 min)
- Discussion (30 min)
- Break (5 min)
- Exercise and further discussion (30 min)
- Rounding off (10 min)

During the Session

Setting the group culture

By 'culture', we mean 'the way we do things'. As the leader in the social context of the group, you have quite a lot of power to set the culture of the group, especially early in the sessions.

Think about things like:

- Your tone of voice and body language
- How friendly and kind you are towards other people in the group
- How enthusiastic you are about EA
- How you behave when you disagree with someone.

In addition, think about what discussion norms you would like for the group to exhibit. We have written a suggested policy [here](#). We suggest discussing these explicitly with participants in week 1, as you will see in the week-by-week guide.

Consider whether you would like to have a hand-raising system for moderating who speaks when. This will depend on group size and dynamics, but we generally recommend using a hand raising system if there are more than 4 participants in the group.

Building a social atmosphere

We think one of the most important things to focus on as a facilitator, especially in early sessions, is to get the participants to know one another socially. This makes people feel more comfortable participating, and have more fun. We suspect that people are more likely to stick around after the program if they feel that they know and like people in the community.

We suggest spending a reasonable amount of time each session (20+ minutes) on letting people get to know each other, at least in the first 2 sessions. We provide icebreakers for the beginnings of the sessions, but there will be opportunities at the start, in the break, and after the session for casual, unstructured conversations. You might want to think in advance about what questions you could ask people. We also recommend leaving the Zoom call running after the sessions, in case people want to continue chatting.

Facilitation tips

We said in the beginning of this document that facilitation is about making the discussion useful and pleasant. Here are some more tips on how you can do that in practice:

- **Explain jargon used by you and the participants**, either giving definitions yourself or asking participants to provide them.
- **Keep track of whether the discussion feels useful for the majority of the group, and move on if not.**
 - **Don't hesitate to interject and change topic** if a conversation gets stuck in an unproductive place ("This seems like an interesting point, but I think we should move on for the sake of time").
 - Conversely, **if the discussion flows naturally and seems useful for most of the participants, feel free to let the group speak uninterrupted.** Straying off-topic to explore interesting side-roads can be fine, as long as the discussion feels "alive" and you're making progress.
- **Ask questions that help further participant's understanding of the key ideas** – this is what the discussion questions in this guide are designed to do, but feel free to complement them with questions you think seem helpful in the context.
- **Use what the participants say themselves to further the discussion**, for example:
 - **Repeat back what people say in your own words** - this is a powerful method for clarifying your understanding and summarising someone's point in a concise way, as well as ensuring that participants feel heard and understood.
 - **Ask participants to clarify** – especially when someone makes a point that seems novel or interesting, but you don't quite understand.
 - **Ask participants to comment on each other's points** – e.g. "What do the rest of you think of this idea, do you agree?"

- **Remind participants of useful framings of ideas**
 - One goal of the program is to help participants feel comfortable with prioritising career paths and donation opportunities, it's sometimes useful to reframe the debate around this.
 - “This has all been quite abstract so far - I’m curious to hear your thoughts on how all of this might affect decisions we might have to make in our careers/lives?”
 - “How does this affect the value of a donation to Y?”
 - “[Give hypothetical decision] How would you choose between these two options?”
- **Promote good discussion norms**, both in the way that you present arguments yourself and by encouraging them in participants.
 - “Should we take that as strong evidence for the claim?”
 - “Seems like we agree on this, what are some reasons we might be wrong?”
 - “This seems like quite an unconventional view; should that make us worried that we’ve gone wrong somewhere?”
 - “A lot of people who have spoken have had x view. Is there anyone who disagrees?”
- **Clarify simple misunderstandings and provide factual information**, for example:
 - “Yeah it’s true that GiveWell does consider more than just overheads”
 - “I think we’re stuck on X - I’ve made a note that we’re confused about this. Let’s move on to the next question and we can do some research on this between sessions”
- **Present arguments** where you find it useful. While the sessions should centre around the participants discussing and presenting arguments themselves, prompting the participants with an argument to respond to or a new idea can sometimes be helpful.
 - If you’re making the case for something, think about who you’re speaking for and try to clearly distinguish between the following:
 - Making the case for something, because it’s a fundamental part of EA and you want the participants to understand it (“a really key point is that...”)
 - Making the case for something, because some people believe it and you want the participants to understand how the argument works, even if you may disagree with it personally. (“some would argue that...”)
 - Playing devil’s advocate for the sake of the discussion
 - Speaking as yourself (“my personal opinion is...”)
- **At the end of the session, encourage participants to reflect on the week.** Either aloud or in writing, you might ask each person to say e.g. one takeaway for them for the week, one confusion or uncertainty they still have, one piece of feedback on the session, or any combination of these.

After the Session

Spend 10+ minutes reflecting on a few key things after each program session:

- Lessons and takeaways to share with others for when they run this session
- Writing down some notes to help build your understanding of the participants
- Reflecting on how you might be able to improve as a facilitator based on the session (see [deliberate performance in people management](#) for some context on this point).

Check in with the program organisers if there is anything you'd like to discuss.

Ending the Program

Much of the last session should be taken up by the participants completing a feedback survey - this is really important to work out how to improve the program in the future. Do also informally chat to the participants to get a sense of any improvements you think should be made.

Have a think about what improvements might well be worth trying in future iterations of the program, and do pass these on to the organisers of the program.

Feel good about your significant contribution to making the program a success!

Week-by-Week Guide

Week 1: The Effectiveness Mindset

Over the course of week 1 and 2 we aim to introduce the core principles of EA. We introduce these ideas in the context of global poverty as we think it's an approachable and intuitive cause area, and makes the tools seem more obviously useful and practical.

Key Points

- We have outstanding opportunities to do good available to us.
- In order to do the most good, we have to make tradeoffs between different outcomes. Because all of these outcomes involve helping someone, it may feel difficult or “wrong” to choose who we help (and thus, who we don't help).
- Aggregation - we can, in principle, compare outcomes of different interventions
- Scientific evidence allows us to quantify the outcomes of different interventions and allows us to make better decisions
- Maximisation - we should choose the best interventions we can
- We should focus on marginal benefit instead of average benefit, making neglectedness an important criteria to look at.

Tips for this session

- The most important goal of this session is to build group rapport so that people are comfortable speaking and listening to each other. It's worth spending a significant amount of time (20-40 minutes) on icebreakers, chatting and other social activities.
- Will's coverage of the case in favour of deworming has been criticised for being overly favourable to the case of deworming, and not adding enough caveats about the uncertainty surrounding this intervention. We don't expect this to come up often, but it probably makes sense to familiarise yourself with the [worm wars](#) in case this comes up as a topic of discussion

Session

Icebreaker (15-20 mins.)

- Everyone introduces themselves
- [Rock paper scissors tournament](#)
- [Would you rather where you go to each side of the room](#)
 - Would you rather live for an extra year, or an extra 14 months but you can't use a screen for that time?
 - Would you rather live for an extra year, or an extra two years but you must be nocturnal?
 - Only need two hours sleep per night, or have a perfect memory?

Give overview of program

It might be helpful to briefly give the high-level details of the program again.

- High-level details of the program:
 - 8 weeks
 - Around 3 hours of personal work including reading each week, to be done before the session
 - 90 minute weekly meeting
 - Aim is to learn about some of the ideas people have had about doing the most good, and work out which of them we think are good ideas, so we can take them into account for our plans to do good in the future.
 - Also to make connections with like-minded people.
- Ask if people have any questions

Set discussion norms

The culture of your group will largely be set in the first session. Try to make it as sociable, fun, thoughtful & considerate as possible - largely by setting this tone yourself.

Discuss the norms that you expect people to follow. Don't forget that especially this week, you help to set the culture!

- Recap [discussion norms](#).
 - Respect
 - Basic stuff; if program is online, perhaps emphasises taking slightly more care not to interrupt, which can be difficult in video calls.
 - Constructive.
 - In the discussions, we'll be clarifying our understanding of the reading, and talking through our perspectives. Naturally, some of us are therefore going to disagree with each other.
 - We want to be approaching these disagreements with a framing of:
 - *Try to figure out what the other person thinks, and why - which differences in worldview or in models generated the disagreement? rather than*
 - *Try to convince everyone that you're right.*
 - Inclusive
 - Explain jargon - there's lots of it around in some of the fields we'll be talking about
 - Do ask for clarification if there is jargon or acronyms you don't know.
 - Be careful of dominating the discussion – it will be most valuable if everyone contributes equally.
 - Explain any hand-raising norms or similar that you'd like to enforce.
- **“It's also important to acknowledge”:**

- Our privilege in being able to discuss these ideas and to not let that overshadow the fact that we are talking about real lives. (Generally being aware that we are lucky to be in a place where we can be able to think about these issues and have a large impact)
- These are very complex issues, and keeping an open mind is important
- Some of these can be uncomfortable to talk about, and some of us may have personal ties to some of these areas. That is part of the reason they are so neglected and so important to talk about.

Discussion questions

Importance of prioritisation/triage

- **Suppose you have been given £100,000 to spend. Charity X will save lives for £20,000, while charity Y will save lives for £15,000. How much would you give to charity X and how much to charity Y?**
- Is it fair to describe James Orbinski as 'cold, calculating' for pragmatically prioritising their patients when faced with overwhelming casualties?
- How fair is it to say that we're in a similar moral situation to James Orbinski, even if much less visceral?
 - Are we in a position to help many others? *Yes*
 - Are we in a position to only help very few of these? *Yes*
 - Does it follow we have a tradeoff, i.e. helping some of these is necessarily at the expense of others? *Yes, since we only have very limited resources*
 - Does it follow that we should pragmatically prioritise in the same way that James Orbinski did? *Yes!*
- What is the importance of an area being neglected – should we (as individuals) prioritise an area more highly if it's neglected?

How to think about our situation

- Do we have a great opportunity or a terrible tragedy? *Both*
- How (if at all) should we take our emotional and/or intuitive judgements into account when deciding where to donate or how to spend our time?

Week 2: Differences in Impact

This week continues in a similar vein to week one, but with more focus on methods of measurement and expected value. We hope that by the end of this week, participants will understand that it is possible in principle (but hard!) to quantify good and adjudicate between causes.

Key Points

- Understand expected value reasoning and how this opens the door to high-risk high pay-off strategies
- **Large differences in cost effectiveness between different interventions mean the very best interventions are much more valuable than typical interventions**
- **We can use quantitative estimates to reduce uncertainty in order to make better decisions**
- (Through discussion) Quantification is a complement to, not a replacement for, other decision-making heuristics.
- (Through discussion) **We should use quantification in cases where the scale of impact may vary a lot, but may not be tracked by our intuitive judgements.**
- [Practicing generating quantitative estimates and comparing outcomes](#)

Tips for this session

Note that the study mentioned in 'Scope Neglect' is actually better designed than the LessWrong post suggests. It actually asks about tax increases, which seems to more clearly get at scope insensitivity. If study participants were being asked to pay themselves, their answers might be correlated with how much spare money they had, or how much the signalling was worth to them.

Also note that this week will often involve discussion of income, which can be a sensitive topic, and should be handled carefully.

Finally, QALYs are a great example of a way that we can begin including quantification in our decisions about how best to help, but it's important that moderators realise that they are far from perfect as a metric.

- QALYs are a measure of health, and as such, can only be said to correlate with individual welfare. They will correlate less, or not at all, with other things we might care about intrinsically, like the environment or justice
- Their weightings are also informed by peoples *expectations* of how bad certain outcomes would be, rather than by sampling their experience in the moment
- Beyond this, as always trying to reduce complex phenomena into a single metric like this will necessarily lose information, and it's important that we acknowledge that QALYs are a useful tool that *adds* to a decision, rather than resolving it immediately.
- See [here](#) for more discussion, and [here](#) for a discussion of what different measures of near-term welfare might imply working on.

Session

In-Session Exercise (10 mins.)

Split into groups of 2-4.

Suppose that a disease, or a war, or something, is killing people. And suppose you only have enough resources to implement one of the following two options:

- 1) Get half of the cohort to decide between:
 - a) Save 400 lives, with certainty.
 - b) Save 500 lives, with 90% probability; save no lives, 10% probability

- 2) Get the other half to decide between
 - a) 100 people die, with certainty.
 - b) 90% chance no one dies; 10% chance 500 people die.

Point out that people often pick 1a) but 2b), even though 1 and 2 are the same gamble. Discuss.

Discussion questions

Discuss exercise

- How did everyone find the exercise?
- If you feel comfortable sharing, what did you estimate your total future income to be? Did this seem surprisingly high or low?
- What did you work out you could achieve with 10% of your future income? Did this seem like more or less than you would have expected?
 - *You may want to compare this to the drowning child/burning building*
- Which charity did you pick to donate to? Why was this, why did the outcomes of donating to that charity seem more valuable than the outcomes of the other charities? Did you find it hard to choose between different outcomes?

Benefits of cost-effectiveness estimates

- How can we go about comparing different interventions/cause areas? Can quantitative estimates of impact be useful even if they're imprecise?
 - How useful are QALYs?
 - How useful is considering importance, tractability, and neglectedness?
 - How useful is EV?

Limitations of cost-effectiveness estimates

- What types of outcomes are particularly hard to measure (or even impossible)? How should we treat such outcomes?

- What kind of problems can we run into when we try quantifying cost-effectiveness? What features are not captured by such estimates? Are there important features of an intervention that are not captured by cost-effectiveness estimates? Which?
 - Things to draw out in discussion:
 - It is particularly important to use explicit quantification in cases where the scale of
 - impact might vary a lot but may not be well tracked by our intuitive judgements
 - Explicit quantification is a complement to, not a replacement for, other decision-making heuristics

Application: Donation decisions

- Discuss GW vs. OPP approaches -
 - Can anyone summarise the difference between the GiveWell and Open Philanthropy approach?
 - If an intervention is not backed by strong evidence, could there still be reasons to pursue it?

Week 3: Expanding Our Compassion

This week is focused on moral circle expansion, with a particular focus on farmed animal welfare as a case example. This week is geared around getting participants to think about their own values and the practical implications that these have.

Key Points

- Awareness of historical moral circle expansion, and that societal morality has often been wrong
- Awareness of various groups which could potentially be within our moral circle
 - There are some strong ethical and empirical arguments to think that animals are morally valuable
- There are vast numbers of animals suffering in factory farms, and we have very cost-effective ways to help them
 - So the choice of where we draw our moral circles can have large practical implications for doing good
 - [Participants reflect on their own personal values, and the practical implications of these values](#)

Tips for this sessions

This week, it's important to consider the tradeoff between exploring more novel or weird-sounding ideas (such as digital sentience or wild animal suffering), and focussing specifically on farmed animal welfare, and farmed animals as valuable beings.

Two contrasting failure modes to avoid:

- The group talks mainly about weird stuff, and spends little time talking about farmed animal welfare and veganism. The vegans in the group feel like EA doesn't care about farmed animals, is too theoretical, weird, and out of touch.
- The group talks too much about farmed animal welfare and veganism. There isn't much time to talk about more speculative things, and the tone around animal welfare is pushy and demanding. Omnivores feel like you have to be vegan or vegetarian to be an EA, EA is unwelcoming and dogmatic, and there isn't space to discuss more speculative ideas
 - Focusing on the effectiveness of ACE top charities, and of systemic changes like clean meat can relieve this somewhat, as it can reduce the focus on individual dietary choice.

Some potential solutions:

- Paying attention to the group and steering the conversation towards their interests
- Making sure all the topics get discussed for some of the session
- Avoid bringing up digital sentience, and discuss it carefully if brought up, acknowledging that it's a very unconventional idea

Discussing animal welfare can be uncomfortable, especially for people who eat meat. It can feel like they are being told they are a bad person which can make them defensive. It's really important to navigate this sensitively, noting that people have different experiences and cultures that influence where they stand on this. We recommend you:

- Reiterate that the purpose of the program is not to convince participants of specific moral convictions but to present them with influential works and arguments they may have not been exposed to before so that they can come to their own conclusions.
- Note that though there are, of course, strong arguments to be vegan or vegetarian, eating meat is normalized in society and changing your diet can be difficult. Remind participants that we want to have an open space where different views are welcome and we respect each other's autonomy.

Session

Icebreaker suggestion (10 mins.)

- Speed friending (expanding your social circle ha ha): Four rounds of 30 seconds conversation with someone you haven't spoken to before or haven't talked to that much. Talk about what the highlight of your semester has been so far.

In-Session Exercise (15 mins.)

Split into groups of 2-3 and do Part 3 of the exercise from the curriculum this week.

Discussion questions

Discuss exercise

- What did you write in your letter to the past?
- How easy did you find it to convince this past person of a moral wrong?
- What sorts of arguments did you use?
- How difficult did you find the exercise from your future self, given that we don't know where we're going wrong?

Moral circle expansion

- Is 'moral circle expansion' a good description of what has happened in the past?
- **Why have we historically failed to recognise the moral importance of others?**
- Do you feel optimistic or pessimistic about society improving its values over time?
- **How likely do you think it is that we currently do not show moral concern to beings that deserve it? Why?**
 - If you think it is likely, then how should we strike the right balance between showing moral concern for beings not widely accepted as morally relevant, and not wasting effort trying to help things that are morally irrelevant?
 - Are there any general techniques or tools we can use to avoid being complicit in atrocities, given that it's hard to know where we're going wrong?
 - If you think it is unlikely, do you accept that the 'moral circles' of every previous generation has been too narrow? If so, isn't it suspicious that you

think we live in the first generation where we have arrived at the correct moral circle?

- What sorts of things make people commit atrocities? Ignorance, selfishness, confusion, malevolence?

Farmed animal welfare

- Do you think farmed animals matter morally?
- Do you think it could even make sense to donate to help farmed animals, instead of donating to help people in extreme poverty, say if the animals are much easier to help?
- If you could magically find out, what information about animals would you most like to know, to help you decide how much they matter morally?
- Is anyone here vegan or vegetarian? If you feel comfortable, would you share what inspired you to become vegan or vegetarian?

Uncertainty and moral concern

- Do you agree that we ought to be open to 'strange' arguments about which beings are of moral concern?
- If you are 99% sure that insects are morally irrelevant, but think there is a 1% chance that they should be considered morally equivalent to a human, how should you treat it?

Week 4: Longtermism

We use this week to introduce people to the philosophical claim of strong longtermism. We wanted people to explore Toby Ord's argument for prioritising existential risk before introducing the philosophical notion of longtermism, to more clearly disentangle the concepts of existential risk and longtermism, which people often conflate, when in fact we think there are strong arguments to work on existential risk reduction that don't depend on longtermism, and longtermism can lead people to prioritise other causes beyond existential risk reduction.

Key Points

- Understand the concept of longtermism: that if you want to help people in general, your key concern should be to ensure that the future goes well in the long term
- There likely will be much more value and disvalue in the future than in the present.
- Influence the future may be tractable eg. through reducing existential risk or trajectory changes
- There are ways to influence the future other than reducing existential risk

Tips for this session

Make a disclaimer along the lines of: "Just like last week, this week we'll be discussing our personal values. This is one of the places where EA definitely doesn't have the answers. Hopefully you'll listen to other people's views, and consider different arguments and perspectives, but in the end you'll have to make up your mind yourself about what kinds of things you value."

Session

Icebreaker (10 min.)

- Ask people to split into groups and share something they're looking forward to in the year 2070. What's a technology they hope will exist? How do they think the world might be better for oppressed people? Do they just look forward to playing with their grandchildren?

Discussion questions

Discuss exercise

- Imagine you could save 100 people today by burying toxic waste that will, in 200 years, leak out and kill thousands (for the purposes of the question, assume you know with an unrealistic level of certainty that thousands will die).
 - Would you choose to save the 100 now and kill the thousands later?
 - Why?
 - Does it make a difference whether the toxic waste leaks out 200 years from now or 2000?
 - Why or why not?

- Imagine you're a wealthy philanthropist, considering how to spend your money. Your first option is to pay for surgeries for blind people in the US. With your donations, you will restore the sight of ten people. You also wanted to consider some nonstandard approaches to philanthropy however, and so your second option is to pay certain couples to have children (who otherwise would not have done so). As a result, ten children with good lives will be born.
 - Which option would you choose?
 - Why did you choose that option? What are the best arguments for and against that option?
- Imagine you donate enough money to the Against Malaria Foundation to save a life. Unfortunately, there's an administrative error with the currency transfer service you used, and AMF aren't able to use your money until 5 years after you donated. Public health experts expect malaria rates to remain high over the next 5 years, so AMF expects your donation will be just as impactful in 5 years time. Many of the lives that the Against Malaria Foundation saves are of children under 5, and so the life your money saves is of someone who hadn't been born yet when you donated.
 - If you had known this at the time, would you have been any less excited about the donation?
 - Why, or why not?

Thinking about long-termism

- Can anyone give a definition of longtermism?
- Sometimes people use an analogy of humanity as a whole being a single person, and longtermism being analogous to someone being prudent about their future - doing things that are good in the long run, rather than only considering the short term effects. To what extent do you think this is a useful way of framing how we should be thinking about humanity as a species?
- Do you think future generations are discriminated against?

What will the future be like?

- How good do you think the future will be? What do you expect to happen in the next 100, 200, or 300 years? What technologies do you think might be developed? What problems do you think we might be able to solve?
 - Dig into what it would be like for people in the future
 - How long would they expect to be living under it
 - Eg. 100 years, millenia, locked in
- What are you afraid that the future might bring? What bad things do you expect to happen in the next 100, 200, or 300 years? What new problems might we face?
 - Dig into what it would be like for people in the future
 - How long would they expect to be living under it
 - Eg. 100 years, millenia, locked in

Our ability to influence the far future

- What do you think is the most tractable way in which we can influence the far future? How tractable is it?
- If you could learn one fact about the future, what would it be?

Arguments against long-termism

- What are the best arguments against longtermism?
 - *The best arguments are made and addressed in [The Case for Strong Longtermism](#), we recommend you read that article to familiarise yourself potential objections including:*
 - *Discount rates*
 - *Person-affecting views*
 - *Intractability*
 - *Risk aversion with respect to welfare*
 - *Non-aggregationism*
 - *Prioritarianism*

●

Week 5: Existential Risk

This week presents the case for existential risk reduction as an important cause area and introduces participants to some of the key anthropogenic risks. While this is a key example when discussing long-termism, this week focuses on x-risk reduction as an independent cause area.

Key Points

- Understand the definition of and what does and doesn't count as an existential catastrophe and an existential risk.
- There are strong reasons to care about reducing existential risks as distinct from other global catastrophes from multiple ethical viewpoints.
- Anthropogenic existential risk over the next two centuries may constitute a significant risk to humanity
- The cause of reducing existential risks is very neglected and is undervalued by markets, nations, and single generations, since existential security is an intergenerational global public good.

Tips for this session

Acknowledge at the beginning of the session that this can be upsetting and difficult to talk about but its importance is why we dedicate time to talking about it.

Session

Discussion questions

Defining existential risk

- Can someone give a definition of existential catastrophe and of existential risk?
- Other than extinction, what other kinds of existential risk might there be?

Importance of reducing existential risk

- So we've been talking about the idea that we and everyone we know could die in a catastrophe. That's pretty intense. How do you feel about it, what are your emotional reactions to it?
- Can anyone explain Ord's reasoning about why an existential catastrophe is so much worse than other global catastrophes?
- What do you think about Toby Ord's estimates of existential risk? What are the implications of existential risk being high or low?
- Do you agree with Toby Ord that existential risks are neglected by society?
- What are the best arguments that we shouldn't be too worried about existential risk?

Other

- Do you think there are other existential risks that we haven't discussed?

- Should we consider doing things that could do a huge amount of good, but don't have lots of supporting evidence, or should we do things that we have strong evidence do a lesser amount of good? How should we decide?

Week 6: Emerging Technologies

This week introduces a different approach to having impact, namely thinking about potential societal developments and transformative events. This big-picture thinking is important in its own right and provides examples of effective altruists thinking about systemic change. It should also help build up to the upcoming discussions of long-termism and existential risk.

Key Points

- One way to look for opportunities to accomplish as much good as possible is to ask “which developments might have an extremely large or irreversible impact on human civilisation?”
- Technological transformations have had a big effect on the world, and this might well continue to be the case.

Session

Icebreaker (10 mins.)

- Groups of 2-4, share your rose, bud, and thorn for the past week
 - Rose = A success, small win, or something positive that happened.
 - Thorn = A challenge you experienced
 - Bud = Something you are looking forward to knowing more about or experiencing

Discussion questions

Past trends and transformative events

- Imagine you’re trying to do as much good as you can. Would you use a time machine to try to influence history? Do you think there’s another time when you could have had more impact than now?
 - Let’s assume that you’re a similar kind of person to now (eg. an educated person living in a developed nation, but you don’t have access to future technologies and knowledge)
 - Things you might want to bring up: agricultural revolution, US constitution, industrial revolution, before/during/after WW1/2
 - Could an altruist in the past have had a greater impact if they knew the transformative events on the way? What would their best strategy have been?
- Imagine you could speak to the past and were speaking to some medieval king or queen who wanted to do as much good as they could. What would you advise them?

Future trends and transformative events

- What important changes do you expect to happen in the future? Why?
- What things do you think will happen in the next 200 years that could have a large impact on the trajectory of human civilisation?
- Is influencing potentially transformative future events a promising approach for an altruist today? Why? What might change your mind?

Week 7: Increasing the Accuracy of our Judgements

The goal of this week is to introduce participants to some key practical skills relevant to cause prioritisation and cost-effectiveness assessment (and useful in most other areas, too). The tools covered are both relevant to the discussions had in every other week and is also itself an example of a cause area many effective altruists consider important.

Key Points

- Cause prioritisation sometimes relies on uncertain, subjective judgements about the future so it's important to improve the accuracy of judgements and forecasts if we can
- There are best practice techniques to improve the reliability of our judgements and predictions about the future
- [Practice calibration, making quantitative estimates, and expressing uncertainty through calibration training](#)

Session

Icebreaker (10 mins.)

The Against Malaria Foundation is consistently a top charity of the independent charity evaluator GiveWell. Let's do some estimates about what the World Health Organisation's estimates for global malaria deaths.

- How many people died from malaria in 2000?
 - Write down your answer (60 sec.)
 - Turn to the person next to you (if online - use breakout rooms), and combine your answers to come up with a joint estimate (2 mins.)
 - Reveal Answer - 840,000 and discuss (2 mins.)
- How many people died from malaria in 2015?
 - Write down your answer (60 sec.)
 - Turn to the **other** person next to you, and combine your answers to come up with a joint estimate (2 mins.)
 - Reveal Answer - 438,000 and discuss (2 mins.)

[Source](#)

If you'd like more questions for people to do estimates of here are a few suggestions:

- In all low income countries across the world today, how many girls finish primary school?
 - 63.2%
- Between 1993 and 2013, the proportion of the world's population living in extreme poverty has gone up or down by how much?
 - The share of people living on less than \$1.9/day fell from 34 percent in 1993 to 10.7 percent in 2013, according to the World Bank.
- What is the average global life expectancy today?

- 72.5 years
- There are 2 billion children in the world today, aged 0 to 15 years old. How many children will there be in the year 2100, according to the United Nations?
 - 2 billion
- How many of the world's one-year-old children today have been vaccinated against some disease?
 - 80%
- What percentage of people in the world have some access to electricity?
 - 85.3%

[Source](#)

In-session exercise (20 mins.)

If you can feel comfortable leading this, this would be a good exercise. But feel free to skip it if you don't feel comfortable.

One of the main reasons the EA community cares about forecasting, is that having a positive impact requires us to reason about the future, which requires us to reason well under uncertainty. One area this is really important is cause prioritisation, so we're going to pick a cause, and generate forecasting questions to help us understand whether to prioritise it.

- Ask the group to suggest a cause they'd find forecasting useful for
 - If people struggle, suggest biosecurity
- As a group, generate questions that'd help understand the cause area more
 - E.g. 'How likely should we think it to be that there is a bioterrorism attack in the next 5 years?'
 - Prompt people to make the questions concrete
 - Prompt people to have clear and precise resolution criteria
 - E.g. [see a concretised question with resolution criteria here](#)
- A good source of questions is Metaculus, there's likely to be a [category](#) for the relevant cause (eg [pandemics](#))
 - If people seem interested, send them a link to Metaculus and encourage them to have a look after the session

Discussion questions

Discuss the exercise

- How did you find the calibration training exercise?
 - Were you under confident or over confident?
 - Do you think people in general tend to be overconfident or under confident?
 - If your calibration improved over the exercises, why do you think that was?
 - Do you have a better sense for what particular probabilities feel like?

What makes a good forecast?

- Can someone summarise the main ideas of forecasting, and what makes a good forecaster?

- Did any of the tips on how to do forecasting surprise you?

How good are we at forecasting?

- Do people agree with Tetlock that ‘political pundits and the media seem to be especially bad at making predictions’? Why?
- Can you think of any examples where you made a bad prediction in the past? Why was this?
- Can you think of any societal or political forecasts that were bad? Why do you think that was?

Implications

- Of the cause areas we’ve discussed so far in the Program, which do you think rely the most on uncertain judgements and forecasts?
- What questions would you be most interested in trying to forecast?
- How can we get better at updating on evidence?
 - What heuristics/techniques can we use?
 - What kinds of evidence should we most pay attention to?

Alternatives to forecasting

- Does it make sense to have probability estimates for everything and anything?
 - What kind of questions might you not be able to give credence in?
- What other societal mechanisms are there for humanity to improve it’s beliefs?
 - For example, how good is academia at this?

Week 8: Putting it into Practice

This week has three aims: wrapping up, evaluating the program and encouraging participants to apply effective altruism to their own life. Make sure you have an overview of the relevant options for staying involved, e.g. groups, mailing lists, programs or events to advertise.

Key Points

This weeks’ content acts as a recap of a lot of the material and concepts explored throughout the program. It makes these ideas action relevant by showing that EA affects people’s career choices. Participants will hopefully feel motivated that they can do lots of good with their lives

- Advice on how to make progress on your own
 - Map of good resources to explore
 - Practice identifying key uncertainties and noticing confusion
- Participants have an opportunity to make a plan to continue making progress on their own

Tips for this session

Emphasise that personal fit considerations could have a major impact on what career path someone chooses, and that this should be factored in as well as cause prioritisation.

Session

Survey

First of all, get the participants to complete your post-program survey in-session. This dramatically increases the number of people who complete the survey, and provides valuable feedback, which is really important.

Discussion questions

Discuss the exercise

- Look at your list of topics or ideas from the program that you felt you didn't understand. Out of this list, which do you feel most excited about learning more about, or which seems most useful to get more clarity on?
- Ask people to share something that they learned in the program and gently prompt them for how it might affect their life and plans

Discuss how to get involved

- What's the best way to have an impact while as a student?
 - Most of your impact will be in your future, when you have more skills and career influence
 - Being a student is a good time to reflect on your values and think about the direction of your life, as well as to network and find other like minded people
 - Learning about the world
 - Deciding which problems you think are most pressing - which problem to work on is a huge and important question.
 - Build skills, explore different career routes, find ways to improve your CV
 - Network with others (again, the EA community can be a great friend here)
 - Networking can feel 'icky' - instead think about it as finding others who you find interesting and get along with
 - Put in time to think and plan your career, while you have more space to think.
- Discuss what your local group has to offer

Shoutouts (15 mins.)

Set aside some time where people can give shoutouts to other participants or organizers or generally remark on their favorite moments. (e.g. I wanted to shoutout X because I thought it was really neat how he changed his opinion on Y)

This is a nice way to wrap up the program on a warm note. Participants can take some time to warm up and think of ideas, so it's useful to think of your own shoutouts to participants in advance.