

# How to run Remote IDF - for Cohort Organisers

## The Point of this Document

The remote In-Depth Fellowship is a big project with a lot of moving parts. This document contains everything you need to know as a group organiser to successfully run your sessions. Practical details are contained in the [onboarding document](#). Message Will Payne if you have any questions or confusions about this document.

## The Point of this Fellowship

The standard in-depth fellowship has a few main goals:

1. Give fellows a deep understanding of some of the most important ideas within EA, and make them aware of the open questions and debates
2. Give fellows a toolkit and the experience necessary to start thinking seriously about these issues themselves, particularly insofar as they relate to career or donation opportunities
3. Fellows feel welcomed into the core of the community, and make genuine connections with others that will sustain past the end of the fellowship

This remote fellowship has similar goals with a few extra benefits, making this a particularly exciting opportunity:

1. We have significantly wider reach than normal, which is exciting in itself, and also provides value to other groups, whose members will hopefully be better placed to run a similar event at their group in the future.
2. Secondly, the lockdown likely gives university students more time to think about EA and more of an inclination to come to sessions.
3. Finally, a remote fellowship seems like a very scalable approach to community building - in particular, it integrates very well with Alex/James/Huw's career mentoring program, meaning that there's a lot of information value in testing out this approach.

The reason I think this is important for session organisers to know is to appreciate this amazing opportunity to develop the global EA community; hopefully the goals also help give some context to the guidelines below.

# Outline of the Fellowship

The fellowship will have 8 weeks of sessions (1.5 hours long) with reading for each session (the curriculum is [here](#)). The current week by week plan is:

1. **Epistemics & Decision Theory**
2. **Moral Uncertainty, and Moral Progress (inc. moral circle expansion)**
3. **\*\*Model building: Cause Prioritisation\*\***
4. **Cause Prioritisation (& Cluelessness revisited)**
5. **Global Health and Animal Welfare**  
**\*\*Grantmaking exercise: GiveWell vs ACE\*\***
6. **Longtermism and Population Ethics**
7. **Grand futures and attendant risks**
8. **Alternative Priorities & Open Questions**

Each group will take a break for a week after session 3, to allow us to reflect on how things are going, but the exact week in which the break happens could vary by cohort.

## Career mentoring

Around 75% of the participants have participated, or will be participating in the career mentoring program. Huw and Alex will be having calls with participants throughout the 8 week period (and likely beyond) and the hope is that this will increase how 'practical' the fellowship feels for the participants.

As a moderator, the main takeaway from this is that you can assume that most people will be working towards planning an EA career, and that bringing the conversation back to careers every couple of sessions (if it seems appropriate) could be a good way to increase the practicality of the fellowship.

## Slack

We're hoping to cultivate a really excellent intellectual community on the fellowship slack - in a similar vein to the EA Fellows FB group. For that reason, we encourage you to post up interesting thoughts or questions that come up (inside or outside of the fellowship) on the different channels, and to engage in any discussions that you find interesting.

# Responsibilities of Session Organisers (and why)

Each week of the curriculum requires 1-5 hours of prep, mostly dependent on how long you need to take on the reading. Each fellowship session lasts 1.5 hours, plus 15 mins for reviewing afterwards.

Session organisers are expected to each week:

- **Prepare for sessions (1-5 hours)**
  - Sessions include a lot of content and reading up beforehand is essential for moderating the discussions
  - If you've already read the relevant resources, you can spend less time on this - use your best judgement
  - The boldened resources are the only resources you absolutely must have read, though you'll provide the most value if you can explore a few of the other resources as well
  - [There are notes on the readings that help you get the main points here](#), compiled by Alex as he goes through each week
  - [There are notes on how to run each session here](#), compiled by Max as he goes through each week
- **Run sessions (1.5 hours per session)**
  - Moderating sessions is needed to keep the discussions on topic
  - If you're unable to make a session one week, post on the cohort leaders channel in the fellowship Slack to ask for someone else to take your session, or swap sessions with you
- **Post-session form (15 mins per session)**
  - This is useful for other organisers to learn from your sessions, for you to improve your moderation, and for you to keep notes on the participants within your session.
  - After each session, fill in this [post-session form](#).
- **Participate in the Slack conversation (NA: ~1 hour is a good amount of time to aim for)**
  - We want participants to get to know each other and other people outside of their group, so we aim to have lots of good discussion on the Slack
  - There's no official requirement to spend this amount of time; we're happy for you to vary this based on how excited you are to interact on the Slack
- **Check-in with Will once a week (~15 mins or less)**
  - I'll be checking in to make sure everyone is feeling good about things and can work through problems as we find them :)
  - This will also be an opportunity for you to give thoughts on how to improve the sessions, or feedback on the overall management of the fellowship

# Tips for Running Sessions

## Preparing for Sessions

The main thing to do in order to prepare for a session is to read the reading list. It's not necessary to feel incredibly resolved or like you can teach the topic in the reading as that's a lot to ask, and anyway the ideal session will have other participants explaining ideas and forming hypotheses way more than you

## Session Norms

- Facilitation of discussion remotely can be more awkward without visual signals from body language, so **more active moderation will likely be required**.
  - At the start of the session, clarify that you'll be moderating. Explain that at any point, people can post the following symbols into the Zoom chat to indicate that they want to speak, and that it'll be up to you as moderator to chime in with how has the next comment or point:
    - “Hand” or 🖐️ to mean “I want to discuss a new Point”
    - “finger” or 📂 to mean “I want to comment on the existing point”
    - “fist” or 👊 to mean “I am confused/jargon/clarification required”
- Ask people to mute their mic when they are not speaking
  - You might want to experiment with this - some groups have reported more lively conversations if mics are left on, so perhaps start with mics on and ask people to switch off if they're creating noise
- If possible, ask people to display their name (Hosts and participants can change names in Zoom)
- Request that people turn on video if possible

## Session structure

- In the first session, give some context on the overall program
  - Introductions:
    - Introduce yourself and ask each of the fellows to introduce themselves
  - High-level details:
    - 8 week program, we'll take a week off a few weeks in
    - Aim is to go deeper than most EA discussion groups are able to go, with a particular eye towards resolving uncertainties within cause prioritisation and inform career decisions

- Also a really great networking opportunity! There are over 100 EAs in the Slack, all of whom are working towards EA careers and keen to connect
    - We'll be running some larger social events, but we're particularly keen for you all to book in chats with people with similar interests - this is a really amazing opportunity to make some lifelong connections!
- Explain how the moderation will work ( 🙌🙌🙌 )
- On resolving disagreements
  - 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.*
- Start with some quick icebreakers - you can keep it simple and run with eg. "[thorn, rose and bud](#)", or vary it by week (here's a [big list](#)), or get inventive and come up with your own question each week. In the first week, it might be good to have everyone talk about themselves for a minute, eg. who they are, what they're up to, and how they became involved with EA.
- Introduce the week's topics briefly, and say how much time we're going to spend on each section
- Discuss!
  - Straying off-topic to explore interesting side-roads is fine, as long as the discussion feels "alive" and you're making progress.
  - If a conversation gets stuck in an unproductive place, don't hesitate to interrupt people.
  - When you've exhausted one topic, you can consult the list you compiled at the beginning. Feel free to dynamically extend the list during discussions.
  - If a disagreement gets heated or otherwise stuck, it's often good to interrupt and reframe the discussion. Some things to try:
    - Ask "What would make you change your mind?" or "Can you imagine making a specific observation (gaining knowledge of a fact about the world) that would change your mind?"
    - Do an intellectual turing test: describe the views of your interlocutor in a way that they themselves endorse.
    - Define your terms to avoid getting stuck in empty semantic misunderstandings.
- Formats - feel free to experiment with discussion formats for the group. We'll sketch out three formats which we think could be worth experimenting with:
  - **Format 1:** Prior to each session, have the group brainstorm topics to discuss. Create a [copy of this template](#) in the "[Shared Notes](#)" folder within the drive, and

call it “Cohort X notes”, where X is the number of your cohort. Consider making notes here as you work through the reading yourself, too.

- Bolden a few topics that you think could be particularly fruitful to work your way through during the session, and go through them one by one.
- **Format 2:** Prior to each session, send a list of questions about this week’s topic (eg. from [this list](#)) that you think people will likely have meaningfully different answers to. Have people write down some answers to the questions prior to the session.
  - Then, in the session, for each question, we go around the circle and give our response to it.
  - “Most of the interesting discussion happens as prompted by things people say when providing their answers. It’s pretty loose: people ask you about your answer, and you respond to them, and the discussion kind of goes where it will. When it’s gone on a long enough tangent, the moderator moves the conversation along by asking the next person to give their answer to the question. This is a great way of ensuring that the conversation can stay roughly on track and everyone gets a chance to speak. Often meetups have a problem where some people don’t know what to say; this goes some of the way to solving that.”
- **Format 3:** This is the simplest way to run the group, and doesn’t require any preparation: simply start the call, and begin [asking questions](#) about each of the different topics - allow the conversation to go where feels natural, and interject with a new question, or a new overall topic, if a conversation thread begins to be stalling or you’re beginning to reach the end of the allocated time for this topic.

## Goals and Heuristics for Session Moderating

Moderating sessions is not about teaching. Cause prioritisation is a complex topic and the syllabus certainly doesn’t contain an absolute answer to the questions raised. The purpose of a moderator is to keep moving the sessions in a useful direction for the fellows. Here are a few goals and heuristics to make this easier:

- **Clarify jargon and definitions** - or better still, have others clarify this for you
- **Minimise teaching**
  - You should be able to summarise the main argument from each article, and be able to define the relevant terms for people etc. - but ideally, you want to minimise the amount of teaching you do.
  - Where possible, remind others in the group to define their terms the first time they use them, or ask others to summarise an article’s argument before stepping in
- **Remind fellows of useful framings of ideas**

- One goal of the fellowship is to help fellows 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?"
- **Remind fellows of useful/relevant concepts**
  - As a cohort leader, you have a wider overview of useful concepts or upcoming topics. Where relevant it can be helpful to bring these up.
    - "If we think the number of animals on the planet is a key factor maybe we can do a quick Fermi estimate on that" [explain what a fermi estimate is]
- **Promote good discussion norms**
  - Having good norms can make a discussion far more useful - see "session norms" below. Enforcing hand signals can reduce interruptions (which may increase due to the virtual format) and being an example for epistemic norms can keep the discussions productive and truth-seeking.
    - "It seems like A has something to add to that point..."
    - "Should we take that as strong evidence for the claim?"
    - "Seems like we all 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?"
- **Resolve factual disagreements where possible, move on otherwise**
  - Often discussions can get lost in easily resolvable disagreements, as a cohort leader you might have access to useful information that can solve these
    - "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"
- **Keep track of whether the discussion feels useful for the majority of the group, and move on if not**
  - "This seems like an interesting point, but I think we should move on for the sake of time"

## Grantmaking and Model Building Sessions

These sessions are designed to stimulate a more active part of the brain with the hope that this will help fellows interrogate their models of EA, see the more practical side of EA, and appreciate why things are the way they are by seeing them in the context of grantmaking and model building. I'll describe the mechanism for each of these in bullet points to help clarify them so you can apply them more easily with moderation. I'll also give moderation tips below. (I'm

talking about “practical EA application” vaguely here but in the specific moderation tips I’ll explain the structure of the sessions.

- Interrogating models of EA
  - A lot of the time when we learn new concepts we place them inside our mental model of the world kind of unconsciously which means they might be in the wrong place. For example when someone learns about “Fermi estimation” you’d want it to occupy a mental space very close to the idea of “needing a quick estimate of something” but this is hard to test just by talking abstractly about the concepts.
  - This is where these sessions are supposed to help. The idea is that by trying to actively solve a problem in an EA context someone can see how these ideas map to problems and exercise those connections more.
  - It’s through this as well that people might appreciate why concepts are the way they are in EA, someone might notice that their understanding of an EA concept isn’t useful in a practical context and have to relearn it to apply it to the problem, e.g they have a vague picture of the INT framework but in practice it’s tough to apply and they might reevaluate the framework or their understanding of it (this happened to me in the sessions)
- Appreciate the practical side of things
  - There is a common objection to EA which I see in new members of the society which goes something like “X seems to be not very useful to doing the most good is EA too philosophical”. I think this objection can be reasonable and nuanced a lot of the time but can also sometimes be given too quickly
  - These exercise based sessions provide a bit of a counterpoint to the naive version of this objection by giving the participants a flavour of EA in practice and allowing them to see how ethical and decision theoretic problems, for example, seem to come up naturally and so may be important to consider
  - Another benefit to seeing the practical side is to give new members of the community the permission to work on tough problems. I think it can be easy to think you’re not ready yet for tough problems when actually these problems are just tough and unapproachable for everyone. By engaging in a practical example the hope is fellows can spot the paths to solving the problem and feel like they have the ownership and agency to take it on themselves.

## Model Building

The model building session involves taking some prompts and attempting to build a model from them. I’ll list some prompts here and plan to have a load of them by the time we run sessions:

- **Define Hinginess**



- This is a quite mathematical problem (and maybe no longer useful as higness hasn't been discussed in weeks 1 or 2) the idea is that fellows try to figure out a model for if a given century is at the hinge of history.
- I think in practice this will look like listing a bunch of considerations and then vaguely discussing how these might be combined and what this would imply about the hingyness of the current century
- **What makes a core member of the EA community**
  - This can be heavily informed by CEA's awareness inclination model for the moderator but I'd prefer fellows not to be aware of it (so they're not too anchored)
  - I think in practice this will look like listing a bunch of considerations and then vaguely discussing how these might be combined and what the best events to set up might be as an EA community builder.
- **How can Fish Welfare Initiative (FWI) best help fish**
  - This is a lot more qualitative than the other two and looks more like a brainstorming session, the goal is to provide a few interventions an imaginary brand new FWI might take and perhaps a vague ranking of those interventions.

So the session involves constructing a model in order to make some decisions as an EA. I think in keeping with the goals of the session it makes more sense to keep thinking and working on the problem than to solve the problem. This means that if at any time the fellows get stuck, or find the problem too easy it makes sense to move on to a different aspect of the problem or reframe the current one to involve more active thinking.

**Script for the session** (although feel free to heavily paraphrase)

- Say hi and do normal session opening
- This session is much more focused on applying our ideas about altruism than developing them.
  - The hope is that it'll inform our understanding of EA and reveal some holes in our knowledge to work on in later sessions and give us a better idea of what all this philosophical stuff is for
- *Group spends some time discussing the idea of a "Model" and how to construct them as well as if they are useful given they're constructed by us.*
- In the session we'll **pick one** of the prompts provided and try to answer it by building an explicit model of our thinking here. The idea is to try to capture as much of our understanding of the problem in explicit considerations as possible whilst keeping this explicit model useful for solving the problem.
- The real point of the session is to practice actually putting the ideas of EA into practice, which can really help challenge and strengthen your understanding. The problems given in the prompts are real and difficult research questions which can't be solved in an hour and a half so we should move on if we get stuck and simplify if we need to. This way we can spend as much time doing what's useful, which is thinking about how to construct models and trying to develop our understanding of EA.

- *Groups should now discuss whichever prompt they find the most appealing*
  - **Prompts the moderator can use to keep discussion useful**
    - “This seems like an important factor here, how can we make this explicit for the model?”
    - “So it seems like our model covers X, Y, and Z, aspect of the problem are we missing anything?”
    - “Yeah this seems like a tough problem, can we make a ballpark or common sense guess here for now and move on?” - Maybe this can be discussed later if there’s time or brought up on Slack
    - “I don’t want to get too bogged down in details here, what’s a good description of this question for our model? Then we can move on to something else.”
    - “What is our goal in making this model?”
    - “How is what we’re doing useful in answering the prompt?”
    - “Does everyone feel confident about the process we’re using, is this the best/only way to construct a model”
    - “It seems like here we could be taking a cluster thinking approach”
    - “Maybe a quick way to solve this sort of problem is to take the epistemically modest view and check for expert consensus”

## Using the Post-Session Form

After each session, spend 10+ minutes filling out [this form](#). This form exists to help you reflect on a few key things after each fellowship session:

1. Lessons and takeaways for others to use when they run this session
2. Writing down some notes to help build your understanding of the participants
3. Reflecting on how you might be able to improve as a moderator based on the session

See [deliberate performance in people management](#) for some context on this last point.

## Norms for Slack Communication

Because we’ll have lots of moderators working on this, and we’ll be taking so many people through the process, we’d like pretty strong norms of communication between moderators in order to make sure we’re able to optimise the sessions, and the overall process.

There’s no official requirement here, but we’d encourage you to post on the Slack whenever you have questions or uncertainties that you’d like other people’s input on, or insights that you’d like to share with the other moderators.