Effective Thesis - Impact Evaluation: July 2021

Goals & Theory of Change

Effective Thesis (ET) aims to create a world where there are a lot more researchers focusing on x-risks, animal sentience and welfare/alternative protein, global development and other priorities that people in EA consider important, such that “not enough knowledge/understanding” or “not knowing how to proceed in solving these problems” is a relatively smaller obstacle compared to today.

ET aims to do this by steering students who will become researchers towards EA-aligned research directions. The ultimate goal of ET is not to influence the focus of individual student theses, but whole research careers. It influences careers via thesis topic choice, since that seems to be a crucial point of research career trajectories (see some evidence that researchers tend not to diverge much from their initial topic choice e.g. in Jia, Wang & Szymanski, 2017). The main focus of the project is on students who are likely to become researchers anyway, not adding more researchers to the system. Chances are that these students will become researchers within academia, rather than within EA research orgs (because of the sheer number of positions available in academia vs in EA), which carries the promise of scalability of this intervention.

Goal 1: Influencing students' long-term research focus via assisting them with thesis topic choice

Other effects that we should track:
Effect 1: allowing students to better test their fit for research jobs (in general as well as specifically at EA orgs)
Effect 2: mitigating the structural bottleneck in EA research space (via helping EAs make use of non-EA resources like supervision and training; and helping them set up their impactful careers in academia)

Goal 2 (potential expansion): bringing new people to EA
Goal 3 (potential expansion): changing trajectories of people who don't aim to become researchers

It is also likely that having a presence in academia gives EA ideas more credibility and is important for the long term movement flourishing, even though independent research in EA orgs might be more directly relevant to EA goals.

However, we are wary about the way EA ideas are promoted, since we assign non-trivial credence to views where EA growth should be slower and low-fidelity transmission is generally negative. We aim to prevent these risks by building high fidelity content and connecting students with mentors experienced in EA to provide space for higher fidelity discussions.

Moreover, we try to prevent potential risks by refraining from using explicit EA labelling and rather focus attention on specific research directions one is interested in. We offer opportunities to get more involved with the EA community and encourage that later on, but this is treated as an additional resource that students can make use of and where they can potentially meet other people interested in the similar research directions, not as an umbrella under which students operate from the beginning. This should prevent risks of diluting the EA brand or harming its credibility by low quality or otherwise harmful work.

How do we track impact?

We are tracking our impact via self-reported testimonies of students who we have served. We ask them about changes in their research focus that they have made and how likely they would be to make such changes if ET didn't exist (trying to get a sense of counterfactual impact that way). We do that for direct effects on thesis topic choices and also for long-term effects on their research careers. We also try to capture other types of value that ET produced for students in a similar way. *For example, we found that network effects (i.e. connecting students with other relevant people) is referred to by approx 20 % of students as the main counterfactual value they got out of the interaction with ET.*

We send students a general feedback form the month after their submission/graduation deadline, and then a long-term feedback form 12 months after their submission/graduation deadline.

Results

Generally, feedback loops are pretty long in this project (up to 2-3 years to get information on how Effective Thesis intervention influenced students' real career choices). It's worth keeping in mind that some data referred to here were created with our 2018 level of services, not 2021 level.

Best case studies

Based on my previous analysis and investigation of current data, I have created an assumption that most of the impact will come from a small number of cases. Summaries of the best case stories should thus be a good proxy for impact. For the full list of the best case studies that Effective Thesis has produced, see this document.

Year to year progress

Since June 2020 (when I did the last collection of best case studies), we have added 11 case studies - an increase of 44% compared to all studies we collected before (n=25). Whilst it is difficult to understand whether the impactfulness/quality threshold for including new studies has gone up, down, or remained the same, we have compared the relative rankings of best case studies across the time span in which the students applied for ET's services. Looking at the impactfulness of our case studies over time (see in-depth below), the quality threshold appears to be remaining stable (with a slight blip in July 2018-2019).

Descriptives of the best cases

→ Best case studies - Google Sheets

Impactfulness of best case studies over time

Long-term effects tracked

Time period applied in	*No of best case studies*	*Ranked case studies in terms of impact (best-ranked case study is 1, the worst ranked case study is 4)*	*Total ranking values*	*Average ranking (lower is better)*
July 2017 - July 2018	1	2	2	2
July 2018 - July 2019	1	3	3	3

July 2019 - July 2020	2	1, 4	5	2.5
July 2020 - Jan 2021	-	-	-	-

Short-term effects tracked

Time period applied in	*No of best case studies*	*Ranked case studies in terms of impact (best ranked case study is 1, worst ranked case study is 31)*	*Total ranking values*	*Average ranking (lower is better)*
July 2017 - July 2018	4	3, 10, 19, 26,	58	14.5
July 2018 - July 2019	13	4, 12, 17, 20, 21, 22, 24, 25, 27, 28, 29, 30, 31,	290	22.3
July 2019 - July 2020	8	2, 5, 6, 7, 11, 15, 16, 18	80	10
July 2020 - Jan 2021	6	1, 8, 9, 13, 14, 23	68	11.3

Types of value generated by ET for our best case studies

Of the 36 best case studies:
50% mentioned value generated by ET in terms of *impact on thesis topic*
39% mentioned value generated by ET in terms of *connections made*
29% mentioned value generated by ET in terms of *research focus trajectory*
22% mentioned value generated by ET in terms of *internship or other opportunities gained*
14% mentioned value generated by ET in terms of *career plan trajectory changes*
8% mentioned value generated by ET in terms of *impact on personal morality*
6% mentioned value generated by ET in terms of *better introduction to EA and the EA community*

Best case studies and their level of study

53%/31%/17% of best case studies were studying at undergraduates/masters/PhD level respectively compared to 39%/32%/28% of 2020 applicants and 35%/42%/23% of all applicants

Best case studies and their study discipline

14% of best case studies' discipline is *philosophy* compared to 10% of 2020 applicants and 7% of all applicants.
8% of best case studies' discipline is *computer science* compared to 6% of 2020 applicants and 7% of all applicants.
8% of best case studies;' discipline is *psychology* compared to 10% of 2020 applicants and 7% of all applicants.

8% of best-case studies' discipline is *economics* compared to 10% of 2020 applicants and 8% of all applicants.
8% of best-case studies' discipline is *politics/international relations* compared to 7% of 2020 applicants and 5% of all applicants.
The sample size is too small to draw any conclusions about the over/under-representation of particular disciplines.

Best case studies and their cause area
28% of best case studies have a cause area of *global priorities research* compared to 11% of 2020 applicants and 7% of all applicants.
14% of best case studies have a cause area of *AI Governance/Policy* compared to 10% of 2020 applicants and 6% of all applicants.
11% of best case studies have a cause area of *animal welfare* compared to 5% of 2020 applicants and 6% of all applicants.
8% of best case studies have a cause area of *biosecurity* compared to 4% of 2020 applicants and 3% of all applicants.
3% of best case studies have a cause area of *AI safety/governance* compared to 13% of 2020 applicants and 9% of all applicants.
6% of best case studies have a cause area of *improving institutional decision making* compared to 6% of 2020 applicants and 3% of all applicants.

Best case studies and their university ranking
Looking at the top 100 universities as ranked by the Times Higher Education 2020 University rankings: 25% of best case studies attend top 100 universities compared to 29% of 2020 applicants and 27% of all applicants.

Best case studies and their country of study
17% of best case studies are from the *UK* compared to 19% of 2020 applicants and 18% of all applicants.
14% of best case studies are from *Germany* compared to 11% of 2020 applicants and 10% of all applicants.
11% of best case studies are from *Australia* compared to 4% of 2020 applicants and 5% of all applicants.
8% of best case studies are from *Czech Republic* compared to 1% of 2020 applicants and 6% of all applicants.
3% of best case studies are from the *US* compared to 13% of 2020 applicants and 13% of all applicants.

DATA: Impact questions July 2021 - Google Sheets

Aggregated descriptive data

Rather than summarizing the impact produced, I believe this section could provide more context on the project operations, which could be useful.

Number and demographics of applicants

In 2020, there were 174 applications. This is almost exactly the same as in 2019. The reason for stagnating in the number of applications is that we were going through organisational changes and did very little proactive outreach in 2020. I expect the number of applications to go fairly up (perhaps 2x or more) once we start with proactive outreach in 2021.

The demographics of these students suggest they come from a very wide range of countries with EU, UK and US amounting to only around 60% of all applications; around 27% of all applications come from one of the top100 universities based on THE ratings; they come from a wide range of disciplines and with interest in a wide range of EA cause areas.

The average time spent on thesis research and writing is approx 600 hours/person, based on data from 11 people from who we collected data. The average grade for their thesis was the best possible (grading systems are very different across countries so it is not easy to convert them to one measure)

DATA:
https://docs.google.com/spreadsheets/u/1/d/1cYus_hpftj9ONw3bsBfktt6LioGvwI5F_MPonHdGzWw/edit#gid=170936743

Effectiveness of our outreach targeting:

Proportion of our students who progressed in research careers (compared to general proportion):
30% of best case studies were still on track for a research career 12 months after their graduation/submission deadline compared to 20% of 2020 applicants and 18% of all applicants. It was unclear whether or not students were still on track for a research career for 20% of best case studies, 49% of 2020 applicants, and 55 % of all applicants.

This suggests that the ET theory of change in targeting students likely to remain in research is somewhat effective, although of course we could further improve our targeted outreach. It is also worth noting that for 20% of our best case studies whether or not they are going to remain in a research career was unclear, due to the aforementioned long feedback loops and length of studies.

Proportion of students who have not identified as EAs at the time of application

Using a ranking of 0-2 to label students prior knowledge of EA with 0=have not heard about EA before applying; 1=have heard about it but have not been highly involved; 2=have been highly involved prior to applying (e.g. reading EA forum and other blogs a lot; regularly attending local EA group meetings; being well-networked in the EA community, etc).

0% of best case studies gave a familiarity with EA ranking of 0 compared to 7% of 2020 applicants and 19% of all applicants
33% of best case studies gave a familiarity with EA ranking of 1 compared to 13% of 2020 applicants and 20% of all applicants
8% of best case studies gave a familiarity with EA ranking of 1-2 compared to 2% of 2020 applicants and 4% of all applicants
44% of best case studies gave a familiarity with EA ranking of 2 compared to 22% of 2020 applicants and 23% of all applicants

8% of best case studies didn't answer this question but found ET through EA related means and therefore had likely at least heard of EA before, compared to 38% of 2020 applicants and 19% of all applicants.
0% of best case studies didn't answer this question and found ET through non-EA related means and therefore probably hadn't heard of EA before, compared to 8% of 2020 applicants and 7% of all applicants.

This suggests that ET is most effective at targeting students with some pre-existing interest in EA perhaps due to the commitment needed towards effective altruism principles in order to reroute one's entire career and research trajectory around such values.

Proportion of PhD student applications
17% of best case studies were studying at PhD level compared to 28% of 2020 applicants and 23% of all applicants.
However, due to the relatively long feedback loops associated with PhD study, especially in comparison to study at the undergraduate and master's levels, it is difficult to draw any conclusions at this point.

Goal 1: influencing theses topics

Influencing research topic choices

About ⅓ of students who filled in our feedback form said that the main value ET generated for them was influencing their topic choice, either by helping them choose a new specific topic or by changing the general direction of their research focus. This service also seems to have the most counterfactual impact, with more than half of these students saying there is only a 40 % or lower chance they would've changed their research direction if Effective Thesis didn't exist.

In the remaining ⅔ of students the impact on their topic choice was either indirect (e.g. via connecting them with relevant people and helping them find supervisors) or the impact was produced in other ways not including topic choice (e.g. helping them with refining idea they originally came up with themselves or giving them useful career advice).

This is based on data from approx 104 feedback forms that we sent out to people who had their thesis submission deadline/graduation before June 2021. We have received 44 answers to date, indicating a response rate of approx 42 %. I assume that mostly people who were more engaged and happier with our services filled in our form, so one should not extrapolate results to the whole sample. Rather, I guess there is little impact hidden in the non-responding cases, and thus the responses should cover most of the impact produced. It is also true that the response rate for long-term feedback forms is generally pretty low, even for some of our best case studies, so it's possible with greater follow-up from our team that we would uncover more best case studies and examples of impact;
Note the mismatch between total number of applications received throughout the existence of the project (approx 500 to date) and number of people who have already finished their thesis and were sent a feedback form (approx 104, again throughout the existence of the project), which suggests that most results will unfold in the future as more students will finish their thesis and will be sent a feedback form. Students who already filled the form also usually applied in 2019 or 2020, so their answers do not track the most recent improvements we made to our services.

DATA: Typeform - Results
ANALYSIS: Evaluation June 2020 - working version - Google Docs

Long term effects

There is still too little data to make conclusions (we have so far received only 15 answers to our long term feedback form which is sent out 12 months after graduation date, which often means almost 24 months after students apply). For this survey, the response rate was 15 % (the form was sent to approx 100 people, excluding the early drop-outs).
Moving forwards we will primarily target people in whom we have identified some impact (i.e. people who have filled in the general feedback form) and since this is going to be a lower number, we will spend more time on eliciting feedback from them - e.g. writing personalized emails or scheduling calls instead of surveys.

Anyway, from that preliminary data it seems that Effective Thesis has some longer term influence on career trajectories (4 out of 12 people who answered that question said that they attribute more than 50 % of the influence over their career trajectory change decision to Effective Thesis).

DATA: Typeform - Results

Goal 2: bringing new people to EA

See the above information on the *"Proportion of students who have not identified as EAs at the time of application".*
17% of 2020 applicants and 26% of all applicants had not heard of EA before applying, and none of these students got into our best case studies list.
In comparison, 44% of best case studies gave a familiarity with EA ranking of 2 compared to 22% of 2020 applicants and 23% of all applicants; and 33% of best case studies gave a familiarity with EA ranking of 1 compared to 13% of 2020 applicants and 20% of all applicants.

How the drop out rate from ET services varies according to prior knowledge of EA
Of the students who rated their prior knowledge of EA as 0, 47% cancelled/dropped out of the process before the first call, compared to 26% of all applicants.
For students who didn't answer the prior knowledge of EA question but came to EA through non-EA related means (and thus probably weren't familiar with EA before) 47% cancelled/dropped out of the process before the first call, compared to 26% of all applicants.

This suggests that ET is most effective at targeting students with some pre-existing interest in EA perhaps due to the commitment needed towards effective altruism principles in order to reroute one's entire career and research trajectory around such values.

DATA: Impact questions July 2021 - Google Sheets

Goal 3: changing trajectories of people who don't aim to become researchers

Interest in research careers at the time of application
42% of best case studies answered *yes* to wanting to become a researcher at the time of application, compared to 18% of 2020 applicants and 26% of all applicants.
17% of best case studies answered *maybe* to wanting to become a researcher at the time of application, compared to 14% of 2020 applicants and 12% of all applicants.
14% of best case studies answered *no* to wanting to become a researcher at the time of application, compared to 6% of 2020 applicants and 13% of all applicants.
28% of best case studies *didn't answer the question* on whether they wanted to become a researcher at the time of application, compared to 62% of 2020 applicants and 49% of all applicants.

Of the students who answered the interested in becoming a researcher question and where more than 12 months had passed since the graduation/submission deadline:
33% of those who answered yes to wanting to become a researcher *were still on track for a researcher career* 12 months after the graduation/submission deadline, compared to 6% of

those who answered no to wanting to become a researcher, and 16% of those who said they were unsure about becoming a researcher.
40% of those who answered yes to wanting to become a researcher *were not on track for a researcher career* 12 months after the graduation/submission deadline, compared to 38% of those who answered no to wanting to become a researcher, and 40% of those who said they were unsure about becoming a researcher.
The *progress of their career (e.g. still in research or not) was unclear* for 28% of those who answered yes to wanting to become a researcher, compared to 54% of those who answered no to wanting to become a researcher, and 44% of those who said they were unsure about becoming a researcher.

Please also see the information above on the *"proportion of our students who progressed in research careers (compared to general proportion)"*. This information is relatively hard to track due to the long time frame of a career and also the low response rate to feedback forms.

Evaluation of specific services

These are questions we are working to answer in the coming months

- How to get a grasp at how much impact is created solely by content vs by coaching vs by opportunity search?
- Should we have some form/expression of interest/some other way to catch people who are not interested in coaching but who have benefitted from our content?
- Should we start offering opportunities search as an independent service for people not interested in coaching?

We are currently working on tracking and analysing the impact of our opportunities search services to understand how useful this service is for our students.