The Effects of Career Advice on Expected Impact for Animals

Abstract

There is limited research on the effects that career advice can have on individuals' expected impact for altruistic causes, especially for helping animals. In two studies, we evaluate whether individuals who receive a one-to-one careers advice call or participate in an online course provided by Animal Advocacy Careers (a nonprofit organisation) seem to perform better on a number of indirect indicators of likely impact for animals, compared to randomly selected control groups that hadn't received these services. The one-to-one calls group had significantly higher self-assessed expected impact for altruistic causes and undertook significantly more career-related behaviours than the control group in the six months after their initial application to the programme. There was no significant difference between the two groups' attitudes related to effective animal advocacy or their career plan changes. In contrast, the online course group made significantly higher levels of career plan changes in the six months after their application than the control group, but there were no significant differences for the other metrics. A number of supplementary analyses were undertaken which support the main conclusion that the one-to-one calls and online course likely each caused meaningful changes on some but not all of the intended outcomes.

Keywords: career decision-making; career decision-making difficulties; career interventions; career behaviors; animals

Introduction

Mentoring and coaching have been found to provide a number of psychological benefits¹ as well as to support positive behavioural changes.² Such interventions can also have positive effects on career-related outcomes. For example, Luzzo and Taylor (1993) found that a one-to-one career counseling session including persuasive messages focused on increasing college students' career decision-making self-efficacy (i.e. their belief that they can successfully complete tasks necessary to making significant career decisions) was successful in increasing participants' self-efficacy at posttest compared to the control group.³ Tarigan and Wimbarti (2011) similarly found that a career planning program was effective at increasing graduates' career search self-efficacy.⁴ Likewise, longitudinal studies have found evidence that careers courses can increase career decision-making self-efficacy or certainty and decrease career decision-making difficulties or indecision.⁵

¹ See, for example, Lillian Turner de Tormes Eby et al., "An Interdisciplinary Meta-Analysis of the Potential Antecedents, Correlates, and Consequences of Protégé Perceptions of Mentoring," *Psychological Bulletin* 139, no. 2 (2013), 5-6 and Tim Theeboom, Bianca Beersma, and Annelies E. M. van Vianen, "Does Coaching Work? A Meta-Analysis on the Effects of Coaching on Individual Level Outcomes in an Organizational context," *The Journal of Positive Psychology* 9, no. 1 (2014), 1-18.

² See "Peer-led interventions and mentoring" and "Brief interventions (BIs)" in "Appendix A: Definitions and Discussion by Intervention Type" of Jamie Harris, Jacy Reese Anthis, and Kelly Anthis "Health Behavior Interventions Literature Review" (July 24, 2020), <https://www.sentienceinstitute.org/health-behavior-appendix-a>.

³ Darrell A. Luzzo and Mary Taylor, "Effects of verbal persuasion on the career self-efficacy of college freshmen," *CACD Journal* 94 (1993), 34.

⁴ Medianta Tarigan and Supra Wimbarti, "Career planning program to increase career search self efficacy in fresh graduates," *Journal of Higher Education Theory and Practice* 11, no. 4 (2011), 75-87.

⁵ See, for example:

- Nadya Fouad, Elizabeth W. Cotter, and Neeta Kantamneni, "The Effectiveness of a Career Decision-Making Course," *Journal of Career Assessment* 17, no. 3 (2009), 338-47,

However, these measures are not sufficient for individuals seeking “to maximize the good with a given unit of resources,” as people involved in the “effective altruism” community are.⁶ For example, if an individual seeks to maximize their positive impact for animals over the course of their career, it is not sufficient for them to have high self-efficacy; they must also apply this to career-related behaviours that enable them to find roles where they can most cost-effectively help animals.

Some studies have evaluated career-related behaviours, but the outcome measures have tended to be generic, rather than impact-focused. For example, Tansley et al. (2007) developed a scale based on questions about career exploration activities like “researching majors, visiting the career development center on campus, completing an interest inventory,” or spending time learning about careers.⁷ Reardon et al. (2020) identified 38 studies of college career courses that

-
- Itamar Gati, Tehila Ryzhik, and Dana Vertsberger, “Preparing Young Veterans for Civilian Life: The Effects of a Workshop on Career Decision-Making Difficulties and Self-Efficacy,” *Journal of Vocational Behavior* 83, no. 3 (2013), 373-85,
 - Caitlin Coyer, Megan Fox, Elena Cantorna, and Lynette Bikos, “Effects of Participation in an Online Course on Undergraduate Career Decision-Making Self-Efficacy” (2018), https://digitalcommons.spu.edu/cgi/viewcontent.cgi?article=1076&context=spfc_research_conference,
 - Michele Lam and Angeli Santos, “The Impact of a College Career Intervention Program on Career Decision Self-Efficacy, Career Indecision, and Decision-Making Difficulties,” *Journal of Career Assessment* 26, no. 3 (2018), 425-44.
 - Diandra Prescod, Beth Gilfillan, Christopher Belser, Robert Orndorff, and Matthew Ishler, “Career Decision-Making for Undergraduates Enrolled in Career Planning Courses,” *College Quarterly* 22, no. 2 (2019), <https://files.eric.ed.gov/fulltext/EJ1221402.pdf>.

For a thorough review of 116 studies of college career courses, see Robert Reardon, Carley Peace, and Ivey Burbrink, “College Career Courses and Learner Outputs and Outcomes, 1976-2019” (2020), <https://www.career.fsu.edu/sites/g/files/upcbnu746/files/TR61.pdf>.

⁶ William MacAskill, “The Definition of Effective Altruism,” *Effective Altruism: Philosophical Issues* (Oxford, UK: Oxford University Press, 2019), 10-28.

⁷ Denny P. Tansley, LaRae M. Jome, Richard F. Haase, and Matthew P. Martens, “The Effects of Message Framing on College Students’ Career Decision Making,” *Journal of Career Assessment* 15, no. 3 (2007), 301-16.

evaluated “job satisfaction, selecting a major, course satisfaction, time taken to graduation from college, [or] cumulative GPA.”⁸ Researchers interested in assessing whether particular career interventions help to “maximize the good” must also assess whether the changes caused by the intervention increase the positive impact on altruistic causes that the participants are expected to have over the course of their career (hereafter abbreviated to “expected impact”). By analogy, experts estimate that, among charities helping the global poor, the most cost-effective charities are 100 times more cost-effective than the average charity;⁹ clearly, not all altruistic actions have equal expected value.

80,000 Hours is the primary organisation providing career support and advice to the effective altruism community, the community striving to help others as much as possible using the best evidence available.¹⁰ 80,000 Hours is optimistic about the usefulness of its advising,¹¹ which it assesses using a system of surveys and interviews to estimate the career plans and trajectory that advisees might have had, were it not for its advising intervention.¹² However, rigorously

⁸ Robert Reardon, Carley Peace, and Ivey Burbrink, “College Career Courses and Learner Outputs and Outcomes, 1976-2019” (2020), <https://www.career.fsu.edu/sites/g/files/upcbnu746/files/TR61.pdf>. Of these, 36 (95%) reported positive gains in the measured outcomes, suggesting that courses can have meaningful impacts on real-world behaviours, not just participants’ attitudes.

⁹ Lucius Caviola, Stefan Schubert, Elliot Teperman, David Moss, Spencer Greenberg, and Nadira S. Faber, “Donors Vastly Underestimate Differences in Charities’ Effectiveness,” *Judgment and Decision Making* 15, no. 4 (2020), 509-16.

¹⁰ 80,000 Hours “help people work out how they can best use this time to help the world, and take action on that basis” (“About us,” 80,000 Hours, accessed 27th April, 2020, <https://80000hours.org/about/>).

¹¹ Benjamin Todd, “80,000 Hours Annual Review – December 2019” (5th April, 2020), <https://80000hours.org/2020/04/annual-review-dec-2019/> and Benjamin Todd, “80,000 Hours Annual Review — November 2020” (May 14th, 2021), <https://80000hours.org/2021/05/80000-hours-annual-review-nov-2020/>.

¹² Benjamin Todd, “80,000 Hours Annual Review – December 2019” (5th April, 2020), <https://80000hours.org/2020/04/annual-review-dec-2019/> notes that “the aim of the criteria-based system is to capture a wider range of changes with less evaluation time, using surveys rather than interviews... From among the criteria-based changes (and other promising changes), we aim to identify the most valuable. If the counterfactual value of a change exceeds a predefined threshold of expected value, we call it a top plan change... We make a judgement call on the value of someone’s plan change based on a 3-30 hour review and (usually) at least one interview.”

evaluating the effects of career advice on participants' expected impact on altruistic causes is difficult. For example, there is evidence constructing counterfactuals from people's self-reports does not produce accurate estimates.¹³ 80,000 Hours has revised its own evaluation methodology in the light of some these difficulties,¹⁴ and bringing additional research methodologies to bear on the question seems valuable.

Until recently, 80,000 Hours did not operate an online course, although they are optimistic about the usefulness of their online content.¹⁵ Others in the effective altruism community have offered online courses, including The Good Food Institute and the philosopher Peter Singer, though neither course is explicitly focused on influencing career outcomes.¹⁶ Our impression is that the only available evidence on whether these courses have successfully influenced career outcomes or not is anecdotal.¹⁷

¹³ See, for example, David McKenzie, "Can business owners form accurate counterfactuals? Eliciting treatment and control beliefs about their outcomes in the alternative treatment status," *Journal of Business & Economic Statistics* 36, no. 4 (2018), 714-22.

¹⁴ See, for example, the section beginning "We were overly credulous..." in "Our mistakes," 80,000 Hours, accessed 19th May, 2020, <https://80000hours.org/about/credibility/evaluations/mistakes/#we-were-overly-credulous-about-how-easy-it-is-to-cause-career-changes-and-our-investigations-into-these-claims-were-insufficiently-skeptical-and-thorough>.

¹⁵ Benjamin Todd, "80,000 Hours Annual Review – December 2019" (5th April, 2020), <https://80000hours.org/2020/04/annual-review-dec-2019/> notes that "well-targeted work on written content is likely higher-impact than work on our other programmes. This is because core online content has played a significant role in driving plan changes in the past (though it's hard to credit share); it continues to create benefits without further investment (unlike headhunting & advising); it supports the other programmes; and it creates more non-plan-change impact, such as introducing people to [effective altruism]."

¹⁶ "A free online course exploring the science behind plant-based and cultivated meat," The Good Food Institute, accessed July 17, 2020, <https://www.gfi.org/OnlineCourses> and Peter Singer, "Effective Altruism," accessed July 17, 2020, <https://www.coursera.org/learn/altruism>.

¹⁷ Peter Singer, "Effective Altruism," accessed July 17, 2020, <https://www.coursera.org/learn/altruism> notes that 41,928 people have already enrolled. "Learner Reviews & Feedback for Effective Altruism by Princeton University," Coursera, accessed July 17, 2020, <https://www.coursera.org/learn/altruism/reviews> includes some reviews specifically referred to "careers," or "jobs":

- "It offered the missing piece and peace of my career puzzle! Thank you!"
- "I particularly enjoyed the lectures on ethical careers because of the guest speakers and was engaged by their experiences and the 'realness' that they were able to present to the students."

There is even less evidence on the impact of careers advice interventions on participants' expected impact for animals specifically.

The main aim of the two studies reported here is to assess whether a one-to-one calls intervention, provided by the nonprofit organisation [Animal Advocacy Careers](#) (AAC), or an online course designed by AAC increase participants' expected impact for animals. The studies also provide evidence for the wider question of whether career advice interventions can improve participants' expected impact on altruistic causes.

Ideally, in order to assess whether a career advice intervention genuinely increases a participant's impact for animals, an evaluation would measure:

1. Whether the intervention causes changes in a participant's career plans,
2. Whether the participant's new intended career plans have higher expected impact for animals than their previous career plans, and
3. Whether the participant actually follows up on those plans and takes actions that have high impact for animals.

-
- “This course is amazing, more than I expected, you can relate it with everything, no matter your career, it's effective.”
 - “While this course is interesting and informative, I was hoping that there would be a larger portion dedicated to jobs, industries, and ideas that help make the world a better place. I was hoping that by the end of this course, I would have learned about many more options on how to act now, today, as opposed to such a strong emphasis on the origin and meaning of ethics.”

In a personal email exchange with an employee of The Good Food Institute, we asked if there were any “stats on participation” or “any form of evaluation of the course's effects.” The employee advised us that the course had had thousands of participants and that “Sri Artham, founder of Hooray Foods, credits GFI's MOOC as being instrumental in his development of a plant-based bacon product,” linking to Julia John, “Hooray for Sri Artham's Plant-Based Bacon!” (November 25 2019), <https://www.gfi.org/blog-hooray-bacon> as evidence.

Unfortunately, there are many complications that would make this ideal evaluation very challenging and expensive.

With regards to the first step, 80,000 Hours has found that there are “long delays between when someone interacts with one of our programmes and when they make a high-value change that we are able to verify and track. The median delay in the current sample [for “top plan changes”] is two years, and only one has ever been tracked within one year.”¹⁸ Hence, to capture all plan changes made as a result of a career advice intervention, and the implementation of those plans, measurement would have to occur after a substantial time delay. With regards to the second and third steps, there are very few clear and rigorous evaluations of interventions that affect animals’ wellbeing, even on short timeframes.¹⁹ After taking longer-term effects into consideration, these evaluations become even more uncertain.²⁰

¹⁸ Benjamin Todd, “80,000 Hours Annual Review – December 2019” (5th April, 2020), <https://80000hours.org/2020/04/annual-review-dec-2019/>.

¹⁹ For a summary of some of the best available evidence for various interventions, see “Briefing Series,” Farmed Animal Funders (February, 2019), https://farmedanimalfunders.org/wp-content/uploads/2019/07/Farmed-Animal-Funders_-2019-Briefing-Series.pdf.

²⁰ The number of sentient beings that could come into existence in the long-term future is astronomically large (see, for example, Nick Bostrom, “Existential Risk Prevention as Global Priority,” *Global Policy* 4, no. 1 (February 2013)). Given this, the impact of any action will presumably be dominated by difficult to measure long-term impacts, which will dwarf the importance of the more measurable impacts on timeframes.

Sentience Institute has undertaken work intended to better understand the long-term effects of interventions that affect animals. See, for example, the evidence summarised in the section on “Momentum vs. complacency from welfare reforms” in “Summary of Evidence for Foundational Questions in Effective Animal Advocacy,” Sentience Institute, accessed 27th April, 2020, <https://www.sentienceinstitute.org/foundational-questions-summaries>.

See also Saulius Šimčikas, “List of ways in which cost-effectiveness estimates can be misleading” (20th Aug 2019), <https://forum.effectivealtruism.org/posts/zdAst6ezi45cChRi6/list-of-ways-in-which-cost-effectiveness-estimates-can-be>.

To address these issues,²¹ the confirmatory analyses in the following studies use a number of outcome measurements that we expect will be correlated with genuine increases in a participant's impact for animals. That is, rather than confirming that the interventions were or were not effective, the studies are only able to provide indirect indications of whether the interventions seem likely to be effective or not. These quantitative, confirmatory analyses are supplemented by an analysis using participants' LinkedIn profiles that takes a more all-things-considered approach but is greatly limited by the above difficulties.

Methodology

The methodology, analysis plans, and predicted results for the two studies were pre-registered on the Open Science Framework.²² A number of small modifications to the methodology and analysis plans were made after the pre-registration which are viewable in the appendices.

Participants and Procedure

The surveys were hosted on Google Forms. Our sample pool was sourced via participants who voluntarily signed up for AAC's one-to-one calls or online course programmes, which were advertised to the effective altruism and animal advocacy communities, such as via relevant newsletters, talks, and Facebook groups. Participants filled out an application form, after which

²¹ AAC is a nonprofit trying to make decisions on the basis of these findings, which made a follow-up period of two years or longer intractable. AAC also expected to offer various other services which would have been accessible to the control groups and intervention groups alike, which would have made effects harder to detect.

²² Jamie Harris, "Pre-registration: The Effects of Career Advising Calls on Expected Impact for Animals" (29 June, 2020), <https://osf.io/pwufc>, and Jamie Harris, "The Effects of an Online Careers Course and Workshop on Expected Impact for Animals" (26 August, 2020), <https://osf.io/cjasf>.

they were randomly assigned to either the intervention group (one-to-one calls or the online course, depending on which they applied for) or a no-intervention control group on a rolling basis. Participants were assigned to the groups using randomised block design that ensured that the intervention groups and their respective control groups did not differ substantially in terms of the extent to which they saw impact for animals over the course of their career as an important priority; randomisation was conducted separately for each possible answer to the relevant question on the application form.²³ They were randomised with an equal allocation ratio.

One-to-one calls were provided by both of the co-founders of AAC,²⁴ with most participants receiving a single session lasting approximately one hour,²⁵ with around one hour's worth of additional time spent on preparation and follow-up support by the advisor. The online course was designed by AAC to focus on what we believed to be the most important content that should affect the decisions that individuals make when seeking to maximise their expected impact for animals over the course of their career. The content draws on research by a number of relevant organisations, most notably 80,000 Hours, Animal Charity Evaluators, and Sentience Institute, as well as AAC's own original research.²⁶ The course design was informed by AAC's research into

²³ The question wording on the application form was: "When deciding on which job to apply for or accept next, how important to you is the potential impact for animals that the job would enable you to have?" The average score for the one-to-one calls intervention group was 4.19 out of 5, compared to 4.24 in the one-to-one calls control group, 4.14 in the online course intervention group, and 4.11 in the online course control group.

²⁴ See "Our Team," Animal Advocacy Careers, accessed 25th June, 2021, <https://www.animaladvocacycareers.org/our-team>. Participants were randomly allocated to have their call with either Lauren Mee or Jamie Harris.

²⁵ A typical one-to-one call would contain: about 5 to 20 minutes of discussion of skills and background; 10 to 50 minutes of discussion of possible pathways the advisee could explore; and 5 to 30 minutes prioritising next steps for the advisee to take.

²⁶ See:

- "A guide to using your career to help solve the world's most pressing problems," 80,000 Hours, accessed 17th July 2020, <https://80000hours.org/key-ideas/>,
- "Advocacy and interventions," Animal Charity Evaluators, accessed 17th July 2020, <https://animalcharityevaluators.org/advocacy-interventions/>,

the “The Characteristics of Effective Training Programmes.”²⁷ For example, the evidence that distance education is often similarly effective or more effective than face-to-face learning²⁸ informed the decision to host the course online (on the platform Thinkific).²⁹ The first cohort of the course culminated in a workshop, which participants were encouraged to attend if they had completed substantial amounts of the online course. The workshop was hosted on Zoom, an online platform supporting group video calls.³⁰ The second cohort of the course did not include a

-
- “Reports,” Sentience Institute, accessed 17th July 2020, <https://www.sentienceinstitute.org/research>, and
 - “Blog,” Animal Advocacy Careers, accessed 17th July 2020, <https://www.animaladvocacycareers.org/blog>.

The sessions in the course were entitled:

- Why farmed animals?
- What does the farmed animal movement look like today?
- Which interventions should we focus on?
- How can you donate effectively?
- How can you have a career that helps animals effectively?
- How can you test your personal fit?
- Can you help animals?
- What are your next steps? Planning your career (as a workshop for the first cohort but an online course session for the second cohort)
- How much have you learned? (Test)

²⁷ “The Characteristics of Effective Training Programmes,” Animal Advocacy Careers (10th January, 2020), <https://www.animaladvocacycareers.org/blog/the-characteristics-of-effective-training-programmes>.

²⁸ See, for example:

- Jorge G. Ruiz, Michael J. Mintzer, and Rosanne M. Leipzig, “The Impact of E-Learning in Medical Education,” *Academic Medicine* 81, no. 3 (2006): 207-12,
- I. Elaine Allen and Jeff Seaman, “Changing Course: Ten Years of Tracking Online Education in the United States” (2013), <https://files.eric.ed.gov/fulltext/ED541571.pdf>,
- Robert M. Bernard, Philip C. Abrami, Yiping Lou, Evgueni Borokhovski, Anne Wade, Lori Wozney, Peter Andrew Wallet, Manon Fiset, and Binru Huang, “How Does Distance Education Compare with Classroom Instruction? A Meta-Analysis of the Empirical Literature,” *Review of Educational Research* 74, no. 3 (2004), 379-439,
- Mary K. Tallent-Runnels, Julie A. Thomas, William Y. Lan, Sandi Cooper, Terence C. Ahern, Shana M. Shaw, and Xiaoming Liu, “Teaching Courses Online: A Review of the Research,” *Review of Educational Research* 76, no. 1 (2006), 93-135,
- Robert M. Bernard, Eugene Borokhovski, Richard F. Schmid, Rana M. Tamim, and Philip C. Abrami, “A Meta-Analysis of Blended Learning and Technology Use in Higher Education: From the General to the Applied,” *Journal of Computing in Higher Education* 26, no. 1 (2014), 87-122,
- David A. Cook, Anthony J. Levinson, Sarah Garside, Denise M. Dupras, Patricia J. Erwin, and Victor M. Montori, “Internet-Based Learning in the Health Professions: A Meta-Analysis,” *JAMA* 300, no. 10 (2008), 1181-96, and
- Traci Sitzmann, Kurt Kraiger, David Stewart, and Robert Wisher, “The Comparative Effectiveness of Web-Based and Classroom Instruction: A Meta-Analysis,” *Personnel Psychology* 59, no. 3 (2006), 623-64.

²⁹ See <https://www.thinkific.com/>.

³⁰ See <https://zoom.us/>.

workshop, though comparable content about career planning was added to the online course.

Participants randomised to participate in the online course were invited to start the course at the same time as the rest of their cohort — 116 participants in the first cohort, 45 in the second — these participants were encouraged to interact and support each other, such as through various channels on Slack, a digital workspace platform.³¹

The application form contained all the same questions as the follow-up form (with minor adjustments to wording so that they made sense in that context), plus some additional questions that were used to help structure the one-to-one calls themselves or to help introduce the online course participants to one another. Hence, the application form serves as a pretest on each of our key outcome measures (see below) for comparison to the follow-up survey (posttest), which was sent approximately six months after each participant submitted their application.³² As an incentive to complete the survey, people who applied to either service were told that they would be entered into a lottery (for \$100 for one-to-one calls, \$200 for the online course) to be sent to one randomly selected respondent (from each service) who completed the follow-up survey.

³¹ See <https://slack.com/>.

³² For various practical reasons, the exact date that they were first sent the follow-up questions varied from about 5.5 to about 6.5 months, then some participants did not reply for several weeks after they were initially sent the follow-up questions. If participants did not respond to the survey within one week of it being sent to them, all except a couple (excluded due to ongoing personal communication) were sent an email reminder from another email address. If they did not reply to that within a week, they were sent a text message, if they provided their mobile phone number in the application form. Participants randomised to the intervention group were invited to their one-to-one call within a week of their application and most took up the offer within a month, although a few delayed for several months for various personal reasons. The start date of the online course may have been as soon as the next day or as distant as about 1.5 months from the time of application, depending on when they applied within the given application windows.

We also sent one additional reminder message on the first online course cohort's Slack channel. We regret this mistake, as it may have increased the difference in response rates between the online course group and the control group, although the response rate difference between intervention and control groups is similar for the first cohort (41% of course participants compared to 33% of control) and the second (36% of course participants, 24% of control).

Meta-analyses tend to find only small effects from interventions intended to modify attitudes or behaviours in general,³³ small effects from advising and mentoring interventions,³⁴ and small effects from career choice interventions,³⁵ though career courses specifically may have slightly larger effects.³⁶ Our calculations suggested that a sample size of 102 would be sufficient to detect medium-sized effects,³⁷ which we used as our target sample size for practical reasons; we did not expect to be able to collect enough participants to detect small effects, besides which an intervention which only caused small effects on our outcome measures would seem less likely to be worth investing in.

There were 134 valid applications to the one-to-one calls intervention; of these applicants, 81 (61%) also completed the follow-up surveys sent approximately six months after their application. There were 321 valid applicants to the online course, 112 (35%) of who completed

³³ See Stephen A. Rains, Timothy R. Levine, and Rene Weber, “Sixty Years of Quantitative Communication Research Summarized: Lessons from 149 Meta-Analyses,” *Annals of the International Communication Association* 42, no. 2 (April 3, 2018), 105–24 and Jamie Harris, Jacy Reese Anthis, and Kelly Anthis, “Health Behavior Interventions Literature Review,” (July 24, 2020), <https://sentienceinstitute.org/health-behavior>.

³⁴ See footnote 1.

³⁵ Susan C. Whiston, Yue Li, Nancy Goodrich Mitts, and Lauren Wright, “Effectiveness of Career Choice Interventions: A Meta-Analytic Replication and Extension,” *Journal of Vocational Behavior* 100 (2017): 175-184.

³⁶ G. Gim, 대학 진로 교과목의 효과에 대한 메타분석 [Meta-analysis of the effectiveness of career college courses] (2015), <http://sspace.snu.ac.kr/handle/10371/127892>, cited and summarised by Robert Reardon, Carley Peace, and Ivey Burbrink, “College Career Courses and Learner Outputs and Outcomes, 1976-2019” (2020), <https://www.career.fsu.edu/sites/g/files/upcbnu746/files/TR61.pdf> as having found “a moderate overall effective size of .556.”

³⁷ G* Power was used to compute the required sample size a priori — see Franz Faul, Edgar Erdfelder, Albert-Georg Lang, and Axel Buchner, “G* Power 3: A Flexible Statistical Power Analysis Program for the Social, Behavioral, and Biomedical Sciences,” *Behavior Research Methods* 39, no. 2 (2007), 175-91.

The option “Means: Difference between two independent means (two groups)” was selected. A one-tailed test was selected. The “Effect size d” was set to 0.5, which is the conventional cutoff for “medium” effect sizes, following J. E. Cohen, *Statistical Power Analysis for the Behavioral Sciences* (Hillsdale, NJ: Lawrence Erlbaum Associates Inc., 1988). Alpha (the risk of false positive results) was set to 0.05, the power (the risk of false negative results) was set to 0.80, and the allocation ratio was set to 1. This suggested that a total sample size of 102 would be needed.

follow-up surveys.³⁸ Although the number of valid applicants allocated to intervention groups were very similar,³⁹ the response rates to the follow-up surveys were different; only 28 of the one-to-ones control group (43%) completed the follow-up survey compared to 53 of the one-to-ones intervention group (78%). For the online course, only 49 of the control group (31%) did so, compared to 63 (39%) of the intervention group. Table 1 shows descriptive statistics (based on questions from the application forms) for participants who completed both surveys.⁴⁰

Table 1: Descriptive statistics

Variable	One-to-one calls		Online course	
	Intervention (<i>n</i> = 53)	Control (<i>n</i> = 28)	Intervention (<i>n</i> = 63)	Control (<i>n</i> = 49)
Based in the Global North	96%	86%	87%	100%
Mean age	31	27	30	29
Female	55%	45%	67%	71%
Self-reported knowledge	3.9	4.0	3.3	3.6

³⁸ For the one-to-one calls, there was one exclusion after a request from the applicant, one exclusion from an individual who applied twice, and seven exclusions for individuals who applied to our online course first. For the online course, there were 4 exclusions after request from the applicants, 2 exclusions from individuals who applied twice, and 13 exclusions for individuals who applied to our one-to-one calls first.

³⁹ For one-to-ones, there were *n* = 68 in the intervention group and *n* = 66 in the control group. For the online course, there were 161 and 160, respectively.

⁴⁰ Two applicants to the online course and two applicants to the one-to-ones had incomplete demographic data — the missing entries have been excluded from the relevant rows. Respondents who selected “other” for their gender were scored as 0.5. For full results on other questions asked in the application form, see the [“Predictions and Results”](#) spreadsheet.

of effective altruism (1 to 5)				
Actively seeking a new role to improve their impact for animals	92%	71%	73%	78%

If certain types of people were more likely to have completed the follow-up questions, then the differing response rates could affect the results. Indeed, comparisons of the pretest answers between those members of the intervention groups and control groups who completed both surveys suggests a number of differences.⁴¹ As such, it may be best to interpret the confirmatory analyses as if from longitudinal, observational studies, rather than randomised controlled trials,⁴² although several supplementary analyses (described below) were undertaken to account for biases potentially introduced by differential attrition.

⁴¹ See the column for “Intervention group and control group pretests for only those who completed both surveys” in the “full results for each group” tabs of the [“Predictions and Results”](#) spreadsheet. Most notably, the one-to-ones intervention group participants were slightly more likely to be female (55% compared to 45%) or from the Global North (96% compared to 86%), and slightly less likely to have changed their career plans in the months prior to the pretest survey (2.08 out of 6 compared to 2.25). The online course intervention group participants had lower average scores for “career-related behaviours” at pretest (2.9 out of 9 compared to 3.5 in the control group) and were less likely to be based in the Global North (87% compared to 100%).

⁴² Janus Christian Jakobsen, Christian Gluud, Jørn Wetterslev, and Per Winkel, “When and how should multiple imputation be used for handling missing data in randomised clinical trials—a practical guide with flowcharts,” *BMC Medical Research Methodology* 17, no. 1 (2017), 1-10 comment that, “[i]f multiple imputations or other methods are used to handle missing data it might indicate that the results of the trial are confirmative, which they are not if the missingness is considerable. If the proportions of missing data are very large (for example, more than 40%) on important variables, then trial results may only be considered as hypothesis generating results... If the [missing at random] assumption seems implausible based on the characteristics of the missing data, then trial results will be at risk of biased results due to ‘incomplete outcome data bias’ and no statistical method can with certainty take account of this potential bias.”

Instruments

The [appendix](#) contains the full text of the main follow-up survey questions. The metrics used in the confirmatory analysis are summarised in table 2 below.

Table 2: Summary of analysis plans

Metric	Components	Components combined how?	Output
Attitudes	Cause prioritisation, importance of impact for animals, confidence that there are viable career paths in effective animal advocacy, and inclination towards effective altruism.	Taking the mean of standardised versions of each component.	A score, usually somewhere between -1 and +1 for the averages of the intervention group and control group.
Career plans	Study plans, internship plans, job plans, and long-term plans.	Multiplying the score for each component question by specific amounts (study plans *1, internship plans *1, job plans *2, long-term plans *2) then summing them together.	A score from 0 to 6.
Career-related behaviours	Secured a new role, applied to one or more positions or programmes, joined the	Multiplying the score for each component question by specific amounts (new role *2,	A score from 0 to 10.

	effective animal advocacy directory, had two or more in-depth conversations about their career plans, changed which nonprofits they donate to or volunteer for, changed substantially and intentionally the amount of money they donate or the time they spend volunteering, and time spent reading content by AAC or 80,000 Hours.	applications *1, directory *1, conversations *1, nonprofits *1, amounts *2, AAC time *1, and EAA time *1) then summing them together.	
Self-assessed expected impact for altruistic causes	NA	NA	A score from 1 to 5, i.e. from "substantially lower impact" to "substantially higher impact."

Where the individual components of a metric are weighted, the weightings are based on AAC's intuitions about the relative importance of these components and the likelihood that these changes will translate into substantial changes in impact for animals.

The attitudes metric includes component questions about:

- The importance to participants that their job has high impact for animals (assessed via a five-point scale from “Irrelevant for you” to “The most important factor for you”).
- Participants’ confidence that there are career paths that are viable for them that will enable them to have a high impact for animals (assessed via a five-point Likert-type scale from “Very low confidence” to “Very high confidence”).
- Whether the participants are focusing on the causes that seem likely to have the best opportunities to do good (scored as a 1 or 0).⁴³
- Whether the participants hold beliefs that seem likely to enable them to focus on the best opportunities within those causes, assessed via seven-point Likert-type scales.⁴⁴

To assess career plans, we asked whether participants had changed their study plans, internship plans, job plans, or long-term plans. The options provided were “changed substantially” (scored as 1), “changed somewhat” (scored as 0.5), “no change” (scored as 0) or “I wasn’t previously planning to do this and still am not” (scored as 0). These questions are not specific to altruistic causes, but are supplemented by a question about whether, all things considered, the participants expect that their impact for altruistic causes has changed (assessed via a five-point Likert-type scale from “substantially lower impact” to “substantially higher impact”).

⁴³ Each participant is assigned a 1 if their answer includes at least one area that we consider to be high priority in the “high interest/priority” column and only includes areas that we consider to be high priority in that column: 1) farmed animals, 2) wild animals (focused on the suffering/welfare of animals), 3) building effective altruism, 4) other causes focused on the long-term future (e.g. AI safety research), 5) climate change (if incorporated alongside “farmed animals”, but otherwise this does not count).

⁴⁴ The questions were about their willingness to use limited time and resources in whichever ways will achieve the most good (scored from 1 to 7, from “strongly disagree” to “strongly agree”) and willingness to think about actions to help others in terms of the numbers of individuals affected (scored from 1 to 7, from “strongly agree” to “strongly disagree”). These two scores were averaged together to form a measure of inclination towards effective altruism.

Given that we did not expect the intended outcome of the intervention — career changes that increase the participants’ expected impact for animals — to occur within the six-month follow-up period employed in the study, we ask about career-related behaviours that seem likely to be useful intermediary steps. For example, rather than evaluating the direct usefulness of the roles that the participants are working in at the time of the posttest survey, we ask them whether they have secured any new role, applied to one or more positions or programmes, or had two or more in-depth conversations about their career plans (each scored as a 1 if they had or a 0 otherwise), since some roles and actions might be beneficial for animals longer-term by helping the participant to develop career capital that they can later apply to helping animals or test their personal fit with certain career paths, rather than being immediately helpful for animals. We also included a number of questions of specific relevance to careers in animal advocacy, such as asking the time they have spent reading content by AAC or 80,000 Hours (scored as a 1 if their answer indicated that they had spent 8 hours or more doing so, otherwise scored as a 0), whether they had changed which nonprofits they donate to or volunteer for (scored as a 1 if so, or 0 otherwise) or changed substantially and intentionally the amount of money they donate or the time they spend volunteering (scored as a 1 if so, or 0 otherwise), and manually checking whether they had joined the “[effective animal advocacy community directory](#)” (scored as a 1 if so, or 0 otherwise).

For each of the main metrics and their component questions, we made predictions about the mean differences that we expected to see between the intervention and control groups, and recorded these in appendices of each pre-registration.⁴⁵

Results

Confirmatory analyses

This section focuses on the pre-registered analysis plans. However, these are arguably not the most appropriate analyses given the differential attrition between the intervention and control groups. Alternative analyses are discussed in subsequent sections.

Tables 3 and 4 shows the pretest and posttest results for participants who completed both surveys for the one-to-one calls and online course interventions, respectively. For full results, including for the components of each metric and for the full set of pretest results (those who did not complete the follow-up survey as well as those who did), see the “[Predictions and Results](#)” spreadsheet.

Table 3: Results at pretest and posttest for the one-to-one calls intervention and control groups

⁴⁵ Jamie Harris, “Pre-registration: The Effects of Career Advising Calls on Expected Impact for Animals” (29 June, 2020), <https://osf.io/pwufc>, and Jamie Harris, “The Effects of an Online Careers Course and Workshop on Expected Impact for Animals” (26 August, 2020), <https://osf.io/cjasf>.

Metric	Pretest, intervention (<i>n</i> = 53)	Pretest, control (<i>n</i> = 28)	Posttest, intervention (<i>n</i> = 53)	Posttest, control (<i>n</i> = 28)
Attitudes ⁴⁶	NA	NA	0.01 (0.60)	-0.02 (0.60)
Career plans	2.08 (1.50)	2.25 (1.43)	2.30 (1.71)	1.82 (1.52)
Career-related behaviours	3.25 (2.20)	3.21 (1.95)	3.66* (2.10)	2.61 (1.95)
Self-assessed expected impact for altruistic causes	3.87 (0.98)	3.71 (1.01)	3.70* (0.82)	3.21 (0.79)

Asterisks () identify significant differences at $p = 0.05$ between the one-to-one calls group and the control group.*

Standard deviations are provided in brackets.

Table 4: Results at pretest and posttest for the online course intervention and control groups

Metric	Pretest, intervention (<i>n</i> = 63)	Pretest, control (<i>n</i> = 49)	Posttest, intervention (<i>n</i> = 63)	Posttest, control (<i>n</i> = 49)
Attitudes ⁴⁷	NA	NA	-0.04 (0.53)	0.06 (0.60)
Career plans	2.30 (1.50)	2.34 (1.63)	2.35 (1.59)*	1.68 (1.78)
Career-related behaviours	2.89 (1.88)	3.53 (2.05)	2.70 (1.90)	2.84 (1.98)

⁴⁶ Note that, since this is a standardised score, the answers at pretest and posttest are not directly comparable. This is because means and standard deviations are calculated from the full sample, which is different at pretest and posttest. E.g. an answer of “High confidence” would have a slightly different value in the context of the pretest to in the context of the posttest.

⁴⁷ Ibid.

Self-assessed expected impact				
for altruistic causes	3.70 (0.84)	3.82 (0.75)	3.79 (0.72)	3.55 (0.84)

Asterisks () identify significant differences at $p = 0.05$ between the online course group and the control group.*

Standard deviations are provided in brackets.

The four metrics were tested for normal distribution and homoscedasticity.⁴⁸ Three of the four metrics were found not to be normally distributed, so one-tailed Wilcoxon rank tests were used for the confirmatory analysis. In each study, a controlled false discovery rate (FDR) of 0.05 was used, meaning we are 95% sure any given significant difference has not occurred just by random chance, for the four tests in the confirmatory analysis.⁴⁹

In both the study of one-to-one calls and the online course, the differences in attitudes between the intervention and control groups at six months' follow-up were not significant ($p = 0.33$ and 0.75 , respectively) and the Mean Difference (MD) on this standardised score was close to zero (0.03 and -0.10 , respectively).⁵⁰

⁴⁸ This was confirmed for the attitudes metric, and so an unpaired one-tailed t-test was conducted between the intervention and control groups, finding no significant differences for either the one-to-one or online course ($p = 0.42$ and 0.82 , respectively, without FDR adjustment).

⁴⁹ The methods described in Yoav Benjamini and Yosef Hochberg, "Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing," *Journal of the Royal Statistical Society: Series B* (Methodological) 57, no. 1 (1995), 289-300 were used to control the FDR.

⁵⁰ The components of this metric were examined as exploratory analysis; none of the differences were significant using one-tailed Wilcoxon tests in either study and the MD s were below our predicted levels in each case (see the "[Predictions and Results](#)" spreadsheet).

Although attitudes appeared to be larger in the online course control group than intervention group, the differences was not significant when the direction of the significance test was reversed ($p = 0.36$). Comparable checks were not conducted for the subcomponents, since this was already a deviation from our pre-registered analysis plans, though some of these may have been found to be significantly different if tested.

The difference in career plans was not significant ($p = .16$) in the study of one-to-one calls at six months' follow-up. However, career plan changes were found to be significantly greater in the online course intervention group than the online course control group ($p = 0.04$). The *MD* was 0.67 on a score from 0 to 6, though this falls short of our predicted *MD* of 1.35. This is roughly equivalent to one in three participants “substantially” changing “the job that you are applying for / plan to apply for next” or their “long-term career plans (including role types, sector, and cause area focus)” compared to what would have occurred without the online course. The *MD* in the one-to-one calls study (0.48) is not far short of the *MD* in the online course study (0.67), and seems potentially meaningful if it is not due to random variation.⁵¹

Career-related behaviours were found to be significantly greater in the intervention than control groups in the one-to-one calls study ($p = 0.03$). The *MD* was 1.05 on a score from 0 to 9, which is about as effective as our predicted *MD* of 1.25. This is roughly equivalent to all participants making a small change such as in “which nonprofits you donate to or volunteer for” or half the participants making a more substantial change such as securing “a new role that you see as being able to facilitate your impact for animals” where such changes would not have occurred without the one-to-one call. There were no significant differences on this metric in the online course study ($p = 0.75$) and the *MD* was -0.14.⁵²

⁵¹ This still falls short of our predicted *MD* of 1.25. One of the components of this metric, relating to job plans, was found to be significantly different ($p = 0.05$, *MD* = 0.15 from a scale with options 0, 0.5, or 1) in exploratory analysis. Recall also that the study is underpowered to detect medium differences.

⁵² Although career-related behaviours appeared to be larger in the online course control group than intervention group, the difference was not significant when the direction of the significance test was reversed ($p = 0.25$).

Self-assessed expected impact for altruistic causes was also found to be significantly greater in the intervention group than the control group in the one-to-one calls study ($p = 0.02$); here, the MD was 0.48 on a scale from 1 to 5 (similar to our predicted MD of 0.6), equivalent to someone moving about half way from an answer of “No change or very similar impact” to an answer of “Somewhat higher impact” to their expected impact due to recent career plan changes from what they would have answered without a one-to-one call. For the online course, before adjustment, there appeared to be a significant difference on this metric ($p = 0.04$), though the difference was not significant after controlling the false discovery rate ($p = 0.09$). Here, the MD was 0.24 on a scale from 1 to 5.⁵³

Exploratory analyses using survey responses

Exploratory linear regressions were carried out with the dependent variables being the participants’ attitudes, career plans, career-related behaviours, and self-assessed expected impact for altruistic causes at follow-up (see the “[Predictions and Results](#)” spreadsheet). This was conducted in order to better understand whether differential attrition rather than the interventions themselves best explain the observed differences on the outcome measures of interest, since it controlled for imbalances in observable characteristics that might have arisen from differential attrition and confounded the relationships between interventions and outcomes.⁵⁴ The 16 predictors included whether the participant was randomised to the intervention or control group

⁵³ This would be equivalent to one in four people moving from an answer of “No change or very similar impact” to an answer of “Somewhat higher impact” to their expected impact due to recent career plan changes from what they would have answered without the online course.

⁵⁴ Inverse probability weighting would likely be a more rigorous test of this. See, for example, Maya Duru and Sarah Kopper, “Data analysis,” accessed 7th September, 2021, <https://www.povertyactionlab.org/resource/data-analysis>.

and the answers to several questions from the application form, such as prior familiarity with effective altruism.

As in the confirmatory analysis for the one-to-one calls study, randomisation to the intervention group was found to have significant effects on self-assessed expected impact for altruistic causes ($p = 0.02$) but not on attitudes ($p = 0.85$) or career plans ($p = 0.36$).⁵⁵ The effect of randomisation to the intervention group on career-related behaviours was no longer significant ($p = 0.08$), though this could be due to the sample size being insufficient to detect medium effects.⁵⁶ For the online course study, the findings were also similar to the confirmatory analysis; randomisation to the intervention group was found to have significant effects on career plans ($p = 0.01$)⁵⁷ but not on attitudes ($p = 0.87$), career-related behaviours ($p = 0.88$), or self-assessed expected impact for altruistic causes ($p = 0.24$).⁵⁸

⁵⁵ In each case, the coefficient identified in the regression was similar to the mean difference observed in simpler comparison between the groups (0.53 compared to 0.48, -0.02 compared to 0.03, and 0.39 compared to 0.48).

⁵⁶ G* Power was used to compute the required sample size a priori — see Franz Faul, Edgar Erdfelder, Albert-Georg Lang, and Axel Buchner, “G* Power 3: A Flexible Statistical Power Analysis Program for the Social, Behavioral, and Biomedical Sciences,” *Behavior Research Methods* 39, no. 2 (2007), 175-91.

The option “Linear multiple regression: Fixed model, R2 deviation from zero” was selected. The f^2 effect size was set to 0.15, which is the conventional cutoff for “medium” effect sizes, following J. E. Cohen, *Statistical Power Analysis for the Behavioral Sciences* (Hillsdale, NJ: Lawrence Erlbaum Associates Inc., 1988). Alpha (the risk of false positive results) was set to 0.05, the power (the risk of false negative results) was set to 0.80, and the number of predictors was set to 16. This suggested that a total sample size of 143 would be needed, whereas our sample size after excluding participants with incomplete demographic data was only 79 (compared to 109 in the online course study). The coefficient of 0.95 in the regression is only slightly lower than the initially observed mean difference of 1.05.

⁵⁷ The coefficient identified in the regression was similar but slightly larger than in simpler comparison between the groups (0.87 compared to 0.67).

⁵⁸ The coefficients were 0.02 for attitudes (now leaning in favour of the online course group, unlike the MD -0.10 in the simpler comparison between groups), 0.20 for self-assessed expected impact for altruistic causes (similar to the MD 0.24 in the simpler comparison) and 0.06 for career-related behaviours (now leaning in favour of the online course group, unlike the MD -0.14 in the simpler comparison).

For the online course study, further regressions were carried out using the same variables except where the online course variable was replaced with participants' percentage completion rate of the course and the population was limited to the 62 participants who completed the follow-up form, had full complete demographic data, and were randomly allocated to the online course group. Comparably to previous tests, the completion rate significantly predicted career plans ($p = 0.02$) but not career-related behaviours ($p = 0.11$) or self-assessed expected impact for altruistic causes ($p = 0.16$).⁵⁹ Surprisingly, the completion rate significantly predicted attitudes ($p = 0.02$).⁶⁰

To check the sensitivity to inclusion criteria, exploratory analysis was undertaken using only the applicants to the first cohort of the online course. None of the differences were significant, though this could be due to the sample size being insufficient to detect medium effects.⁶¹ Additionally, exploratory analysis was undertaken where individuals randomised to the online course who completed less than 50% of the course were excluded from the analysis; the results were similar to the confirmatory analysis, albeit suggesting slightly more positive effects of participation in the course.⁶² A comparable analysis was not necessary for the one-to-one calls study, since no one who skipped the call after being invited to participate actually completed the follow-up survey.

⁵⁹ The coefficients were 0.015, 0.014, and 0.005, respectively.

⁶⁰ The coefficient was 0.005. This could potentially just indicate that attitudes influenced the completion rate, rather than vice versa, although the inclusion of attitudinal variables from the pretest in the regression should have reduced the importance of such effects.

⁶¹ See footnote 37. This suggested that a total sample size of 102 would be needed, whereas our sample size was only 85. Attitudes: $p = 0.817$, $MD = -0.13$. Career plans: $p = 0.111$, $MD = 0.50$. Career-related behaviours: $p = 0.679$, $MD = -0.05$. Self-assessed expected impact for altruistic causes: $p = 0.111$, $MD = 0.26$.

⁶² Attitudes and career-related behaviours were not significantly different ($p = 0.37$, $MD = 0.02$ and $p = 0.55$, $MD = -0.03$, respectively). Career plans and self-assessed expected impact for altruistic causes were significantly larger in the online course group ($p = 0.00$, $MD = 0.86$ and $p = 0.03$, $MD = 0.28$, respectively).

As a robustness check, non parametric two-samples Wilcoxon rank tests were carried out comparing the intervention group's follow-up survey answers and their answers in the application form itself.⁶³ None of the primary metrics described above were found to be significantly different.⁶⁴

Rethink Priorities conducted re-analysis of the data, using methods of their own choosing that they expected to be most appropriate, though they didn't expect that this analysis could overcome some of the difficulties with the data such as differential attrition. They carried out linear mixed modelling regressions that included both pretest and posttest data, allowed for random intercepts for each participant, and used bootstrapping to obtain robust standard errors. Randomisation to the intervention groups was not found to have significant effects on any of the four main metrics for either of the two services.⁶⁵

⁶³ Unlike in the confirmatory analysis, two-tailed tests were used, since initial review of the data suggested that some of the pretest scores were notably higher than the posttest scores.

⁶⁴ For the online course, importance of impact for animals (a component of the attitudes metric) was found to be significantly larger at pretest ($p = 0.03$), which is disappointing. However, the number of people who had joined the effective animal advocacy directory in the last six months was significantly larger at posttest ($p = 0.00$). None of the components were significantly different in the study of one-to-one careers calls.

⁶⁵ However, in each case, the coefficient identified was similar to the mean difference observed in simpler comparison between the groups. Career-related behaviours in the online course was a notable exception; the coefficient was 0.50, compared to -0.14 previously.

The p values were also close to significance for two of the comparisons that had been found to be significant in the confirmatory analyses: career-related behaviours in the study of one-to-one calls ($p = 0.074$) and career plans in the study of the online course ($p = 0.069$). However, for self-assessed expected impact for altruistic causes in the study of one-to-one calls was now quite far from the conventional cutoff for significance ($p = 0.234$).

See [here](#) for the full results and confidence intervals.

Exploratory analyses using participants' LinkedIn profiles

Supplementary analyses were conducted, comparing the roles that the participants (regardless of whether or not they completed the follow-up survey) appeared to be in at the time of their application to their roles in late July or early August 2021, i.e. 8 to 13 months after their application to a one-to-one careers call or 7.5 to 10.5 months after their application to the online course. This information was gathered from the participants' LinkedIns, where this information was available.⁶⁶ The main results are reported in Table 6 below.⁶⁷

Table 6: Analysis of role changes according to LinkedIn profiles in the 7.5 to 13 months after application

Subjective categorisation	One-to-one calls		Online course	
	Intervention (<i>n</i> = 54)	Control (<i>n</i> = 55)	Intervention (<i>n</i> = 100)	Control (<i>n</i> = 93)
No or minor change	43%	40%	54%	65%
Negative-seeming change	7%	7%	0%	6%

⁶⁶ In the one-to-one calls application form, we had asked participants: "Please link to your LinkedIn profile so we can learn more about your background. If you don't have LinkedIn, a CV in a Google Doc would be fine." In the online course application form, we had asked participants: "Please provide details on any other contact details you would like to provide to other participants." We did not make provision of a LinkedIn account obligatory in either study, but some participants provided their LinkedIn in answer to these questions; this was true for a larger proportion of the one-to-one calls applicants (74%) than the online course applicants (60%), presumably due to differences in the question wording. For all those who did not, we checked whether the email that they had used to apply was affiliated with a LinkedIn account. We did not seek other LinkedIn profiles, since it is often hard to tell whether the identified profile belongs to a particular applicant or not. When an individual could be put into one of two categories (e.g. both "Started new non-EA(A) position (no clear EA(A) relevance)" and "Started new EA(A) volunteering"), they were placed into the more optimistic, EAA-relevant category.

⁶⁷ Narrower groupings can be seen at the "Qualitative analysis" tab of the "[Predictions and Results](#)" spreadsheet.

Change with unclear EA(A) implications	17%	27%	29%	23%
EA(A)-relevant, positive-seeming change	33%	25%	17%	6%

EA(A) = effective animal advocacy and/or effective altruism.

In the online course study, a smaller proportion of the intervention group than the control group had apparently not changed their role much or at all (54% compared to 65%).⁶⁸ Unlike in the control group (6%), none of the intervention group had undergone negative-seeming changes such as becoming unemployed or not making any change apart from stopping EA(A) volunteering. A larger proportion of the intervention group had made some change with unclear implications, such as starting a new position with no clear EA(A) relevance (29% compared to 23% in the control group). There did not seem to be comparable differences between the intervention and control groups for any of these categories in the one-to-one calls study.

Interestingly, in both studies, a larger proportion of the intervention group than the control group had made some sort of positive-seeming change relevant to EA(A); either starting a new EA(A) volunteering position (9% compared to 7% in the one-to-one calls study, 7% compared to 1% in the online course study), starting a new EA(A) internship (4%, 5%, 1%, 1%), starting a more substantial EA(A) position (11%, 6%, 7%, 2%), or starting in a new role in another path that

⁶⁸ This method may exaggerate the number of individuals with no change, since it seems likely that some people simply had not updated their LinkedIns. In one or two cases, there were individuals who we knew from personal contact had changed roles in the past few months, but where their LinkedIn remained unchanged. (In this analysis, we relied exclusively on the information on their LinkedIn, to avoid biasing the results due to having greater information about the online course group.)

seemed potentially impactful for animals, e.g. in policy, politics, or academia (9%, 7%, 2%, 2%). Logistic regression found that the difference in EA(A)-relevant, positive-seeming changes between the intervention and control groups was significant for the online course study ($p = 0.018$) but not the one-to-one calls study ($p = 0.479$).⁶⁹

Discussion

The findings from comparisons between the intervention groups and the control groups (our pre-registered confirmatory analyses) provide evidence that career advice interventions can increase participants' expected impact for animals. Although the two studies do not suggest that either one-to-one calls or an online course successfully caused change in all our intended outcomes — attitudes, career plans, career-related behaviours, and self-assessed expected impact for animals — the studies do suggest that three out of four of these outcomes were significantly altered in the expected direction by one intervention or the other.

The study of one-to-one calls provides evidence that the participants' own self-assessed expected impact for altruistic causes will tend to be higher than it would be without a one-to-one call and that participants are likely to engage in more career-related behaviours than they otherwise would. Given that previous meta-analyses have found only small effects from advising and mentoring interventions⁷⁰ and our study had insufficient sample size to detect small effects, these findings are arguably quite impressive. Indeed, the mean differences (our best guess of the effect

⁶⁹ The estimates were 0.068 and 0.052, respectively. No controls were included in these regressions.

⁷⁰ See footnote 1.

sizes of the intervention) suggest changes that seem intuitively meaningful. There is some suggestive evidence that the one-to-one calls caused career plan changes, but the difference between the intervention and control groups was not statistically significant. In contrast, the second study provides evidence that online courses may cause individuals to change their career plans. The identified effects fell short of our predictions for each of the main outcome metrics for both interventions. This is disappointing, but the *MDs* on several of the outcome metrics still suggest changes that seem intuitively meaningful.

Differences in the design and marketing of the services mean that the average scores from the participants in the two studies cannot easily be compared directly. Nevertheless, the two studies suggest different sorts of effects from the interventions. This weakly suggests that these services do not have equivalent, interchangeable effects and could be mutually complementary.

Quantitative analyses such as this are limited in a number of ways and one might reasonably believe that, even if the tests we designed failed to find any evidence of positive (or negative) effects, a one-to-one calls intervention could still be impactful (or not) for one reason or another. Although we carefully created measures that we expected would provide indirect indications of whether the interventions seem likely to be effective or not, our measures may have been imperfect indicators.⁷¹

⁷¹ For example, our questions about career plans simply asked participants whether their plans had changed somewhat, substantially, or not at all, and we have no systematic and quantitative method of assessing the quality of those plan changes. It seems plausible that even if only a handful of participants changed their long-term career plans “somewhat” as a result of their one-to-one calls, this could make the intervention worthwhile if those changes were meaningful enough or the expected impact of those individuals was sufficiently high.

The high attrition rates from pretest to posttest and differential attrition between the intervention groups and the control groups in both studies also makes interpretation of the results difficult; arguably these are fatal flaws in the study, since they remove many of the benefits of randomisation between intervention and control groups and mean that the uncertainty about the effects of the interventions is actually higher than the uncertainty that is implied by the 95% confidence intervals implied in the “[Predictions and Results](#)” spreadsheet. The mechanism that seems most plausible to us for explaining the differential attrition is that the intervention group participants felt (on average) more obliged to respond to our request for them to fill out a survey, given that we had provided them a one-to-one call or an online course free of charge. The responses we received from the control group might be disproportionately likely to only come from individuals who are especially mission-aligned with AAC, or especially likely to have made changes that they felt proud of (and wanted to report in a survey). This would presumably make the interventions seem *less* effective than it would have seemed if we had had 100% response rates, and the findings from the confirmatory analyses would be conservative estimates of effects.

One could imagine plausible alternative explanations for the differing response rates which might push in the opposite direction. For example, involvement in the service could have magnified social desirability bias, relative to the control groups.⁷² However, our exploratory analyses using linear regression and controlling for other factors mostly had coefficients suggesting similar or more positive effects from the interventions on attitudes, career plans,

⁷² Individuals in the intervention groups who made changes in relevant outcome metrics may have been especially likely to respond to the follow-up survey and individuals who did not make such changes may have been especially unlikely to respond.

career-related behaviours, and self-assessed expected impact for altruistic causes than those suggested by a simple comparison of the means of the online course and control group posttest surveys.

Furthermore, the exploratory analyses using participants' LinkedIn profiles does not rely on judgments by the participants themselves or have differential attrition rates that could plausibly have been caused by involvement in the online course itself,⁷³ yet also found evidence (stronger for the online course than one-to-one advising and fairly independently of the confirmatory analyses) that the interventions studied here can have positive effects on career outcomes, increasing their expected impact for animals. The LinkedIn analysis suffers from the limitations described in the introduction which encouraged us to use proxy metrics like “career-related behaviours” in the confirmatory analyses. It is also very subjective, and difficult to independently verify without compromising participants' anonymity.⁷⁴

Finally, there was a small difference in favour of the intervention groups on the only truly “objective” measure of their effects — whether the participant had joined the effective animal

⁷³ Relatedly, there was a small difference in favour of the online course group on the only truly “objective” measure of the course's effects — whether the participant had joined the effective animal advocacy community directory. 7% of the online course group (11 out of 161) had joined in between the pretest and posttest surveys, compared to 2% of the control group (3 out of 160) and 1% of the online course group had joined in the 6 months before the pretest compared to 0% of the control group. This was used as part of the “career-related behaviours” metric (for the 112 participants who completed follow-up surveys).

⁷⁴ In the online course study, we had planned (but not pre-registered our intention) to use a qualitative analysis method which relied on self-report by the participants; comparing their answers to the question “In one or two sentences, briefly describe what your current role or study programme is” from the pretest and posttest. However, we decided to use the LinkedIn method instead in order to avoid the risks of bias caused by self-report and differential attrition.

advocacy community directory.⁷⁵ Unfortunately, this metric is likely not a very accurate indicator of expected impact for animals, since it has very low barriers to entry and someone could in theory join it even if they had no intention of pursuing a career that helped animals.

Each of these analyses has different risks of bias and readers may reasonably have different views about which analysis is most informative, though each of them points towards at least some positive effects from the tested interventions.

At first glance, it seems surprising that comparisons between the intervention group posttest and pretest answers found no significant differences. However, it should be borne in mind that the career plans, career-related behaviours, and self-assessed expected impact for altruistic causes questions (but not the attitudes questions) in both the pretest and posttest surveys asked about changes over the previous six months. Hence, nonsignificant differences between the posttest and pretest do not provide very strong evidence that the interventions did not have positive effects, they simply suggest that, on average, the recipients of the interventions changed their plans, behaviours, and expectations a similar amount both before and after participating. For example, mean scores of 2.08 and 2.30 on the career plans metric indicate that an average participant in the one-to-one calls intervention might have “changed substantially” their “long-term career plans (including role types, sector, and cause area focus)” or “changed somewhat” each of the studies programme, volunteering, and job that they were planning to

⁷⁵ 7% of the online course intervention group and 17% of the one-to-one calls intervention group had joined in between the pretest and posttest surveys, compared to 2% and 7% of their respective control groups. As a further comparison, 1% of the online course intervention group and 6% of The one-to-one calls intervention group had joined in the 6 months prior to their application. This was used as part of the “career-related behaviours” metric (for the participants who completed follow-up surveys).

apply for next — and they might have made changes of this magnitude both in the 6 months prior to the one-to-one call and the 6 months following it.⁷⁶ This seemingly high degree of change may be because people tended to apply for the interventions during periods where they were making important career decisions. For example, the proportion of people actively seeking a new role to improve their impact for animals fell from 93% and 71% to 50% and 57% in the one-to-one calls intervention group and control group, respectively, from pretest to posttest; for the online course, the comparable figures are 73% and 78% to 59% and 65%.

If we assume that the confirmatory analysis is correct — that the one-to-one calls successfully encouraged more career-related behaviours and more positive self assessments of expected impact for animals, whereas the online course successfully changed participants' career plans, and neither intervention altered attitudes — which of these proxy outcome metrics is most important for increasing expected impact for animals? All these metrics seem like useful indicators, but our guess is that “career plans” is the only metric where (in most cases) we *need* to see meaningful changes within six months of the intervention in order to believe that the intervention has had positive effects.⁷⁷ One can imagine that an individual would change their career plans after an intervention, but only actually put these plans into action (i.e. undertake more “career-related behaviours”) after some time delay; 80,000 Hours have commonly found

⁷⁶ Of course, in practice, some people made substantial changes before but not after the intervention, some made substantial changes after but not before, and some people made similar changes before and after.

⁷⁷ It seems possible that an individual could end up following a more promising career pathway without explicitly changing their plans if the intervention made certain options more salient to them or otherwise nudged them to take some options over others. Additionally, communications studies have found that sometimes persuasive interventions can have a “sleeper effect,” where “a message initially discounted by message receivers comes to be accepted over time,” but this only tends to happen in certain circumstances, such as when “messages disseminated by low-credibility communicators can come to be viewed as true over time, particularly if they are memorable” (see Richard M. Perloff, *The Dynamics of Persuasion: Communication and Attitudes in the Twenty-First Century*, 6th ed. (New York, NY: Routledge, 2017), 374-7).

this sort of delayed implementation following their services (see introduction). It seems plausible that participants' assessments of their own expected impact would be more accurate after an intervention, so even if the average score does not change much, it could be that the plan changes they make will be better plan changes than they otherwise would have made. And attitudes do not necessarily need to change for expected impact to increase. An individual might already hold attitudes conducive to careers that are highly impactful for animals, just not have identified the best career pathways. For example, they might not have thought about certain promising options before.

There is some weak evidence for this idea from the exploratory analyses using participants' LinkedIn profiles. Although the confirmatory analyses in the online course study only provides evidence that that intervention successfully encouraged changes in participants' career plans, the results from the LinkedIn analysis for this intervention seem more promising than the results from the equivalent analysis of the online course.⁷⁸ It suggests that a higher proportion of people in the online course intervention group than the control group had made meaningful changes to their current roles and responsibilities, including a larger proportion of individuals taking on new responsibilities directly relating to effective animal advocacy or effective altruism. In fact, the proportion of people making such EA(A)-relevant, positive-seeming changes was more than three times higher in the intervention group than the control group. By comparison, The one-to-one calls intervention group also had a larger proportion of individuals making such

⁷⁸ Recall that, by comparison, the one-to-one calls study did not find evidence of significant differences on the career plans metric between the intervention and control groups but found stronger evidence of changes in career-related behaviours and self-assessed expected impact for animals than the study of the online course did.

EA(A)-relevant, positive-seeming changes than its control group, but the increase in the proportion of people doing so was much smaller.

These two studies have a number of limitations that prevent a simple interpretation of whether the interventions succeeded or failed. Nevertheless, they bring important new evidence to the question of whether career advice interventions such as one-to-one calls and online courses can increase participants' expected impact for altruistic causes, finding positive (albeit mixed) results.

References

Allen, I. Elaine, and Jeff Seaman, "Changing Course: Ten Years of Tracking Online Education in the United States" (2013), <https://files.eric.ed.gov/fulltext/ED541571.pdf>.

"Blog," Animal Advocacy Careers, accessed 17th July 2020, <https://www.animaladvocacycareers.org/blog>.

"Our Team," Animal Advocacy Careers, accessed 25th June, 2021, <https://www.animaladvocacycareers.org/our-team>.

"The Characteristics of Effective Training Programmes," Animal Advocacy Careers (10th January, 2020),

<https://www.animaladvocacycareers.org/blog/the-characteristics-of-effective-training-programme>
[S.](#)

“Advocacy and interventions,” Animal Charity Evaluators, accessed 17th July 2020,
<https://animalcharityevaluators.org/advocacy-interventions/>.

Benjamini, Yoav and Yosef Hochberg, “Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing,” *Journal of the Royal Statistical Society: Series B* (Methodological) 57, no. 1 (1995), 289-300.

Bernard, Robert M., Philip C. Abrami, Yiping Lou, Evgueni Borokhovski, Anne Wade, Lori Wozney, Peter Andrew Wallet, Manon Fiset, and Binru Huang, “How Does Distance Education Compare with Classroom Instruction? A Meta-Analysis of the Empirical Literature,” *Review of Educational Research* 74, no. 3 (2004), 379-439.

Bernard, Robert M., Eugene Borokhovski, Richard F. Schmid, Rana M. Tamim, and Philip C. Abrami, “A Meta-Analysis of Blended Learning and Technology Use in Higher Education: From the General to the Applied,” *Journal of Computing in Higher Education* 26, no. 1 (2014), 87-122.

Bostrom, Nick, “Existential Risk Prevention as Global Priority,” *Global Policy* 4, no. 1 (February 2013).

Caviola, Lucius, Stefan Schubert, Elliot Teperman, David Moss, Spencer Greenberg, and Nadira S. Faber, “Donors Vastly Underestimate Differences in Charities’ Effectiveness,” *Judgment and Decision Making* 15, no. 4 (2020), 509-16.

Cohen, J. E., *Statistical Power Analysis for the Behavioral Sciences* (Hillsdale, NJ: Lawrence Erlbaum Associates Inc., 1988).

Cook, David A., Anthony J. Levinson, Sarah Garside, Denise M. Dupras, Patricia J. Erwin, and Victor M. Montori, “Internet-Based Learning in the Health Professions: A Meta-Analysis,” *JAMA* 300, no. 10 (2008), 1181-96.

Coyer, Caitlin, Megan Fox, Elena Cantorna, and Lynette Bikos, “Effects of Participation in an Online Course on Undergraduate Career Decision-Making Self-Efficacy” (2018), https://digitalcommons.spu.edu/cgi/viewcontent.cgi?article=1076&context=spfc_research_conference.

“Learner Reviews & Feedback for Effective Altruism by Princeton University,” Coursera, accessed July 17, 2020, <https://www.coursera.org/learn/altruism/reviews>.

Duru, Maya, and Sarah Kopper, “Data analysis,” accessed 7th September, 2021, <https://www.povertyactionlab.org/resource/data-analysis>.

Faul, Franz, Edgar Erdfelder, Albert-Georg Lang, and Axel Buchner, “G* Power 3: A Flexible Statistical Power Analysis Program for the Social, Behavioral, and Biomedical Sciences,” *Behavior Research Methods* 39, no. 2 (2007), 175-91.

Gati, Itamar, Tehila Ryzhik, and Dana Vertsberger, “Preparing Young Veterans for Civilian Life: The Effects of a Workshop on Career Decision-Making Difficulties and Self-Efficacy,” *Journal of Vocational Behavior* 83, no. 3 (2013), 373-85.

Gim, G., 대학 진로 교과목의 효과에 대한 메타분석 [Meta-analysis of the effectiveness of career college courses] (2015), <http://sspace.snu.ac.kr/handle/10371/127892>.

“Briefing Series,” Farmed Animal Funders (February, 2019), https://farmedanimalfunders.org/wp-content/uploads/2019/07/Farmed-Animal-Funders_-2019-Briefing-Series.pdf.

Fouad, Nadya, Elizabeth W. Cotter, and Neeta Kantamneni, “The Effectiveness of a Career Decision-Making Course,” *Journal of Career Assessment* 17, no. 3 (2009), 338-47.

Harris, Jamie, Jacy Reese Anthis, and Kelly Anthis, “Health Behavior Interventions Literature Review,” (July 24, 2020), <https://sentienceinstitute.org/health-behavior>.

Harris, Jamie, “Pre-registration: The Effects of Career Advising Calls on Expected Impact for Animals” (29 June, 2020), <https://osf.io/pwufc>.

Harris, Jamie, “The Effects of an Online Careers Course and Workshop on Expected Impact for Animals” (26 August, 2020), <https://osf.io/cjasf>.

Jakobsen, Janus Christian, Christian Gluud, Jørn Wetterslev, and Per Winkel, “When and how should multiple imputation be used for handling missing data in randomised clinical trials—a practical guide with flowcharts,” *BMC Medical Research Methodology* 17, no. 1 (2017), 1-10.

John, Julia, “Hooray for Sri Artham’s Plant-Based Bacon!” (November 25 2019), <https://www.gfi.org/blog-hooray-bacon>.

Lam, Michele, and Angeli Santos, “The Impact of a College Career Intervention Program on Career Decision Self-Efficacy, Career Indecision, and Decision-Making Difficulties,” *Journal of Career Assessment* 26, no. 3 (2018), 425-44.

Luzzo, Darrell A., and Mary Taylor, “Effects of verbal persuasion on the career self-efficacy of college freshmen,” *CACD Journal* 94 (1993), 34.

MacAskill, William, “The Definition of Effective Altruism,” *Effective Altruism: Philosophical Issues* (Oxford, UK: Oxford University Press, 2019).

McKenzie, David, “Can business owners form accurate counterfactuals? Eliciting treatment and control beliefs about their outcomes in the alternative treatment status,” *Journal of Business & Economic Statistics* 36, no. 4 (2018), 714-22.

Prescod, Diandra, Beth Gilfillan, Christopher Belser, Robert Orndorff, and Matthew Ishler, “Career Decision-Making for Undergraduates Enrolled in Career Planning Courses,” *College Quarterly* 22, no. 2 (2019), <https://files.eric.ed.gov/fulltext/EJ1221402.pdf>.

“Reports,” Sentience Institute, accessed 17th July 2020, <https://www.sentienceinstitute.org/research>.

“Summary of Evidence for Foundational Questions in Effective Animal Advocacy,” Sentience Institute, accessed 27th April, 2020, <https://www.sentienceinstitute.org/foundational-questions-summaries>.

Šimčikas, Saulius, “List of ways in which cost-effectiveness estimates can be misleading” (20th Aug 2019), <https://forum.effectivealtruism.org/posts/zdAst6ezi45cChRi6/list-of-ways-in-which-cost-effectiveness-estimates-can-be>.

Singer, Peter, “Effective Altruism,” accessed July 17, 2020,

<https://www.coursera.org/learn/altruism>.

Sitzmann, Traci, Kurt Kraiger, David Stewart, and Robert Wisher, “The Comparative Effectiveness of Web-Based and Classroom Instruction: A Meta-Analysis,” *Personnel Psychology* 59, no. 3 (2006), 623-64.

Rains, Stephen A., Timothy R. Levine, and Rene Weber, “Sixty Years of Quantitative Communication Research Summarized: Lessons from 149 Meta-Analyses,” *Annals of the International Communication Association* 42, no. 2 (April 3, 2018), 105–24.

Reardon, Robert, Carley Peace, and Ivey Burbrink, “College Career Courses and Learner Outputs and Outcomes, 1976-2019” (2020),

<https://www.career.fsu.edu/sites/g/files/upcbnu746/files/TR61.pdf>.

Ruiz, Jorge G., Michael J. Mintzer, and Rosanne M. Leipzig, “The Impact of E-Learning in Medical Education,” *Academic Medicine* 81, no. 3 (2006): 207-12.

Perloff, Richard M., *The Dynamics of Persuasion: Communication and Attitudes in the Twenty-First Century*, 6th ed. (New York, NY: Routledge, 2017).

Tallent-Runnels, Mary K., Julie A. Thomas, William Y. Lan, Sandi Cooper, Terence C. Ahern, Shana M. Shaw, and Xiaoming Liu, "Teaching Courses Online: A Review of the Research," *Review of Educational Research* 76, no. 1 (2006), 93-135.

Tansley, Denny P., LaRae M. Jome, Richard F. Haase, and Matthew P. Martens, "The Effects of Message Framing on College Students' Career Decision Making," *Journal of Career Assessment* 15, no. 3 (2007), 301-16.

Tarigan, Medianta, and Supra Wimbarti, "Career planning program to increase career search self efficacy in fresh graduates," *Journal of Higher Education Theory and Practice* 11, no. 4 (2011), 75-87.

"A free online course exploring the science behind plant-based and cultivated meat," The Good Food Institute, accessed July 17, 2020, <https://www.gfi.org/OnlineCourses>.

Theeboom, Tim, Bianca Beersma, and Annelies E. M. van Vianen, "Does Coaching Work? A Meta-Analysis on the Effects of Coaching on Individual Level Outcomes in an Organizational context," *The Journal of Positive Psychology* 9, no. 1 (2014), 1-18.

Todd, Benjamin, "80,000 Hours Annual Review — November 2020" (May 14th, 2021), <https://80000hours.org/2021/05/80000-hours-annual-review-nov-2020/>.

Todd, Benjamin, “80,000 Hours Annual Review – December 2019” (5th April, 2020),
<https://80000hours.org/2020/04/annual-review-dec-2019/>.

Turner de Tormes Eby, Lillian et al., “An Interdisciplinary Meta-Analysis of the Potential Antecedents, Correlates, and Consequences of Protégé Perceptions of Mentoring,” *Psychological Bulletin* 139, no. 2 (2013).

Whiston, Susan C., Yue Li, Nancy Goodrich Mitts, and Lauren Wright, “Effectiveness of Career Choice Interventions: A Meta-Analytic Replication and Extension,” *Journal of Vocational Behavior* 100 (2017): 175-184.

“A guide to using your career to help solve the world’s most pressing problems,” 80,000 Hours, accessed 17th July 2020, <https://80000hours.org/key-ideas/>.

“About us,” 80,000 Hours, accessed 27th April, 2020, <https://80000hours.org/about/>.

“Our mistakes,” 80,000 Hours, accessed 19th May, 2020,
<https://80000hours.org/about/credibility/evaluations/mistakes/#we-were-overly-credulous-about-how-easy-it-is-to-cause-career-changes-and-our-investigations-into-these-claims-were-insufficiently-skeptical-and-thorough>.

Acknowledgements

Many thanks to Lauren Mee, Oska Fentem, Mattie Toma, David Reinstein, David Moss, David Rhys Bernard, Brenton Mayer, Karolina Sarek, Erik Hausen, Johanne Krog, Pegah Maham, Jonah Goldberg, Tom Beggs, Thej Kiran, and Juliette Finetti for providing feedback.