Our Dec 2018 annual review is published online. The following document is best thought of as a long appendix of notes to the review, which goes into more detail on many of the key questions our stakeholders have. It's not designed to be read all the way through, so we recommend skipping ahead to the parts you're most interested in.

It's also been subjected to less checking than the main review document, and it was mainly written in Dec 2018, so some parts are out of date. We decided to err on the side of publishing these additional notes despite their relatively rough condition in order to uphold our norm of defaulting to transparency.

See all our previous reviews and evaluations.

Glossary of key terms used in this review

- **Significant plan change** - Someone who told us they shifted their career plans due to 80,000 Hours, and expect to have more impact as a result.

- **Rated-N plan change** - We categorise the plan changes into buckets, depending on how much counterfactual impact we expect will result from them. This depends on (1) how likely the change was to happen without involvement from 80,000 Hours (2) the amount by which their impact increased (or decreased). Currently, we use buckets of 1/10/100/1000, which can also be negative. If a change is 'rated-10' that means that it falls into the '10' bucket. The aim is that we should very approximately be indifferent between ten rated-10 plan changes and one rated-100 plan change, but in practice the ratios don't perfectly line up.

- **Impact-adjusted significant plan change points** (IASPC) - A significant plan change that's rated-N is worth N points. So, if we record two rated-10 plan changes and one rated-1 plan change, we'd record 21 IASPC points.

- **Tier 1 advising applicant** - 'Tier 1' applicants are roughly the top 20% of coaching applications in terms of chances of entering a priority path, focus on social impact, and how much we can help them. For instance, a typical tier 1 applicant might have studied a quantitative subject at a top 20 university, have become interested in AI safety after reading our online materials, and be looking for advice and connections to enter the field. Another might be a medical doctor who read our materials on problem selection and is wondering whether to exit medicine or enter public health.

- **Donor dollars** - What it would cost for a donor to cause an approximately equally valuable outcome in dollars at the current margin. One donor dollar can be roughly approximated as $1 donated to the EA Long-Term Fund. There are many challenges in making donor dollar estimates, such as diminishing returns, changes in the margin over time, and hugely different views about the value of additional funding in the community. We hope to create a more precise alternative next year.

- **Passive target** - What performance we expect on a metric if we dramatically cut back effort invested in 80,000 Hours (e.g. to one person who just does the most low hanging fruit activities, like keeping the website running, and continues tracking plan changes).

- **Baseline budget** - Our future costs if we fulfill existing commitments, including staff salaries for current staff with increases according to our policy, increases in software costs to keep pace with our audience, and our office lease. The baseline budget doesn't include any hiring or significant new projects.

- **An advisor** - Someone who gives advice to our users within the in-person team. Formerly called a 'career coach'.

- **Priority area** - the global problems we most prioritise solving, listed here.

- **Priority path** - the roles we most prioritise filling within our priority areas, listed here.

What does 80,000 Hours do?

80,000 Hours aims to solve the most pressing skill bottlenecks in the world's most pressing problems.

We do this by carrying out research to identify the careers that best solve these problems, and using this research to provide free online content and in-person support, aimed at talented graduates aged 20-40.

The content aims to attract people who might be able to solve these bottlenecks and help them find new high-impact options. The in-person support aims to identify promising people and help them enter new paths that are a good fit for them.

Currently, the problems we're most focused on are our those listed here as 'especially promising'. These concern issues that are important for improving the long-term future, such as reducing catastrophic risks and 'meta' ways of helping. Within the problem areas, we're most focused on enabling people to enter our 'priority paths', such as AI policy, AI safety research, global priorities research and operations management at top organisations.

How our programmes work

The aim of the online content is to tell our readers about new high-impact career options and how they can contribute to them, especially our priority paths. The in-person advice helps our most promising readers decide where to focus and enter a high-impact role.

We attract our readers through the effective altruism community, promoting online content and referrals.

The online content breaks down into the following categories, which are roughly in order of engagement:

- The career guide and advanced articles, which introduce and justify the key concepts of effective altruism applied to careers.
- The problem profiles, which introduce our priority problems and explain how they can be solved most effectively.
- The career reviews and priority path write-ups, which review different careers, and explain who is a good fit and how to enter them.
- The podcast, which goes into more depth on these topics.
- The job board, which recommends specific positions to apply to.

The in-person advice breaks down into:

- ‘Generalist’ advice, which helps people narrow down their options, and provides introductions to mentors, jobs and funding.
- ‘Specialist’ advice, where a staff member provides advice, headhunting and coordination to solve skill bottlenecks in a single priority path.
- Headhunting, in which we aim to fill specific high-priority vacancies.

The in-person advice also helps to get people involved with relevant communities, especially the effective altruism community, which helps them coordinate and stay motivated over the long term.

The online content, in-person advice and broader community all play a role in most plan changes. Typically, someone will read our materials for several years and gradually become convinced they should make a shift. But, almost none of the most valuable changes we've tracked occurred without some kind of in-person contact, whether with the team or the community more broadly, so this appears to be necessary too.

Both types of programmes complement each other. The online content makes the advising more efficient by improving our understanding, finding more people to advise, and getting the people we advise more up to speed. The in-person advice helps us focus our research on the questions that are most relevant to our readers.

We give more detail on what the plan changes look like and how they come about in the next section.

2018 plan changes

How we track plan changes

We quantitatively track our impact by measuring the number of ‘impact-adjusted significant plan changes’ we cause (IASPC). There are many challenges with this approach, but we find it helps us to stay focused and get a rough idea of our impact over time.

Someone is counted as a significant plan change if they (1) tell us they shifted their career plans due to us (2) expect to have more impact as a result. We identify these reports through online surveys and by asking people we advise.

We then classify these plan changes into buckets, depending on how much counterfactual impact we expect will result from them. This depends on (1) how likely the change was to happen otherwise (2) the amount by which their impact increased (or decreased).[1] Currently, we use buckets of 1/10/100/1000, which can also be negative. The idea is that we should approximately be indifferent between ten rated-10 plan changes and one rated-100 plan change, but in practice the ratios don't perfectly line up (something we'd like to fix in 2019).

We also generally try to take a sceptical view on the likelihood that plan changes will be successful if they haven't yet led to concrete contributions to their field so we tend to rate particularly uncertain or promissory plan changes less than we think they're actually worth. For this reason, the average rating of our plan changes tend to increase over time.

For the highest-value plan changes, we speak to the person concerned and write up a case study of what happened. We also typically ask relevant organisations and our donors to provide estimates of the value of the increase in impact. For instance, if someone started working on AI safety research, we'd ask a couple of researchers in the field to guess their potential. We update the ratings as we learn more. For instance, if someone doesn't follow through with an intended shift, we'd down-rate the change.

Example plan changes

In addition to the three short examples in the review (Cullen, Cassidy and Michael), below you can find four more examples that were newly recorded in 2018. You can also see some older examples in our 2016 annual review.

Habiba Islam

---

[1] We are aware of some issues with using 'counterfactual impact' as our impact measure. For example, let's say Sally makes a plan change worth 50 utils. Group A spent 49 hours helping Sally to make the change. Group B then spent just 1 hour helping Sally. Both A and B were *necessary* for Sally to make the change (i.e. if only A or only B had helped Sally, she would not have made the change at all). In this case, both groups were 100% counterfactually necessary for Sally's change so Group B would be counterfactually responsible for 50 utils. In this case, using 'counterfactual impact' overestimates the credit we'd ideally give to them for deciding how much funding they should get. By contrast, let's say either Group A's work or Group B's work would have been sufficient for Sally to make the change (i.e. if only Group A or only Group B had helped Sally, she would have made the change). In this case, neither group was counterfactually responsible for the change and neither group would receive credit.

We know that this system is not ideal but unfortunately have not come up with a better system for allocating credit than attempting to track our counterfactual impact. We hope to write in the future about other methods of credit sharing that we've considered, each of which had major complications or drawbacks.

*In 2011, Habiba was a Politics, Philosophy and Economics (PPE) student at the University of Oxford. She attended the first 80,000 Hours presentation and talked to Will MacAskill following it. This led led her to take the [Giving What We Can further pledge](), *with which she committed to giving away everything she earned above £25,000*.*

After graduating with first class honours, she trained as a barrister and became a consultant at PwC. Over the next 5 years, she kept to her pledge and passively followed the progress of the effective altruism movement.

In 2018, both Will MacAskill and 80,000 Hours independently reached out to encourage her to apply to be Senior Administrator for the Future of Humanity Institute (FHI) and the Global Priorities Institute (GPI). After talking more to us and listening to the 80,000 Hours Podcast (especially the episode with Toby Ord), she decided she was happy to work in a role focused on improving the long-run future and that this particular job was a good fit for her.

Since taking this role in the Oxford University central Philosophy Faculty, she has been able to unblock FHI and GPI's hiring, finances and office move. Without 80,000 Hours, it's unlikely she would have become involved in effective altruism or considered this or similar roles.

Caleb W

In 2015 Caleb finished his economics and information systems undergraduate degree and started working in the New Zealand government. He was a policy advisor on issues such as housing and transport. When he applied for one-on-one advice in 2017, he was already committed to working in existential risk reduction but unsure what direction to take.

We spent most of our advising time with him looking into policy options outside New Zealand, which we thought would set him up better to work on reducing existential risk. This included discussing postgraduate programs in the US and UK, introducing him to people who were applying for those programs, directing him to our research on the UK civil service and linking him up with several people currently working there.

Caleb is now doing a highly ranked security studies masters degree at Georgetown University and aiming to engage with the US policy community after he graduates. He thinks there's a <50% chance he'd be on this path if not for 80,000 Hours. Instead, he'd likely be working on a less prestigious strategic studies masters degree in New Zealand or perhaps still be working in government there, though still aiming towards existential risk reduction and open to moving overseas in the coming years.

Carina Prunkl

Carina holds a PhD at Oxford University in the philosophy of quantum mechanics, and planned to become an academic in philosophy of physics prior to reading 80,000 Hours' research. Doing so made her aware of the importance of AI governance and led her to apply for advising. The advising gave her a deeper understanding of effective altruism and pointed her towards the Future of Humanity Institute at Oxford's research scholars program. She was accepted into it, and is now a Senior Research Scholar working on the Ethics and Governance of Artificial Intelligence.

Zac Kenton

In 2016 Zac was planning on being an academic in theoretical physics. He was a PhD student at Queen Mary University of London and had previously studied maths at Cambridge, and graduating distinction from the Masters of Maths. However, due to coming across effective altruism in the 80,000 Hours career guide he joined Giving What We Can, became involved in the community and decided to switch towards AI safety.

After leaving his PhD Zac worked as a data scientist at ASI (where some of his colleagues were also interested in AI safety) and had an 80,000 Hours advising session during which we put him in touch with a number of AI safety researchers. We also helped him craft a plan for publishing in ML and applying to roles that would lead to technical AI safety work. He is now a post-doc in machine learning at Oxford, where he works with Prof Yarin Gal and is based at the Future of Humanity Institute. You can see his publications here.

Description of typical plan changes

Plan change rating	Description
Rated-100	Someone switching into a promising position in one of our priority paths, in significant part due to us e.g. becoming an AI strategy researcher at FHI or OpenAI; taking a management position at an 'effective altruism' non-profit, earning to give with a realistic prospect of earning over $1m per year while planning to donate to the EA Long-term Fund or equivalent.
Rated-10	Often someone with roughly a 10% chance (in practice a bit higher) of succeeding in a rated-100 option or someone with a less convincing

	counterfactual story. Otherwise they often involve something like starting in a junior role at an 'effective altruism' non-profit, entry level policy roles aiming at a pressing problem area, or starting graduate school aiming at AI safety. A shift to earning to give while on track to earn about $200k would also count.
Rated-1	We initially aimed to make these roughly equivalent to an additional median Giving What We Can Pledger, and about 25% of the rated-1s are people doing exactly this due to 80,000 Hours. Otherwise they often involve people who have a 10% chance of making a rated-10 plan change; shifts to gain better career capital where we're not sure how it will work out (e.g. switching major) or by people who are less focused on social impact.

When setting up this system we initially aimed for a rated-100 plan change to very roughly correspond to $1 - $10m of counterfactual value created for our priority problem areas. You could roughly think of this as a donation of $1 - $10m to the EA Long-term Fund at the margin. However, there are a lot of issues with this quantification into 'donor dollars' and we want to develop a better system next year.

The scale is unbounded. We've recorded one rated-1000 plan change so far, which is similar to the rated-100 category but doing something roughly 10 times higher impact.

We roughly categorised the top plan changes by priority path:

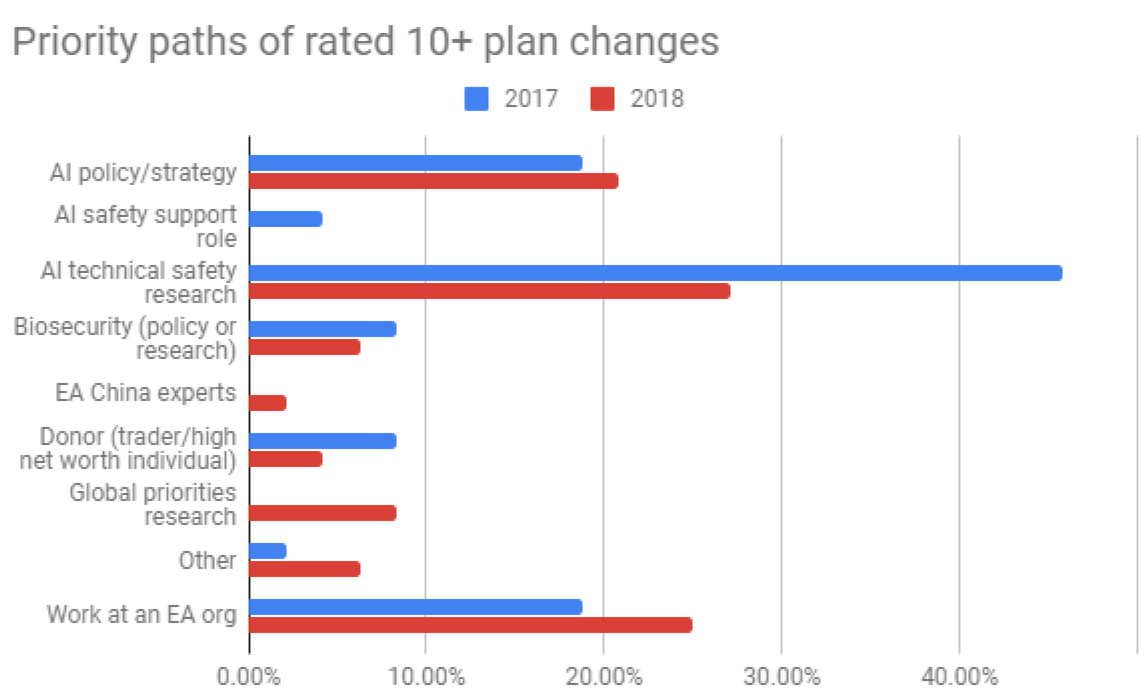

We raised the bar at which we recommend entry into AI technical research, leading to a decline in these plan changes.

We also categorised the plan changes by how much they've followed through with the intended shift. A 'milestone' would be something like applying to graduate school but not yet getting an offer. The 2017 plan changes were assessed as of the end of 2017.

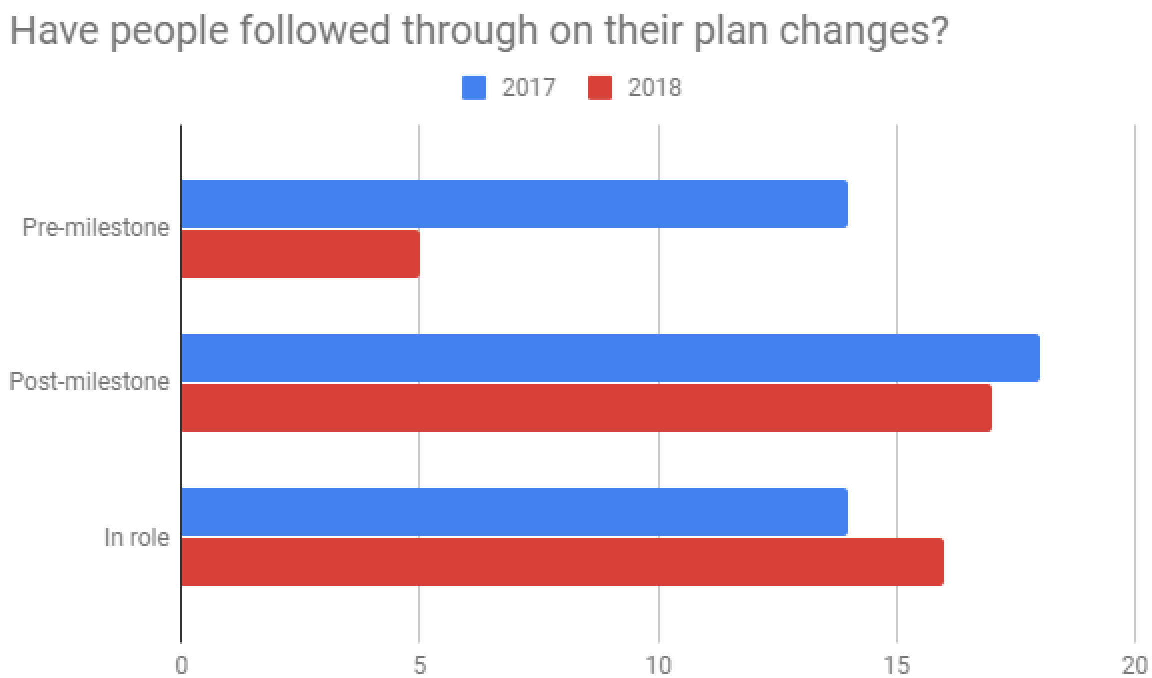

This shows that most of the 2018 plan changes have followed through to some degree. In 2018, the fraction that are pre-milestone dropped, potentially reflecting more stringent standards for confirming rated-10 plan changes.

How do the plan changes come about?

Typically, someone will engage with 80,000 Hours over several years in several different ways before being tracked as a highly-rated plan change.

When we ask people who make rated-10 plan changes, they say they most commonly find out about 80,000 Hours through another effective altruist organisation, following by coming across our online promotion, referral from a friend, and Peter Singer.

How did you first find out about 80,000 Hours? Answers given by people who reported a rated-10 plan change.	**2017**	**2018**	**2017**	**2018**
Through the effective altruism community	34%	36%	10	8
I don't remember	17%	18%	5	4
Recommendation from a friend	17%	14%	5	3
Online promotion	21%	14%	6	3
Link on social media	7%	5%	2	1
Search engine	7%	5%	2	1
Link on another website	3%	5%	1	1
TEDx talk by Benjamin Todd	3%	0%	1	0
Peter Singer	10%	14%	3	3
University activities fairs	0%	5%	0	1
Grand Total			**29**	**22**

For rated-100 plan changes, online promotion turns out to be more important, but the sample is small.

After first finding out about us, they'll typically explore the website, perhaps over several years. Some of the key shifts this might cause include:

- They hear about and become convinced by one of our priority areas, such as AI safety.
- They identify some particular careers or jobs in which they might contribute that they hadn't considered before e.g. become interested in studying ML, entering government or operations management, or they might see a particular job listing.
- If they didn't already know about effective altruism, they might explore the rest of the community online.
- They come to take our general approach to career choice, and use concepts like our career framework in comparing their options.

From there, they start to get involved in-person. Our current core programme here is advising, which helps in a couple of important ways:

- We help them analyse their options and become confident enough in their top choice to pursue it.
- We make introductions to people in the field to further help them gain confidence and enter.
- We sometimes also spot mistakes with their plan e.g. highlight a promising option they've missed.

Alternatively, our content might cause them to have discussions with other people in the relevant community if they already have connections. Otherwise they might apply to an EAG conference or local group to get involved with the effective altruism community.

As explained earlier, for most people the online content and in-person contact play a complementary role; though some kind of in-person contact with 80,000 Hours or with the effective altruism community seems to be close to necessary for a large proportion of people who make rated-10+ plan changes.

To find out about the plan changes, we follow up with people we've advised and survey our readers. We stay in touch over time to check that they followed through.

When we ask the plan changes what was most important in causing the change, they typically report a mixture of factors, but if we focus on the top one, we find the following:

What was most significant in causing these plan changes? Answers given by people who reported a rated-10 plan change.	**2017**	**2018**	**2017**	**2018**
Reading our online research	38%	54%	16	21
One-on-one advising	33%	31%	14	12
Talking to someone in our community	14%	5%	6	2
Speaking informally to a staff member	7%	3%	3	1
Attending an event	7%	0%	3	0
Using our make a decision tool	0%	5%	0	2
Using our job board	0%	3%	0	1
Total			42	39

For rated-100 plan changes, online research is less important and informal engagement with the community is more important.

What was most significant in causing these plan changes? Answers given by people who reported a **rated-100** plan change.	**2017 - 2018**	**2017 - 2018**
One-on-one advising	38%	5
Talking to someone in our community	30%	4
Reading our online research	15%	2
Speaking informally to a staff member	8%	1
Event	8%	1
Total		**13**

In the past, we've found that it takes about a year from someone joining the newsletter to reporting some kind of plan change, and then it takes a further two years on average to be tracked as a rated-100 plan change. This is because it takes time and multiple interactions for people to absorb the ideas and become convinced to make a shift. Then it can easily take 6 - 24 months to follow through with the shift and for us to learn about it. If they report a change in intentions, we might record that as a rated-1 or rated-10 plan change, but we typically require evidence of follow through to be rated-100.

You can see more statistics on each stage of engagement in the 'funnel metrics' section later.

Changes to how we track plan changes in 2018

This year we decided to stop tracking 0.1-rated plan changes, because they account for an overly small fraction of our total impact.

For the plan change surveys on our website which have a lot of responses, we stopped carefully rating each response. Instead, we quickly read every response and marked potential rated-10 plan changes. We then additionally chose a random sample of the remainder to look at more closely, and used the number of rated-1s in the random sample to estimate the rated-1s in the remaining portion.

We started formally recording negative plan changes for the first time this year. These are cases where we think we caused someone to have less positive impact (including causing them to have a negative impact). Currently, we've only recorded five of these and don't want to give details in case they become identifiable. There are likely more that we've failed to track, so we added a section to our impact survey this year that asks directly about ways we've made the respondents' careers worse.

We also became more rigorous about derating previous plan changes that didn't work out. We also think we became more stringent about what's counted as a rated-10 plan change. We tried to correct for this effect, but think our rated-10 plan changes may be understated by about 20% compared to 2017 (see the issues section later).

All these corrections mean that the figures below don't fully line up with our last annual review.

There are many problems with the metric, and we plan to make a round of revisions early in 2019.

Total plan changes and IASPC recorded

The following table shows the total number of plan changes recorded in each category by year.

Let's say someone was coached in 2013, made a career change that we tracked and rated-1 in 2014, and then, upon learning more, we upgraded their plan change to rated-10 in 2015. In that case, in the below, their plan change would add one rated-1 plan change in 2014 and one rated-10 plan change in 2015. In 2015, the IASPC total would go up 9 points.

	2013	**2014**	**2015**	**2016**	**2017**	**2018**	**Total**
Rated 1	13	57	144	653	632	487	1986
	NA	338%	153%	353%	-3%	-23%	
Rated 10	8	11	10	16	50	38	133
	NA	38%	-9%	60%	213%	-24%	
Rated 100	0	1	4	3	2	10	20
	NA	NA	300%	-25%	-33%	400%	
Rated 1000	0	0	0	0	1	0	1
	NA	NA	NA	NA	NA	-100%	
Rated 10 in previous years, but downgraded	0	0	0	-2	0	-7	-9
Rated -1	0	0	0	0	0	-4	-4
Rated -10	0	0	0	-1	0	0	-1

If we naively weight the categories 1/10/100 etc, then we find:

	2013	2014	2015	2016	2017	2018	Total
IASPC change from previous year	93	266	614	1084	2211	1713	5981
		186%	131%	77%	104%	-23%	

(Note that if a change is recorded as 10 in 2015 and then re-rated 100 in 2016, it'll appear in both years as +10 and +90, which is why the numbers in the graph often aren't multiples of 10 or 100.)

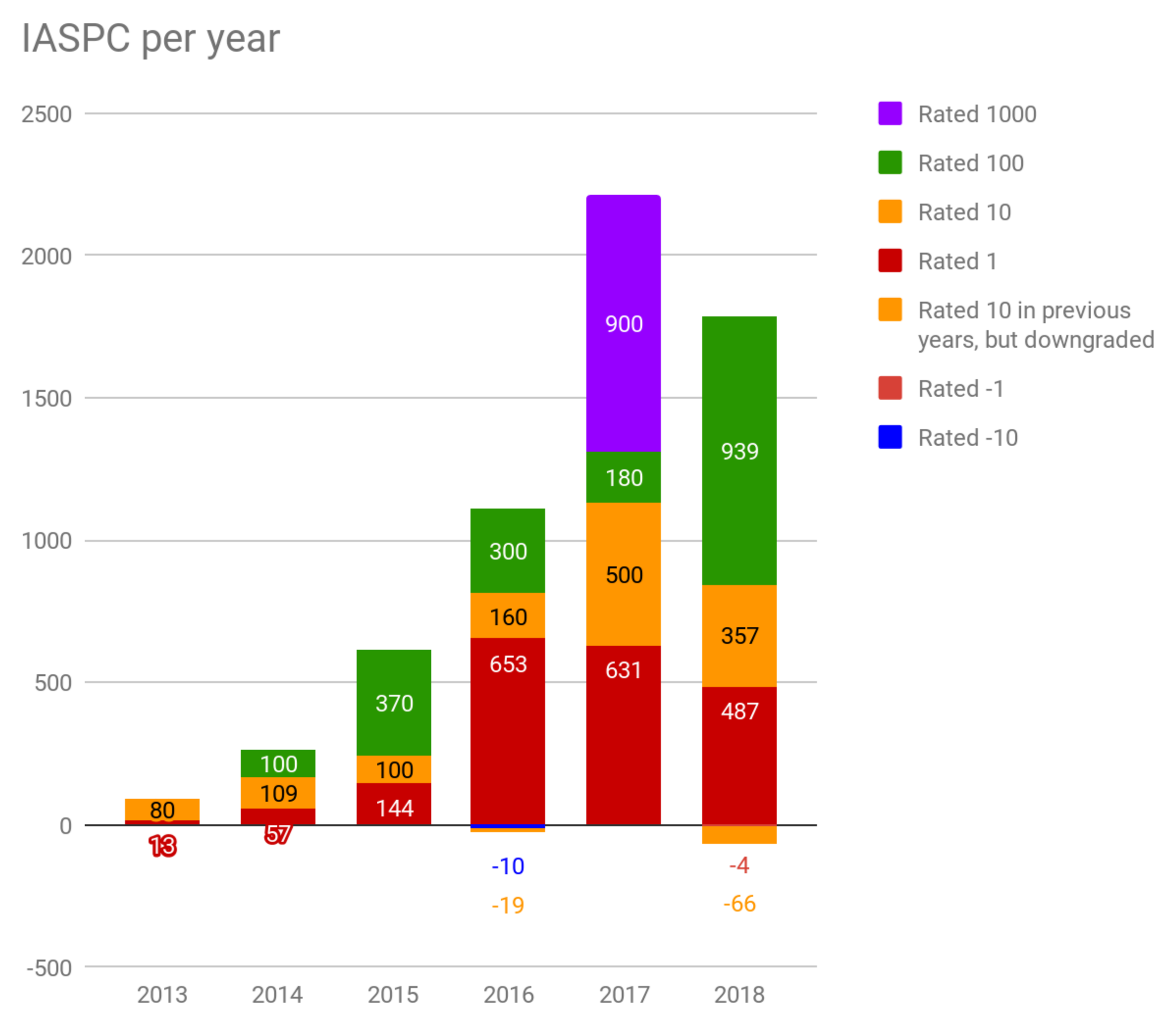

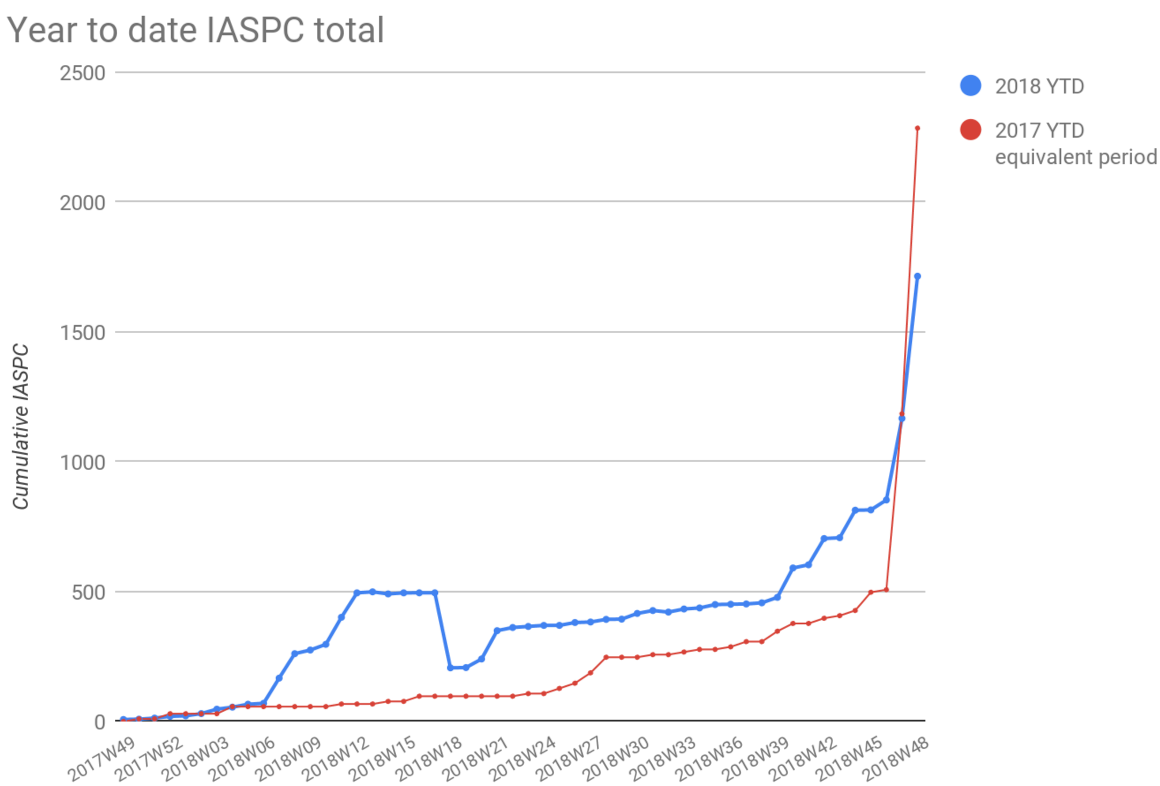

Total IASPC points with updated weightings

When we make a best guess at the expected value of the plan changes, we don’t think the average rated-1 plan change is actually 1/100th the value of the average rated-100 plan change. We think a more accurate weighting is something like the following, which we’ll refer to throughout this document as the ‘updated weighting’:

Updated weighting	
Rated 1	0.7
Rated 10	10
Rated 100	150
Rated 1000	2000

If we use these figures, then we get the following change in total IASPC per year, which we think better reflects how our tracked counterfactual impact changed each year.

	2013	2014	2015	2016	2017	2018	Total
IASPC change from previous year (using updated weightings)	89	299	771	888	3072	2068	7187
		235%	158%	15%	246%	-33%	

This means that:

- Over our entire history, we've recorded 7161 IASPC points.
- Since 2013, we've grown the annual number of IASPC points 23-fold, 88% p.a.
- Since 2015, we've grown the annual number of IASPC points 2.7-fold, 39% p.a.

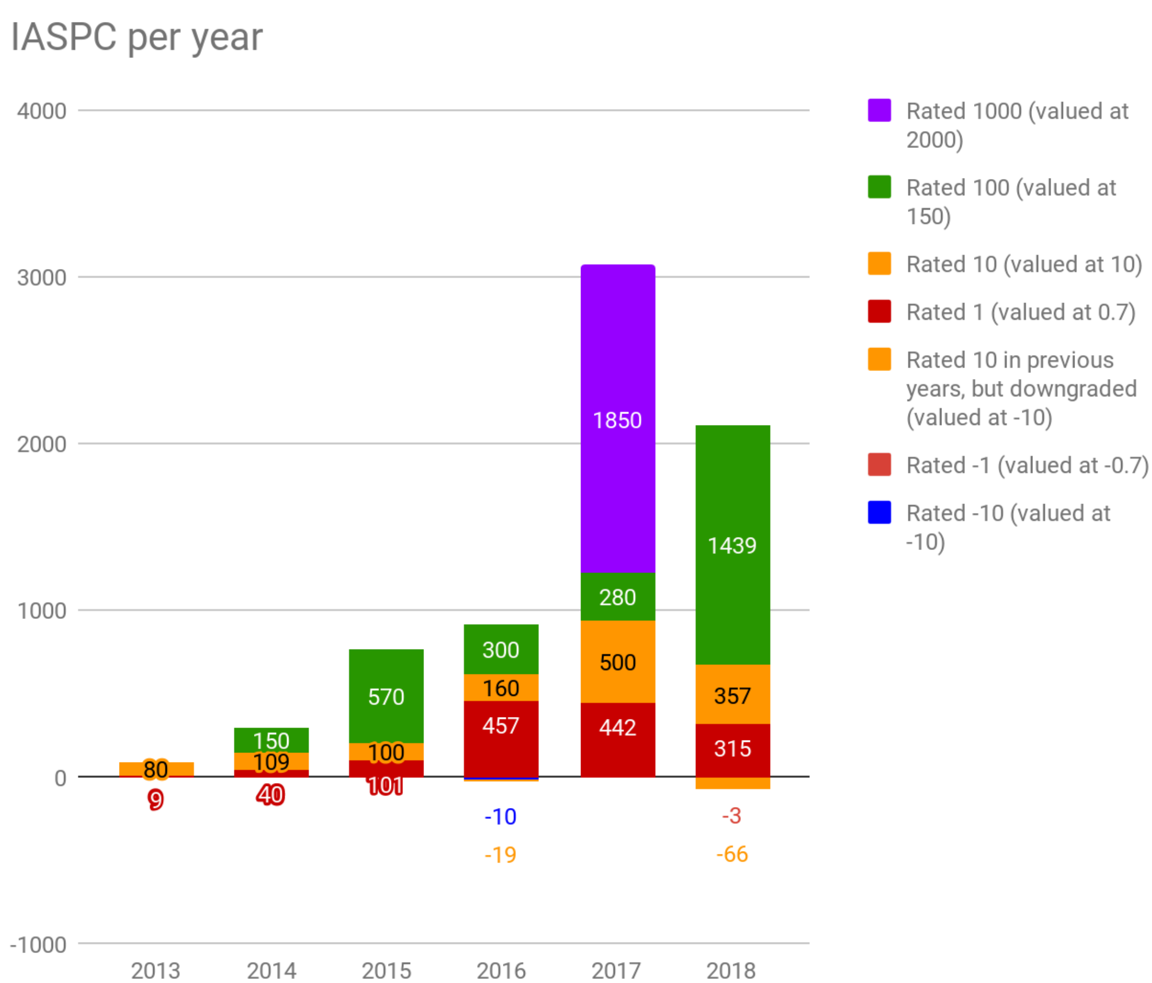

Why did IASPC fall in 2018?

There are a couple of potentially worrying trends in our IASPC metrics this year:

- We generated fewer IASPC points in 2018 than we did in 2017.
- We missed our growth target of 2200 IASPC points (old weightings).

We said that we expected a passive target of the following (i.e. the plan changes we'd get if we radically scaled back 80,000 Hours to one person who just does the most low hanging fruit activities and continues tracking plan changes):

- 600 rated-1 plan changes
- 25 rated-10
- 3 rated-100

This adds up to 1100 points. Our growth target was 2200 points, a doubling of the passive level.

We actually tracked 1713 points over 2018. This means we exceeded our passive target by 613, but missed our growth target by 487.

If we use the updated weightings, then the growth target would have been 2130 points, compared to the 2068 that we recorded, an 62 point shortfall.

What drove each type of plan change, and why did total IASPC points fall?

The biggest reason is that we didn't record another rated-1000 plan change. If we exclude the rated-1000 change in 2017, then the total would have grown 28%, or 73% with the updated weightings.

We estimated there was only a 25-50% chance of recording a rated-1000 plan change in 2018, so this was as expected.

Our key focus is growing the number of rated-100 plan changes, and the number grew 400% over the year. We take this as encouraging evidence that we have a reproducible process for creating rated-100 plan changes.

However, rated-100 plan changes typically take many years, so the rated-100 plan changes we recorded this year were mainly due to efforts in earlier years. It doesn't tell us much about the success of our efforts in 2018. (Though, as covered in the progress section, we expect our 2018 efforts will lead to more highly rated plan changes in later years.)

We recorded 145 fewer rated-1 plan changes (-24%). We think the main reason for this decrease was that we put less effort into tracking the rated-1 plan changes we caused, because they're not our key focus and the surveys take up time from users. For instance, we removed several prominent survey appeals on the website that were generating a lot of plan change reports with a low average rating. This meant that only 3909 people filled out an impact survey in 2018 (0.24% of traffic) compared to 6680 in 2017 (0.43% of traffic).

We also recorded 12 fewer rated-10 plan changes in 2018 relative to 2017, a decline of 24%. There are several potential explanations for this. One is that we think we probably applied a more stringent standard to rated-10 plan changes, perhaps leading to a 20% underestimate of their number. We will make another round of changes to the process we use to rate plan changes next year to try to avoid this problem.

A more concerning potential explanation would be that the effectiveness of in-person advice has gone down. Our best guess is that we spent somewhere between the same number of hours and 20% fewer hours advising in 2018. So, one possibility is that we invested less time in advising, leading to fewer plan changes. Another possibility is that we spent the same amount of time advising, but the 6-month conversion rate into rated-10 plan changes declined about 24%.

However, this data is very noisy, and the 6-month rate doesn't tell us much about the long-term conversion rate, or the overall case for the cost-effectiveness of advising. We think there's a good chance that the long-term conversion rate of 2018 advising will actually be higher, because we focused more on people who might make rated-100 plan changes over the coming years. Nevertheless, we'll closely monitor this trend next year.

Overall, the decline in rated-10 and rated-1 plan changes was more than offset by the increase in rated-100 plan changes. However, we're disappointed by the decrease in the number of rated-10 plan changes, because it will make it harder to get rated-100 plan changes in future years.

What explains why we missed our growth target?

We had a short-fall of 580 IASPC compared to our growth target, or 127 using the updated weightings. With the updated weightings, this can be explained by any of the following:

1. Not recording a rated-1000 plan change.
2. The decline in rated-1 plan changes tracked.
3. Having 0.5 less full-time equivalents on advising, due to unexpectedly losing a staff member for over six months. This could have likely generated 150 IASPC.
4. Not having any other major positive surprises.

2018 Funnel metrics

Below, we sketch out the typical stages of our engagement 'funnel', from first landing on the site, to subscribing to the newsletter or podcast, applying to advising, and making a plan change that we track.

Reporting period is Dec-Nov for all rows apart from Financial and Labour costs									
Year (Dec-Nov)	**2011**	**2012**	**2013**	**2014**	**2015**	**2016**	**2017**	**2018**	All-time total
Reach: Unique visitors to site	2,870	42,761	93,610	136,525	475,487	889,630	1,583,322	1,634,192	4,858,397
Year on year growth rate	NA	1390%	119%	46%	248%	87%	78%	3%	NA
New newsletter subscribers added	703	1,612	1,930	2,125	20,810	58,815	104,428	62,730	253,153
Year on year growth rate	NA	129%	20%	10%	879%	183%	78%	-40%	NA
New podcast subscribers added	NA	NA	NA	NA	NA	NA	4,473	10,481	14,954
Year on year growth rate	NA	NA	NA	NA	NA	NA	NA		
Total engagement hours (website + podcast)	NA	1496	3,407	2,793	11,998	28,633	86,179	141,030	275,536
Year on year growth rate			128%	-18%	330%	139%	201%	64%	
New tier 1 coaching leads identified	NA	NA	NA	NA	NA	NA	296	131	427
Year on year growth rate	NA	NA	NA	NA	NA	NA	NA	-56%	
People coached	Informal only	34	78	28	139	100	278	223	880
Year on year growth rate	NA	NA	129%	-64%	396%	-28%	178%	-20%	NA
New significant plan changes recorded	NA	NA	21	69	158	669	685	524	2,126
Year on year growth rate	NA	NA	NA	229%	129%	323%	2%	-24%	NA
New impact-adjusted significant plan changes recorded	NA	NA	93	266	614	1,084	2,211	1,713	5,981
Year on year growth rate	NA	NA	NA	186%	131%	77%	104%	-23%	NA
New rated-100 significant plan changes recorded	NA	NA	0	1	4	3	2	10	20
Year on year growth rate	NA	NA	NA	NA	300%	-25%	-33%	400%	NA
Financial costs converted to USD, calendar year (estimate)	0	$57,721	$163,108	$229,133	$347,950	$470,574	$751,480	$1,430,394	$3,450,360
Labour costs (in person-years, core paid staff only)	0	1.5	2.2	2.1	3.2	4.1	6.3	7.9	27
Labour costs (in person-years, including freelancers, interns and volunteers)	1.7	3.4	7.5	4.9	5	5.5	7.8	10.1	46
All time total costs divided by impact-adjusted plan changes, as of this date			$2,375	$1,253	$820	$617	$473	$577	NA

[See larger version.](#)

Tier 1 applicants are defined in the glossary. We tracked fewer of these tier 1 applicants this year because our capacity was filled, so we closed the coaching application process for most of the year. Instead, we collected names in a waitlist, which received over 3800 sign-ups. Of the 2098 people on the waitlist who we invited to apply, 23% applied. 16% of these applications were marked as tier 1 leads -- a similar fraction to previous years.

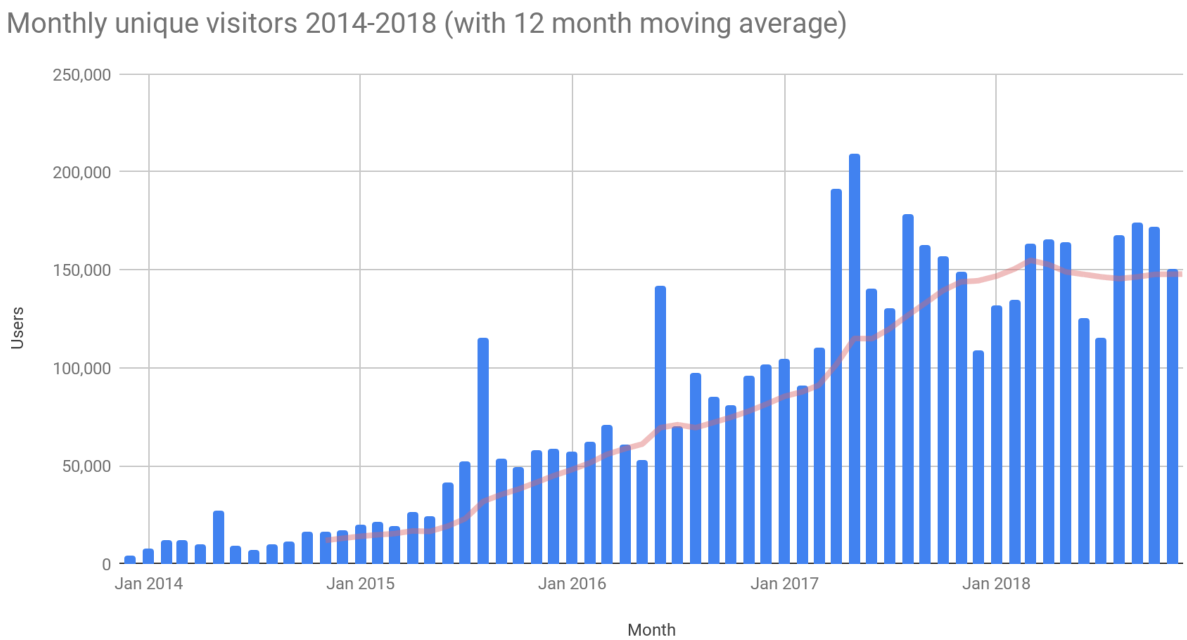

Web traffic flattened off this year. This was because we released fewer articles in order to focus on the podcast, which grew from 4,000 to 15,000 subscribers. The articles we did release were also more aimed at informing our existing audience rather than gaining traffic.

If we estimate the total engagement time with the podcast, and add it to engagement time on the website, total engagement grew 60% from 2017 to 2018. This is because podcast listeners tend to engage for much longer than users of written content.

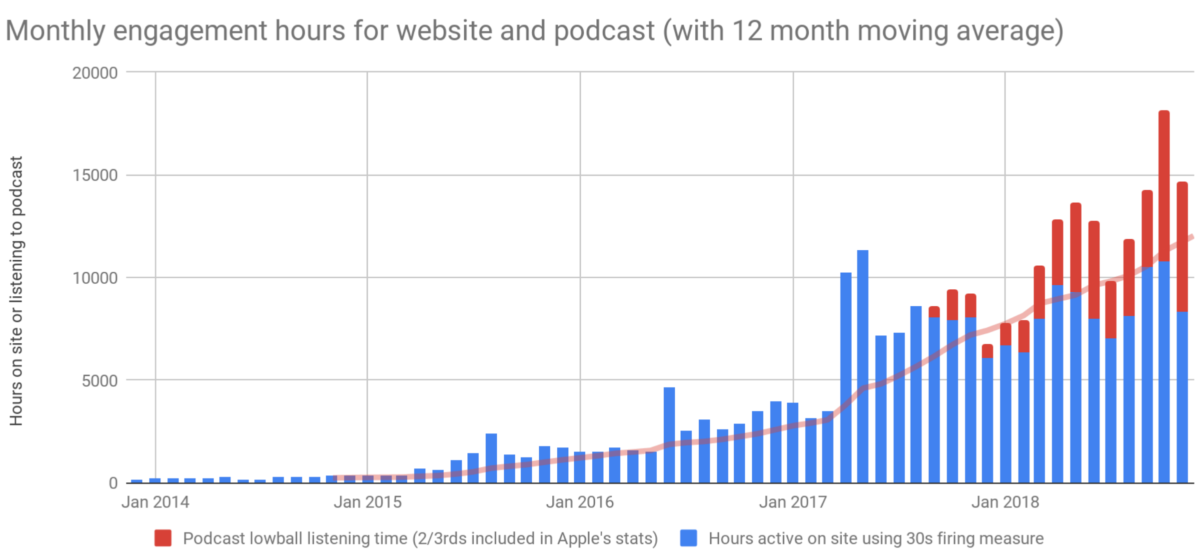

New newsletter subscribers per month declined since we stopped promoting the newsletter as heavily. For instance, we removed most links to the career quiz, dramatically cutting its traffic, and it was one of the largest sources of newsletter subscribers. Nevertheless, we still added over 60,000 subscribers.

Historical one-on-one advising efficiency

In this section, we present some preliminary data on the number of plan changes produced by one hour of advising time to date.

In the following figures, we count a plan change when it was discovered or partially caused by advising. This means the following do not represent the marginal counterfactual impact due to advising alone. Rather, they measure something closer to the impact of advisors *when combined with* the rest of 80,000 Hours. (To estimate the impact of advising alone, we'd roughly divide the figures by 2 or 3.)

Unfortunately, the data can't tell us all that much about the efficiency of coaching in 2017 and later because the sample sizes are small and more time needs to pass before we can see how well the more recent cohorts have done (more detail below).

That said, these are the figures:

Advising cohort	**2014-2016**	**2017 June - Dec**	**2018**
Number of rated-100 plan changes	4	4	1
Number of IASPC to date (old weightings)	570	645.9	187.2
Number of IASPC to date (updated weightings for the rated-100s)	770	845.9	237.2
Number advised	267	278	219
Hours spent	596	889	560
IASPC per hour (old weightings)	1.0	0.7	0.3
IASPC per hour (updated weightings)	1.3	1	0.4

Other notes: (1) The estimate of 2014-2016 total IASPC is inferred from the number of rated-100 plan changes, using our historical average of 70% of impact being from rated-100+ plan changes. (2) We count the hours that are spent directly on advising, including selecting applicants, preparing for meetings, taking meetings and following up. A full-time advisor has about 900 'advising hours' per year. (3) We don't include the first half of 2017 because we didn't analyse that time period yet. We had worse tracking for this period, which makes the analysis more difficult.

We think the main reason the efficiency seems to decline over time is because it takes a long time for plan changes to get rated-100. The average time from receiving advice and being rated-100 is 1.2 years, and it can take more than three years. This means that we expect many of the newer cohorts' plan changes have yet to come.

To accurately compare efficiency over time, you need to compare total IASPC after *the same amount of time* has passed since the coaching session. We've collected data that allows us to do this since summer 2017, in which time we've seen no discernible trend.

In part, this is because we also have very small samples that are dominated by outliers, also making it challenging to compare efficiency over time.

Our best guess is that our most recent advising will end up being similarly efficient as the 2014-2016 advising, and has a reasonable chance of being better. This is because our impression is that today we're able to work with people who are more qualified and more focused on top priority areas than in the past. However, we won't have the data to show this for several more years.

We also tried to carry out a similar analysis for headhunting and specialist advising. The sample sizes are even smaller here, and it also becomes hard to tell which changes were due to these programmes compared to our regular advising. However, the initial indication is that the efficiency is similar or higher than regular advising.

Also note that these programmes have additional benefits over the immediate plan changes they cause. For example, advising organisations on their recruiting strategy may end up being the most valuable benefit of headhunting.

Advising cohort	**Headhunting second half 2018**	**Previous headhunting 2017-2018**	**2018 AI policy specialist advising**
IASPC per hour (old weightings)	0.55	0.5	0.3
IASPC per hour (updated weightings)	0.55	0.5	0.7

See an estimate of the marginal effectiveness of advising later.

Main activities and achievements over 2018

We'll break our core activities into: the big picture; online content; research; in-person advice; internal systems and operations; hiring; and introductory content.

Our core progress in 2018 was:

1. We feel more confident we have a reproducible process for producing rated-100 plan changes, though we would still like to increase our confidence in it.
2. We improved our content to build our 'tier 1' audience and get them closer to making a plan change, especially through the podcast, growing engagement 60%.
3. We continued to experiment with how to make the in-person team more efficient per hour through headhunting and specialist advising.
4. Capacity building -- we appointed five team leads, made two hires and improved processes, putting us in a better position to grow the team next year.

The big picture

The main focus this year was creating an effective and scalable 'machine' for producing rated-100 plan changes.

We've already recorded over 100 rated-10 plan changes, and many of these appear to be driven directly by our online content, advising and headhunting, so we feel confident we can produce these reproducibly.

However, last year we had only recorded ten rated-100 plan changes over our entire history, so it was less clear we have a process for producing these.

We've now recorded ten more, so we feel significantly more confident that we have a reproducible process.

From looking at the case studies, we think about half of these changes follow fairly directly from our current core programmes -- online content and one-on-one advice.

We gave three examples in the review, Cullen, Cassidy and Michael.

Our evidence that we can produce these rated-100+ plan changes is not strong -- we ultimately only have a handful of examples -- but it seems worth making a larger bet on.

Online content

See a list of all the content we released this year. The aim of the new content was to better help tier 1 applicants and attract them to advising, which we think we're improving at doing. The following contains some more details.

The podcast was a major focus. We released 34 episodes (66 hours) and created a new landing page, growing subscribers from 4,000 to 15,000, meeting our targets.

The aim of the podcast is to provide in-depth content at the right level for potential tier 1 applicants, and we think it's succeeding.

Many of our rated-10 plan changes and key people in the effective altruism community say they are fans and find it highly useful. We receive a notable amount of positive feedback about the podcast, including these five examples:

> "Thanks for making an incredibly informative podcast. ...this has been the most intellectually stimulating audio that my eardrums have ever registered. You truly live up to your "unusually in-depth" description. ... A bit of background on myself, if you're curious. I am a recent MIT graduate in math+CS, and I had the good fortune of learning about EA through friends..."

> "Brilliant. Ops roles are alien to many people, including myself. I had never considered it a suitable role for me, but this episode changed that view."

> "I have often felt very nihilistic and fatalistic in the face of the problems that humanity faces in the coming decades. The voices and ideas in the 80,000 Hours podcasts have given me a lot of reason to not despair and just get to work on trying to make the world a better place. I also had heard and read a lot on AI safety in the past, but there was something in the Paul Christiano interview that really made the ideas click for me and helped me better understand the urgency of the risk."

> "Initially I suspected that this podcast would be a little dry; how to use your career to benefit others? But it's actually just fascinating discussions with brilliant people, and I've found the topics genuinely fun to hear about… it reminds me a lot of the Sam Harris podcast, but with short segments of practice advice for people who are thinking about these topics in the context of a career decision. So good."

> “Your podcast has substantially enriched my life. The format is really engaging for me (more so than any other 80000 Hours or EA content) and I think this has improved over time. ”

We also actively seek out constructive criticism, and here are five examples of negative feedback:

> “I found Rob not so inspiring this episode, for the same reasons that Tara mentions about what she thinks of 80k's career profile. The framing is too much "great person theory of history" and too little "standing on the shoulders of giants" (the giants here are capable ops people)."

> “I would like to hear more about the background of people working at GiveWell. I would like to be able to guess whether my background would make me a suitable candidate to apply for a position with GiveWell.”

> “Ng had this annoying habit of just saying what he felt like saying in response to your questions rather than answering them, which meant that I didn't learn stuff that I wanted to know (which was normally your questions).”

> “The section on applying Bayes rule to Trump, the China-US conflict and North Korea felt messy. It is difficult to do justice to such complex issues in a short amount of time; it may have been better to spend more time on fewer examples.”

> “I wish less time was spent on the survey results and spent more time on the career decisions of Katja and her co-workers/collaborators, the logistics of running an org like AI Impacts, and the prospects or desirability of many young EAs pursuing such a path and starting their own version of AI Impacts, as opposed to trying to get high-status positions in gov't or prominent EA orgs like OP and FHI.”

The average recent episode has been downloaded about 14,000 times and people typically listened to more than half of each. This meant that despite annual traffic remaining roughly unchanged at 1.6 million unique visitors, total monthly engagement hours grew 60%.

We would be surprised if this didn’t convert into more highly rated plan changes over the coming years. There are now probably thousands of people who have listened to tens of hours of this in-depth content. We started to see readers mention the podcast in our annual impact survey this November.

The most downloaded episodes were:

- Tyler Cowen on long-termism and economic growth - 13,900 on our feed and 60,000 on Tyler's feed.
- Bryan Caplan on whether higher education is mostly signalling - 16,600.
- Robin Hanson on The Elephant in the Brain - 15,800
- Paul Christiano on AI safety - 15,100

The episodes that were most highly rated on enjoyment and usefulness by our committee of about 10 volunteer raters drawn from our audience were:

- Tara Mac Aulay on operations management - 4.4/5
- Allan Dafoe on AI strategy - 4.3/5
- Holden Karnofsky on the Open Philanthropy Project - 4.3/5

Our other main focus was improving the website so that it better appeals to the readers who have the best chances of entering a priority path, which we operationalise as 'tier 1' applications to advising. As explained later, we think the current career guide is probably better at attracting undergraduates rather than people in their late 20s; not focused enough on our top priority areas; and not in-depth and rigorous enough for this audience. If we fix this problem, we can make the advising more effective.

Within this focus, we've **started redesigning the website**, doing user interviews and sketching out a new key ideas page. We released an initial round of improvements to the homepage to better highlight our content, and a new version of the problem profiles page (our second most viewed page at 25,000 pageviews per month). We think the new page does a much better job at conveying our views on which problems are most pressing, which is one of our most crucial positions.

We've also **drafted about 50% of our new key ideas series**. Most importantly, we released an article explaining our views on which career paths are highest-impact, including an updated list of priority paths. This was probably the most important single piece of content we released this year -- our views about which paths are best were not clear before, but it's perhaps the most informative advice we can give.

To further support the new list, we released a supplementary article on suggestions by skill type, articles on operations management and China specialists as new priority paths, and drafted an in-depth article on AI policy careers. We think these articles have helped to shift our readers and the broader effective altruism community to focus more on these options, and we've especially noticed many more people talking about careers in operations management.

We finished an article on coordination and released articles on accidental harm and comparative advantage. The rest of the key ideas series covers how we choose problems, our moral values, how to weigh the importance of career capital and exploration, and our process for choosing a career. We've prepared drafts and improved our understanding of each topic.

We also released several major updates to the job board. We now list over 50 new vacancies each month, updated every two weeks. Due to this, the board receives over 20,000 views per month, making it our 5th most viewed page. This results in more than 5,000 clicks through to the external sites listing the jobs each month. We've received reports from several organisations that the board accounted for a significant fraction of their applicants, and from people who secured jobs at DeepMind, GovAI and the Open Philanthropy Project that they discovered from the board. This could mean we're now able to cause rated-10+ plan changes with only online content, which we've not been able to do in the past.

Some other notable content we released this year include:

- A blog post on the Freakonomics experiment about the status quo bias, which received over 50,000 views.
- The psychology replication quiz - 37,000 views.
- How to transition into ML engineering - 19,000 views.
- Think twice before talking about 'talent gaps' - an important clarification to our advice.
- New career reviews: academic research, being a congressional staffer, commercial law
- Why not to rush to translate effective altruism in new languages - one of the most upvoted posts on the EA Forum.

In total, we released 24 new articles or pages. Combined with the podcast, we released 57 pieces of content, compared to 48 last year. See a list of every release.

Research

By research, we mean improving our understanding rather than communicating it. We didn't have any staff focus directly on research this year (something we'd like to change), but we learned a significant amount in the course of producing the content above and in our advising.

Some examples are:

- **We learned lots of details about how to enter the priority paths**, who is best suited to them, and which roles are most valuable. This is especially true within the paths we wrote up this year - AI policy careers, operations management and China specialists - as well as the guide to entering ML engineering. For instance, within AI policy, we learned that studying ML and working in an AI lab may be a good entry route, which we write about in our upcoming article.

- While writing the key ideas series, **we thought about key career strategy questions**, such as the importance of career capital and discount rates. We started the year with very high discount rates (following our talent survey), but then realised there were extra reasons they should be low, and now put some credence on both very high and low or even negative discount rates. We also thought about how to best characterise our key advice, and now put more emphasis on entering a priority path directly, rather than taking a detour to gain flexible career capital or broadly explore. We hope to write more about this next year.

- We found many of the podcasts improved our understanding, such as Paul Christiano's interview on his AI safety agenda, Brian Christian on the explore-exploit tradeoff and Tom Inglesby on preventing global catastrophic biological risks.

- We redid our survey of leaders in the effective altruism community about talent needs. The main thing we learned is that the results were similar to last year, suggesting some stability on how people respond. However, after discussing the results in more depth, we became more sceptical about how much to rely on the estimates of the value of recent hires and discount rates. Next year, we'd like to do more in-depth interviews with managers to understand what's driving their estimates.

In-person

We gave one-on-one advice to 223 people this year. The advising continues to receive very positive feedback:

We think the advising was similarly effective as last year. Our impression was that we advised more talented people, which could result in more rated-100 plan changes over the coming years. However, as discussed, there is some evidence that the conversion rate into rated-10 plan changes declined up to 24% in the short-term. Overall, the conversion rate is very noisy so it's hard to draw firm conclusions.

We advised fewer people than intended, mainly due to unexpectedly losing a staff member for more than six months, but also from spending more time investing in hiring and online content. We may have also spent up to 20% longer per person compared to 2017. We discuss the lack of coaching capacity further in the setbacks section.

Niel continued with AI policy specialist advising, as a test of the specialist model in one of our highest-priority areas. We think that more experienced people have the potential to make higher-rated plan changes, but they're hard to advise without more domain expertise. This suggests we should focus on specialist advisors rather than generalists. If the model succeeds, we could imagine hiring 3-5 specialists covering our highest-priority areas in the coming years.

In order to test out the model Niel surveyed 134 people in the community about what prevents them from entering the path -- we found that it's a lack of information about concrete options, followed by a lack of credentials and funding. To help with this, Niel wrote an article making the case for the path and how to enter it, which involved several significant changes in how we frame the path. We also significantly improved our understanding of the path over the year. However, we've only advised 46 people within it (included in the 216 figure above). It's too early to tell whether specialist coaching is more effective than generalist coaching, but our initial estimates suggest that to date specialist coaching is somewhere between about as good and a fair bit more effective than generalist coaching, and might be many times more effective in the long-term.

In 2017 and early 2018, we carried out several short trials of **headhunting**, which suggested it might be more effective than our existing advising, especially from the perspective of getting highly rated plan changes. Due to this, Peter McIntyre started to work full-time developing the headhunting programme in September.

His main focus so far was developing our process and systems for headhunting. For instance, he recruited a professional headhunter as an advisor, and started working with several organisations. Most importantly, he ran an intensive process to help the Global Priorities Institute to find a new Head of Research Operations, turning up several promising candidates and leading to a successful placement.

Given results so far, we think it's likely we can add value to many high impact organisations both by finding leads and advising them on recruitment best practices, and that headhunting could be more effective than advising. However, we've only done about five months of headhunting to date, which resulted in only one or two plan changes, so we don't yet have a good measurement.

Headhunting is also promising as a programme because it's an existing skill-set and doesn't require as in-depth knowledge of effective altruism, which makes it easier to hire for. There also seems to be easily two to three times as many roles we could usefully headhunt for than we currently have the capacity to fill, and as we expand we can get more efficient from network effects. However, we're not sure we'll have enough management capacity to oversee this hiring over the next year.

We also reviewed several other ways to improve the in-person advice, such as setting up an AI policy scholarship, but decided to deprioritise them.

Internal systems

We made a **significant upgrade to our systems for tracking plan changes**. We now record when the plan change was first recorded, then update its rating over time, recording all the changes. We also added tracking of negative plan changes. This became necessary after adding more categories of plan changes at the end of 2017.

We also made lots of improvements to coaching systems; setup a new headhunting CRM; migrated to G Suite; automated our metrics sheets; and dealt with GDPR compliance.

CEA operations

We cover 30% of the costs of CEA operations team, who provide our essential systems in finance, legal compliance and HR.

Despite major changes in the team, they kept all of these systems running, and also significantly improved many of them.

For instance, **CEA US passed its first ever US audit with an 'unqualified opinion'**, the best rating the auditors can deliver, and the team set up processes to make this even faster.

We also made significant updates to the way we create and track spending & budgets, including replacing our bookkeeper, so that that our budgets are far more fine grained, more accurate, and more predictive than they were previously.

Hiring and capacity building

Hiring went better than our expectations in 2018. We ran two main hiring processes, and as a result hired:

- Michelle Hutchinson as a full-time advisor. She was one of the first staff hired by CEA in 2012, formerly the Executive Director of Giving What We Can. After that, she helped to launch the Global Priorities Institute as their Operations Director. She makes a great advisor, and was able to increase our capacity within a month of starting.
- Howie Lempel on research & strategy. He started at GiveWell in 2013, and later became a program officer at the Open Philanthropy Project, with a focus on funding global catastrophic risks and biosecurity. Before that, he went to Yale Law for two years, where he spent his summers as a public defender in Louisiana and suing prisons with poor conditions at the ACLU. Howie has also worked on white collar crime at the Manhattan DA's office and was a research assistant at the Brookings Institution for two years. After he started in September, we've already found his input extremely helpful.

We also:

- Retained all our core staff.
- Increased Keiran Harris to ⅔ time as podcast producer, which enabled us to consistently release episodes with only about 2 days per week from Rob. Keiran was previously a professional poker player and writer.

- Increased Maria Gutierrez to half time as job board manager. She has worked with us for years on graphic design, and studied animation at NYU.
- The CEA operations team hired Josh Axford, an experienced operations manager, to lead on UK operations, as well as Chloe Malone as an operations specialist in the USA.

Roman Duda successfully transitioned from the research team to the web and tech team. He has completed valuable work on the systems that support the in-person team, his web development skills are improving rapidly, and he's already making changes to the site. He has also taken over responsibility for several internal systems, and continues to oversee our metrics.

We also revamped our hiring process, ran hiring rounds for an advisor and researcher and ran five work trials that did not lead to permanent hires. We have two more trials scheduled for early 2019.

The other major progress we made on capacity building this year was **developing five team leads**. I have handed over most responsibility for hiring and product design to the leads within their area. Given the success in hiring, this seems to have gone well. If each team lead can manage 3-4 people, this now means we have room to hire about 10 more, whereas before we were bottlenecked by my management capacity.

Team org chart

The following lists everyone working on the team over half time, and indicates their key responsibilities (though their official job titles are different):

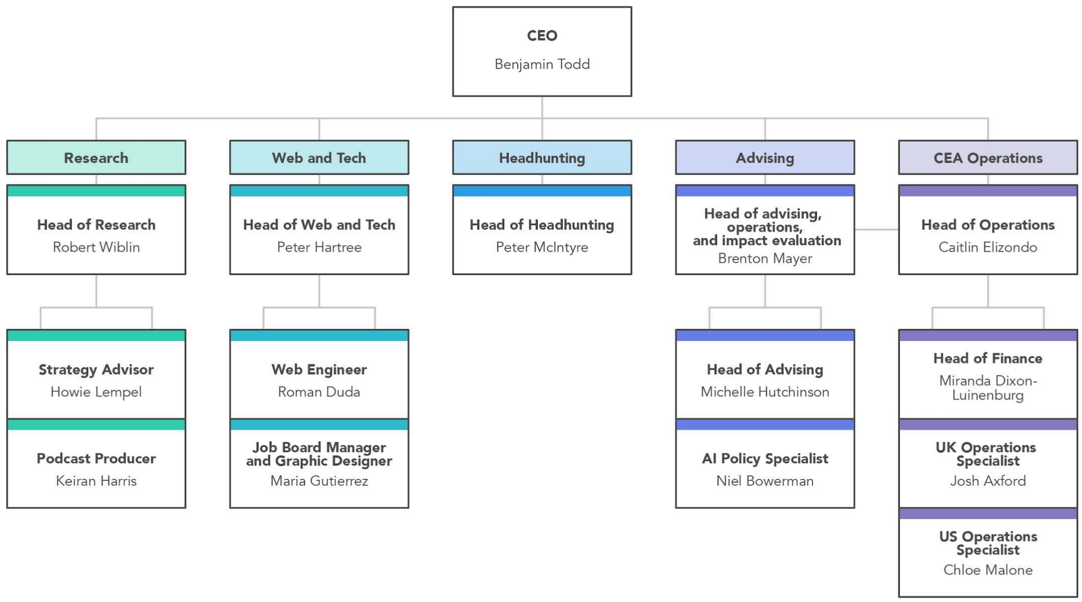

We also have about 11 freelancers, many of whom act as mentors to the people we advise, but also do editing, marketing and operations. In total, they account for about about one full-time equivalent.

All considered, the team grew about 35%.

	Jan 2018	**Dec 2018 (expected)**	**Average over the year**
Full-time staff (inc. central)	7	10	7.9
Freelancers (FTE)	2.3	2.6	2.2
Total	9.3	12.6	10.1

Introductory content

Besides getting a small number of high-value plan changes, 80,000 Hours also introduces a broader audience to the ideas of effective altruism through our website. We didn't focus directly on this broader aim, but the website continues to deliver value of this kind.

First, our overall traffic held steady at 1.6 million unique visitors, of which 80% were new according to Google Analytics.

In addition, traffic held steady to many of the career guide pages, despite a lack of promotion. These are our 12 most viewed pages in 2018:

	Views in 2018	**All time views**
Home page	455,354	1,344,515
Problem profile landing page	326,795	512,527
Career quiz	196,457	687,861
Career guide part 1: job satisfaction	189,180	521,956
Job board	161,774	235,906
Career guide home page	153,960	490,521
Career guide part 8: personal fit	105,813	213,929
Career guide part 5: the world's biggest problems	105,663	193,056
List of the most urgent global issues	87,523	229,317
Career guide part 6: Highest impact jobs	72,946	186,358
Career guide part 9: how to be successful	71,485	251,198

The pages where the largest number of people initially landed on the site over 2018 were:

Landing Page	**Sessions**

1. /	**342,983** (12.31%)
2. (not set)	**259,073** (9.3%)
3. Problem profile landing page	**246,185** (8.84%)
4. Career quiz	**147,455** (5.29%)
5. Career guide part 1: job satisfaction	**118,695** (4.26%)
6. Job board	**93,355** (3.35%)
7. Career guide home page	**58,818** (2.11%)
8. Career guide part 5: the world's biggest problems	**56,146** (2.02%)
9. Randomised experiment: If you're genuinely unsure whether to quit your job or break up, then you probably should	**50,030** (1.8%)
10. Career guide part 8: personal fit	**49,977** (1.79%)

The website appears to be succeeding in introducing people to the effective altruism community. We argued last year that 5-15% of the membership of the community first found out about it through 80,000 Hours.

The case seems stronger this year. The 2018 effective altruism survey suggests that 80,000 Hours is becoming an increasingly important source. Over the entire sample, only 8% said they first found out about effective altruism through 80,000 Hours, but:

> "The numbers of EAs reporting first hearing about EA from 80,000 Hours explodes from 9% to 13% to 25%, making them the largest single source of EA recruitment in 2018."

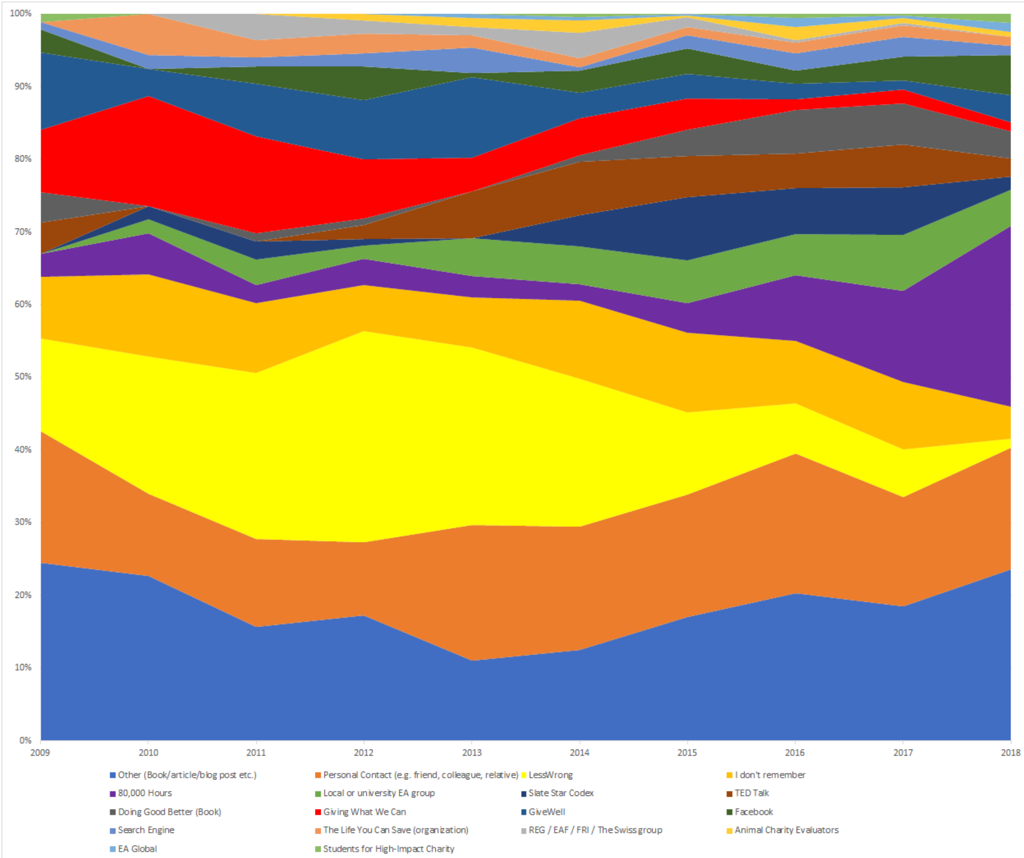

(Note that the figures aren't reliable in the early years due to the small sample. For instance, 80,000 Hours didn't exist until 2011, but several people say we told them about effective altruism before this date.)

80,000 Hours seems to be even more important in how people say they first 'get involved' in effective altruism. In the survey, 80,000 Hours was the most commonly cited factor, and from 2017 - 2018, it was mentioned by over half of respondents (though people could give multiple answers).

We'd like to see more analysis of how the source of involvement changes by level of contribution to the community. The EA Survey team looked at their data but didn't find a clear difference. In our survey of leaders in the community, 2 out of 29 (7%) said that 80,000 Hours helped to get them involved.

The increasing role of 80,000 Hours almost exactly lines up with our growth in web traffic, which started in 2015 from the release of the career quiz and promotion of Doing Good Better, and continued in 2016 due to the release of the career guide. It may also be because other sources have declined, such as GWWC and LessWrong.

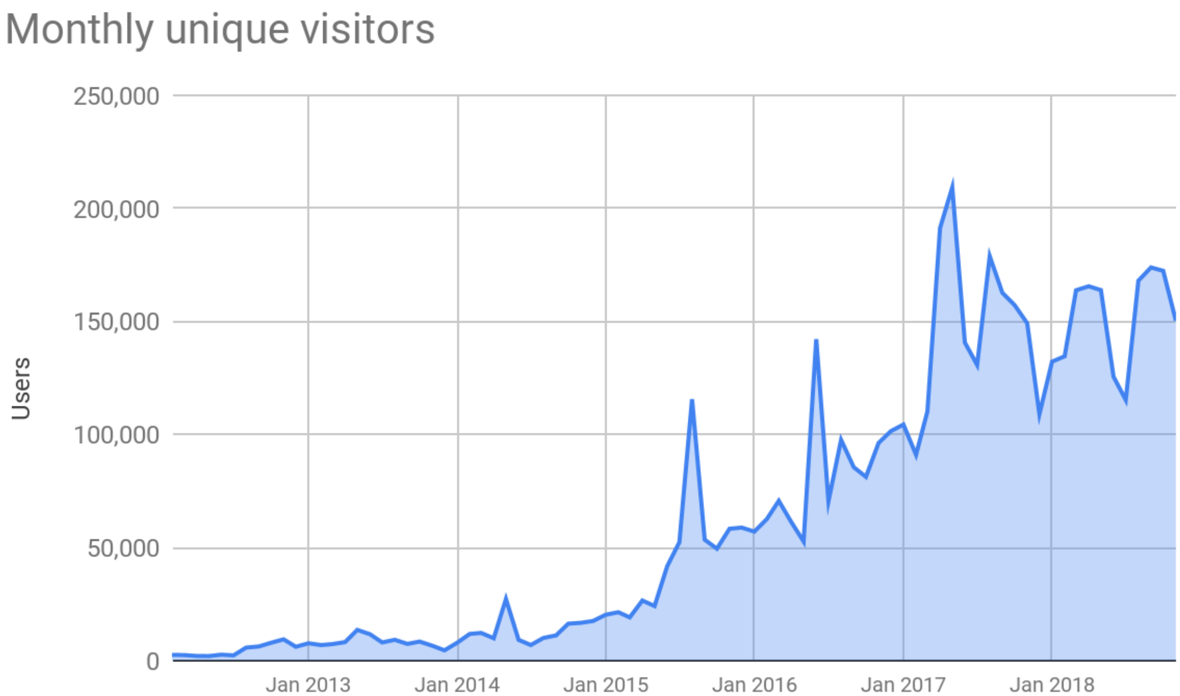

In 2018, our traffic didn't grow, so we wouldn't be surprised if the trend flattens off.

The fact that 80,000 Hours is introducing so many people to effective altruism is both a source of impact and risk -- if we do it badly, then we might put people off or cause other problems in the community. The two main issues we're worried about here are putting off certain groups who are well placed to enter our priority paths (e.g. mid-career policy professionals and academics), and demoralising those who aren't able to enter the paths, which are covered in the setbacks section.

How our plans changed over the year

We stuck with the focus we set out in the [last annual review](#) of designing a 'machine' that produces rated-100 plan changes.

We also delivered on most of the [concrete projects](#) we listed, though some projects were deprioritised. For instance, we swapped updating the biorisk and nuclear security profiles for write-ups of new priority paths, such as operations management.

The main strategic shifts we made during the year were:

1. We made hiring a greater priority, since we wanted to increase the diversity and robustness of the team, and we realised there were some areas where hiring is likely the key bottleneck, such as generalist advising, web programming capacity and certain types of research.

2. We decided to focus the in-person team on advising and matchmaking, rather than other ideas such as a scholarship or fellowship. This was because: (i) these alternatives didn't seem obviously more effective (ii) we want to stay focused (iii) we have a comparative advantage in our existing programmes.

Problems we face

Weaknesses of 80,000 Hours and risks to growth

We think these represent the most significant drawbacks of 80,000 Hours to donors and stakeholders. We covered these in the review.

Below we also list some 'mistakes' from 2018 – things that went badly but which seemed like they could have reasonably been prevented – and some 'issues' from 2018 – difficulties we struggled with and held up progress.

Mistakes in 2018

We still haven't updated our writing on career capital

In our last review, we noted that we hadn't done a good enough job on communicating our views on career capital, and especially how they changed.

We intended to write more about this issue, then update the career guide. We have drafted two major articles on career capital, but didn't publish any of them due to a shortage of my time for writing. In the meantime, our old article was still mis-representing our views, and this was rightly criticised on the EA Forum.

We've now added a note to the top of the career capital page, but we should have done this a year ago. We'll continue to prioritise writing about this topic next year.

This mistake also points towards a broader issue. We have hundreds of pages of old content, but only 2 full-time staff working on content, so it's easy for our old advice to get out of sync with our views. For instance, the career quiz often doesn't return useful results, so we've added a disclaimer to it and removed most links to it.

Going forward, we're considering creating an automatic warning that will appear on old articles saying they might not represent our views. The warning will have to be actively overridden when we think the old article is still accurate. There are also some other areas where our core advice is not presented clearly enough on the site, and we intend to fix that with the key ideas series.

Not working on a summary of the key ideas series earlier

One of our main aims over 2018 was to make the site better appeal to our tier 1 audience. In particular, we think the career guide is not at the right level for this audience (as explained elsewhere). To fix this, we've been writing a 'key ideas' series to replace the career guide.

However, in the autumn we realised we could start by writing a summary instead, and that this, combined with existing articles, would be good enough to replace the career guide right away.

We could have probably realised this earlier in the year if we had spent more time thinking about how to minimise new writing, which we have very little capacity to do right now. This could have meant fixing the main problem with our site a year earlier, perhaps attracting hundreds of additional tier 1 applicants.

We intend to prioritise this project at the start of 2019.

We're planning to relocate from the San Francisco Bay Area to London

We recently announced this intention and gave some brief reasoning here.

Although we've received many benefits from being based in the Bay Area, our intention to move suggests that we may have made a mistake to move from the UK in 2016. Hiring over the last two years hasn't been obviously better (contrary to our expectations) and we should have probably put more weight on the long-term preferences of the senior staff in deciding where to be based.

Another upshot is that finding an office in SF was one of my major projects over the year, and if we'd realised sooner that we wanted to move, I could have saved a significant amount of time and attention, which might have been enough to publish several more articles in the key ideas series.

In particular, one major element in our decision to move to London was concluding that having offices in both cities would be a bad idea. If I had looked into this question earlier in the year, I could have cut the SF office search short earlier.

Smaller mistakes and issues

(Not in order)

- We're unsure what level of ability is needed in computer science to contribute to be an AI safety researcher. In early 2018 we spoke to some key people in AI safety who think that people should be able to get into top ~3 world-ranked machine learning PhD programs in order to pursue this option. Our AI safety career review now suggests that people should be able to get into a top 10 program, but we're still unsure of the ideal recommendation for us to make. This means there were some people we spoke to in 2017 who are following old recommendations, but we're insufficiently sure of our new recommendations to suggest they switch. Fortunately, ML graduate study (what most of these people are doing) is good preparation for AI policy, ML engineering and earning to give, so they still have good back-up options. We could have avoided this if we'd done more to understand the right profile of person for this path ahead of our 2017 advising. We intend to prioritise this kind of information within the research we do over 2019.
- We contributed to confusion in the community by using the term 'talent gaps'. See our article explaining the misunderstanding and our proposed solution (talk about specific skill bottlenecks instead).
- Whenever we wanted to write about AI policy publicly, we ran into difficult and controversial issues about how it should be framed in the community. Although these discussions resulted in progress on framing, it meant we had less output in this area than expected. This is an update against trying to invest heavily in a priority path before it has been clearly framed.
- The AI policy talent survey was only moderately useful and also took longer due to the issue above, so in hindsight, it might have been better to spend the time learning more about the path by interviewing experts than conducting the survey.
- Typeform, a third party service that hosts some of our web forms, suffered a data breach which affected some of our users.
- After more discussion, we think we should have been more sceptical about the results in our talent survey, especially concerning the value of recent hires and discount rates. We don't think this changes our core advice, but we should rely less on the figures in justifying our positions, and make it clearer that others shouldn't rely on them either. We've updated the survey results blog post to be clearer about their weaknesses. For instance, we gave too little thought to the scenario in which hiring and managing new recruits absorbs a lot of senior staff time, which significantly offsets its benefits.

- We think we let our standard for what qualifies as a rated-10 plan change drift up this year, which led to an artificial reduction in our number of plan changes reported from this group in 2018. After we realised this error we re-scored several plan changes, but expect that the number reported in this review is too low by around 20%. We're working on writing up a detailed guide on how to score plan changes with lots of examples, which should increase our consistency in future.
- On two occasions we failed to respond to applicants for advising when we said we would. We'll aim to be more conservative with these commitments in future.

Issues and bottlenecks faced in 2018

In this section, we list some other issues we especially struggled with in 2018. These do not necessarily represent mistakes or key weaknesses.

Lack of growth in IASPC points

As discussed earlier, this year we missed our IASPC growth target by about 580 points (original weightings), and recorded fewer IASPC points than 2017. We explained some potential reasons for this in the review, which we don't repeat here.

We also think there's a significant chance that the IASPC total does not grow next year, as explained in the later section on plans.

Only moderate progress on team diversity

Last year, we highlighted team diversity as one of our most significant issues.

We succeeded in our core goal of hiring one person from an underrepresented background. We also greatly increased the diversity of our hiring pipeline, and implemented a generous maternity leave policy.

However, there is still a long way to go until we reach parity with the rest of the effective altruism community. As of the end of 2018, among our full-time staff we have 8 white men and 1 white woman (among our freelancers, we have 6 women and 6 men, some of whom are people of colour).

For these reasons, we intend to continue to prioritise improving diversity as laid out last year. (Update in April: we've already had 3 candidates from under-represented backgrounds accept offers.)

Limited senior management time

Our team leads often face a tough tradeoff between high-level work to make our programmes more effective (e.g. redesigning the website) and hiring new staff who might increase our future capacity, because both of these can only be done by senior staff. What's more, if we focus on hiring, we won't see much increase in IASPC over the next year or two. This is similar to the situation described by GiveWell here, and is perhaps the key dynamic limiting our growth.

This is especially pronounced in my role, since I'm perhaps in the best position to write the key ideas series and work on strategy, but also have plenty of management responsibilities as CEO.

One way to see the tradeoff that I face is that 80,000 Hours needs to do two very different things:

- Improve our ideas: continue to improve our core advice and how it's framed.
- Company-building: scale up delivery of our more proven programmes, such as one-on-one advising.

These both require full attention and involve different styles of work. This suggests I should work towards handing off one of these to another senior staff member.

We made progress handing off some of the company-building responsibilities to the team leads this year, and I plan to continue in 2019.

Hiring a full-time executive assistant or other operations staff might be another option, which could help the other team leads as well.

(Update in April: We've appointed Brenton head of internal systems, and he'll aim to take on more responsibility for operations, impact evaluation, and other projects to give me more capacity.)

We're not focused enough on reaching and helping potential tier 1 applicants

Focusing on this audience was one of our major aims for 2018, but we still think there's much more we could do.

The problem is that often the ideal person we'd like to advise could do something like contribute to research, management or policy-setting within an area like AI safety or global priorities research as soon as possible. For this profile of person, much of our online advice seems:

- Overly aimed at university students rather than people age 25-35.
- Too simple and not rigorous enough.
- Not focused enough on explaining our key ideas about maximising impact. For instance, the career guide opens with articles on job satisfaction and donating 10%, but we should probably open with an introduction to problem selection, extinction risk or longtermism.

These problems have been mentioned in user interviews with the types of people we'd like to appeal to.

This means we fail to attract some of the readers who are the best fit for our priority paths, and in the worst case, might put them off the broader ideas of effective altruism. Attracting an overly broad audience also results in the demoralisation problem listed later.

We think we made significant progress on this problem this year by developing:

- The podcast, which we think is at the right level for this audience.
- The headhunting and job board programmes, which are useful for this audience.
- The specialist advising in AI policy, which can make it easier to add value to a more experienced audience.
- New articles and pages such as the one outlining our list of priority paths and updated problem profiles page.

However, we would have liked to make more progress on the website. The majority of the traffic ends up in the career guide section of the website, which still has these problems. We list this as a mistake.

Ultimately, it remains to be seen how much we'll be able to appeal to and add value to this narrower audience.

It's hard to know whether we have a reproducible process for creating rated-100+ plan changes

We find it very challenging to measure our impact. It takes significant time from the team, and we became aware of even more problems with our process this year.

This makes it hard to know whether our programmes work. It's especially hard for rated-100 plan changes because there are so few of them.

This situation will naturally improve as we grow, however, we also intend to keep investing in improving our impact evaluation process. Unfortunately this comes at the cost of delivering and improving our programmes.

Demoralising some readers

By focusing on our priority paths, we've made the advice less suitable for a broader audience. Several people and this [EA Forum] post have raised concerns that by mentioning high-impact but extremely competitive options, we might be putting off and demoralising the broader audience. It's difficult to tell how concerned to be about this, since impact survey respondents overwhelmingly reported finding our content motivating (despite our explicitly asking for negative feedback in addition to positive). We discuss how broad our audience should be later.

Lack of advising capacity

We had less advising capacity than we expected this year, which contributed to missing our IASPC growth targets, due to:

- We lost one member of the in-person team unexpectedly for six months.
- Our AI policy specialist decided to focus more on research and content.
- We spent more time on hiring than planned. This led to hiring Michelle Hutchinson, but she only started at the end of the year.

Costs

Time allocation in 2018

Among the core team, we roughly allocated time across our core functions as follows, which is roughly 35% online, 35% in-person and 30% support:

	Full-time equivalent	**Fraction of total**
Research	1.8	23%
Web	0.93	12%

In-person	2.8	35%
Internal systems	0.62	8%
Operations (via CEA)	1.1	13%
Senior management	0.75	10%
Total	7.9	

2018 financial costs

Spending was broadly in line with our projections.

In December 2017 we produced a ‘baseline budget’ that had us spending $1,083,000 in 2018 if we didn’t launch any new projects or hire any new staff.

In addition, we raised enough for an expansion budget of $1,425,000 -- the increase was to hire 2.5 extra full-time equivalents; raise salaries and spend an additional $25,000 on marketing.

As of November, we expect to spend $1,360,000 over 2018, close to our expansion budget.

We increased salaries and hired two extra full-time staff, in line with our expansion plan. However the staff started in September, and so didn’t increase costs as much as budgeted. We also spent less on marketing since it didn’t seem like a top priority, and saved on legal fees by not applying for US Green Cards. On the other hand, Peter McIntyre returned to the team in the summer, and we made a larger contribution to CEA operations than budgeted. These various corrections roughly cancelled out. Most other line items came in similar to what we projected.

	2018 baseline spending ($) projected in Dec 2017	**2018 spending ($) projected in Dec 2018**
Staff and contractors salaries, payroll tax and benefits	695,856	957,159
Non-salary staff expenses (e.g. travel, conferences and food)	71,860	95,176

Office rent, supplies and utilities	132,900	119,090
Contribution to CEA's expenses for operations	50,000	90,000
Legal and immigration	37,200	4,530
Marketing	27,000	13,746
Computer Software & Hardware	13,200	15,797
Internet/Web/Hosting Fees	50,520	46,117
Non-employee insurance	600	113
Books/subscriptions/reference	1,440	5,346
Other	0	6,943
Uncategorized Expenditure	2,400	2,017
Total	1,082,976	1,356,034

Donations in kind

We also receive discounted services from companies (listed here), though for the most part, the value of these only accounts for a few percent of the budget.

One exception is that Google gives us $40,000 of free AdWords per month as part of their Grantspro program (which can be spent subject to various restrictions). If valued at their nominal rate, this would be 35% of our budget, though we would not spend this much if they weren't free.

Bellroy also donate hosting services for our website, which would otherwise cost us approximately $10,000 per year.

We also receive free advice from many in the effective altruism community (many of whom are listed on our acknowledgements page), and in the past we've used the time of student group leaders and other volunteers. We haven't tried to quantify these costs, and ignore donations in kind in the rest of the estimates.

Summary of historical costs

	Spending ($)	Income ($)	Number of full-time staff (excl. central CEA)	Total staff FTE (inc. interns & volunteers & share of central & freelancers)
2011	0	0	0	1.7
2012	57,721	70,285	1.5	3.4
2013	163,108	237,074	2.2	7.5
2014	229,133	325,686	2.1	4.9
2015	283,297	355,689	3.2	5
2016	406,074	388,020	4.1	5.5
2017	752,888	2,060,938	6.3	7.8
2018 (projected)	1,356,000	1,370,000 (to date)	6.8	10.1
Total	3,248,221		27.1	45.8

Opportunity costs of staff time

We think our largest costs are not financial, but are rather the opportunity costs of staff time i.e. our staff could have worked on something else high-impact if they didn't work at 80,000 Hours.

We mainly leave it to individual stakeholders to compare these costs to our impact, because estimating them mainly comes down to a judgement call based in large part on information about the effectiveness of other organisations.

In the past, we've also provided a very rough estimate of the financial value of these opportunity costs assuming that some of our staff would have worked at other organisations in the community otherwise, and using the value of staff estimates from our talent survey. However, we've found out about conceptual issues with using that data as a measure of our staff's opportunity costs and no longer believe those estimates are reliable enough to report.

Update on historical cost-effectiveness

See the review.

Marginal cost-effectiveness

People often assume that our marginal cost-effectiveness is going down -- that we'll have less impact per dollar in 2018 compared to 2017 -- but we don't think that's obvious. We're probably not close to saturating our market, we see benefit from economies of scale, and each year we learn how to be more efficient. Overall, we think it's hard to know whether marginal returns will go up or down, and our rough estimate is that marginal returns are similar to recent history. For more detail, see the discussion from last year.

One update from this year is that the all-time ratio of financial costs per IASPC rose by 27%, when it has fallen the last four years, which could suggest we've hit diminishing returns.

However, we don't think this increase is surprising. In 2018, we nearly doubled financial costs, but this investment went into programmes aimed at generating rated-100 plan changes, which takes several years to pay off. This will increase the ratio of costs per plan change until the returns arrive. Our plan change metrics are also volatile -- if we had recorded an extra rated-1000 plan change, the ratio would have been flat.

Future advising cost-effectiveness

Another way to estimate marginal cost-effectiveness is to analyse specific programmes we could pursue at the margin. We would expand several of our programmes with additional funding and don't have a strong view on which is best at the margin but it's easiest to estimate the impact of additional advising capacity.

In the following, we roughly estimate the impact of an additional advisor supplied with tier 1 advising leads at the current margin compared to costs.

We start with costs:

- A full-time advisor uses one year of labour per year.
- We also need to include the costs of finding the tier 1 applicants for them to advise. At the current margin, we have a surplus of leads, so this cost is small. In the long-term, however, we'll need to continue to grow our audience to continue expanding the number

of advisors. This requires adding additional staff to the web and research teams (or staff working on other ways of finding leads).

- We roughly estimate that 1 person on the web and research teams can produce an extra 300 tier 1 applicants per year of work. This is partly based on our current ratio of web/research staff to in-person advisors, which is roughly 1:1. It also lines up with our longer-term estimates of what audience growth is possible outlined in the 3-year plan later. Alternatively, we think a full-time person focused on actively reaching out to potential leads could find enough to fill up one advisor.
- Each person we hire requires about 10% of an operations staff member to run the organisation, so to cover 2 people, we need an extra 20% of someone working on operations, bringing the total up to 2.2 years.
- Each person we hire also requires about 10% of manager's time, so we would need another 0.22 years of management.
- So, in total, the costs of a year of advising time are 2.20 years of labour and 0.2 years of management time.
- The financial costs of salaries, expenses and overhead would be roughly $350,000 per year.

What would the impact of this be? Over a year of work, here is how we might estimate the impact of an additional advisor supplied with tier 1 applicants.

- A full-time advisor has about 900 hours of 'delivery' time free per year, which is enough to advise 300 tier 1 applications with our current approach.
- Advising over 2014-2016 has produced 1.9 IASPC per hour (updated weightings), which would imply 900 hours produces 1710 IASPC in the long-term. This might be something like 8 rated-100 plan changes and 50 rated-10 plan changes.
- These returns would make marginal advisors highly effective. To illustrate, if a rated-100 plan change is worth over $1m, then the combined returns would be over $8m. Given that we've only tracked 20 rated-100 plan changes to date, only about half of these were influenced by advising, and our current advising programme differs in some ways from previous iterations, we can't confidently extrapolate forward such high returns, and should take regression to the mean into account.
- On the other hand, if the advising finds an extra rated-1000 plan change, then the returns would be much higher.
- We think our 'tier 1' applicants are significantly more talented and up-to-speed than the people we worked with in 2014-2016. Though, it may have been that advising in the early days of effective altruism was much more effective than it is today.
- Overall, our estimate is highly uncertain, ranging from perhaps 300 to 3000 IASPC, with a central estimate of about 1000. However, our best guess is that it's cost-effective, and it's hard to tell whether returns are diminishing or increasing.

One of the greatest uncertainties facing the effectiveness of marginal advising is whether we will find enough tier 1 applicants to fill up their capacity. We're confident we can do this if we hire 1-2 additional advisors, but it's hard to know what would happen if we hired over 4. We don't take this into account above – the estimate is for an advisor filled with tier 1 applicants.

Plans for 2019

Key focus for 2019

Broadly, we expect to keep our focus unchanged on building a machine to produce rated-100+ plan changes, and scaling it up. If we do this, we'll be able to make a major contribution to some of the key skill bottlenecks facing our top priority problems.

To get these plan changes, we'll continue to attract and introduce people to our priority paths with online content, and help them enter through in-person advice.

Compared to last year, we think it's more likely we have a reproducible process, so intend to focus more on making it more efficient and hiring to scale it up – working towards the 3-year plan sketched later – while also addressing our most pressing weaknesses

How we set strategy and key strategic decisions

In 2018, we clarified the four key strategic frameworks we use. The two most important approach planning from opposing perspectives:

1. **Long-term growth approach** - what's our long-term vision and what step best takes us towards that vision? This involves aiming to follow our optimal growth trajectory, and resolve key strategic issues in the right order.
2. **Marginal growth approach** - what are the key bottlenecks preventing growth next year and how can we resolve them? This involves trying to maximise our growth over the next year.

We also consider:

3. **Capacity building** - how can we best increase the productivity of the team? We expect team capacity to be one of our biggest bottlenecks over the long-term, so it deserves extra attention.
4. **Community** - what are the most pressing needs in the effective altruism community, and which are we best placed to help with?

Here are some of the key strategic questions we need to settle with these frameworks:

Decision	Position	Quick justification
Should we change our vision & mission?	Unchanged. Focus on filling the most pressing skill bottlenecks in the most pressing problems.	The same reasons for focusing on rated-100 plan changes given last year.
What should our key impact metric be?	Stick with IASPC, but update how we make the impact ratings in several ways.	We think IASPC gives a reasonable proxy for our historical impact, though there are some problems, such as the ratios being out of line.
Should we focus on getting plan changes rated 0.1/1/10/100/1000?	Rated-100.	Same as last year. As we get larger, it might make sense to move our focus to rated-1000+ plan changes.
Which target market?	Same as last year, but tilt even more towards people age 25-35 who have experience relevant to a priority path, rather than recent graduates.	It's hard for younger people to make rated-100 plan changes and the community is also short of these skills.
Which problem areas and career paths should we focus on?	Our views are here and here.	Justification for why they're high-impact is sketched in the articles. With the priority paths, we also try to assess the paths we're best able to help with given our audience.
Do we have a reproducible process for getting plan changes?	Yes when it comes to getting rated-1/10 plan changes, and probably when it comes to getting rated-100 plan changes as well. But not when it comes to getting rated-1000 changes.	We only have 20 examples of rated-100 plan changes, so it's hard to be confident. We also have a lot of uncertainties about how the programmes should work.
How quickly should we hire?	Medium. Last year, we said we'd focus on essential hires and people who can act	Some of our work has become more systematised, the hiring pool seems better,

	autonomously immediately. We would still hire anyone like this, but we now think it's more likely we should hire 2-7 people into more clearly defined roles.	and it seems more likely that hiring is the key bottleneck preventing growth (despite its large costs). See more explanation later.
What are the most powerful drivers of top plan changes?	We usually find that all the main drivers (online content, coaching, community connections) are important, and their impact is difficult to disentangle, because they serve different purposes. Though, some kind of in-person interaction is basically necessary for many of our users.	This comes from doing case studies of our top plan changes. That said, there may still be ways to generate more 'online only' top plan changes, and there's some evidence the job board could achieve this.
How much to focus on **research** vs. **providing advice**? (i.e. improving the accuracy of our advice vs. telling people our findings)	We're currently spending under 5% of time on research; but would like to increase towards 10%.	10% seems like a reasonable long-term ratio, and there are lots of concrete topics that seem useful.
How much to focus on **online content** vs. **in-person advice**, such as coaching?	Roughly 1:1, allocate staff mainly based on fit.	Both seem important for plan changes, and have different, complementary benefits.
How much to focus on **improving our programmes** vs. **outreach**?	Mostly focus on improving programmes. (Though if we increase headhunting, that involves more outreach.)	We already reach lots of people from our online content, the EA community, and word-of-mouth. Improving programmes also increases each of these.

How broad should our target audience be?

There was recently a thread on the Effective Altruism Forum about whether 80,000 Hours is focusing too narrowly on people who can enter its priority paths, and not enough on the broader audience of graduates interested in effective altruist career advice. 80,000 Hours is one of the most common ways people are introduced to effective altruism, but our narrower focus can be off-putting and demoralising to some people who do not believe they are able to enter these paths. A large fraction of our online content, especially what appears in the career guide, is relevant to the broader audience, but even mentioning the priority paths sometimes causes this problem (though we don't have much data to quantify how serious the problem is).

We feel fairly confident that focusing on the narrower audience will have more impact for the coming years (though we may return to a broader audience in the future), so we think this is what we should do. We also don't think it's possible to focus on both the narrower and broader visions right now, since we don't have enough capacity, and it would be hard to separate the audiences.

However, we deeply regret any negative impact we might have on our readers, so we will look for ways to improve. For instance, in the site redesign, we'll try to clarify who our target audience is to reduce the chance that people get given the wrong recommendations, and consider doing more to be up-front about the competitiveness of our top suggestions. We'll also consider releasing more content on less competitive options.

We also think much of the community is mistaken about just how competitive some of our priority paths are. For example, we think government jobs are some of the highest-impact options right now, and that some of these, while still quite competitive, are achievable for a significant fraction of graduates.

If it's important for the effective altruism community to also provide broader advice, then we think it should be addressed by other groups. Here are some projects in this area we don't intend to do in the near-term but others could:

- One-on-one advising, headhunting and specialist content within global health and factory farming.
- Greater support for people focused on donating and earning to give (though we're very concerned that effective altruism doesn't become even more associated with donating, so such a project would have to be done carefully).
- Material on how to contribute to useful political advocacy, which could be a way that many people could make high-impact part-time contributions.

Key priorities and projects for 2019

We currently intend to continue with the core categories of projects from 2018. The following are examples of projects we might pursue. What we prioritise might change during our coming strategy review and over the year.

Continue to improve the site to better appeal to and help ‘tier 1’ applicants, by replacing the career guide and continuing with the podcast and job board

We think we can substantially improve the site’s ability to attract the readers most able to switch into our priority paths, and be more useful to them.

Doing this makes the in-person advice more effective in the short-term -- since the people coming into it are a better fit and more up-to-speed -- and it also builds our long-term audience, which needs to grow if we want to maintain our growth rate beyond a few more years.

To this end, we’ll do the following, roughly in order of priority.

- **Continue with the podcast.** We think we can continue to get at least 300 subscribers per new episode, so if we release another 30 episodes, we’ll hit 23,000 subscribers next year (though we might release more or fewer episodes depending on how we prioritise the podcast compared to articles). We’ll likely focus more on episodes that deliver useful advice rather than attract new listeners, such as the recent episode on ML engineering.

- **Replace the career guide with the key ideas series.** The current site doesn’t do a good job of explaining our key views and attracting tier 1 applicants. We intend to replace the career guide with a new ‘key ideas’ series as the main introduction to our advice. We’ll start by producing a new summary of our advice, and updating the front page and site structure. Eventually, we’d like to finish part drafted articles on our values, how we choose problems, career capital, and exploration. This could be turned into a book, however, this project is bottlenecked by my time as the primary author, so may not be finished next year.

- **Continue with and perhaps double down on the job board.** We were happy with the performance of the job board this year, which seems to have already resulted in several plan changes, and would like to keep it running. We currently have a job board manager who works about 50% time sourcing, filtering and adding jobs. We could bring them up to full-time and roughly double the jobs listed each month from 100 to 200. We’d also like to add automatic, customised email alerts, which are by far the most requested feature. However, we’re not yet sure how to prioritise this against the projects above.

- **Continue to add articles on why and how to enter the priority paths.** We think that developing a package of content on policy careers is a top priority, since we think it’s one of the highest-impact paths many more of our readers could enter, but they often don’t due to lack of knowledge about how to get started. We’d like to create a landing page that makes the case for policy careers and gives an overview of how to enter. Then, we’d like to add or improve our career reviews on key sub-options, such as working in a campaign or joining the executive branch. We think this content could help to shift our audience and the community to focus much more on policy careers. There are many other articles we could write about our priority paths, for instance, our AI safety research profile is significantly out of date but is one of the most important in generating plan changes. However, we won't be able to make much progress unless we can hire an additional writer / researcher.

In the above, Rob would focus on the podcast, Peter Hartree would focus on the site redesign, and I would focus on the key ideas series.

Continue to improve and scale up the in-person advice

Some kind of personal contact seems to have been necessary for almost all of our past plan changes, so it seems necessary if we want to grow quickly. As noted, advising is also heavily oversubscribed. This means that increasing advising capacity by hiring is likely a key priority, which is covered in the next section.

Besides hiring, next year we intend to:

- **Advise over 100-400 people** (compared to 211 in 2018). We're unsure how much we'll prioritise delivering advising against hiring and other ways of improving it. We're also unsure when new hires will join and increase capacity.

- **Continue with the AI policy specialist experiment.** This year we'd like to focus on delivering the advice while learning as we go, rather than producing content and research, so we can better measure its effectiveness compared to regular advising.

- **Continue to develop the headhunting programme.** We continue to be pleased with the results of headhunting over 2018, and would like to continue to develop the programme to evaluate its effectiveness compared to advising.

In the above, Michelle would focus on delivering and improving the core advising, Niel would focus on the AI policy specialist experiment, and Peter McIntyre would focus on the headhunting.

Increase how much we prioritise hiring

Last year, we said we'd only focus on hiring for essential roles or people who can be highly autonomous right away. We'd still like to hire anyone like this, but we think it's more likely we should hire for a wider range of roles at a faster rate.

Hiring seems more attractive, because:

- We seem to have a greater volume of useful 'systematised' work, where it's clearer what's needed and performance can be more easily measured, which makes it easier to hire for.
- We have more management capacity. Previously, I was managing almost everyone on the team. In April, we made five people team leads, and now each of them has some experience in hiring and management. If each team lead could manage 3-4 people, then we could absorb about 10 extra people before needing to add another manager.
- The hiring pool might be better. Several hiring processes in the community seem to have gone better than expected, and our headhunting has turned up more promising people than expected. In particular, there seems to be lots of people who are talented but have 0-2 years' experience, who could be a fit for some of the roles that have opened up.

For all these reasons, it seems more likely that hiring could be one of the best ways to grow our impact over the next year (rather than developing better programmes or improving efficiency), and it also helps to build the capacity of the team in the long-term.

This said, we still think hiring has a lot of hidden costs and risks, so we will remain wary of hiring too quickly, and maintain a high bar. We're also very uncertain about the ideal rate of hiring.

You can see a list of roles we'd like to fill and why we think we can fill them in the review.

Improve systems and operations

Some of the key projects over 2019 will include:

- **Move to London.** We've decided to move our office to London (announcement and explanation here). This will require finding an office in London, getting visas for several staff, switching our budget to the UK and coordinating the move. This will likely cost several weeks from each staff member, as well as several months from me and our operations staff.

- **Update our plan change rating system.** There are many problems with the existing system we'd like to fix. We'd especially like to improve how we write up case studies of plan changes, how we estimate the extent of our counterfactual impact on plan changes, and the rubric we use to rate them 1/10/100/1000 to make them more consistent.

- **Continue to improve core financial and legal systems.** We made good progress on these last year, and passed our audits in the US and UK without significant problems. The next stage is to expand capacity, and proactively set up the systems that will let CEA and 80,000 Hours scale up their team over the coming years.

2019 metrics

IASPC growth might be flat next year

The majority of the plan changes we'll record in 2019 will be the result of efforts made in previous years, so predicting how many we'll get is mainly a matter of predicting how much these efforts will pay off. These predictions are also difficult, because a small number of large plan changes play a big role.

Our expected levels for 2019 are approximately as follows (updated weightings, rounded, brackets show 80% CI):

- Passive: 950 (400 - 1600)
- Growth: 1800 (1000 - 4000)

This compares to about 2000 IASPC recorded in 2018.

This is disappointing because it means there's a good chance we don't grow the IASPC total next year. (Though there's also a good chance it grows a lot.)

The main reason we are not more confident in growth is that we're still bottlenecked by advising capacity. We expect to be able to advise 100-400 people next year compared to 217 in 2018, but that's not large enough to cause a major increase in the bottom line. This is in large part because it takes several years to record all of the high value plan changes from advising, so even if we expand the team in 2019, the payoff will be delayed until 2020 and beyond.

We have other pathways to growth besides increasing advising capacity (sketched below), but it's hard to predict them with confidence.

However, as we sketch in the next section, if we are able to increase advising capacity, then we think it's likely we can grow significantly over a 3+ year period.

What are these IASPC projections based on?

We expect our passive level of plan changes (if we heavily cut back our activities) will be something like:

- 5 rated-100s (80% CI: 1-8)
- 20 rated-10s (80% CI: 10-40)
- 400 rated-1s

Using the updated weightings, this would give us 940 IASPC points with the updated weightings (400 - 1800).

The passive level could be higher than the 1100 passive level we projected last year if plan changes keep upgrading to rated-100 at a higher rate.

Otherwise, the passive level hasn't grown more compared to last year because we didn't significantly grow the pool of rated-10 plan changes that can upgrade (130 vs. 100 at the end of 2017), while an unusually large number of plan changes upgraded to rated-100 in 2018 and we're unsure whether this will continue.

We've also reduced how many rated-1 plan changes we expect to record, though this is mainly due to making less effort to track them. We may stop tracking rated-1 plan changes entirely next year, in which case we'll drop them from the targets.

On top of the passive level, we think it's likely we can advise 100-400 people next year. This assumes no-one leaves the team.

We expect that advising 200 people will lead to about 300 IASPC over the year, raising the total to 1240.

How might we grow more than this next year? Here are some options:

- We record a rated-1000 plan change, which we estimate has about a 50% probability.
- If we can hire an extra advisor early in the year, we'll be able to advise an extra ~200 people over the rest of the year, which would get us another 300 IASPC points (and much more the year after).
- If the podcast and other content has brought our audience more up to speed in our priority paths, the conversion rate of the advising might be significantly higher, or we might start to get more online-only plan changes.
- If we can continue the headhunting, it should start to pay off, which could lead to several placements in top roles.
- If we can continue the job board, that might produce several extra large plan changes.

If several of these happen, that might grow the total by another ~800 points, with a range of something like 100 - 3000. This is an especially rough estimate.

2019 leading metrics

We roughly expect the following measurable outputs (assuming no-one leaves the team):

- Web traffic will remain roughly constant at 1.6 million unique users, unless we redirect attention from the podcast to writing articles and promotion.
- We'll continue to gain about 300 subscribers per new podcast episode on average, unless we focus on more shareable episodes.
- We'll able to release about 50 new pieces of content again this year, assuming a 3:2 ratio of podcasts to long articles. If we can hire a researcher, we'll be able to exceed this amount in 2020 but not 2019.
- We'll be able to advise over 200 people. If we hire early in the year, this might be higher.
- The job board will continue sending at least 5000 clicks through to high impact job listings each month.

The figures above assume a moderate amount of effort invested in hiring (e.g. about 6 months spread across the team leads). However, if we invest heavily in hiring, then we will release 25-50% less content and do a similarly reduced amount of advising.

What might we optimistically be able to achieve in 3 years?

In the following, we sketch out how it might look if we focus on scaling up our existing approach. We assume we don't hit major diminishing returns to our audience (which we think is plausible at this scale, but might not work out) and that we face no major setbacks. Overall, this trajectory is optimistic but achievable if things go well.

In brief, we expand the team from about 10 to 25 people: 10 on the web & research teams; 10 on the in-person team; and 5 in operations and management. The web and research team aims to expand the audience, developing more tier 1 applicants. The in-person team helps them become plan changes.

Starting with the research team, we would grow it by three, reaching 1 manager, 1 podcast producer, 3 researcher / writers, 1 marketer. We could also expand the web team by three, reaching 1 manager, 2 engineers, 1 designer, and 1 job board manager.

Collectively, we think with this expansion we could likely continue to grow audience engagement at at our recent rate of ~60% per year. This would result in a 4-fold increase in three years. Our aim would be to focus on developing tier 1 applicants (or better), so our hope is that this would lead to 4-fold growth in how many we identify each year.

It's possible we could grow the audience this much without hiring, or just from growth in the broader community, but it'll be easier with a larger volume of higher-quality content. Having the larger team will also make the research on which our advice is based more thorough, guiding people towards better paths.

Our advising is already oversubscribed with promising applicants, despite little promotion this year. If we ramped up promotion, we expect we could double the number of tier 1 advising applicants to about 500 (or an equivalent number of even better leads).

If we also grew online engagement by 4-fold over three years, then we'd probably also have about 4-times as many new 'tier 1' advising applications per year, which would be about 2000.

We also think up to half of our advising capacity can be usefully filled with people we've advised before. Moreover, about one third of tier 1 applicants come through other sources, such as referrals. This could likely be further ramped up if we started actively searching for people to work with. So, this means the total number of people worth advising each year would likely be more than 2000, but we'll ignore that in what follows.

To advise 2000 tier 1 applicants per year, we'd need to hire 4 more full-time advisors. The team might then consist of 1 manager, 4 'generalist' advisors, and 3 specialists in AI policy, AI technical and biorisk, who would become key point people for coordinating to solve skill bottlenecks in these areas. (In practice the specialists might do a mixture of research, headhunting and advising, but we're splitting out the roles for simplicity. We will also adjust the balance of generalist advising, specialist advising and headhunting and other potential in-person programmes depending on what's most effective.)

In the section on marginal cost-effectiveness, we estimated that each year of advisor time with tier 1 applicants can produce 1000 IASPC (range 300 - 3000). (Note that over 1 year it would be more like 400 IASPC, which is why the 2019 projections are not higher.)

In addition, if we build a headhunting team of 3, we think they could each produce a similar or higher number of IASPC. And it may be possible to produce more plan changes that are mainly online-driven, in addition to the existing passive level.

So, in this scenario at equilibrium, in total we'd track 3000 - 30,000 IASPC per year, with a target of around 10,000.

The optimistic trajectory could look like this (bearing in mind the huge confidence intervals we'd need to place on these figures):

Year	IASPC points (updated weightings)	Number of advisors & headhunters
2018	2000	1.5
2019	2500	5
2020	4000	7.5
2021	6000	10
2022	8000	10
2023	10,000	10

10,000 IASPC would be something like 40 rated-100 plan changes per year and another 4 rated-1000. This would be much more than today, but it would only require a growth rate of about 50% per year over four years.

To support all of this, we'd need about 5 more operations and management staff, bringing the total team to 27, an increase of 5.6 per year for three years. We have five main teams, so this is about 1 hire per team per year, which seems like a manageable rate of hiring.

This would mean the total team would grow by about 2.5 times, so our financial costs would grow about the same. This would mean the cost per IASPC ratio is about the same as today.

It's possible we can achieve this growth without hiring as many people by working out how to make our programmes more efficient -- and we will try to do that -- but it's less predictable.

If we succeed in this growth, then we'll be filling tens of the most important roles in our priority problems each year. The fields we're focused on, such as long-term AI safety, effective altruism organisations and global priorities research have under 100 people focused on them full-time, so if we did this then we'd be responsible significant fraction of the people working on the issues, and therefore play a major role in solving them.

In five years, this could mean outcomes like:

- 100 people have embarked on promising careers in emerging technology policy.
- Our headhunters are one of the key sources of hires for the DeepMind and OpenAI safety and policy teams.
- A new $1bn+ foundation has been established following the ideas of effective altruism, and we find half of the initial operations and research staff for it.

- Global priorities research has expanded to academic centres at Harvard, Cambridge and Stanford, and half of the new junior researchers say they became interested through 80,000 Hours.

This trajectory would also have spillover benefits for the effective altruism community in general, such as much more thoroughly worked out research and content, and more concrete successes.

At this scale, we'd be advising the equivalent of 0.3% of graduates of top 30 global universities per year. This is a lot, but it seems likely our audience could eventually be larger than that. Alternatively, it may be possible to further increase the impact of our existing audience by coordinating them better, such as by creating a fellowship or incubator for new projects.

Budgets and fundraising targets

We start by outlining a 'baseline' budget, in which we continue operating with our current scale and policies. We're confident this will be cost-effective.

We then outline an expansion target based on following the growth trajectory we just sketched for two years. We think the expected value of this plan is high, but there's a greater risk it doesn't work out.

Projected baseline budgets

In the baseline budget, we assume that we fulfill all our existing commitments, including retaining all our 12.6 full-time-equivalent existing staff with salary increases in-line with our current policies (about 10% per year on average), and increase our marketing and online services with our expected growth in audience. During 2019, we've also already made offers that will increase our full-time equivalent staff by about 3.7, and we've included these in the baseline budget.

Staff costs are significantly higher because we averaged only 9.3 FTE over 2018, and we've also increased average salaries and health benefits by about 20%.

The other major increase compared to 2018 are the costs for an office. We had intended to set up our long-term office during 2018 but it took longer to find than we hoped. However, we started renting in March 2019 in London. Office-related costs will be higher in 2019 while we set up the space.

	2018 ($)	2019 ($)	2020 ($)
Staff and contractors salaries, payroll and benefits	957,159	1,522,885	1,759,613
Non-salary staff expenses (e.g. travel, conferences and food)	95,176	165,355	194,632
Office rent, supplies, utilities and services	119,090	387,251	584,974
Office and country moving expenses	0	303,869	0
Contribution to CEA's expenses for operations	90,000	210,400	253,200
Legal and immigration	4,530	55,758	32,448
Marketing	13,746	23,650	21,310
Computer Software & Hardware	15,797	132,013	35,674
Internet/Web/Hosting Fees	46,117	73,149	75,795
Books/subscriptions/reference	5,346	9,737	10,672
Other	6,943	2,629	3,964
Uncategorized Expenditure	2,017	2,472	2,584
Total	1,355,921	2,889,167	2,974,866

We have excluded a 2019 £217,500 rental deposit which we expect to recoup. The 2018 figures are an estimate made in December.

After 2020, we budget for 10% increases per year to the baseline.

Why raise multi-year commitments?

We would like to raise enough to let us commit to two years of the 3-year growth trajectory sketched earlier.

Growing 80,000 Hours from a team of about 10 to a team of 25 is a multiyear project. It's valuable to have multi-year commitments because in order to be on this expansion trajectory, we will need a committed management team and to make long-term investments, such as in our office and internal systems. For instance, it's easier to hire a great head of operations if we can say that we've raised commitments to let us expand to 25 people.

It will then take us several years to build the team and see the expected increase in IASPC points.

For these reasons, our donors often make multi-year commitments to give us further stability. In particular, the Open Philanthropy Project gave us a 2-year grant in 2016.

This said, we recognise the risks to growth flagged in the review. In particular, we think there are two major factors that could limit our growth:

- We may not be able to hire sufficiently skilled staff as quickly as projected in the plan (or we find a way to grow without hiring). In this case, we would accumulate reserves, which would let us fundraise less in later years. If our reserves' exceed 30 months, then we could cut the commitments in future years.

- Another risk is that the additional advisors are not able to generate enough plan changes to be highly cost-effective. If we run out of tier 1 applicants, we can stop hiring, mitigating this risk. However, because it takes a long time to measure the impact of advising, there's a chance that our estimates of effectiveness turn out to be wrong.

Due to these risks, we think that focusing on raising two years of funding makes most sense right now. However, we think that contributions to our Dec 2021 target would have benefits, and could be worth considering for a donor keen to support 80,000 Hours.

How much funding is required for the expansion budget?

To be on our expansion trajectory, we need to hire about 5 full-time staff per year.

The starting salary of these staff would be about $80,000 on average. We predict the marginal overhead would be about $40,000 (this includes payroll tax, benefits, and expenses), so the total cost per person would be $120,000. These costs would increase about 10% per year after someone starts.

They would work in the roles described earlier, such as advising, web engineering, design, operations, or headhunting. We might hire a smaller number of senior staff, or a larger number of junior staff or freelancers.

We've already made offers to 3.4 full-time equivalent staff in 2019, and would like to hire another 2.5.[2] We would then like to have the funds to hire 4.5 full-time-equivalents in 2020, bringing the total over the period up to 10.4.

This would have the following impact on costs:

	2019 ($)	2020 ($)	2021 ($)
Budget increase by year	150,000	600,000	957,000
Updated total spending that year	3,041,917	3,574,866	4,229,353

This is about 45% growth in costs per year from 2018, though only 18% per year 2019-2021.

Reserves policy

We aim to hold more than 12 months' reserves for greater financial security -- we have to make multi-year commitments to our staff and office rent. We think there are benefits up to around 24 months' reserves.

---

[2] These are not whole numbers because in some cases we've increased the hours of part-time freelancers rather than hire full-time staff, and because the forward looking numbers are expected figures.

We also aim to fundraise only once a year to allow us to focus on execution the rest of the time. In combination with the above, this means at the start of the year we aim to accumulate 24 months' reserves, which we then spend down to 12 at the end.

Fundraising targets

In order to maintain 12 months of reserves, our Dec 2018 fundraising target needs to take into account our expenses for 2019 and 2020. Specifically, we need to raise enough to cover (1) all our spending over 2019 (2) enough to end the period with 12 months' reserves. If we already had 12 months' reserves, then the second part would just be our 2020 budget. In fact, we only have 7 months' reserves.

Due to our greater maturity, we'd like to at least raise 2-year commitments this year, especially in the expansion scenario.

Therefore, we have calculated our fundraising targets for Dec 2018 and Dec 2019.

Dec 2018 target	
Budget 2019-2020	$6,616,783
Expected cash end Dec 2018	$1,631,640
Projected passive income 2019-2020	$300,000
Additional amount needed (to end 2019 with 12 months)	**$4,685,143**

The Dec 2019 target would then be the amount needed to have 12 months' reserves at the end of 2020, which is just the 2021 expansion budget minus the $150,000 of passive income we expect that year.

Dec 2019 target	
2021 budget	$4,229,353
Projected passive income	$150,000

Additional needed	**$4,079,353**

What would we do with further funding?

We think we could usefully accept more funds than these targets. They would be put towards our Dec 2020 target, which would be approximately $4.7m at baseline (maintaining the team of 20 staff), or about $6m if we continue to hire up to 25 staff on the expansion trajectory.

As further expansion options, we also considered investing more in marketing. It's likely that we spend more than the $24,000 in the baseline budget on marketing, but moderate expansion from this level wouldn't be a material part of the budget. We considered more aggressive expansion in marketing spending than this, but overseeing that spending doesn't seem like a top use of staff time over the next year given the returns from past campaigns (though it is still likely cost-effective).

The other major way we could increase our budget is by running a scholarship, but (1) we currently think this is less effective than advising and headhunting, (2) if we did go ahead, the funds would likely be provided by another group (3) or we could fundraise separately for it.

Current funding situation

We have already raised $4.3m towards the current round, and expect a further $2.7m in Dec 2019, leaving the following gaps:

- Now: approximate gap: $400,000
- Dec 2019: approximate gap: $1,300,000

We think providing funding to expand over the two years, and especially next year, is highly cost-effective in expectation for someone fairly aligned with our approach.

Donors who want to support 80,000 Hours more aggressively could also commit to our Dec 2020 target, which will be over $5m.

One risk of giving us funds to cover expenses beyond the next year is that we won't be able to, or it won't be prudent to, grow our spending as quickly as the expansion trajectory sketched out. For instance, we may not find enough tier 1 applicants or qualified staff.

If this turns out to be the case, then we will accumulate reserves. We propose a policy that if we exceed 30 months' reserves during our next fundraising round, we will not collect all of the funds committed. For instance, if we have $6m of committed funds for Dec 2020, but at that point have $1.5m more than is required to have 30 months' reserves, we will reduce commitments by 25%. (We did exactly this in 2017 with several donors.)

If you might be interested in covering our current funding gap, you can donate directly through the Effective Altruism Funds platform. After making an account, you can assign your donation to 80,000 Hours. If you're already logged in, you can access our page here.

The platform accepts all major credit & debit cards, monthly direct debits and bank transfers. For large donors we can also accept direct bank transfers and donations of assets, including cryptocurrencies.

If you're interested in providing a larger donation and have questions, or would like to make commitments to donate in future years, please contact me: direct.ben at 80000hours.org.

If you think our research and advice is useful, and want to support our vision to expand and help solve the biggest skill bottlenecks facing the most pressing problems, we really appreciate your donations.