

How valuable is movement growth?

Owen Cotton-Barratt

Executive Summary

Movement growth may be very important for young social movements. It's obvious that movement building cannot be always better than direct work, but knowing how to compare them presents a challenge.

In this article I introduce and explore a model of movement growth which tracks individuals' awareness of and inclination towards the movement. I aim to understand when movement building activities are better or worse than direct work, and apply the model to give my views on movement growth in the effective altruist and related communities.

Part 1: Theory

In the first half of this paper I introduce a model for thinking about movement growth, and terminology to refer to critical concepts. We model individuals as having varying levels of *awareness* about the movement, and varying *inclinations* towards it. We assume that these two characteristics can represent the major drivers of interaction with the movement. We explore the consequences this Awareness/Inclination Model (AIM), particularly looking at the long-term counterfactual effects of direct work compared to '*publicity*', which aims at increasing awareness of the movement, and '*advocacy*', which aims at improving inclinations towards the movement. This involves analysing different possible long-term trajectories the movement may be on.

If we accept the model, this has some general implications:

- For early-stage movements, the effects on movement growth are a key consideration in deciding between different activities. For relatively mature movements, direct work is usually better than movement growth.
- It is more important to focus on increasing awareness than improving inclination, if:
 - the movement has a natural maximum size that we cannot change; or
 - essentially everyone will join the movement after they know enough about it; or
 - direct work earlier is much more important than direct work later; or
 - it is very hard to change inclination relative to awareness.
- Otherwise improving inclination may often be better than increasing awareness (this is sensitive to beliefs about some parameters).

- It is particularly key to avoid being controversial and focus on improving inclination rather than increasing awareness, if:
 - the views of people around them have a significant effect on the inclinations of people towards the movement; or
 - the movement might plateau at a wide range of sizes, depending on how well-perceived it is; or
 - building political consensus will be useful for the direct work.

Part 2: Application

In the second half of the paper, I apply the conceptual tools developed in Part 1 to answer questions about how to find the best work for the young effective altruism movement and related areas. The conclusions here are not certain, but represent my informed best judgement. Some of them are driven purely by qualitative considerations, and some are based in part on numerical estimates.

My conclusions are:

- Getting movement growth right is extremely important for effective altruism. Which activities to pursue should perhaps be governed even more by their effects on movement growth than by their direct effects.
- Increasing awareness of the movement is important, but increasing positive inclination is at least comparably important. Therefore we should generally:
 - prefer *advocacy* to *publicity*;
 - strive to take acts which are seen as good by societal standards as well as for the movement;
 - avoid hostility or needless controversy.
- Direct work has a very important role to play in movement building. It is likely to increase positive inclination, by:
 - Demonstrating commitment, and showing that the people engaged in the movement think the work is valuable;
 - Increasing the credibility of the area by demonstrating that there is productive and valuable direct work that can be done.
- Within global poverty, work focusing on movement growth may be more effective than direct work for most people (at the margin today).
- Within areas that seem promising but do not have an established track record, direct work aimed at demonstrating that there are credible interventions may be one of the most effective forms of movement building.

Acknowledgements: Thanks to Tom Ash, Daniel Dewey, Seb Farquhar, Michelle Hutchinson, Josh Jacobson, Toby Ord, and Kerry Vaughan for helpful conversations and comments.

0. Introduction

Suppose you have discovered a great way to help the world. Now you face a trade-off: how should you balance pursuing this directly against spreading the word and convincing others to help pursue it?

The basic problem is similar for many possible ways of helping the world. Perhaps you've worked out that washing your hands before surgery helps your patients' survival rate. Perhaps you've realised that treating men and women equally will help welfare levels and productivity. Perhaps you think that giving money to charities that have been shown to do a great deal with it is an effective way to help people. In each case you might have more impact by *growing the movement* of people who care about the issue than just doing your part of the *direct work*¹, but it may be challenging to see how to compare these against each other.

In order to make the best decisions, it's useful to think about how valuable growing the movement is compared to the direct work. This is very hard to know for sure, since we never get to see how all the counterfactuals play out. But I think we can improve our intuitions by understanding some simple models and getting a clearer picture of the kind of dynamics that may be at play.

Section 1, the majority of the paper, aims to give the reader a set of conceptual tools for thinking about movement growth. We introduce the Awareness/Inclination model, and use it to explore the possible life-cycles of movements without intervention, and the counterfactual effects of interventions.

In Section 2 we combine these insights with personal impressions of the effective altruism movement to provide recommendations.

1. Modelling movement growth

In order to understand the counterfactual impact of a marginal intervention to foster movement growth, it's helpful to first have a picture of what we expect to happen without that intervention. Generally this will be of the form: (i) the movement grows as more people become aware of it; (ii) as it grows more direct work is done; (iii) at some point the movement size plateaus.

I'm going to introduce a simple model of the space in which this is happening. This will allow us to analyse the effects of different kinds of intervention.

¹ In this paper I count earning to give to direct work as direct work. It is of course indirect in an important way, but in a direct work / movement growth split, it fits with the direct work. In one sense, the idea of 'earning to give' is one example of a way to help the world that a person might have discovered.

1.1 The Awareness/Inclination Model (AIM)

People will know a variety of different things about the movement, and have different attitudes towards it. As a modelling assumption we will compress these down to two dimensions: *awareness* (how much they know about the movement) and *inclination* (how favourably disposed they are towards the movement, or how favourably they would be if they knew more).

Then a person (at a fixed point in time) will be represented by a point in this space. Some examples are shown in Figure 1:

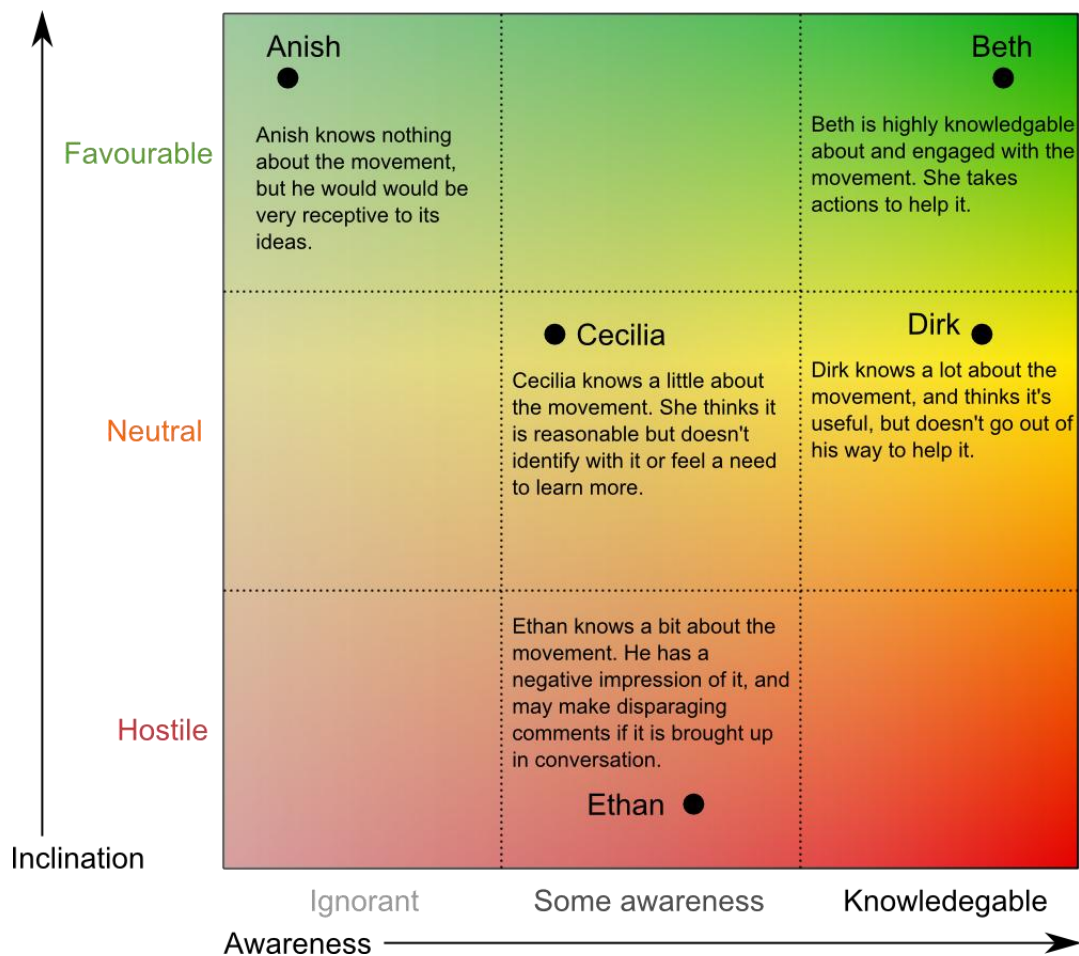


Figure 1: Individuals in the AIM space

Of course over time people may move around in this space. This will usually be driven by finding out more about the movement. So when movement occurs, we should expect it to be primarily in the left-to-right direction of the diagram. As they learn more, people may also become a bit more favourable or hostile to the movement. It seems to me that significant changes in inclination without significant increases in awareness are likely to be relatively rare during the growth phase. This is a substantive assumption which could be worth exploring, but the conclusions of the model do not require it.

Examples of some trajectories individuals could take through this space are shown in Figure 2.

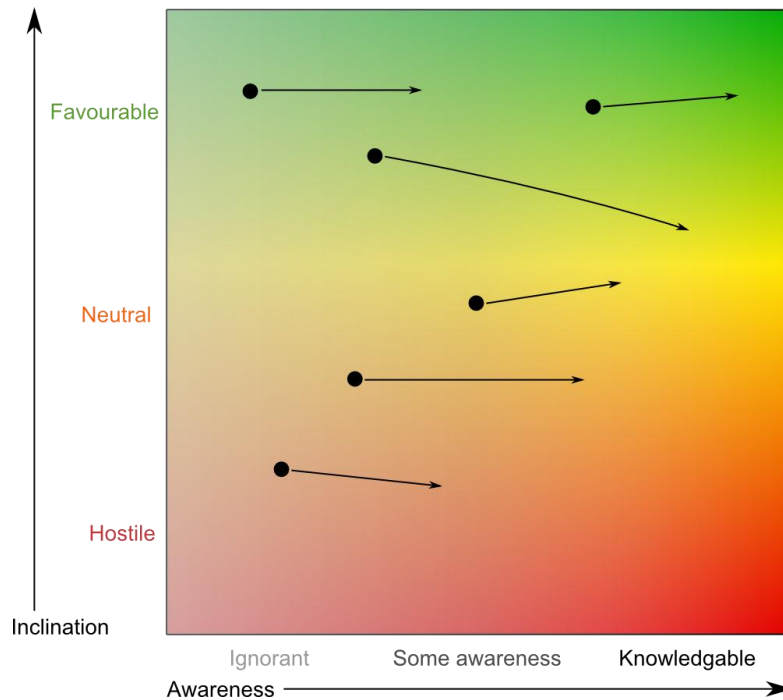


Figure 2: Example trajectories as individuals become more aware of the movement

1.2 Life-cycle of a movement

Early on in the life of an idea a large majority of the population will be ignorant of it. We cannot in general say how favourable they will be. Movements where the people on the left are clustered towards the top and away from the bottom have more potential for growth, and less potential to lead to a vocal opposition. It is the people in the top-right of our diagram who are both knowledgeable and favourably inclined who will actively pursue and promote the idea, and may be regarded as the movement built around that idea. This is illustrated in Figure 3:

How valuable is movement growth?

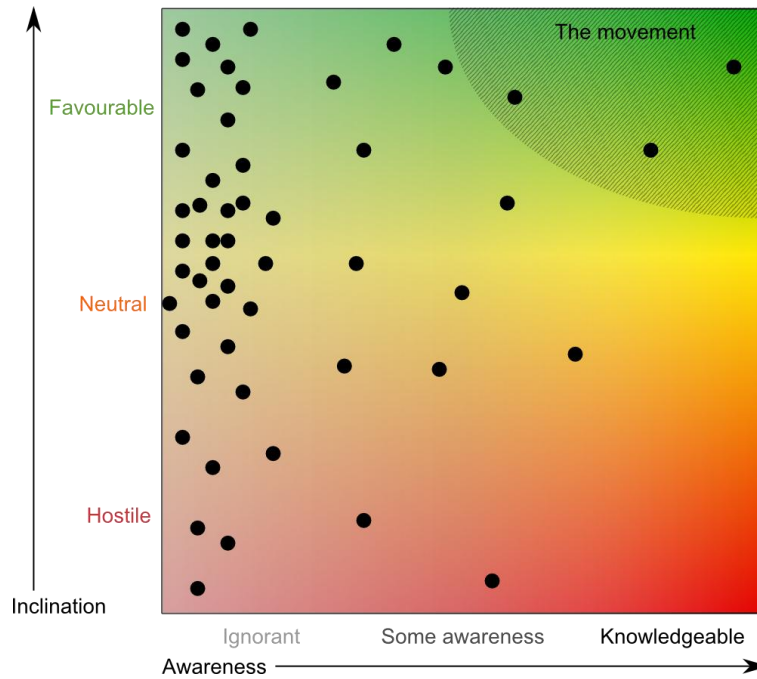


Figure 3: Early in the life of a movement

As the movement starts to grow, some people will become more aware of it, moving along trajectories to the right and perhaps changing inclination slightly (Figure 4).




Figure 4: Growth of a movement

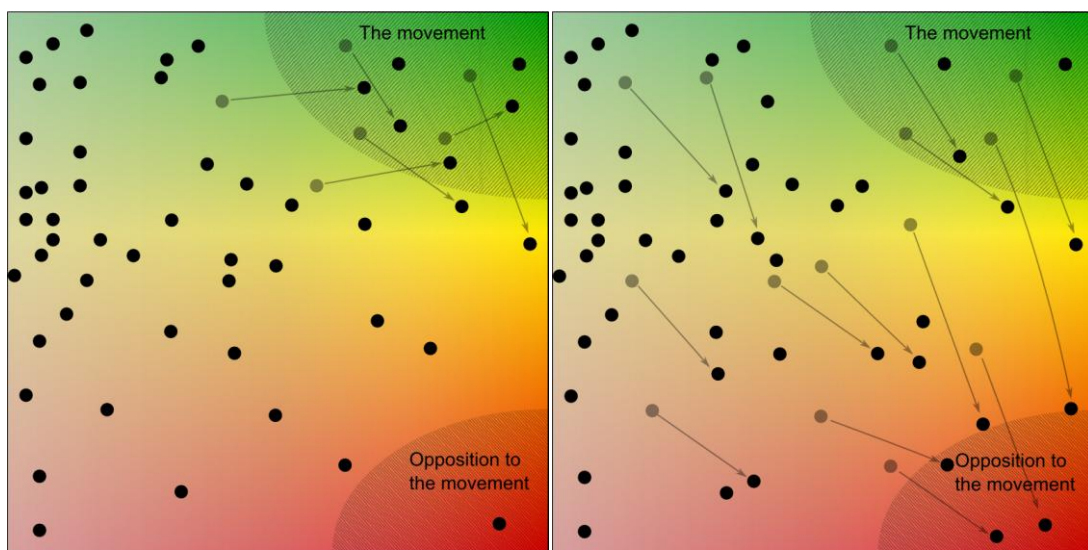
In our example, the growing awareness resulted in a substantially larger movement, as well as someone moving into the bottom-right corner of the diagram, becoming knowledgeable about the movement and also hostile to it. Why should this section matter? Consider what drives the increases in awareness and changes in inclination.

At the level of the model it seems that more people being knowledgeable and favourable to the movement will help to increase the spread of awareness, and that more of the knowledgeable people being favourable rather than hostile will help to nudge people’s inclination in a positive direction. So having a larger, more active opposition to the movement may decrease the frequency of positive inclination-shifts as people learn more (and increase frequency of negative inclination-shifts). It may also affect how the movement is able to influence public policy: a small highly favourable group with no opposition may be better able to shift policy than a larger group with a vocal opposition.

What is the long-term picture? There are four stable situations that can come about in the mid- to long-term as the number of people informed about the movement grows:

Scenario A – Saturation: There is a certain amount of work that should usefully be done on the issue at hand (perhaps each year, or perhaps in total). When the movement gets large enough, this work gets done, and there is no need to make further people aware; the movement is the correct size and more growth might even be bad. Existing people drift away from the movement as they see less need to be involved. *e.g. fundraising for specific projects.*

Scenario B – Collapse: Something happens to cause negative inclination towards the movement. This could be an inevitable consequence of increased awareness, or could be caused by some particular event. This affects both existing supporters and people who do not know about it. A subcategory of *collapse* is *stagnation*, where there is no real opposition but people lose their strong positive inclination, and nobody works to bring new people into the movement. In either case the negative inclinations stops the movement from growing, and it eventually collapses. *e.g. Marxism*



A - Saturation
B - Collapse

Figure 5: Scenarios where many people never become aware of the movement

How valuable is movement growth?

Scenario C - Controversy: In this scenario almost everyone knows a reasonable amount about the movement. A fraction of society is favourable enough to be actively involved, but many people are not in support. Most large movements can expect to be here. *e.g. abortion*

Scenario D - Acceptance: It becomes socially normal to support the movement. Almost everyone knows something about it and agrees.² *e.g. ending the use of lead in consumer products*

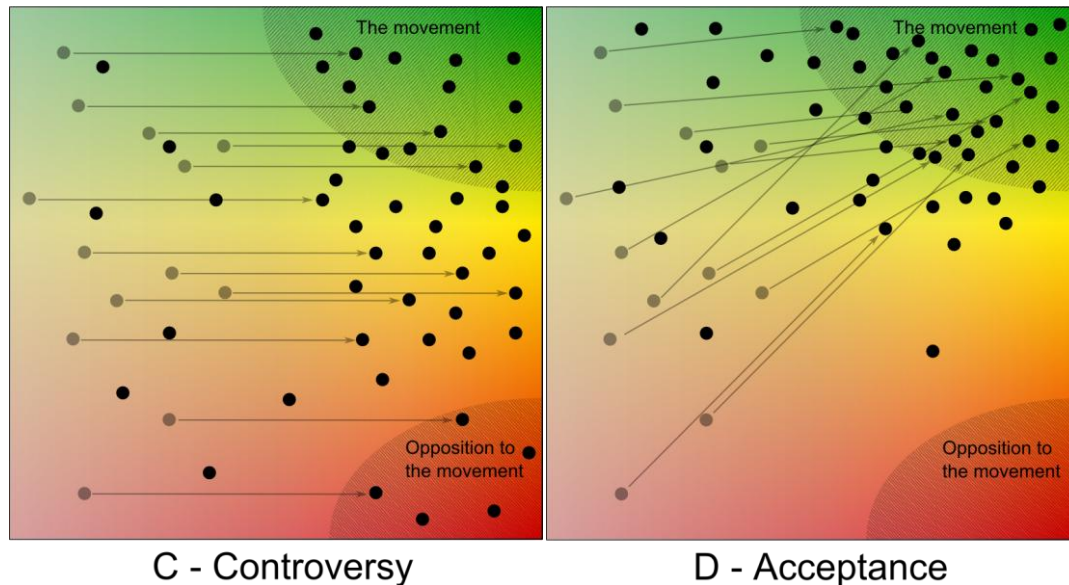


Figure 6: Scenarios where virtually everyone becomes aware of the movement

1.3 Effects of interventions

The marginal effect of an intervention may be quite different depending on which long-term scenario we are in. For this reason it's important to have a good understanding of which are the most likely. To help think about the effects of interventions, we'll consider the two intervention archetypes illustrated in Figure 7.

Intervention X - Publicity: This intervention focuses on raising awareness of the movement. It is consistent with the view that all publicity is good publicity. *e.g. performing stunts that are widely reported.*

² The *Acceptance* scenario has a couple of variations: one where there is a positive feedback loop which increases inclinations towards the movement; a second where people's natural inclinations are positive enough that shifting right in the diagram is enough, and the movement is only limited by awareness. In the very long term *Controversy* could shift to *Acceptance*, for example as in the abolition of slavery.

Intervention Y - Advocacy: This reaches fewer people than *Publicity*, but is concerned chiefly with improving their inclination towards the movement. It is closer to brand management. *e.g. writing a persuasive opinion piece for a newspaper.*

Both of these interventions affect the population roughly uniformly; in practice many interventions may target subgroups.

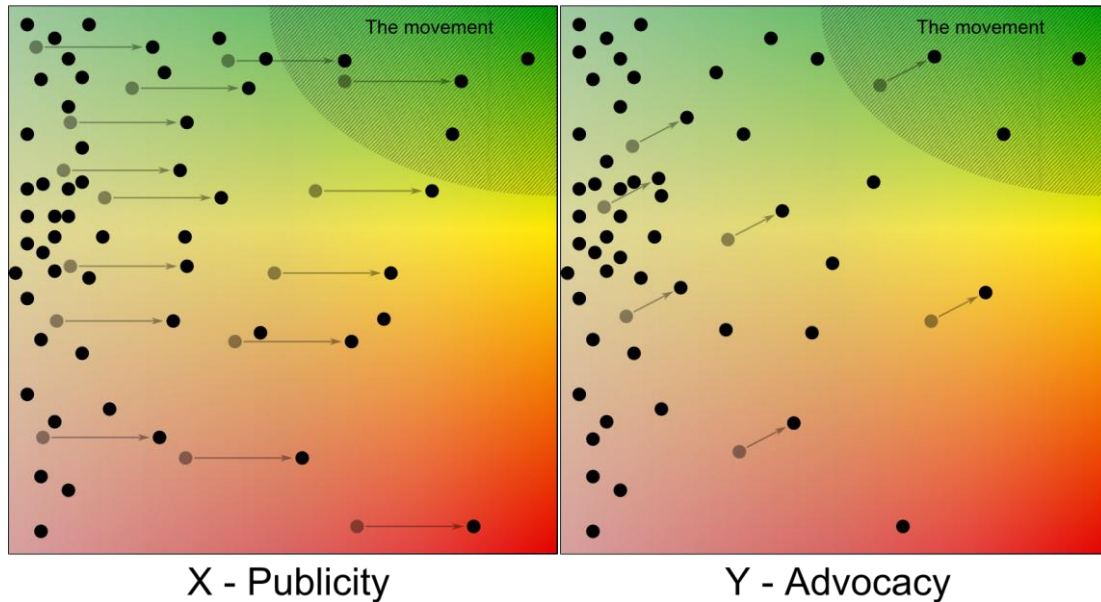


Figure 7: Two intervention archetypes.

1.4 Shifting timelines

The *Saturation* and *Acceptance* scenarios are perhaps the simplest to analyse. In both cases there is a natural eventual size of the movement, so marginal work on movement growth will not change that size, but change when it is realised. This suggests that the focus should be on increasing awareness rather than improving inclination: *Publicity* has a bigger effect than *Advocacy* on reaching the end state quickly. There is still a question of how good increasing awareness is compared to direct work.

The curve of movement size against time before intervention (black curve in Figure 8) thus has a horizontal asymptote. Suppose that we intervene to make the movement grow by k people. We will assume for now that our intervention just moves the entire curve forward in time (red intervention and blue curve in Figure 8).

How valuable is movement growth?




Figure 8: Intervention to accelerate movement growth

What is the total effect of this? The growth of the movement is advanced in time by a certain amount a (see Figure 9). This will produce extra direct work on corresponding to the area between the blue and black curves.³

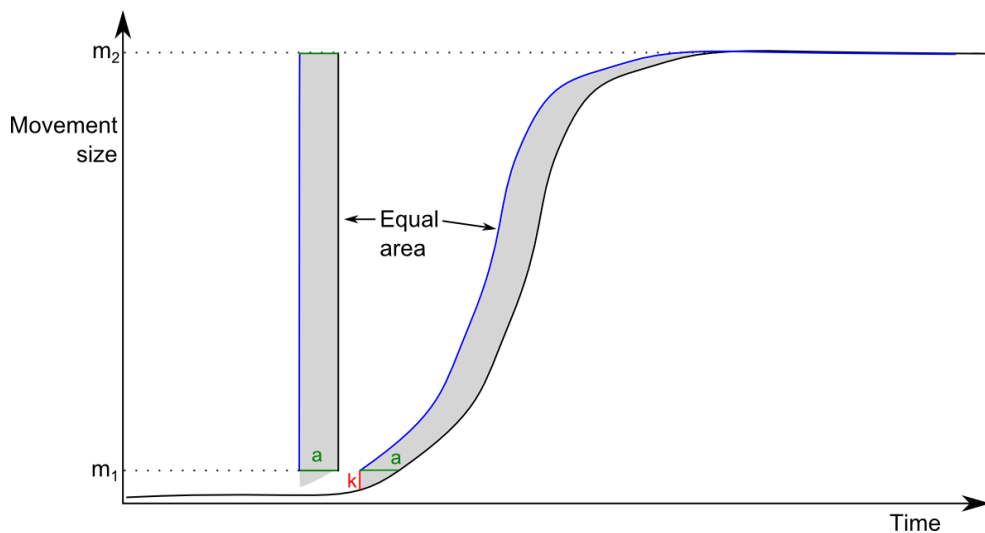


Figure 9: Assessing the impact of accelerating movement growth

This area can be shown to be equal to a times the difference between current movement size and eventual movement size (see Figure 9 – so long as k is small the area of the extra grey triangle is negligible). If we value direct work equally regardless of when it happens, this gives a simple way of understanding the value of acceleration: it is only necessary to estimate the current rate of growth and the eventual movement size. Then:

³ Note that area in this diagram represents person-years in the movement; we are assuming that the direct work done is in proportion to this.

$$\text{Benefit} = (\text{eventual movement size} - \text{current movement size}) * \text{advancement period}$$

For example suppose the movement currently has 10,000 people in, is growing at 10 people per day, and is projected to reach an eventual size of 100,000 people. Then an intervention which gains 10 extra new members beyond what would have happened anyway for the movement accelerates the process by a day, *so long as it also advances the background awareness of the movement at the same rate.*⁴ This means the value of this advancement is the same as the direct work produced by 90,000 people in the movement for a day – which is the same as one person for 250 years. So the value of just one person joining can be estimated as having their future direct work for around 25 years.

Some caveats. A person being worth about 25 years seems reasonable, but this is a bit of a coincidence of the numbers. The value is quite sensitive to eventual movement size. If the movement will only reach 20,000 people, each person added only equates to 2.5 years of direct work. If it will reach 10 million, each person added (while it's currently only 10,000 people) produces a whopping 2,500 years of direct work, because they pull forward the future hump of the distribution by a macroscopic amount. Uncertainty about eventual movement size will complicate this, but much of the expected value may be driven by scenarios where the movement gets very large, if these are realistic. In practice interventions won't advance all parts of the awareness process equally, so just counting new members is likely to overestimate the value (though probably by less than an order of magnitude). And if we value direct work differently at different times (or differently according to how much has been done), that would change the analysis, most likely in the direction of making direct work better relative to movement building.

For these scenarios it seems that movement growth can compare very favourably to direct work, and that *Publicity* is often better than *Advocacy*.⁵ The major exceptions

⁴ This corresponds to the assumption that we are simply moving the curve forwards. An example of an intervention which does not satisfy this assumption might be persuading 10 people who already think of themselves as almost members of Giving What We Can to fill in the membership form. However doing that together with sufficient outreach to refresh the pool of such people probably would satisfy the assumption. This means that outreach which only shifts people along at one part of the curve should be discounted by up to an order of magnitude compared to the direct measured shift.

⁵ However, there is a mechanism which could make *Advocacy* crucial in these situations too: relative to *Publicity* it is probably slower at creating movement members, but it is also much slower at creating opponents of the movement. If having opponents could significantly slow movement growth, this indirect effect may mean that in the long term *Advocacy* speeds getting the movement to full size more than *Publicity*.

How valuable is movement growth?

are when the movement has already done most of its growth, or if the movement is simply never expected to get large.

1.5 Shifting the long-term state

The *Controversy* scenario in some ways admits a similar analysis. Treating the eventual size of the movement as known and fixed, we can estimate how good accelerating movement growth is in terms of direct work. But an important difference is that we may be able to affect the eventual size of the movement. In *Saturation* and *Acceptance*, movement growth can be very important, and it is generally worthwhile to speed it up as much as possible by whatever means are most effective. In *Controversy*, long-term stability is achieved when everyone knows about the movement but people have different opinions about it. It may be possible to change this eventual distribution of opinions, and this could have a large effect.

For example suppose we have a movement of 10,000 people growing at 10 people per day, and projected to reach a size of one million people. Suppose there is an opportunity which would bring in 300 new people - accelerating growth by a month - but by being controversial, reduce the pool of people interested by 1%. At the current movement size, putting off 1% is just 100 people, so attracting an additional 300 seems like a great deal. But if the change persists and affects the stable level, this would be reduced by 10,000. Projecting twenty years past the plateau of movement size, the intervention will destroy more than twice as much direct work as it creates in this model.

Since we hypothesise that a major driver of improving inclination towards the movement is having lots of people with high inclination and few people who are opposed, in *Controversy* it may be more important to focus on improving inclination towards the movement as a first priority, and raising awareness only as a second priority. *Advocacy* is probably a preferable intervention to *Publicity*.




Figure 10: Distribution of level of inclination in a population of 1 billion people

In Figure 10 we suppose we have a metric for inclination towards the movement such that the population is normally distributed.⁶ The horizontal bands show standard deviations (z-scores), and the dots and numbers show how many people lie in each band, from a population of around a billion people.

Changing the threshold required for eventual participation in the movement can have a large effect size. For example suppose in one scenario the movement only catches on in those who are 3.8 standard deviations towards favourable or higher, whereas in another the movement catches on in those who are 3.7 standard deviations towards favourable or higher. Shifting from the first to the second scenario seems like a small change in required attitude, which might be achievable by shaping the right positive view of the movement. But its effect is to add 35,000 people to the eventual size of the movement, on a base of 72,000 – almost a 50% increase! So if how positive people will eventually be towards the movement can be adjusted, that could be crucial (see Figure 11).

⁶ If necessary this might be a synthetic metric, in the same way that IQ is.

How valuable is movement growth?

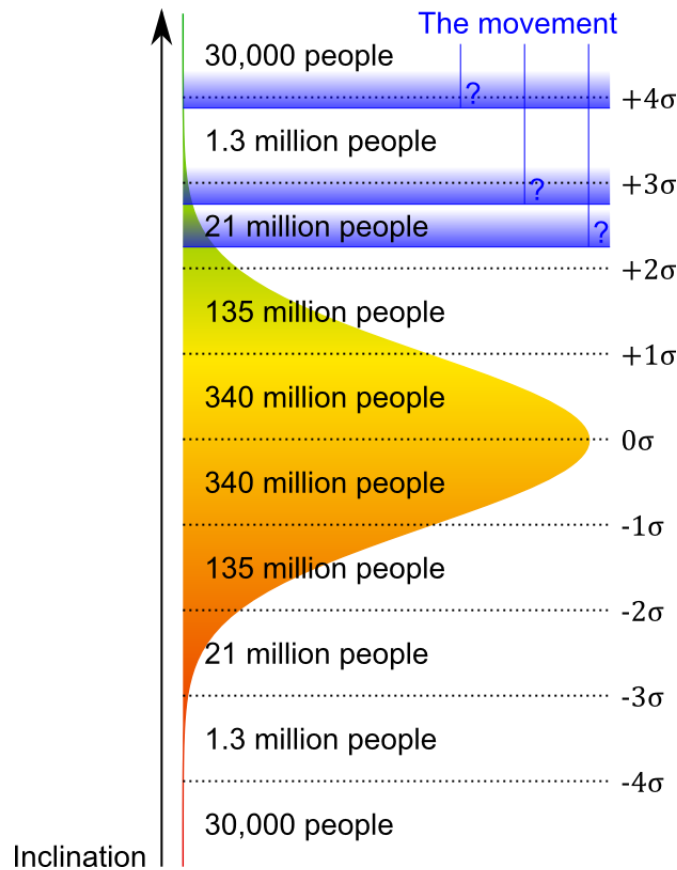


Figure 11: Possible eventual movement sizes – depending on how extremely favourable people have to be in order to become involved

Interventions which change the likelihoods of different long-term scenarios may also be crucial. Changing scenarios is most plausible between *Collapse* and the other scenarios. This could have a very large total effect. Avoiding movement collapse may be difficult to do with positive interventions, but it becomes extra important to avoid events which could catalyse the creation of opponents of the movement and drive movement collapse.

1.6 Changing the movement

In the above, we've imagined that changing the framing around the movement can affect the size it grows to, but that the underlying message and the direct work associated are essentially the same. It's also possible that by choosing a slightly different movement to promote, for example by excising less popular ideas, we could change the growth.

This may be a powerful tool. Because the eventual movement size may be so sensitive to – among other factors – how appealing the ideas are (as in Figure 11), it could be valuable to look for opportunities to adjust and improve this. It's also a tool that should be used very carefully, however – the value of spreading a different movement may be substantially different, and it could be too easy to throw the baby out with the bathwater by removing unpopular components.

1.7 Limitations and extensions

The model we have used necessarily makes a number of simplifications. I think these have generally been worthwhile, in allowing us to go further and faster through the analysis without distorting the picture too much. Nonetheless, it's worth being aware of possible limitations or shortcomings, and it may be that future work could explore relaxing these assumptions.

1.7.1 *Compressed space of attitudes*

Restricting to the 2-dimensional AIM space meant that we could not represent more complex attitudes, such as distinguishing between someone who has strong positive affect towards the movement but is not currently taking action on this, and someone who has much more mixed feelings but is currently pursuing some direct work. Much of the time I think these two dimensions suffice to capture the important distinctions, but there may be cases where they fail.

1.7.2 *No loss of awareness*

We assumed that people tend to move in AIM space only when presented with something external, and cannot move to the left. In fact people may drift around without interacting with others, and in particular may forget and so slowly lose awareness. It seems likely to me that the pictures presented are still roughly correct *on average* for growing movements, but further consideration of this point might show extra benefits to some interventions.

1.7.3 *No different behaviour for distinct subgroups*

Perhaps it is easier to affect inclination by large amounts in some groups such as students. The current model doesn't make any assumptions about this, but doesn't make it easy to model. An extension of the model could perhaps provide easier ways to think about heterogeneity in the population beyond the position in AIM space.

1.7.4 *Growth is primarily endogenous*

When we looked at the marginal effect of a larger movement today, we assumed that this would accelerate the movement growth process fairly uniformly. This is reasonable if growth of the movement is driven primarily by people in the movement. If growth is driven significantly by some exogenous factor (for example uptake of a new technology, reaction to political events, or flow from a related movement) then the long-term effect of growing the movement today could be much smaller. Changing the dynamics so that the long-term state alters could still be very important.

2. Applications

2.1 Types of intervention

Our example interventions *Publicity* and *Advocacy* were quite a long way removed from concrete actions. What might they look like in reality? Since the effects are all about changing minds (via awareness and inclination), interventions will tend to take the form of communication – although they could also include taking actions that can usefully be reported.

Example intervention might include:

- Writing an introductory article in a popular media venue. This would increase awareness across the population. Depending on style, it might also increase inclination, or be more controversial and sacrifice inclination for extra awareness.
- Writing a detailed articulation of an idea important of the movement, to publish on a website associated with the movement. This would increase awareness and inclination, primarily in a group near the top-centre of AIM space. The article you are currently reading is an example of an intervention which intends to be of this type.
- Direct work on the problem the movement is aimed at. This will not affect awareness, but improve inclination as people become more aware.

The last example may be surprising: we started out setting movement building against direct work, and are now using direct work as an example of movement building! What's going on here?

The answer is that while direct work has some movement growth effects, it is usually not a very effective way of movement building, so there is a meaningful choice between direct work and (effective) movement growth. However if the ratio of direct work to movement growth activities is too low, this could hinder the growth of the movement, either because it gives a natural line of criticism, or because it is harder to tell a clear story of what the movement is about. In such cases direct work may become the most effective form of movement growth at the margin – albeit with a slant towards direct work that can be communicated clearly or sets a good example.

As a final note, when choosing what to do, it's important to keep the concept of comparative advantage in mind. Even if you think a bit more movement growth is the most effective thing to be doing at the margin, if you have much better opportunities to pursue direct work, that could be better. Think about the portfolio of work done by those involved in the movement. Direct work should generally be done by those best at it, and movement building by those best at that. The importance of movement building relative to direct work sets the location of the split

between these activities and shows what people with intermediate skill sets should be working at.

2.2 Effective altruism

The effective altruism movement is concerned with using evidence and reason to do the most good it can in the world. It is currently quite small (in the thousands of people, probably the low thousands) and growing rapidly (for example membership of Giving What We Can has roughly doubled each year).

2.2.1 Decision-relevant parameters

A key decision-relevant parameter is how big the movement will get. It seems that it would be hard to run out of effective things to do, so *Saturation* is impossible. *Collapse*, *Controversy*, and *Acceptance* are all possibilities. Movement collapse is probably unlikely, but it is very important to avoid.

Knowing how big the movement might get would tell us whether we're likely to end up in scenario *Controversy* or *Acceptance*, and if *Controversy*, what the scale of the eventual plateau might be. I think there is a lot of uncertainty around the eventual size. It could be valuable to do market research to estimate what the uptake rates could get to in different slices of society, although this is not straightforward.

If, as seems to me most likely, we are heading towards *Controversy*, it is also crucial how much we could hope to affect the eventual size of the movement. Again, this is hard to estimate, but might be done with market research, seeing how much effect different framings have on people's willingness to engage.

2.2.2 Implications

The higher we think the plateau is, the more we should care about movement growth relative to direct work. The more we think we can affect the level of the plateau, the more we should prefer quality to speed of growth, focusing more on inclination than awareness. And the higher the discount rate we place on direct work, the more we should care about direct work compared to movement growth, and the more about spreading awareness than improving inclinations.

My impressions are that it is possible the movement could get very large – millions of people within twenty to fifty years. Until we have ruled this out, it is important to keep movement growth a priority, and to weight improving inclinations more than increasing awareness in many trade-offs. Focusing on improving inclinations means that we should be careful with our framing, trying to explain how we help people to do things they already wanted to do.

Effective altruism has several facets, as people pursue direct work in multiple areas. Dealing with developing-world poverty has been perhaps the largest focus area of the movement. I think that at present movement growth is not significantly limited

How valuable is movement growth?

by a lack of direct work in poverty, and this work is not dramatically more useful now than in a few years. The implication of this is that it may be better at the margin if we put a few more resources into movement growth rather than direct work in poverty. Of course this statement has several caveats:

- This is by no means a claim that poverty isn't a hugely important issue, and one with great opportunities for direct work available. It may well be that the big success of effective altruism will be to help eliminate poverty. The claim is just that extra direct work today may have less total impact on poverty than extra movement growth today.
- This doesn't mean that we should collectively stop doing direct work in poverty altogether. This is important to continue: because it would be faithless to abandon the work; because it helps attract more people to such work; and because it helps us learn how to do good such work.
- Some individuals may be best increasing their direct work in poverty, because they have strong comparative advantage there or because they could be more effective advocates for effectively alleviating poverty if they could talk about their direct work in poverty.

Does the same argument apply to say that effective altruism should reduce its focus on direct work altogether to prefer movement building? Not entirely. I think that the movement's growth may be limited by credibility, by ability to assess cause areas or interventions, or by direct work in other areas such as engagement with policy. In general the direct work that is done will have an effect on the way that people perceive the movement, so may affect awareness and will certainly affect inclination.

2.3 Very approximate numerical estimates

Suppose the effective altruism movement currently has 3,000 people doing an average of \$5,000 direct work per person per year (present movement size equivalent \$15M/yr), and is growing by 10 people per day.

Great uncertainty over eventual movement size. Possibilities, excluding collapse:

- 10th percentile: \$250 M/yr (50,000 people doing \$5,000 per person)
- 50th percentile: \$3B/yr (1M people doing \$3,000 per person)
- 90th percentile: \$50B / yr (50M people doing \$1,000 per person)

Mean approx \$10B/yr

Adding a person to the movement ~ bringing this forward by 0.1 days. Adjust down to 0.025 days to account for the fact that there are lots of steps in getting people towards movement.

0.025 days * \$10B/yr = \$700,000 worth of direct work (excluding movement building effects from direct work)

Perhaps adjust this down a bit for future work holding less value than current work, and down a bit more for people doing otherwise useful work. *Say \$300k for getting a person into the movement.*

That was for spreading awareness. What about improving inclination? (Numbers for this are if anything even more crude.)

Suppose increasing inclination of everyone who currently knows something substantial about the movement (say 10,000 people, which includes many people not in the movement) by 1σ leads to opinions being 0.1σ higher when movement reaches plateau, and that that change lasts 20 years.

Then causing a noticeable (0.2σ , corresponds to $\sim 1.5\text{cm}$ height) shift in one person's opinion leads in expectation to $2 \times 10^{-6} \sigma$ shift in population opinion. That leads to an expected ~ 200 people at 90th percentile, and around 25 people in expectation total, causing about \$1.2M extra direct work. This should be adjusted down for later work being less valuable than early work – perhaps to \$600k.

This is a large effect for improving opinions. It's probably most important for improving the opinions of people near the top and bottom of the inclination spectrum, as those are the ones who are likely to communicate most about the movement; therefore I suggest increasing this estimate for those people as well as for influential people, and decreasing it for those in the middle.

It might be better to model shifts in opinion as decaying with time and when spreading to new people. This will probably lead to a lower overall estimate of the value of changing opinions, since the decay will be biggest in scenarios where the movement gets very big and inclination is particularly important (mock-up in Excel suggests reduces impact by a factor of 4). *Say \$150k for a noticeable shift in an opinion.*

All of these numbers are basically pulled out of my head and are neither carefully modelled nor terribly robust. Nevertheless, they are suggestive, and I'd be mildly surprised if they were wrong by much over an order of magnitude. It's important to remember that we excluded the movement growth effects of direct work in order to make the comparison, so direct work will tend to compare more favourably in practice than these suggest.

Key empirical questions for better numerical estimates:

1. How big is the movement likely to get?
2. How many people know about the movement now?
3. How fast is the movement currently growing?

How valuable is movement growth?

4. How much do attitude shifts decay with time?
5. How much are attitude shifts transferred across people?

2.4 Practical upshots

Here I'll try a little more interpretation, to discuss what this means when deciding between actions. These views are informed by the theory but not entailed by it:

- We should be very careful to avoid actions which could cause negative inclinations towards the movement for little gain. We should therefore avoid hostility where possible, and should try to avoid expressing opinions on controversial topics when we may be perceived as affiliated with the movement.
- When allocating resources to direct work, we should do so in such a way that we can demonstrate what we've done, and we can learn from the experience, even if there are associated overheads.
- We should communicate our successes clearly. If a significant amount of the value of direct work today comes from demonstrating credibility and persuading others, it's important to be able to draw on this easily when relevant. (Note this is not the same as trumpeting our successes widely.)
- We should probably avoid moralising where possible, or doing anything else that might accidentally turn people off. The goal should be to present ourselves as something society obviously regards as good, so we should generally conform to social norms. Similarly presenting ourselves as on the side of the person addressed may help to avoid creating negative inclination shifts.
- We should try to reach good communicators and thought leaders early and get them onside. This increases the chance that when someone first hears about us, it is from a source which is positive, high-status, and eloquent.

2.5 Other areas and movements

How does this analysis translate over to other young movements, or other focus areas of effective altruism? For such linked movements, it generally seems that *collapse* and *controversy* are the most likely outcomes, that many people are relatively ignorant of it, and that the community of people who might care about it is much larger than the community that currently does. So for the same reasons as before it seems that movement growth is important, and that inclination may be more important than awareness. In cases where many of the eventual routes to impact are likely to involve policy, focusing on inclination and avoiding making opponents could be particularly important.

However, I think there is a significant disanalogy between effective work in global poverty and some more speculative areas such as existential risk reduction. Effective

interventions in global poverty are well-demonstrated, so the growth of attention is not significantly limited by direct work. In contrast, the growth of attention for existential risk reduction is substantially limited by a lack of credible direct work in the area. Therefore, in order to foster movement growth, one of the highest value activities is to do direct work, and to set an agenda for direct work, that clearly helps with the problem. Choosing work that demonstrates credible ways to help, and explaining the value of this work, is probably more important than choosing the direct work with the biggest actual effect in the short term.⁷ There may also be risks that need to be addressed in the short term, which could increase the discount rate on direct work, and so make direct work more valuable relative to movement building.⁸

⁷ See Ord: <http://www.fhi.ox.ac.uk/the-timing-of-labour-aimed-at-reducing-existential-risk/>

⁸ See <http://www.fhi.ox.ac.uk/Allocating-risk-mitigation.pdf>