Introduction

- I have only spent a very limited amount of time on these questions. I anticipate that I could change my values quite a lot by next year.
- Overall, I benchmark the value of 1 to equal a healthy life year.

Moral Weights

1. Health versus consumption

Choice:

- Value assigned to doubling consumption for one person for one year = 0.5.
- However, I will update this (potentially quite substantially) based on IDinsight's work on beneficiary preferences in Kenya. Depending on the extent of methodological limitations

in that work, I will place between a 0% and 50% weight on that study relative to my current best guess.

Explanation:

- GiveWell has previously done some work reviewing the existing literature on the value of a statistical life in this report. It summarises that studies often conclude that "a year of healthy life is roughly as valuable as 2-3x gross domestic product (GDP) per capita". I doubt I can add much to that literature review in a short space of time.
- This figure is mostly based on studies in developed countries. I believe that this trade-off will be highly dependent on an individual's relative levels of income and health. This might differ substantially between the target populations of GiveWell recommended top charities and the populations in these studies. I have a weak prior that the target populations of GiveWell recommended top charities are relatively more income scarce than health scarce as compared to developed country study populations, and so would place a slightly lower value on health relative to income. But I think this is ultimately an empirical question. I therefore choose the lower end of this range: 2x GDP per capita.

- I plan to take the results from IDinsight's study very seriously, conditional on the quality of the eventual experimental design. I will take a simple weighted average of my existing best guess (based on GiveWell's literature review) and the IDinsight study. If I believe the IDinsight study is very well done, with no serious methodological limitations, I am willing to place up to a 50% weight on it (even though it is only a single study), for the following reasons:

  1) As mentioned above, I believe this trade-off is highly dependent on an individual's initial level of income relative to health, and so getting information from beneficiaries directly is very important given the lack of existing literature from LMICs.

  2) I'm quite skeptical of revealed preference approaches (which I think form a lot of the existing literature) relative to stated preference approaches for this kind of question. I am not convinced that, in revealed preference studies, individuals really have the risk of death top of mind when weighing up labour supply/consumption decisions. Stated preference studies directly prime people to make trade-offs between health and income.

Future steps to potentially update my view:

- Read and analyse IDinsight's work when it is completed.
- Literature review to try to understand whether the target populations of GiveWell's recommended charities are likely to be relatively more income or health scarce as compared to developed country study populations. This is probably a very difficult question to answer.

2. Value of averting death at different ages

a) Acquired life potential

Choice:

- ALP rises linearly from 0 at birth to 1 at age 13.

Explanation:

- I find the Time Relative Interest Account as useful as any approach to think about the badness of death at different ages (after reading James' thoughts on this question, and this and this). To use that approach, I need to estimate an acquired life potential function, representing an individual's interest in her future years of life.
- The main way I am thinking about ALP is in terms of the level of engagement a person has with the world. The more engaged a person is with the world - e.g. the more social relationships they have, the more projects or activities they are working on, the more hobbies or interests they have; and the more interest I believe they will have in continued years of life.
- It therefore seems intuitive to me that ALP would rise quite slowly during early childhood. [1] It feels sensible to me that this would accumulate something like linearly until an individual perhaps enters their teenage years (though I'm aware that this might be anchored on my own experiences).
- From that age onwards, I find it difficult to argue that a person would be lacking 'full' interest in continued years of life (i.e. I certainly believe that children can be as engaged in the world as adults and so have an ALP = 1).

b) Life expectancy

Choice:

- See my calculations spreadsheet.

Explanation:

- I use 2016 WHO estimates of life expectancy at age x, available here.
- Estimates are available for men and women separately. I take a simple average assuming gender parity.
- I then take a weighted average across countries to which GW-directed funds (to charities that avert death - i.e. AMF, MC and HKI) were allocated, where the weights are the

[1] I.e. relative to some other GW staff members (e.g. James, Amar), I feel more comfortable with a slower accumulating ALP function.

proportion of GW-directed funds to charities averting death that are allocated to each country. I calculate country-weights in this document.

c) Emotional costs of death to others (immediate family)

Choice:

- Emotional cost after loss of a child (aged 0-17) = 0.396 x 3.7 x 5 = 7.33.
- Emotional cost after loss of an adult (aged 18+) = 0.2 x 3.7 = 0.74.

Explanation:

- I follow Olivia's approach to this, which is to capture emotional cost as the Grief Severity (G) x Number of Grievers (N).
- Like Olivia, I base the estimate of G for the loss of a child (aged 0-17) in the GBD disability weight for moderate major depressive disorder, which is 0.396.
- Like James in his model, I count only grief felt by immediate family members. Assuming an average household size of 4.7, this implies N = 3.7.
- I believe that the level of grief described as moderate major depressive disorder could last for many years after the loss of a child. I assume it lasts for five years (though it could last for much longer, perhaps at a lower severity).
- I make a distinction between grief after the loss of a child (aged 0-17) and an adult. I think this is the most relevant age at which to make a distinction in terms of the severity and duration of grief felt. The death of a child of any age seems particularly unnatural and likely to induce more severe and longer-lasting grief than the death of an adult.
- After the loss of an adult, I assume that G is half as bad (G = 0.2), and lasts for one year.

d) Economic costs of death to others (immediate family)

Choice:

- I follow James' model for this.

Explanation:

- With the time I have, I doubt I can improve on that model. This is something I may come back to and update my belief on.

e) Replaceability

Choice:

- Apply a 92% discount factor to averting deaths of children aged 2 or younger (including unborn children).
- I apply this discount factor to the direct badness of death.[2]

Explanation:

- There seems to be reasonable evidence that in some contexts averting the death of a child can decrease subsequent fertility (see this post by David Roodman for a review).
- My gut instinct (which I guess is a total utilitarian view) is that we should place some weight on the value of the hypothetical child's life who would have otherwise been born had a given child (whose death is averted) died. It seems in-keeping with the way I think about counterfactuals generally - trying to reason out the likely causal effects of a given event and factor those in to the analysis.
- To be clear, I am not in any way suggesting that parents can deal with emotional losses of a given child's death by 'replacing' that child with another. The word 'replaceability' is pretty awful and I think confuses this discussion.
- According to Roodman's article, the effect of infant death on subsequent fertility is highly dependent on total fertility in any given context. In contexts where total fertility has declined towards 2, family planning is practised and we might expect something close to full replaceability. In contexts where total fertility is high and not declining by much, we might expect a mortality drop-fertility drop ratio of 1:0.5 to 1:0.33.[3]
- I use the country-weights in this document, reflecting which countries GiveWell directed funds to charities averting death were spent in ~2017. Bottom line (AMF, HKI and MC): Burkina Faso (26%), Nigeria (18%), Uganda (13%) and the remaining 43% is spread across many countries, all but one (PNG) in SSA.
- I then look at recent fertility data (available up until 2016) from the World Bank (downloadable from this website) for these countries. I look at total births per woman

[2] I have so far only thought about this as an adjustment to the direct badness of death. Subsequent fertility does not reduce the grief felt by family members, and so this discount factor should not apply to the grief component of the overall calculation. In future, I will make an adjustment to the economic costs of a death of a child aged 2 and under to account for replaceability, because presumably the subsequently born child would make a similar economic contribution to the child who has died. However, so far I have borrowed James' model for economic costs and not dug deeply into it, so I will make this adjustment at a later date. I do not expect it to affect the bottom-line much.

[3] "In places where lifetime births/woman has been converging to 2 or lower, family size is largely a conscious choice, made with an ideal family size in mind, and achieved in part by access to modern contraception. In those contexts, saving one child's life should lead parents to avert a birth they would otherwise have. The impact of mortality drops on fertility will be nearly 1:1, so population growth will hardly change.
In the increasingly exceptional locales where couples appear not to limit fertility much, such as Niger and Mali (Bongaarts 2013, Pg 3), the impact of saving a life on total births will be smaller, and may come about mainly through the biological channel of lactational amenorrhea. Here, mortality-drop-fertility-drop ratios of 1:0.5 and 1:0.33 appear more plausible", see this post.

between 1996 and 2016:

Burkina Faso - decrease from 6.80 to 5.35
Nigeria - decrease from 6.2 to 5.53
Uganda - decrease from 6.99 to 5.59
SSA overall - decrease from 5.97 to 4.84.

- I then classify Burkina Faso, Nigeria, Uganda and other SSA countries as all being of the 1:0.5 to 1:0.33 mortality drop-fertility drop category.
- My prior is that within all of these countries, GW probably operates in regions with higher than average fertility for the country as a whole (based on the assumption that fertility declines with 'development', and the areas that GW-directed funds go have achieved a less than average level of 'development'). So I use the lower bounds. I assume that for every death averted in Nigeria, Burkina Faso, Uganda and other SSA countries, 0.33 fewer births occur.
- Then the weighted average decrease in births for a death averted in areas receiving GW-directed funds = 26%(0.33) + 18%(0.33) + 13%(0.33) + 42%(0.33) = 0.33
- However, I want to discount this because I value hypothetical children's lives less than children who are actually born (haven't thought deeply about philosophical basis for this, just my gut instinct). I place a one quarter weight on a hypothetical child (who is never born because a death is averted) compared to the child who actually lives and whose death is averted. Then the discount = (¼) x 0.33 = 0.08.
- So I then apply a 92% discount factor to child deaths averted (due to the lost value of life from the child who would have otherwise been born due to replacement).
- Remaining question is over what age-group of averted deaths this discount factor should be applied. Given that a key mechanism for replacement in high total fertility contexts is biological (lactational amenorrhea), I consider only young children - children aged 2 years or younger (including unborn).
- I'm worried that something about this sits uneasily with the way I am thinking about ALP. It's unclear how a hypothetical child who is never even conceived, because a death is averted, could have developed any interest in their future years of life, which would suggest that their ALP should = 0. At the moment I am thinking about this as a somewhat separate consideration, but I think there is a reasonable chance I will remove this from my estimates in future if I believe this really is inconsistent after more thought.

f) Putting these considerations together.

- I put this all together in this spreadsheet.
- I follow James' model, taking an average of DALYs lost across age groups based on the proportion of global malaria deaths that occur at different ages.

Future steps to potentially update my views:

- Spend more time modelling the economic costs of death for others (I have followed James approach so far, but would like to spend more time on it).
- Refine the model of country-weights for GiveWell-directed funds (this exercise was very rough - see the document for details). And update based on future allocations.
- More thought about replaceability (something I feel quite uncomfortable about in my current approach). (If I keep the replaceability adjustment in, I will apply the concept to the economic costs of the death of a child aged 2 and under also).
- Adjust the calculation used to estimate average DALYs lost across age groups to be based on the proportion of malaria deaths at different ages in the countries where GW-directed funds are allocated, rather than on the global burden. I might also consider mortality rates of diseases averted by HKI's vitamin A supplements (particularly diarrhoea and measles), though HKI accounts for a much smaller proportion of GW-directed funds than AMF and MC combined so I don't think this would make a huge difference.

User Inputs

1. Discount Rate

James' framework seems like a good place to start answering this question. He outlines five considerations that GW staff members have factored in when choosing a discount rate. I also found Caitlin's notes on this very useful.

a) Improving circumstances

Choice:

- My best guess is that GW beneficiaries increase real consumption by 2.57% each year.
- Combining this with James' assumption that $\eta = 1.59$ in the utility function he outlines in his Google doc, this gives an implied discount rate of 1.5%.

Explanation:

Growth in real consumption:

- In this spreadsheet, I calculate the compound annual growth rate in real GDP per capita for the major country recipients of GW-directed funds for years (X to 2017) from X = 1960 to X = 2016, using World Bank data that James downloaded in this spreadsheet.
- I am interested in understanding long-term growth rates in these countries. CAGRs calculated over short-time periods (e.g. 2014-2017) are probably quite noisy, reflecting short-term economic changes. I therefore only consider CAGRs for X as recent as 2010.

- I'm not sure that I want to use data as far back as 1960. I'm unconvinced that economic growth during the 1960s, say, is that informative about likely future growth in these countries. I consider CAGRs only as far back as X = 1980.
- For each country, I take the average CAGR from X = 1980 to X = 2010. (In Uganda, data is only available as far back as 1982, and so I just take the average from X = 1982 to X = 2010).
- I then take a weighted average of these CAGRs across countries, where the weights reflect where GW-directed funds are allocated (as calculated in this document).
- Note that for these country weights I consider all GW-directed funds, not just GW-directed funds to charities averting death, since the discount rate is not just applicable to charities averting death.
- Amar also did this exercise, using the same country-weights but using IMF growth forecasts for real GDP per capita until 2023 instead of the long-term growth rates based on past World Bank data. Reassuringly he came out with an estimate of 2.63% (very close to my estimate of 2.57%)

Shape of the utility function of ln(consumption):

- I intuitively agree with James' belief that increasing ~1.5 individuals' consumption from $1,000/year to $2,000/year is as good as increasing 1 individual's consumption from $500/year to $1,000/year.
- I have therefore followed James' assumption that $\eta = 1.59$ in the utility function he presents in his Google doc. I have spent next to no time reviewing this, and this is something I will come back to in future.

b) Individual reinvestment

Choice:

- 1.9%

Explanation:

- I have used James' best guess as a placeholder, due to lack of time.
- This is something I will revisit in future. In particular, even if I stick with James' approach, I will use the country-weights I calculated for GW-directed funds to take a weighted average of the real deposit interest rates across countries.

c) Uncertainty

Choice:

- 0.75%.

Explanation:

- I think that other staff members have considered two main factors here:

  1) There is some risk that beneficiaries in programs with long-term consumption benefits (particularly deworming charities) die/become disabled, and so do not experience that long-term stream of consumption benefits.

  2) Major negative external effects, for example natural disasters or political unrest, cause our programs to become less effective in the future. (The external events we are thinking about here are new shocks that are not already factored into the long-term real GDP per capita growth rate estimate that we have based the "Improving circumstances" discount on[4]).

- I have a fairly different view on these two factors to some other staff members, as follows:

  1) We already build the risk that death or disability into our CEA models for the benefits of deworming charities. We estimate the duration of benefits (using an estimate for the length of a full career) based in part on an adjustment for the risk of death and injury.[5] I don't have a strong reason to believe that the size of the adjustment we have made is too small, and so I do not add a discount for this factor. However, we have not researched that question in much detail, and so this might be something to come back to.[6]

  2) If major negative external events like natural disasters hit the areas where our charities work, making their programs less effective in the future, then those same events will presumably also lower future beneficiaries' consumption. In that case, future beneficiaries' marginal utility from consumption gains will be higher. In other words, it seems to me like discounting for this is inconsistent with also adding a discount for 'Improving circumstances', described above.

---

[4] I.e. the estimate for growth in real GDP per capita already incorporates many downside and upside shocks that have occurred in the past. Here we are talking about historically unprecedented shocks that cause major deviations from past long-term growth trends and therefore haven't already been factored in to the analysis.

[5] We assume 40 years of benefits, based on the reasoning laid out in this document. "...the benefits likely last for a full career, which I'd guess for the average person will be ~40 years. (My true best guess of a career for a fully healthy individual is ~45 years, but I cut that estimate by ~5 years for the considerations that Natalie raised.)", p.2. Natalie's considerations were: "death, injury, obligations to work in or near the home caring for family", p.1.

[6] Josh suggests that in future work we could: "Better quantify the magnitude of effects that would cut careers short", p.1.

There are perhaps two reasons for making this discount whilst also discounting for 'improving circumstances': a) If an external event which makes the program less effective kills beneficiaries, then there is no offsetting effect on the marginal utility of consumption for them. b) It is possible that the external event has a (in some loosely defined sense) larger effect on program effectiveness than the offsetting impact on the marginal utility of consumption.

As a separate point, whilst there are downside risks, there are also upside risks to future program effectiveness from these external events. E.g. a major technological improvement could increase the cost-effectiveness of our charities' programs in future; as could the settling down of political unrest (if our current estimates of program effectiveness are based on studies done during times of political unrest). My gut instinct is that the downside risks are slightly greater than the upside risks, but I don't have a strong opinion on that.

- For the reasons outlined in paragraphs 2 and 3 of the bullet point above (2), I make a small discount for this factor of 0.75%.

d) Pure time preference (beneficiary)

Choice:

- 0%.

Explanation:

- My gut instinct is that pure time preference is generally irrational (in the sense that today's self would not want tomorrow's self to be present biased). However, I do find the empirical lack of take-up of commitment mechanisms (for their own sake) quite puzzling. Could just be that people are naive about their present bias, or it could be a sign that there is something more intentional about this preference. This is something I will spend more time thinking about it future.

e) Pure time preference (donor)

Choice:

- 0%.

Explanation:

- Intuitively doesn't feel right to me that we should philosophically value people now more than people later.

f) Summing these values

1.5% + 1.9% + 0.75% = 4.15%

Future steps to potentially update my views:

- More time thinking about the shape of the utility function of ln(consumption).
- Update my value for individuals' reinvestment (at the very least to try to better reflect where GW-directed funds are going using the country weights).
- Refine country-weights and update for future allocations.

2. Deworming replicability adjustment

Choice:

- 11%.

Explanation:

- This is something I know very little about and have spent very little time on. Reading Josh's thoughts on this, it seems like lots of staff members have thought about this parameter seriously and provided inputs for his final suggested change. As a placeholder, I trust that Josh has weighed this all up sensibly, and I think it's preferable for me to use his suggested input than rush a decision based on my own quick review of all that evidence.

3. VAS development benefits

Choice:

- 6%.

Explanation:

- I have spent very little time on this. I have read through James' and Chris' arguments (and their Slack discussion), and both seem quite persuasive. I don't have a strong prior for whether development benefits are likely to be more or less important when tackling VAD as against malaria. But development benefits generally have less evidence, and there is particularly less evidence in the case of tackling VAD relative to malaria. So I'm going to discount a bit for this - I think the proportion of benefits from VAS which come from development should be roughly 80% of the proportion of benefits from SMC/AMF which come from development. For that to be (approximately) the case (which depends on all my other inputs to the respective CEAs), it turns out that I need to choose a value

for "Value of development benefits from a year of VAS coverage relative to a year of deworming in Miguel and Kremer" equal to 6%.

4. Other user inputs

- I have not had a chance to think about the other user inputs. I have therefore used the rule that: I defer to the parameter owner's view where it exists, and if not I choose the median view of other GW staff members.