Evaluating Effective Altruism and its Implications on the Fight Against Malaria

Samuel M. Williams

TC 660H

Plan II Honors Program

The University of Texas at Austin

May 10, 2018

_____

Ian Proops, Ph.D.

Department of Philosophy

Supervising Professor

_____

Christopher Meakin, J.D., M.A.

Department of Business, Government and Society

Second Reader

Abstract

Author: Samuel Williams

Title: Evaluating Effective Altruism and its Implications on the Fight Against Malaria

Supervising Professors: Ian Proops, Christopher Meakin

While malaria cases and fatalities have fallen since 2000, there were still an estimated 216 million cases and 445,000 deaths worldwide in 2016 -- comprising 0.8% of global deaths. One movement calling great attention and action to the global burden of malaria, particularly in sub-Saharan Africa, is effective altruism. Effective altruism is focused on prioritizing donations to causes and organizations with the greatest impact and highest cost-effectiveness. It hinges on the argument that when faced with cost-equivalent decision alternatives, we should choose those that do more good than less.

The first task is to evaluate whether effective altruism is correct in its assertions and aims. This will rely upon philosophical discussion centered on the works of a few key effective altruists, namely Peter Singer, and various objections raised against the movement. The second task will be to understand the current state of antimalarial efforts in prevention, treatment, and elimination. Last, I will apply the valuable principles of effective altruism to these efforts to determine what methods the ideal fight against malaria would utilize.

Acknowledgements

I’d like to share my gratitude to those who helped me in my efforts.

To Professor Proops, your wealth of knowledge and precise critique of my work were invaluable resources without which I could not have undertaken this task. I am immensely grateful for the amount of time and energy you invested in me and my work.

To Professor Meakin, the impact of your ethics class on my own moral viewpoints was influential in helping me reach this topic. Your open mind and thoughtfulness helped continually drive me in my research and conclusions.

To my friends and family, thank you for challenging my viewpoints and their impact on my work.

Table of Contents

Chapter I: Effective Altruism

I.a. Background

I recently received one-thousand dollars I didn't deserve. This came in the form of a scholarship for my membership in an organization which exists primarily for social reasons. There were roughly sixty of these scholarships available to the organization's members which were funded by alumni donations and distributed through my university. Since there are currently fewer than sixty members in the organization, it basically guaranteed that any member could earn the $1000 simply by applying. The dues for this organization were less than $1000, meaning members could effectively be paid for their membership. While admission is extremely selective, this notion was still unsettling to me. So I asked myself, *what should I do with this money*? (When I use *should* here, I more properly mean "what I morally ought to do.")

The two important words here are: *what* and *should*. What we *should* do (particularly with money) is a topic of fierce philosophical debate ranging from extreme altruistic stances of self-imposed poverty to ethical egoist stances that we *should* only spend money to further our own personal pursuits. The same question divides countries like the United States on partisan lines about the morality of uses of taxpayer money. This ponderance falls within the scope of normative ethics, the ethical study of what we *ought* to do. Depending on the answer to *should,* the question of *what* actually had a surprisingly defined answer. This was how I discovered effective altruism.

I.b. What is Effective Altruism?

Effective altruism is a guide for how we[1] should consider where and what we donate to. As defined by William MacAskill and the Centre for Effective Altruism, "[it] is the project of using evidence and reason to figure out how to benefit others as much as possible, and taking action on that basis."[2] In *The Most Good You Can Do*, Peter Singer extends this definition as "a philosophy and social movement."[3] On the underlying principles of value theory the philosophical aspect focuses on comparing how much good is accomplished by different donations. This essentially takes form in the argument that if, all other things being equal, there are two donation options with the same cost but with different amounts of resultant *good*, it is morally obligatory to choose the option that does more good for the same cost.

In reality, effective altruism has two meanings. As discussed above, it is in itself a thesis in ethical theory. The other side, though, is a pragmatic, activist vision of what it means to employ this theory in the real world. From this social aspect, effective altruism has grown into a diverse community with a number of notable characteristics. As laid out by Peter Singer, effective altruists *tend* to (1) live modestly and donate a large part of their income to the most effective charities, (2) research and discuss charity evaluation, (3) choose a career in which they can earn the most in order to give, (4) talk to others about effective altruism in order to spread it, (5) be willing to give part of their body, such as blood or bone marrow, to a complete stranger, (6) have equal concern for global causes as local causes, and (7) aren't "warm glow givers" -- people who give very small amounts to lots of causes for personal satisfaction. These

[1] "We," in this context, is referring to those of us in positions of financial security to donate without substantially harming our own ability to live.
[2] William MacAskill, *Effective Altruism: Introduction*, 2.
[3] Peter Singer, *The Most Good You Can Do: How Effective Altruism Is Changing Ideas About Living Ethically*, 5.

characteristics are not required to be part of the effective altruist community, nor must participants meet all of them. For example, an effective altruist can reasonably choose a career that pays $180k annually over a career that pays $200k because of the personal value it provides. An effective altruist can also rely on reputable organizations to inform his or her giving rather than conducting their own extensive research. An effective altruist, even, can meet none of these criteria while still believing in the value of the movement.

Singer frames effective altruism as "the most good we can do," and subsequently attaches a normative claim that it is also what we ought to do. His book makes persuasive arguments, and in itself could be considered an act of effective altruism. To convince others to join the philosophical and social movement which he has helped champion, writing *The Most Good You Can Do* is a justifiable use of time and funds for the impact it could create. In that sense, we do have to be slightly wary of the bias Singer may hold in favor of his own goals. Contrary to Singer, MacAskill asserts that effective altruism is not itself a normative stance, but, rather, that it contains normative views within it. In other words, MacAskill does not go so far as to say effective altruism is the most good we can do, or that it is what we ought to do, but rather establishes that it is an approach to altruism that focuses on utilizing quantifiable data and rational decision making to do the most good. It is within effective altruism, MacAskill claims, that there would be a normative stance that one donation is morally obligatory over another. In order to best determine the application of such a philosophy, we will evaluate the tenants of these stances and break the argument down to its roots.

I.c. Breaking Down the Argument

In order to best understand the implications of effective altruism and evaluate it properly, we must ascertain the components of the argument. First, let us apply working definitions to moral permissibility, impermissibility, and obligation. Actions are morally permissible according to a certain moral code if, and only if, (iff) they are actions that that moral code would deem acceptable. Actions that are morally obligatory are actions that a moral code would deem to be required to perform or else violate the moral code. For example, wearing a green shirt is morally permissible, but not morally obligatory. Actions that are morally obligatory are also morally permissible. For example, saving a child, when you have all the resources and there is no risk to oneself, would usually be regarded as morally obligatory (and therefore also permissible). Actions that are morally impermissible are actions that violate a moral code. Murder is a simple, common example of a morally impermissible action. We will avoid the terms *right* and *wrong* because their descriptive power is not specific enough for detailed ethical analysis.

These working definitions are dependent on a moral code itself. Since establishing a unified moral code or theory falls outside the aim of our scope, we will utilize "commonsense morality." Commonsense morality, per Shelly Kagan, is a shared "common moral outlook," in which "people may differ about the details, but at least the broad features are familiar and widely accepted."[4] This code may include some simple, or commonsense, moral claims such as *murder is morally impermissible* or, in our case of interest, that *donation is good*. Of course there are obvious exceptions to these, such as killing in self-defense or donating to morally dubious organizations, but we can still start our evaluation from the basic outlook. MacAskill lays out the

[4] Shelly Kagan, *Normative Ethics*, 25.

path that effective altruism is supported by a wide variety of moral views, saying that, "All plausible moral views care about making the world better, impartially speaking."[5]

Another key component of effective altruism is value theory. Value theory is focused on identifying how to quantify and compare abstract concepts such as *good*. Much like right and wrong, good and bad are vague terms that can be used many ways; we need to narrow to a working understanding for our purposes. In his paper "Value Theory," Mark Schroeder lays out four kinds of good: (1) value claims, (2) goodness *simpliciter*, (3) *good for*, and (4) attributive uses.[6] The following sentences help exemplify their uses:

1. Pleasure is good.
2. It is good that you contributed.
3. It is good for Jack to talk to Jill.
4. That is a good knife.

For effective altruism, the focus primarily lies on value claims and goodness *simpliciter*, as Schroeder calls it. Goodness *simpliciter* claims are "the 'good' claims that consequentialists hold to have a bearing on what we ought to do."[7] Relationally, this understanding of good has comparative power. For instance, one could say, "It is good that you did that. It is better that he did this." There are a few theories for how to best understand good *simpliciter*, namely: the *point of view* theory and the *agglomerative* theory. In the way that sentence 3 shows a *good for* from the point of view for Jack, goodness *simpliciter* is good from the point of view of the universe. This theory, though, can have some undesirable conclusions such that, for the universe, it may be

[5] MacAskill, *Effective Altruism*, 2.
[6] Mark Schroeder, *Value Theory*.
[7] Schroeder.

better that humans not exist at all. Under *agglomerative* theory goodness *simpliciter* is the sum of what is good for all the various people that there are. This is often treated as (one version of) the utilitarian view of good. For the sake of limiting the scope of this discussion, we will mostly adhere to this view of *good* in order to evaluate effective altruism.

Value claims, like the first sentence, attempt to show that things that are *good* are things that have value. In other words, their being good is what gives them value. This particular approach is referred to as the *good-first* theory. For effective altruism, this may look something like "donation is good" or "altruism is good." This means we derive value from things that would be considered good by goodness *simpliciter*. The other side of this coin is the *value-first* theory which conversely says that things have value and that is what makes them good. There is also a difference between "intrinsic" value and "instrumental" value. Things with intrinsic value are good in virtue of their intrinsic properties. There is significant debate over what things are intrinsically good that falls outside of our scope. Things with instrumental value are good because they lead to other good things. For our purposes, money has instrumental value because of its ability to facilitate other good.

Value claims are inherently difficult to compare. Take for example, "knowledge is good" and "pleasure is good." Which is better? It is hard to determine a correct answer, or if there is an answer at all. Because of the difficulty to make comparisons with value claims, there is good rationale to avoid Singer's normative stance that effective altruism is what we ought to do, and stay with the second-stage that MacAskill suggests. In other words, we can make the value claim that altruism is good, but we should avoid prescribing it as the best course of action and stick to the comparative nature of goodness simpliciter which can be used to evaluate donation alternatives. This, of course, weakens the stature of the effective altruist movement, but is a

necessary decision in order not to invalidate its self-contained claims. More on this in the Evaluation section.

Effective altruism relies on the principle that donations and outcomes can be compared. Value theory lays out a theoretical approach to quantifying and comparing good. Pragmatically speaking, a number of criteria are utilized in the evaluation process, narrowing down to a singular scaling metric for comparison. In turn, it actually becomes more feasible to measure how much an action/condition damages the resting state of good. In other words, it is easier to say how bad things are than how good all things are. The quantifiable scaling factor for this burden is often called a quality-adjusted life-year (QALY) or a disability-adjusted life-year (DALY). Singer often uses the QALY, but based on the World Health Organization's (WHO) use of the DALY, that is the terminology we will stick with. One DALY represents one lost year of life of full health. The WHO lays out the following formula:

> DALY = Years of Life Lost (YLL) + Years Lost due to Disability (YLD).
>
> YLL = N x L where N is number of deaths and L is standard life expectancy at age of death in years.
>
> YLD = I x DW x L where I is number of incident cases, DW is disability weight, and L is average duration of case until remission or death

This metric is applied in the WHO's Global Burden of Diseases for 2010 report:

> Researchers conducted nearly 14,000 face-to-face interviews in several countries and supplemented these findings with a web survey. The researchers found generally consistent results across distinct cultures. For blindness, it indicated a discount of 0.2. In other words, 1 year when blind is equivalent to 0.8 years of healthy life, or curing a person of blindness

> for 5 years is equivalent to extending a healthy person's life by 1 year. At that discount rate… in the population we could help for $100,000, untreated blindness causes the loss of 1000 x 0.2 = 200 DALYs per year, while starvation threatens to cause the loss of 500 DALYs per year. On these figures, we should feed the starving.[8]

There are many noteworthy reasons to be skeptical of this approach, namely in the difficulty to assign consistent quantitative figures to every "burden" people face. We will discuss these issues in depth under Challenges and Criticisms.

The final term of the argument is "cost." When we discuss cost, the pragmatic focus is largely financial. Having said that, the cost we consider ought to be inclusive of all relevant factors in making donation decisions. For example, making a $100 donation has an opportunity cost of any other action that could have been made with the $100. We ought to also consider any other consequences, such as the probability that our donation could have a negative impact on the environment or negative economic impacts such as state governments providing less aid. Often, these consequences are unforeseen, though, and speak to our inability to properly evaluate decision alternatives. More on this in Challenges and Criticisms.

[8] Singer, *The Most Good*, 132.

I.d. Applied Examples

Considering a few examples of the application of the theory can help bring this topic to life. Singer lays out an example comparing curing blindness and funding seeing-guide dogs. In total, supplying one person with a guide dog in the U.S. costs about $40,000. Trachoma -- the most common cause of preventable blindness -- can be prevented after identification for a cost ranging from $20-$100 per person. Per Singer, "the choice we face is to provide one person with a guide dog or prevent anywhere between four hundred and two thousand cases of blindness in developing countries."[9]

In 1998, Ted Turner gave one billion dollars to the United Nations to scale up already-proven health programs focused on the world's deadliest diseases, which mostly kill children. Singer evaluates its impact as such:

> Since 2000, 1.1 billion children have been given a combined vaccine that prevents measles and rubella. The vaccine now reaches 84 percent of the world's children. Between 2000 and 2012, worldwide deaths from measles have fallen 78 percent, with a total 13.8 million deaths averted. The cost per vaccination is estimated to be $1. If that figure is correct, the estimated cost per life saved would be just under $80.

This example exhibits the thought process of effective altruism. The approach is largely systematic and rational to create the greatest positive impact, rather than utilizing a decision process that relies on emotion.

Turner's example could stand in contrast, say, to the donation of Lucile Packard to establish the Lucile Packard Children's Hospital in Palo Alto, California. This hospital is well-

[9] Singer, *The Most Good*, 111.

known for its work with separating conjoined twins; a notable nine-hour, twenty-two-person surgery costing an estimated $1 million to $2 million grabbed national headlines. Effective altruists would look to this example as a case of an ineffective use of resources that could have done significantly more good elsewhere. There are a number of difficult implications in reaching these conclusions, though, that we will evaluate in the coming sections.

I.e. Organizations and Current Impact

There are a number of existing organizations employing effective altruism to help find and implement solutions to the world's most pressing problems. GiveWell and GoodVentures are two of the better known organizations. GiveWell focuses on researching and publishing details on the most effective giving opportunities and directing donations that are "evidence-based, cost-effective and scalable." GoodVentures is a $10 billion foundation that focuses on grantmaking based on three criteria: "importance, neglectedness, and tractability." The two organizations partnered together and helped start the Open Philanthropy Project to advise major donors on maximizing their giving. [10]

In 2016, GiveWell moved over $100 million for the most effective charities; GoodVentures funded $126 million in grants recommended by the Open Philanthropy Project;[11] Giving What We Can Members have collectively made over $1.5 billion in lifetime pledges to charity and already reported over $25 million in donations.[12]

Rather than evaluating specific organizations, the Abdul Latif Jameel Poverty Action Lab (J-PAL) and Innovations for Poverty Action (IPA) evaluate specific interventions, such as the effect of price on coverage of insecticide-treated bednets.[13] J-PAL has a core staff of over 300 research, policy, education, and training professionals worldwide and receives funding from organizations such as the MacArthur Foundation, the Gates Foundation, the Australian Department of Foreign Affairs and Trade, and the UK Department for International Development. J-PAL has evaluated methods such as school-based deworming programs in Kenya, microfinancing in India, and

[10] See GoodVenture's "Grantmaking Approach" page.
[11] MacAskill, *Effective Altruism*, 3.
[12] This figure it tracked and displayed on GivingWhatWeCan's homepage.
[13] Singer, *The Most Good*, 15.

improving immunization rates through regular camps. IPA has found insecticidal bednets have significantly better coverage when provided for free than when any cost, no matter how small, is distributed to recipients. The impact of its work can be seen in the WHO's endorsement of free nets, in certain organizations' movement to provide free nets in endemic areas, and Population Services International's (PSI) decision to rapidly scale up free nets for high risk groups such as pregnant women.[14]

The organizations detailed here name a few notable examples, but only serve to show the growing trend of effective altruism and the scale of impact it is creating.

[14] Per IPA's summary of results on providing free bed nets. https://www.poverty-action.org/impact/free-malaria-bednets

I.f. Challenges and Criticisms

Denouncing Good for Better

Within its root ambitions, effective altruism presents us with a number of challenges. First, it lays out a view for how to give effectively for those individuals who are committed to donation. This is formulated in a two-step structure: Step One being the decision to give and Step Two being the decision of what to give to. This could be a decision to donate a percentage of one's income every year, meaning the dollar amount is fixed and there is a second decision of where to donate the money. While this is an accurate description of some individuals' decision-making process, many individuals' charitable choices are motivated emotionally and carried out in a single-step procedure. In other words, sometimes one doesn't decide just to give but rather decides to give to a certain cause. This is often done on the basis of some emotional or personal connection to a cause.

Within the normative stance of moral obligation to choose efficient options, this leaves us in the uncomfortable position of criticizing/denouncing actions that are still charitable but to lesser degrees. It is hard to foresee the overall impact of such a stance, but one possibility is a decrease in overall donation amounts. Take for example someone who is interested in helping maintain the public library system in his or her city by donating money. If effective altruism were widespread and accepted, this person may feel too much shame or judgement for making that donation and instead do nothing altruistic with that money. That would result in a net negative under almost any working application of *good*. Take another situation in which your neighbors approach you that their child has been diagnosed with a form of cancer and asks for any donation that you could manage in order to help pay for treatment. Even knowing that you could save more children elsewhere in the world with the same money, would you really feel

good turning them down on such a basis? You would probably seek to help your neighbor first, meaning this money may come funds you had intended to do *more* with. If not, it could condition us to earmark less to charity in advance because we can foresee events like this.

Melissa Berman, the president and CEO of Rockefeller Philanthropy Advisors, an organization Singer critiques, "explain[ed] why it is best to allow potential donors to follow their 'personal convictions' in choosing a charity. Doing so leads, she observes, to their giving more and more consistently."[15] Even if this is true, it does not necessarily mean that more good cannot still be done with less. For example, an effective $50 donation may help more people more significantly (think DALYs) than an ineffective $100 donation. Having said that, since this figure is difficult to quantify precisely we cannot rule it out of consideration against effective altruism.

Terms of the Argument

Challenges also lie in the difficulty of precisely defining terms within the argument and how they should be applied in the real world. In the real world moral choices do not always fall into an easy bucket of morally obligatory, morally permissible, and morally impermissible. We disagree amongst ourselves on a number of moral topics such as abortion, capital punishment, the role of governments, and the "costs" of certain freedoms like the First and Second Amendment. Categorizing our decisions within these buckets is, in itself, an entire debate, of which effective altruism makes some presuppositions. Defining and quantifying *good* may be the most difficult. Effective altruism relies on the work of value theory to establish value claims and comparative good through goodness *simpliciter*, but there is fair disagreement with these

[15] Singer, *The Most Good*, 125.

principles entirely in the philosophy community. For example, can we really say that some acts are "better" than others or just that they are "good?"

Quantifying and Comparing Good

The effort to quantify that decision with metrics such as a DALY presents problems as well. Inherently, there can be no DALY that is truly correct. We often must use averages based on interviews and surveys based on how people think they would make tradeoffs between different conditions such as blindness and deafness. For every person the answer is different, though. More so, we have significant evidence from psychological research pointing to the unreliability of healthy individuals making judgments about what it would be like to suffer from certain conditions.[16] What this leaves us with, often, is precise results from imprecise inputs. For example, for the global disease burden, the DALYs for a certain disease can shift drastically on small changes in the discount rate. For example, certain blind individuals saying that one year of living with blindness is worth 0.8 years of healthy life helps us build a guideline but can it be compared to, hypothetically, different individuals suffering from paralysis saying it is worth 0.75 years of healthy life? Is there not a significant margin of error here for people's inability to accurately pinpoint such a precise figure? These judgements are largely emotional and subject to bias. A more reliable way to build out comparisons would be to find single individuals suffering from more than one disease in order to compare from a single standpoint, but finding such individuals can be quite difficult -- if they exist at all -- for certain comparisons. Additionally, this fails to consider whether we place any significance on an individual having a causal relationship with their burden. A commonly-analyzed population is the inmate population in the

[16] Donald Redelmeier and Daniel Kahneman, "Patients' Memories of Painful Medical Treatments: Real-time and Retrospective Evaluations of Two Minimally Invasive Procedures," *Pain* 66:1 (1996): 3-8, and Donald Redelmeier, Joel Katz, and Daniel Kahneman, "Memories of Colonoscopy: A Randomized Trial," *Pain* 104 (2003): 187-194.

U.S. and the DALYs lost every year to the prison system. While there are certainly non-violent offenses that are over-punished and innocent individuals who are wrongly imprisoned, it would make sense to very few people to suggest releasing all prisoners back into society to reduce the collective DALYs lost.

Focusing on evaluating and producing quantifiable metrics for comparison can also have negative consequences what we give to. Steven Brown provides a critique of the implications of focusing on metrics:

> GiveWell's methodology prioritizes easily measurable metrics, and those outcomes that can also be researched and tracked in a particularly rigorous way. This leads to their most serious drawback: overlooking projects that should be a high priority, but are difficult to measure. When one takes a step back and asks what it would take to better a place that is not doing well, one will surely come across many difficult to measure answers. For example, it is striking that since it began giving ratings in 2008, GiveWell has only recommended a single international educational charity… education is certainly one of the most important things a community must have if it is to rise out of poverty and not merely survive."[17]

Metaphorically, Brown's point suggests that effective altruism is applying bandages to wounds but not taking necessary action to heal the underlying injury. The focus on metrics and evidence points to an issue voiced by Emily Clough in "Effective Altruism's Political Blind Spot." She criticizes effective altruist groups' reliance on randomized control trials (RCTs) saying that an RCT "might determine whether a bed net distribution program lowered the incidence of malaria

[17] Steven G. Brown, "Supporting the Best Charities is Harder than it Seems." Journal of Global Ethics 12, no. 2 (2016): 242, 240-244.

among its target population. But it would be less likely to capture whether the program unintentionally demobilized political pressures on the government to build a more effective malaria eradication program, one that would ultimately affect more people.[18]

In face of all the criticism presented of using RCTs and the DALY as a metric, its theoretical and pragmatic benefits still make sense to Toby Ord, the founder of Giving What We Can, to whom it represents the "best effort so far" to make a comparative metric for healthcare intervention. Additionally, it still remains the WHO method for estimating the global disease burden which provides participants in this realm a consistent basis for data.

Unforeseen Results

One criticism to the highly-targeted donation of effective altruism is the difficulty to account for unforeseen results. As Brown mentioned, focusing on DALYs and RCTs can create a lack of focus on non-quantifiable risks. For example, as Emily Clough mentioned, "The presence of NGOs induces exit from the state sector… the pressure on the government to maintain and improve services eases, and the quality of government provisions is likely to fall."[19] This effect, oftentimes called skimming, can have consequences for the poorest individuals who lack awareness of non-government organizations' (NGOs) efforts to provide services. As a result of the reduced pressure by citizens on the state to provide quality healthcare, education, and protection, the quality of government provision is likely to fall. Clough concludes that:

> [Effective altruists] must contend with the fact that the state remains the primary provider of basic social welfare for most poor citizens in most poor countries, and that pumping

[18] Emily Clough, *Effective Altruism's Political Blind Spot*, Boston Review, 14 Jul 2015.
[19] Clough.

> money into a parallel set of providers—even good ones—without a plan for reaching the coverage or scale of a state may do serious harm to the poor who are left in the state system.[20]

Clough suggests not only that effective altruists be more considerate in their factors, but specifically that they look towards funding advocacy and watchdog groups in certain countries. Particularly in countries where governments don't allow advocacy groups to receive funds from foreign governments but which can accept contributions from private organizations.

Another, more directly attributable risk has been realized in the last several years: Insecticide-treated nets used to combat malaria were also being used by fishermen across the world, particularly in tropical areas such as East Africa where malaria is the most prevalent. There isn't causal data to show whether this is a result of the growing availability of free nets, in particular. Regardless, there is sufficient information to know that this poses a health risk to both humans and the fish stocks. The latter is predominantly because the nets are so finely woven that they catch even the small, young fish. The exact impact of these findings isn't yet known, but it calls for greater research in this area.[21] This particular unforeseen consequence may stand out to effective altruists given the popular ranking of malaria nets as a cost-effective intervention method. GiveWell, for example, ranks the Against Malaria Foundation -- a distributor of long-lasting insecticidal nets -- in its top tier of charities.

---

[20] Clough.

[21] Damian Carrington, *Global use of mosquito nets for fishing 'endangering humans and wildlife,* (The Guardian, 2018).

Personal Moral Judgements

Limiting effective altruism to the second step of the donation process (deciding what to donate to after already deciding to donate to something) still requires us to make some moral judgements in order to compare causes. In a blog post, Holden Karnofsky -- a founder of GiveWell -- put forth three scenarios:

> 1. Prevent 100 deaths-in-infancy, knowing that in all likelihood these 100 people will grow up to have consistently low income and poor health for a 40 year lifetime
>
> 2. Provide consistent nutrition and health care to 100 people, such that instead of growing up malnourished they spend their lives healthy. (Assume for simplicity that their lifespan is also 40 years.)
>
> 3. Prevent one case of relatively mild non-fatal malaria (say a fever that lasts a few days) for 10,000 people without having a significant impact on the rest of their lives

Holden chooses (2) because he is "very excited by the idea of changing someone's life in a lasting and significant way." He rejects (3) because he doesn't think quality of life consists simply in the sum of the quality of days in it. He rejects (1) because he does not put much value on "potential lives," especially when the lives may be filled with health problems.

Holden's decision shows how different people may reach different conclusions based on personal moral judgements. Even if DALYs were assigned to all of these (say option 3 was the highest for example), he may still opt for (2) because of his own moral and logical beliefs. This points to a characteristic of effective altruism -- the lack of a consistent, underlying moral code. On one hand, this leaves individuals with room to apply their own beliefs and become more impassioned in their giving. It also provides a diverse approach of methods and causes. On the

other hand, the lack of a consistent moral philosophy can weaken the effectiveness of the movement and the organizations that represent it. If effective altruist organizations want to be a recipient of donation in order to reroute funds effectively, they must show that they function under some consistent moral code.

One of the large divides within effective altruism lies within how to view animal suffering. While most would agree that animal suffering matters, comparing it to human suffering is a difficult task. This opens up questions such as: Should animal suffering count the same? Do animals consciously process suffering the way that people do? Do they have a capacity for suffering? Do humans suffer differently from one another? The difficulty of answering these underlying moral questions of causes is evidenced by various effective altruist organizations ranking different causes and charities at different levels.

Another significant issue which effective altruists cannot come to agree on is whether to donate or to invest (for future donation). In other words, is it better to donate $1 million today or invest the $1 million to let it grow and be able to donate more money later? The primary arguments for investment are: (1) more money can have a greater impact later, (2) it provides more time to determine what the most efficient investments are, (3) greater ability to fund specific projects -- i.e. if a new intervention is uncovered for a region and needs a certain amount of funding, you can contribute more -- and (4) the bargaining power of more funding allows us to demand greater transparency.

The arguments in favor of donating now are primarily: (1) compounding benefits of current investment -- for example, the benefit of opening a school today compounds as time progresses and the benefit at year five is not equivalent to the benefit if it was opened at year five, (2) network effect -- seeing donation may encourage others to give going forward, (3) the

best donation opportunities may disappear, (4) reduces the risk that you may become less altruistic and not donate all the money later, and (5) deferring donation for investment may become indefinite -- the next year it will always be more money, so when do we actually use it? There are some hybrid solutions to this issue such as giving now to meta-charities like GiveWell so they can appropriately distribute funding in the most effective amounts (due to some of the compounding benefits mentioned earlier).[22]

Normative Incompletion

Stepping away from only evaluating the second step of donation, there is some criticism that effective altruism is incomplete as a normative movement. This is considered in the sense that the same actions and lessons of effective altruism should be applied in other aspects of life. A few notable movements are selective asceticism, ethical consumerism, and the moral market. While effective altruism deems donation the best way to intervene in solving the world's pressing problems, these other movements focus on the non-donation-based actions we can take.

Selective asceticism calls for individuals to take strategic inaction to combat the world's issues. In her paper on selective asceticism as "the other half of effective altruism," Kathryn Muyskens says that, "effective altruists need to take seriously the ways in which their actions contribute to systemic inequality and structural violence. Donation is not enough to create a paradigm shift or to stop systemic injustice."[23] This could include inactions like not eating meat; which could logically be paired with the effective altruist position to support charities against animal cruelty.

[22] Bastian Stern, *Donating vs Investing*, (GivingWhatWeCan, 2012).
[23] Kathryn Muyskens, *The Other Half of Effective Altruism: Selective Asceticism*, (Essays in Philosophy, 2017), 2.

Ethical consumerism would call for individuals to consume things (food, entertainment, etc.) that are ethically sourced. On the vegetarian example, buying vegetables that were grown on farms with fair wages and taking measures to protect the environment.

The moral market is a largely conceptual market (mostly being tested at universities) in which our entitlements can be bought. Entitlements are actions we are permitted to do or not do. From a normative point of view, a moral market could suggest that we ought to pay others to cease behavior that we believe is morally impermissible which they are morally indifferent towards. While this could conflict with a normative view of effective altruism because it calls for us to put our money towards buying entitlements rather than toward donation, I think it can work together with effective altruism in the cases that the moral market is more efficient to accomplish good. In the vegetarian example, this would involve paying people to not eat meat. There are numerous reasons a moral market may be problematic and not pragmatic, but that evaluation falls outside the scope of what we want to accomplish here with effective altruism.

Is it Inhumane?

The concern that there may be something inhumane about effective altruism is a pressing issue for both the philosophy and movement. In "Effective Altruism and the Altruistic Repugnant Conclusion," Gianfranco Pellegrino argues:

> Compared to a possible charitable action ***a***, whereby a given amount of benefit is given to a great number of people, there must be some alternative charitable action ***b*** that, at the same cost, can give a much greater amount of very tiny benefits to a much larger number of people and that, if other things are equal, should be preferred to ***a***.[24]

[24] Gianfranco Pellegrino, *Effective Altruism and the Altruistic Repugnant Conclusion*, (Essays in Philosophy, 2017).

Pellegrino refers to this as the Altruistic Repugnant Conclusion (ARC). His argument against effective altruism is complex -- I will try to simplify it while making an active effort not to create a strawman. At its root, Pellegrino argues that effective altruism, particularly based on Holden's decision (which I have provided above under *Personal Moral Judgements*), can lead us to make choices that would favor small benefits towards many rather than significant benefits to a few. This decision, he argues, is repugnant in that puts minor, unnecessary benefits before the true needs of others. ARC, he argues, is not an unrealistic occurrence for effective altruism either:

> It is perfectly possible that current aid strategies will lead to an overcrowded world of illiterate people, living very poor lives, even though being saved from some days of avoidable morbidity. If this scenario appears to us repugnant, this is not due to its being unrealistic.[25]

Singer briefly addresses this notion when he discusses Thomas Scanlon's argument of 'justificatory weight.' An example is proposed in which a technician for a televised event suffers a serious injury and the only way to help him is to stop the broadcast. Scanlon asserts that it doesn't matter if a billion people are watching; the sum of the smaller pleasures of the many have no 'justificatory weight' compared to the needs of the severely burdened.[26] This seems like a promising workaround to Pellegrino's repugnant conclusion, but it contains some problems in itself. First, defining the 'severely burdened' may be difficult. Are only fatal conditions severe? It also does not address comparing some burdens to others (such as deafness vs blindness) rather than to an absence of pleasure. What both Scanlon and Pellegrino seem to be angling at, though,

[25] Gianfranco.
[26] Singer, *The Most Good*, 120.

is that there are certain utilitarian conclusions that may be associated with effective altruism that defy some intuitive or commonsense morality. More on this in the Evaluation section.

I.g. Evaluation

"It was impermissible to donate to that children's hospital." Framed like this, there is a significant part of me that has a negative gut reaction to some of the conclusions of effective altruism. It seems incredibly difficult (emotionally speaking) to reach some of the conclusions which prioritize certain lives over others based on financial viability. It seems to me that there should be a valid effort to save every life. In a world where many people fail to contribute anything to those in need, saying that someone's decision to do good, such as donating to a children's hospital focused on separating conjoined twins, is morally impermissible seems counterintuitive -- or at minimum unforgiving. This is why it's important to sort out what the philosophy *actually* says as opposed to strawman arguments that overextend the principles of the movement. The real basis of the effective altruist argument is that, all else being equal, we should do more good than less when given the opportunity. There is nothing wrong with making emotional donations in the one-step process we have detailed earlier. Effective altruism, in the form we have laid out, should come into effect if there is a two-step structure and it only says that if we are looking at various recipients for donations, we should pick those that are the most efficient.

I'll elaborate. Having laid out the first-level normative stance of Singer and the second-level stance of MacAskill, it seems wiser to apply the tenets of the second level. It may seem puzzling to say that it is morally permissible to not donate anything (and do no good) yet it is not morally permissible to donate to less efficient causes, but Derek Parfit provides a good example to why this can be true:

> "Suppose that I have three alternatives:
>
> A: at some great cost to myself, saving a stranger's right arm;

B: doing nothing;

C: at the same cost to myself, saving both the arms of this stranger."[27]

Parfit's claim here is that, while one would not be morally obligated to take the great cost to oneself to save a stranger's arm, if one did decide to take on the cost, he or she would be morally obligated to save both arms of the stranger since it poses no additional cost. In light of criticism that this argument is ineffective because there is no tradeoff cost to do C, Shelly Kagan offered another scenario. Facing a burning building, if one made the decision to enter it and found a trapped bird and child (and only one could be saved), one would certainly be obligated to save the child. In this sense, the tradeoffs are made at the second level. The decision to not enter the building is not itself morally impermissible based on the potential harm, but if one undertook the risk, saving the bird would be an impermissible choice.

In this sense, we should implement at the second order of, "I have decided I'm going to give… Now what should I give to?" Isolating to this level allows us greater freedom to apply the principles of effective altruism without having to adhere to a number of other movements and conclusions outside the scope of effective altruism, such as selective asceticism and ethical consumerism.

Let me take this opportunity to offer a realistic donation case that supports this point. If a homeless individual approaches me and asks for money, it is not impermissible to give him money so long as that money doesn't come from funds I would have donated elsewhere. If that meant forsaking my morning coffee when I otherwise wouldn't have, then the net altruistic impact is still positive. One could reasonably argue that that act was still a good act, and not

[27] Derek Parfit, *Future Generations: Further Problems*, (Philosophy and Public Affairs 11, 1982), 131.

wrong at all. Now, if while sitting at home I made the choice that I would go search for a homeless man to give $10, we would be critical of this decision considering the comparative impact of $10 elsewhere in the world -- this money may only buy the man a single lunch whereas it could provide a child immunization to a number of diseases elsewhere in the world. By utilizing effective altruism at the second level, we can avoid making judgements on good acts with emotional motivations (that would not have happened without an emotional stimulus) and focus on the best ways to apply money that has been earmarked for altruistic purposes.

In concluding my evaluation here, it should be said that a true form of effective altruism is effectively impossible and only exists in theory. With the theoretical form of effective altruism we would always know which decisions bring the most good and what their consequences are; every alternative would be quantifiable and wouldn't use flawed metrics like the DALY; and our underlying moral judgements and values wouldn't have a significant impact on which causes we prioritize. Unsurprisingly, we do not live in a world in which we always make the best decisions, let alone a world in which we are able to quantify every aspect of said decisions. In spite of these challenges, there are benefits to a pragmatic form of effective altruism. When confronted with a set of choices, very few would argue that it is wrong to seek to do more good than less. The idea of comparing donation options to do the most good, no matter how imperfect, is better than the alternative of trying not do more good when possible.

I.h. How We'll Apply It

Looking to apply effective altruism means looking to the areas of the world most in need of aid. Per WHO estimates, the global DALYs per 100,000 population sat at 36,331 DALYs in 2015. The WHO Africa region had an estimated 63,349 DALYs per 100,000 in 2015 which was significantly higher than any other region. (The WHO Eastern Mediterranean region ranked second at 38,211 per 100,000.) This figure for Africa is down significantly from 107,065 in 2000 which suggests that actions taken in the last few decades have made significant progress in bettering lives in that region. It also signals that we can continue to target certain areas in order to reduce Africa's disease burden to levels in line with the rest of the world.[28]

Again, I must re-mention that DALYs are a flawed metric for a number of previously-detailed reasons, but I think they still merit some pragmatic use as a guide to identifying which areas carry high burdens relative to the rest of the world and how we can act.

The WHO Global Health Estimates 2015 report shows increases in deaths caused by most cancers globally from 2000 to 2015. Malignant neoplasms, as the WHO classifies them, has grown from roughly 7 million deaths in 2010 to nearly 8.8 million deaths in 2015. Crude death rates (CDRs) and age-specific death rates (ASDRs) are relatively flat or slightly decreasing for most cancers over the same time period. The fight against cancer is certainly an important one, but it is not one to which I can significantly contribute to here. Death counts continue to rise in spite of growth in funding.[29] This isn't to say the funding isn't effective or worthwhile, but only

[28] See IHME-GHE 2016 Data
[29] Eckhouse et al., *Trends in the global funding and activity of cancer research*, (Molecular Oncology, 2008).

to say that cancer research has a step-function structure in which meaningful developments come from breakthroughs rather than the marginal effects of day-to-day activities.

Malaria, meanwhile, has seen a 49% reduction in attributable deaths from 2000 to 2015 (down from 859,000 to 439,000). As a percent of total global deaths, malaria's share has fallen from 1.7% to 0.8%. Accordingly, the CDR and ASDR of malaria have both fallen by more than half. While it is difficult to accurately quantify the exact amount of total funding going to these causes, we can feasibly conclude that malaria has a high level of preventability and reducibility. In the following section we will summarize the attributions of the falling mortality rates of malaria and seek to apply tenets of effective altruism to our findings. The purpose of such an effort is to determine which efforts in the fight against malaria are the most effective and deserve the most backing. This is not to say that there aren't other areas worth consideration. To name just two, the number of DALYs lost to the HIV/AIDS epidemic (45,870 in Africa) or diarrhoeal diseases (44,483 in Africa) merit attention and action.

As we continue to battle the great afflictions facing humanity, such as cancer, I can't help but feel that we've left behind a significant portion of the world that is still afflicted with very pervasive, and very preventable, health threats. For the scope of this essay, I wanted to spend time in the following sections exploring malaria for two primary reasons: (1) it significantly affects young children without the agency to defend themselves -- age is a large reason why the DALY count is so high for malaria -- and (2) there are a large variety of methods being employed that we can look to evaluate.

Chapter II: Malaria

II.a. Disease Background

Epidemiology

"Malaria," as we commonly know it, is a disease caused by plasmodium parasites which are spread through the bites of infected female *Anepholes* mosquitoes, which act as "malaria vectors." Female mosquitoes bite humans to take "blood meals" in order to gain nutrients to develop eggs; male mosquitoes do not bite humans. The plasmodium parasites successfully infect two hosts: humans and female *Anepholes* mosquitoes. There are five parasite species that cause malaria in humans. The two in particular that the World Health Organization (WHO) recognizes as posing the greatest threat are *Plasmodium falciparum* and *Plasmodium vivax*. *P. falciparum*, for short, is the most prevalent malaria parasite on the African continent and responsible for the most malaria-related deaths globally. *P. Vivax* is the most dominant malaria parasite in most countries outside sub-Saharan Africa.[30]

In humans, the plasmodium parasites initially grow and multiply in the liver cells before expanding into red blood cells.[31] In the blood, the parasites destroy the red blood cells and release "meterozoites," that continue the cycle of infecting and destroying other red blood cells. The parasite burden expands logarithmically by approximately ten-fold per 48-hour cycle.[32] The blood stage parasites are responsible for the symptoms associated with malaria. When these "gamocytes" are picked up by the female *Anopheles* mosquito, they start a new cycle within the mosquito. After 10-18 days, the parasites are considered "sporozoites" and found in the

[30] These claims are all detailed in the WHO Malaria Factsheet
[31] Per the CDC's "Biology of Malaria" page
[32] This figure is specifically based on studies of P. falciparum. See Simpson et al., 2002; White, 2004

mosquito's salivary glands. When the mosquito takes a "blood meal" on a human, these sporozoites are injected into the humans which then target the liver cells. And so ensues the cycle. Notably, the mosquito acts solely as a vector and does not suffer from the parasites.[33]

Symptoms

Malaria is an acute febrile illness that carries with it a number of symptoms which can be different for children and adults. In general, symptoms appear 10-15 days after infection. Symptoms of malaria "include fever and flu-like illness… shaking chills, headache, muscle aches, and tiredness. Nausea, vomiting, and diarrhea may also occur."[34] Malaria may also lead to anemia and jaundice as a result of the diminished red blood cell count. If not treated within 24 hours, malaria can progress to severe illness "including kidney failure, seizures, mental confusion, coma, and death."[35] Children, particularly, experience severe anemia, respiratory distress in relation to metabolic acidosis, and/or cerebral malaria. *P. vivax and P. ovale* have potential to relapse up to 4 years later as dormant parasites in red blood cells reengage.[36]

It is not uncommon for individuals in malaria endemic areas to develop partial immunity. This means that infection can be harder to determine based on asymptomatic infections (infections that don't manifest in the usual symptoms).

Impact and Costs

According to the 2017 *World Malaria Report*, an annual document prepared by the WHO Global Malaria Program, nearly half the world's population is at risk of contracting malaria. In

[33] See Appendix A for a detailed illustration of the cycle.
[34] Per the CDC's "About Malaria" page
[35] Per the CDC's "About Malaria" page
[36] Per the CDC's "About Malaria" page

2016, 91 countries reported 216 million cases of malaria and 445,000 deaths due to the disease. 285,000 of these deaths were attributable to children under the age of 5. This figure is down from 440,000 deaths of children under the age of 5 in 2010. The United States saw 1700 cases of malaria, attributed mostly to travelers returning from overseas.

The WHO African Region accounts for roughly 90% of malaria cases and deaths worldwide, with 15 countries carrying 80% of the global malaria burden. All but one of those 15 countries is in sub-Saharan Africa. According to the CDC, "in Africa alone, costs of illness treatment, and premature death from malaria are at least $12 billion per year."[37] Additionally, malaria has been shown to reduce the gross domestic product of endemic countries by several percentage points.[38]

Information regarding impact and costs is primarily obtained through reports from national malaria control programs (NMCPs) in the 94 countries with malaria transmission in 2000. Supplemental data is attained from nationally representative household surveys and databases held by other organizations participating in the Alliance for Malaria Prevention; the Global Fund to Fight AIDS, Tuberculosis, and Malaria (Global Fund); the Organization for Economic Co-operation and Development; Policy Cures; the US President's Malaria Initiative; and WHO. Due to the depth of sources utilized by the WHO in the *World Malaria Report*, this document will serve as a principal source of data and figures for estimating impact, scope, and outlook of various aspects of the fight against malaria.

[37] Per CDC estimates
[38] Per the WHO Malaria Vaccine Position Paper

II.b. Aims

Goals and Outlook

The *World Malaria Report* sets out a number of goals and targets, specifically in the "Global technical strategy for malaria 2016-2030 (GTS)" and the "Roll Back Malaria advocacy plan, *Action and investment to defeat malaria 2016-2030* (AIM)." These goals align with the United Nations' "Sustainable Development Goals (SDGs)" which entails a "plan of action for people, the planet and prosperity."[39] The goals set in the GTS and AIM are targeted for 2020, 2025, and 2030 with a baseline comparison to 2015.

These goals and targets can be found in Appendix B. The most notable targets are to reduce both malaria mortality rates and malaria case incidence globally by at least 40% in 2020, 75% in 2025, and 90% in 2030 relative to 2015 figures.[40] According to the CDC, scale-up of malaria prevention and treatment interventions has saved 6.8 million lives globally from 2000-2015 and mortality rates due to malaria in Africa over that period were cut by more than half.[41]

The goals of GTS and AIM face a number of self-identified threats and pressures including: inadequate funding, evolution of resistance of parasites to drugs and vectors to insecticides, and the interruption of inventions due to complex situations such as insecurity.

With data only going through 2016, it is too early to say whether the targets established in 2015 are on course. Looking back at the period of 2010-2016 can indicate how the significance of malaria is already trending. In that period, the world saw a 32% decrease in the malaria mortality rate (deaths per 100,000 population at risk). In Africa, the mortality rate has

[39] World Malaria Report, 2.
[40] World Malaria Report, 3.
[41] Per the CDC's "CDC and Malaria" brief

declined every year since 2010 and sits at 43% as of 2016. The trend for the world malaria mortality rate isn't as strong, but currently sits at 12.9%. Of the 15 countries with nearly 80% of the malaria burden, Nigeria accounts for 30% of estimated malaria deaths followed by the Democratic Republic of the Congo which accounts for 14%.[42]

Investment

In 2016, an estimated $2.7 billion was invested in malaria control and elimination efforts -- short of the WHO minimum target investment of $6.5 billion. Per the WHO, investment would need to reach the level of $6.5 billion annually by 2020 in order to meet the long-term GTS targets for the fight against malaria. Outside of the affected regions, the United States contributed the most funding: $1 billion (38%), followed by the United Kingdom and other donors including France, Germany and Japan. Governments of endemic countries contributed $800 million (31%) of the $2.7 billion in funding. 57% of international funding was channeled through the Global Fund in 2016.[43]

Total research and development (R&D) funding was estimated at $575 million in 2015 which accounts for 83% of the estimated $686 million annual funding required for R&D to meet GTS targets. 2014 funding was higher than usual due to a large disbursement from the Bill and Melinda Gates Foundation for vaccines. Funding decreased across certain groups, namely: the Australian National Health and Medical Research Council, the Wellcome Trust and the European Union. Among the private sector and US Government Agencies, R&D funding increased primarily in the area of drug development. Over the last 3 years, the three main funding channels have been the US Government National Institutes of Health (27% of total

[42] World Malaria Report, 41.
[43] World Malaria Report, 4.

funding), the Bill & Melinda Gates Foundation (22%), and biotechnology companies (21% collectively). Appendix C provides a thorough breakdown of sources of funding and research areas.

Methods to combatting malaria generally fall in one of three categories: Prevention, Treatment, or Elimination. Prevention and treatment are primarily aimed at reducing incidences of malaria, whereas elimination is focused on methods to eliminate risk of the disease entirely.

II.c. Prevention Methods

The primary methods of malaria prevention currently employed are: insecticide-treated bed nets (ITNs), which includes long-lasting insecticidal nets (LLINs); indoor residual spraying (IRS); vaccines -- particularly the PATH Malaria Vaccine Initiative; and seasonal malaria chemoprevention (SMC).

Bed Nets: ITNs and LLINs

The first major prevention method, ITNs and LLINs, have shown in trials to reduce the deaths of children under 5 by approximately 20%.[44] These bed nets repel and kill mosquitoes and other insects, protecting sleeping individuals from mosquitoes carrying malaria. Per the WHO, 54% of people at risk of malaria in sub-Saharan Africa were sleeping under an ITN in 2016. This figure has increased substantially from 2010 to 2016.[45] Nets are currently the primary method of prevention in the form of ITNs and LLINs, which differ from ITNs in that they are designed to remain effective for multiple years without retreatment. Studies show that LLINs can also save up to $3.8 billion over 10 years by reducing the need for retreatment from 3 to 5 years.[46] Per the CDC, ITNs compound into a community-coverage effect which can protect all members of a community (even those not using nets) when over half of the community uses nets.[47] Net-based prevention methods are a primary driver of the cost-effectiveness of treating malaria and according to GiveWell's analysis of the Against Malaria Foundation (AMF), AMF can produce and distribute LLINs for an estimated $4.22 in most regions.[48] Per the WHO, 582 million ITNs

[44] Per the CDC page on ITNs
[45] See Appendix D
[46] Per the CDC page on ITNs
[47] Per the CDC page on ITNs
[48] Per GIveWell's cost analysis of the Against Malaria Foundation.

were delivered globally between 2014 and 2016. This marks significant growth from the previous three year period (2011-2014) in which 301 million ITNs were delivered. 16 countries in sub-Saharan Africa accounted for over 80% of deliveries in the most recent three year period. ITNs are primarily distributed through mass campaigns (75% of distribution) with additional distribution occurring through antenatal care clinics (13%) and alongside child immunization (5%). The latter two allow for nets to be distributed to the two most malaria-susceptible groups: pregnant women and children.

Indoor Residual Spraying

While the coverage of nets is up significantly from 2010, indoor residual spraying (IRS) has declined significantly over the same period. See Appendix D for the contrasting trend lines. Because mosquitoes have shown tendencies to rest inside domiciles post blood meal, IRS has potential to be effective at controlling indoor mosquito populations. Because IRS kills mosquitoes that come in contact with walls and other sprayed surfaces, it is not a direct method of prevention. Rather, IRS is effective at preventing mosquitoes from transmitting plasmodium parasites to other humans because of the tendency to rest inside buildings after biting individuals. To produce effective results, IRS must be implemented in a high percentage of households, usually over 80%. Indoor residual spraying fell out of popularity following concerns over the environmental impact of DDT, but interest has renewed with the success of IRS in South Africa at reducing malaria incidences by more than 80%.[49]

[49] Per the CDC page on Indoor Residual Spraying.

Vaccinations

Another prevention method receiving funding with high aspirations is malaria vaccination. The only currently approved vaccine is RTS,S which requires four injections and has a relatively low efficacy ranging from 26-50%. RTS,S was developed by PATH Malaria Vaccine Initiative (MVI) and GlaxoSmithKline (GSK) with aid from the Bill and Melinda Gates Foundation. Per the WHO, the median of 4 models of cost effectiveness for the RTS,S vaccination in a 4-dose schedule is $87 per DALY averted (assuming $5 per vaccine).[50] Per the WHO, "These estimates are consistent with the cost per DALY averted for other vaccines in a broad range of developing countries."[51] On the basis of DALYs as a metric, WHO estimates the RTS,S vaccine as an effective method after LLINs and seasonal malaria chemoprevention (escalating preventative measures during malaria season) achieve high usage and coverage. Based on shifting efficacy among age groups, WHO particularly does not recommend the RTS,S vaccine among children of age 6-12 weeks.[52]

Seasonal Malaria Chemoprevention

Season malaria chemoprevention (SMC) targets *P. falciparum* in highly seasonal transmission areas in the Sahel sub-region of sub-Saharan Africa. Following WHO recommendations in 2012, SMC implementation was scaled up in 2015 and 2016 via door-to-door distribution which has proved to be more effective than fixed distribution points.[53] The seasonal malaria chemoprevention method entails "maintaining therapeutic antimalarial drug

[50] Per the WHO Malaria Vaccine Position Paper.
[51] Per the WHO Malaria Vaccine Position Paper, p49.
[52] Per the WHO Malaria Vaccine Position Paper, p50.
[53] World Malaria Report, 19.

concentrations in the blood throughout the period of greatest malarial risk."[54] This is specifically targeted towards children and achieved by using antimalarial drug treatments containing sulfadoxine-pyrimethamine and amodiaquine. Effective antimalarial treatment in monthly intervals "has been shown to be 75% protective against uncomplicated and severe malaria in children under 5 years of age… [and] is cost-effective and safe and can be administered by community-health workers."[55]

Approximately 15 million children were included in SMC programs in 2016 with 60 million monthly treatments administered. This represented a 91% inclusion rate for children in areas where SMC programs were established.[56] This means 13 million children were left uncovered who could have benefitted from this prevention method, with the identifiable reason being lack of funding.[57]

Rapid Diagnostic Tests

Rapid diagnostic tests (RDTs) have helped bridge the gap between prevention and treatment of malaria. RDTs provide a cost-effective alternative to diagnosis from clinical grounds or microscopy which are not always available in remote regions. Usually using a finger prick, RDTs detect specific proteins produced by malaria parasites in the blood of infected individuals.[58] This allows individuals to seek out treatment when effected. From 2010-2016, 1.66 billion RDTs were sold globally by manufacturers.

[54] Per the WHO's position of preventative therapies for Malaria for children
[55] Per the WHO's position of preventative therapies for Malaria for children
[56] World Malaria Report, 20.
[57] World Malaria Report, xiv.
[58] Per WHO page *How Malaria RDTs work*

II.d. Treatment Methods

Artemisinin-Based Combination Therapies

For the treatment of uncomplicated malaria (malaria for which there are symptoms present but lacking in severity and organ dysfunction) from *P. falciparum* and *P. vivax*, the World Health Organization recommends artemisinin-based combination therapies (ACTs). WHO currently recommends 5 different ACTs -- the difference in treatment is dependent on the local strain that is present.[59] ACTs involve the combination of "2 active ingredients with different mechanisms of action… [and] are the most effective antimalarial medicines available today."[60] Proportionate use of ACTs has risen from a median of 39% in 2010-2012 to 54% in 2014-2016.[61]

For the treatment of severe malaria, injectable artesunate should be used for at least 24 hours before administering a 3-day ACT treatment course (if the patient can ingest oral medicines).[62] For children under 6, rectal artesunate should be given when injectable artesunate cannot be. The effectiveness of these treatments is contingent upon the artesunate administration being followed by a 3-day course of ACTs. As ACTs became the first-line treatment in 2014, "the number of courses of ACTs procured from manufacturers increased from 187 million in 2010 to a peak of 393 million in 2013."[63]

The WHO has documented plasmodium resistance to antimalarial drugs in 3 of the 5 human-affecting species: *P. falciparum*, *P. vivax* and *P. malariae*. "Resistance" means a delayed

[59] Per the WHO malaria treatment overview
[60] Per the WHO malaria treatment overview
[61] World Malaria Report, 27.
[62] Per the WHO malaria treatment overview
[63] Per the WHO malaria treatment overview

or incomplete clearing of the parasites from human blood that has been treated with antimalarial drugs. Resistance can include *antimalarial resistance*, *artemisinin resistance*, *multidrug resistance*, and *treatment failure* (improper dosage or poor drug quality, for example). While measuring the efficacy of these drugs is complex, antimalarial drugs have proven to be a highly effective method for treating malaria. Median treatment failure rates for effected and studied countries are usually below 10% with a few exceptions.[64] As mentioned earlier, ACTs also function as a targeted preventative measure against malaria through SMC programs.

Issues

While insecticides and ACTs have proven to be effective measures in the past to curb the mortality rate of malaria, there is growing concern with the long-term effectiveness of these methods due to their selective pressures. This is largely because humans serve as the only significant infectious reservoirs for human malaria parasites.[65] In other words, it does not take long for insecticide-resistant mosquitoes and drug-resistant parasites to reproduce and occupy greater proportions of the mosquito and parasite population.[66] Karen Barnes and Nicholas White explain in 2005 in the journal Acta Tropica:

> After a steady decline in malaria deaths, the malaria mortality rate in eastern and southern Africa has increased dramatically in the last two decades, despite a drop in all-cause deaths among children (Snow et al., 2001; Korenromp et al., 2003). This increase is largely accounted for by the continued use of cheap and widely available drugs, chloroquine and

[64] Per WHO's *Summary of treatment failure rates among patients infected with P. vivax, grouped by treatment and country* and *Summary of treatment failure rates among patients infected with P. falciparum, grouped by treatment and country*

[65] Killen et al., *Going beyond personal protection against mosquito bites to eliminate malaria transmission: population suppression of malaria vectors that exploit both human and animal blood*, (BMJ Global Health, 2017).

[66] Koekemoer et al., *Characterization of multiple insecticide resistance in Anopheles gambiae (Diptera: Culicidae) from Pointe Noire, Republic of the Congo*, (Malaria Journal, 2010).

> sulfadoxine-pyrimethamine, which have become progressively ineffective (Attaran et al., 2004). Antimalarial resistance in Plasmodium falciparum parasites results in an enormous public health burden because of prolonged or recurrent illness, and progression to severe malaria, which is associated with increased hospitalisation and death… For example in Siaya, Western Kenya, 69% of malaria deaths are attributed to resistance to chloroquine treatment (Zucker et al., 2003).[67]

It is difficult to assign efficacies and failure rates to any drug or method at large because of the geographical dependencies of such issues. Subsequently, these methods are still highly effective in certain regions and less effective in others. All things considered, though, and in light of these growing problems with prevention and treatment, it makes sense to look for a permanent solution. This has led to a surge in research towards malaria *elimination*. Our use of the term "elimination" will entail measures taken to eliminate the occurrence of malaria completely which would reduce and, ideally, eliminate the need for prevention and treatment methods.

[67] Karen I. Barnes, Nicholas J. White, *Population biology and antimalarial resistance: The transmission of antimalarial drug resistance in Plasmodium falciparum*, Acta Tropica, Volume 94, Issue 3, 2005, Pages 230-240.

II.f. Elimination

Gene Drives

For 60 years, "gene drives" have served as a beacon of hope (and controversy) for antimalarial efforts. Now, they are becoming reality. A gene drive involves altering the genetic makeup of an organism to include a genetic characteristic not naturally found and utilizing various forms of technology to propagate that characteristic through a population.[68] Unlike a typical genetic characteristic which would be passed down 50% of the time, gene drives are capable of forcing their way into 99% of offspring.[69] One such example is the recently-developed CRISPR--Cas9 -- a bacterial nuclease system that can be directed to use guide RNA to cut almost any DNA sequence with high specificity. The exact mechanisms by which gene drives like CRISPR (clustered, regularly interspaced, short palindromic repeats) operate and can be manipulated biologically are outside of the scope of what needs to be, and can be, covered here. The focus, rather, is on the potential application and risks.

Oxitec's Method

Oxitec, an Oxford-based biotech startup, is utilizing gene drives in the attempt to eliminate the *Aedes aegypti* mosquito in certain areas of the world. Oxitec has developed genetically modified *Aedes aegypti* male mosquitoes whose offspring don't develop past the larval stage.[70] In other words, these mosquitoes occupy the reproductive capacity of a mosquito population and begin to quickly reduce the population due to the short lifespan of mosquitoes. This works via a "self-limiting gene" called the tTAV (tetracycline repressible transactivator

[68] Bull & Malik, *The gene drive bubble: New realities*, (PLOS Genetics, 2017).
[69] Antonio Regaldo, *The Extinction Intervention*, (MIT Technology Review, 2016).
[70] Flavio Devienne Ferriera, *Inside the Mosquito Factory That Could Stop Dengue and Zika*, (MIT Technology Review, 2016).

variant) that is inserted into male mosquitoes. This gene variant, which only works in insect cells, causes offspring to create a non-toxic protein that "ties up the cell's machinery so its other genes aren't expressed and the insect dies."[71] See Appendix E for details. Oxitec is able to breed these mosquitoes in large quantities by feeding mosquitoes an antidote in the breeding facilities. This allows them to create effective breeding centers in local areas where the trials are conducted.

Previous trials in Brazil, Panama, and the Cayman Islands have reduced the population of *Aedes aegypti* by more than 90% – a significant result compared to conventional mosquito-killing methods, such as insecticides which only have a 30-50 per cent efficiency rate.[72] Use of the method in Brazil, in particular, can be a good example. In April 2014, a year after the epidemic of dengue fever (transmitted by *Aedes aegypti* mosquitoes), Oxitec was brought in to run a limited trial. In 10 months of testing in two small neighborhoods of 5,600 residents, the number of cases fell from 133 annually to one case. In regards to ecology, the FDA published "a final finding of no significant impact (FONSI)" for Oxitec's "self-limiting mosquito" methods.[73]

It should be noted that the *Aedes aegypti* is not a primary carrier of malaria, especially in Africa -- that would be *Anopheles* mosquitoes. But, if this method continues to prove successful there would certainly be greater efforts to implement a similar method for *Anopheles* mosquitoes. In its current state, Oxitec's method mirrors preventative measures in that it still requires reapplication. If, and when, it is developed to the point that it can be implemented once to kill a mosquito population and left be, then it can be fully considered an elimination method. This fits

[71] Per Oxitec's "Friendly Mosquitoes" page, see http://www.oxitec.com/our-solution/technology/the-science/
[72] Madhumita Murgia, *Genetically-modified mosquitos released into the wild to wipe out offspring*, (The Telegraph, 2016).
[73] Per Oxitec's program page, see http://www.oxitec.com/programmes/united-states/

into a greater narrative for gene drives to target and eliminate mosquitoes globally, though we are not quite at a point yet when that is feasible.

Target Malaria

Feasibility may be closer than previously thought, though. Target Malaria, an Imperial College project with 16 other institutions, has developed GM *Anopheles gambiae* mosquitoes which are used to implement two strategies. *Anopheles gambiae* is the most prominent malaria vector in Africa. One strategy focuses on using nucleases which target reproductive genes and cause female mosquitoes to become sterile. The other strategy is to skew mosquito reproduction to produce males.

The Fertility strategy, as we'll deem it, works as such: the nuclease identifies fertility genes, cuts through them, and offers itself to be used to repair the stretch of DNA. The altered gene does not affect fertility in females unless there are two modified genes present. Target Malaria's modified mosquitoes, both male and female, each carry one copy of a chromosome with the nuclease and are able to reproduce normally. As the modified gene spreads through the population, females will naturally pick up two sets -- thus preventing their ability to reproduce. This allows a time-delay effect that can last up to 11 generations in order to let the gene drive spread through the population.[74]

The Sex Biasing strategy, as we'll call it, is based on reducing the number of female *Anopheles gambiae* which transmit malaria. Reducing the number of female mosquitoes in populations lowers transmission rates but also reduces the size of mosquito populations by constraining their reproductive capacity. It works as such: the nuclease is attached to the Y

[74] See the Target Malaria page "Our Work," https://targetmalaria.org/our-work/

chromosome in the sperm of male *Anopheles gambiae* and targets unique sections of the X chromosome to cut and fragment. In reproduction, this means only the in-tact Y chromosome can be passed on and the offspring will all be male. (Males have XY chromosomes and females have XX chromosomes.) Attaching the nuclease to the Y chromosome also makes the gene drive self-sustaining since it stays attached to all the male offspring too. Target Malaria's lab experiments show the ability to effectively alter the reproductive output to be 95% male and 5% female for a population.

Target Malaria is backed by some big names. Notably, it has received $44 million in funding from the Bill and Melinda Gates Foundation, the largest investment into gene-drive research as of 2016.[75] The Open Philanthropy Project also rewarded Target Malaria with a $17.5 million grant.

The implementation of Target Malaria's methods may still be a ways down the road. According to a Gates Foundation business plan, the project could be ready to launch by 2029. Fil Randazzo, a deputy director at the Gates Foundation, outlined potential implementation where "buckets of mosquitoes" would be released every 50 kilometers which would begin the desired geographical and genetic chain reaction in a timeframe of roughly two years. Austin Burt, an evolutionary theorist and leader of Target Malaria, believes a drive could spread 5-20 kilometers per year from a release point and could be done with fewer than 500 mosquitoes. [76] From that point, there would remain only 1 percent of mosquitoes (based on their projections). Using nets and sprays while implementing a campaign of drug treatment to kill the parasitical reservoir would, in theory, break the malaria transmission cycle.

[75] Antonio Regaldo, *The Extinction Intervention*
[76] Antonio Regaldo, *The Extinction Intervention*

Other Potential Elimination Methods

There are a number of other new, promising methods from recent research that should be considered and potentially implemented when evidence from research and experimentation deems fit. One method is to disrupt mosquitoes' ability to 'smell' humans. This method can be achieved through gene editing by disabling mosquitoes ability to detect humans via smell, though this has shown to have its difficulties as mosquitoes seem to have other biological work-arounds in place for finding humans. Other versions of this bear more resemblance to insecticidal methods in terms of implementation and usage. Certain molecules and substances have been found to activate mosquitoes Orco coreceptors which can have the effect of flooding a mosquito's olfactory system and disrupting its ability to detect humans. These methods aren't currently ready for widespread implementation but they offer a targeted approach to preventing human infection without taking measures to necessarily kill mosquitoes. Helen Shen reviews it as such:

> Elucidating and learning to hack this complex system could pay off in a big way, as traditional insecticidal approaches have proven to be a "zero sum game," Vosshall explained last November at the Society for Neuroscience annual meeting in San Diego. Multiple efforts since the 1960s have shown that under such extreme selective pressure, the animals rapidly evolve resistance to different pesticides.[77]

Unfortunately, like other current methods, the selective pressure of this method would still favor the mosquitoes that are able to find humans in spite of the chemical effect. The resistant female

[77] Helen Shen, *How do mosquitoes smell us? The answers could help eradicate disease*, (PNAS, 2017).

mosquitoes would lay the most eggs because they would be able to obtain the most blood to develop and nourish their eggs.

Another method is to breed mosquitoes with specific bacteria that prevent it from carrying disease. A notable example is the Wolbachia bacteria, which, if present in *Aedes aegypti* mosquitoes, prevents the mosquito from carrying dengue, yellow fever, chikungunya, and Zika. The bacteria is passed between generations and, like other gene drive methods, creates a crowding-out effect which reduces the number of dangerous mosquitoes.[78] The comparative benefit, though, is that it does not kill the second generation and therefore continues to exponentially crowd out dangerous mosquitoes since the bacteria is passed on. Conversely, this does little to eliminate mosquito populations if that is the overall goal.

So what should the ultimate goal be? This is where we will dive into applying effective altruism to the antimalarial efforts we have laid out.

[78] Michelle Innis, *In Australia, a New Tactic in Battle Against Zika Virus: Mosquito Breeding*, New York Times, 2016.

Chapter III: Application

Prevention and Treatment Analysis

Background

As was mentioned in Part 2, preventative SMC programs did not reach the full population at risk of contracting malaria primarily because of a shortage of funding. Meanwhile, the Open Philanthropy Project also believes Target Malaria has funding bottlenecks are slowing progress. These issues hit the heart of what we hope to accomplish and the questions we seek to answer. Is there a way in which to make the funding go farther and better reduce Africa's burden? Are certain antimalarial methods more viable and efficient than others -- and therefore more deserving of our funding than others? Are certain methods more ethical and does it matter?

At the start of this project, I set out to build a unified model which integrated the relevant costs, efficacies, resistances, DALYs, and other factors of the different antimalarial methods. What I quickly discovered, though, is that building a model to deliver precise judgements meant I needed precise inputs. The inputs we currently have, sadly, are not sufficient to make such precise judgments; doing so, I would argue, would be irresponsible considering the ramifications of such decisions. For example, a misestimate of the rate at which parasites develop resistance could render certain prevention and treatment methods ineffective in parts of the world.

Even beyond the ethical ramifications of building such a model, the complexity of this topic goes far beyond what is feasible to overcome at this level. Geographically, even small communities can have mosquito and parasite populations with drastically different levels of resistance to various methods. Risks associated with certain methods also vary with factors such as education. Collecting dynamic, up-to-date information on education in some of these regions

is frankly infeasible. Measuring the dollar return on R&D is also an extremely challenging prospect given the need to evaluate on an individual project level. Its results often go beyond the scope of one disease or method, too, as they can offer solutions for other areas of research. This doesn't mean that comparing methods is hopeless, but, rather, that we will have to utilize quantitative analysis critically and rely on a heavier mix of qualitative analysis and expert opinion. We can start by evaluating some of the methods we detailed in Part 2. After evaluating the methods, I will make a recommendation on where we should continue to focus.

Insecticidal Methods -- Nets and Spraying

Insecticide-treated bed nets (ITNs), long-lasting insecticidal nets (LLINs), and indoor-residual spraying (IRS) utilize insecticides in order to repel and/or kill mosquitoes. These methods are non-targeted, meaning they are not designed such that their effects are only limited to mosquitoes, or specific mosquito species.

There are a number of reasons that these methods, particularly nets, are commonly recommended by effective altruist organizations like GiveWell. First, insecticidal methods are cheap. As previously mentioned, the Against Malaria Foundation (AMF) can produce and distribute LLINs for an estimated $4.22 in most regions.[79] With the WHO putting an estimated half of the world's population at risk of contracting malaria, not to mention other mosquito-delivered diseases, nets are an extremely effective way to start to help a significant amount of people. The WHO estimates that, in Sub-Saharran Africa alone, 54% of at-risk individuals are sleeping under an ITN. That means there is still significant capacity to distribute nets effectively before they become redundant. With an effective lifetime of 5 years, LLINs also can cover a

[79] Per GiveWell's evaluation of AMF, see https://www.givewell.org/charities/amf

single child for the amount of time at which they are at the highest risk (0-5 years old). The distribution process is also very cost-efficient. Nets are lightweight and space-efficient (collapse to reduce wasted air space), and can be targeted to high-risk individuals by being given out at pregnancy and health clinics. In GiveWell's 2018 Cost-Effective Analysis, the median estimate for cost per death averted (after accounting for leverage and funging) is $3,280 and $4,237 for children under-5. This is based specifically on the Against Malaria Foundation's (AMF) pricing of distribution and efficacy. "Leverage," in GiveWell's usage, refers to "charities causing other entities to spend more on the program than they otherwise would have." "Funging," is the inverse of leverage; it is when they cause less spending from other entities. Consideration of these factors is a response to the "political blind spot" suggested by Emily Clough which we discussed in Part 1. GiveWell's analysts estimated that donating to AMF nets has an 18% funging effect -- meaning the donation is 18% less cost effective than a stand-alone analysis. This is based on consideration of potential reactions of governments, the Global Fund, and unfunded distributions.

These figures can help illustrate the impact of donations and intervention as it can be compared to other interventions, but we must still be skeptical of a unified quantitative metric. In November of 2016, GiveWell's cost per death averted using AMF was $7,500. The last year and a half have not seen a significant change in pricing or efficacy that should drive this change -- especially in consideration that resistant is likely growing.[80]

[80] See GiveWell 2018 cost-effectiveness analysis ('Bed nets' sheet, cells B91:92): https://docs.google.com/spreadsheets/d/1FApXxsBziIo2bcCFBMp9G3F7sDk0VRjzCETbP9J84TE/edit?usp=sharing
And 2016 GiveWell cost-effectiveness analysis ('Bed nets' sheet, cells B78:80): https://docs.google.com/spreadsheets/d/1gVbwg7Og9ClekG-9r2J0lyolNILQ_yDWbEmR7OimCKs/edit?usp=sharing

Nets and spraying both have communal compounding effects, too. When 50% of a community utilizes nets, it has been found to create a coverage effect for the entire community. For IRS, the number is higher: 80% of households need to be sprayed to achieve the communal coverage effect. Spraying can be effectively implemented in public settings, though, to reduce the risk of contraction in higher-traffic areas. Compared to some other methods, insecticidal methods are also easier to understand. This means it is easier to educate communities and local health workers about how to properly utilize these methods and understand their risks. Nets and sprays also give people some agency the fight against malaria. Teaching a child the importance of using their bed nets at nights can reinforce at an early age how important it is to take action to protect oneself and to consider health risks.

Insecticidal methods have their downsides, though. While using nets targets a time when people are most at risk (while they're sleeping), it still doesn't do much to protect children when they play outside, explore, and do all the things children enjoy doing. Nets also cannot protect women and mothers, among all individuals, as they go to get water for the family (keeping in mind that mosquitoes breed around stagnant water). IRS can help protect homes but has similar issues protecting people when they are outside. Nets are also not a targeted method -- i.e. they can affect other species in unintended ways. This could include flies and other insects that don't pose a threat to humans; it could even include fish. As mentioned in Chapter 1 under Unintended Consequences, fishermen across the world, particularly in tropical areas like East Africa, have been using antimalarial bed nets to fish. Again, there isn't causal data to show whether this is a

result of the growing availability of free nets, in particular. Regardless, there is sufficient information to know that this poses a health risk to both humans and the fish stocks.[81]

There is significant concern and growing research showing the growing resistance of mosquito species to insecticidal methods. Menno Smit, MD, a Ph.D. candidate at the Liverpool School of Tropical Medicine who led a new study on resistance, summarizes, "Since 2015, the number of annual deaths from malaria has stabilized. We're not making any more progress. We need new tools…"[82] As my insights have largely been driven by WHO data from 2000-2015, this is a huge insight for recent trends. Smit points to a large issue with insecticidal methods: when these methods are implemented but mosquito populations aren't eliminated, the selective pressures can quickly result in the proportionate population growth of resistant mosquitoes. In other words, if insecticides are prevalent, only the mosquitoes that are resistant will continue to take blood meals on humans and carry the population forward. Nets do have the benefit of still acting as a physical barrier between mosquitoes and humans at night, whereas resistance can make IRS a wasteful effort.

Seasonal Malaria Chemoprevention

Seasonal Malaria Chemoprevention (SMC) is one of the most cost-effective methods we have available to combat malaria. To quickly restate from earlier, SMC is the practice of providing children with full treatment courses of antimalarial drugs during malaria season in high transmission areas. This consists of four separate rounds of treatment. In 2016, 91% of children in areas where SMC programs were established received at least one round of treatment.

[81] Damian Carrington, *Global use of mosquito nets for fishing 'endangering humans and wildlife'*, The Guardian, 2018.
[82] Nadia Whitehead, *What If A Drug Could Make Your Blood Deadly To Mosquitoes?*, NPR, 2018.

81% received two rounds, 69% received three rounds, and 51% received all four rounds.[83] Per the WHO, each round protects a child for roughly one month. That left 13 million children completely uncovered on the basis of a lack of funding to scale it out.[84]

GiveWell's 2018 cost-effectiveness analysis puts the median cost per death averted via SMC at $2,280. This figure is excellent compared to other methods; it comes in lower than GiveWell's estimated cost to save a life using AMF nets. Comparing these figures exactly isn't a great idea, though, based on the amount of uncertainty of some figures. This is why it is more appropriate to organize methods into tiers, and undoubtedly SMC would join AMF nets in the top tier of options.

As a method, SMC has a lot of advantages that allow it to be cost-effective. First, the drugs themselves are relatively cheap (we will discuss this in *Treatment* in the next section). Distributing them during malaria season allows their implementation to be limited to only times at which the method would be effective. This is the opposite of, say, bed nets which are present constantly throughout the year and are constantly degrading and wearing down. SMC also can be paired with almost any other preventative measure. One downside, though, is that it doesn't necessarily reduce the need to implement nets and other prevention methods. For example, providing SMC for a region would not also preclude it from being an effective place to distribute nets in order to cover the contraction risk for the rest of the year, even though it is lower. Another reason for cost-effectiveness is that this method can be provided by community health-workers after a small amount of training. GiveWell particularly recommends SMC programs run by the Malaria Consortium for their ability to actively scale up funding effectively in regards to

[83] World Malaria Report, 20.
[84] World Malaria Report, xiv.

distribution and education. Additional funding is currently needed to expand programs to reach children in Nigeria, Chad, Burkina Faso, and Guinea Bissau.[85]

There are few downsides, though, that are worth considering. Like many other methods, it contributes to the increase in plasmodium resistance. This works differently than the insecticidal methods which favor the resistance of mosquitoes. For SMC, only the parasites that are resistant to the drugs will survive inside of the human host. This means that the total population of people carrying malarial parasites would be carrying a strain with higher resistance that could then be spread. This is why it would be important to pair SMC with methods such as nets in order to reduce transmission.

Treatment

The position on treatment is ethically much simpler from that of prevention. There is a substantial difference between the obligation to take preventative measures against any disease that doesn't have an extremely high contraction rate and the obligation to treat people who have already contracted an illness. In the U.S. it would not be obligatory to extend me significant preventative measures against me contracting, say, a rhinovirus (common cause of the common cold), but if I showed up at a hospital seeking treatment, there would be an obligation to provide a best-effort treatment. (This doesn't mean the U.S. would be unwise to take greater preventative measures that could reduce stress on the healthcare system, but that is another topic entirely.)

Excluding cases of misdiagnosis, funds spent on treatment are theoretically 100% effective. This is in the context that all funding for treatment is applied to malaria cases whereas

[85] Per GiveWell's evaluation of the Malaria Consortium, see https://www.givewell.org/charities/malaria-consortium#Roomformorefunding

funding for prevention may be spent on someone who never would have contracted malaria. It also, as mentioned before, may go towards distributing nets that are misused for other purposes such as fishing and don't even contribute to the communal coverage effect. While it is important to scale up prevention methods, it shouldn't be at the expense of treatment, even if prevention is cheaper. There may be a point at which this would not be true; say, for example, it cost $100,000 to treat one case of malaria but only $1 to prevent it with 90% certainty. In these cases, if there was a possibility of shortage of funds for treatment, prevention may be the ethically acceptable decision. Thankfully, that is not the position we are in for malaria.

Artemisinin-based combination therapy (ACT) dominates as the most prevalent and effective malaria treatment method and is the consensus recommendation of treatment among health organizations. ACTs are fairly cost-effective. There are lots of different price points based on the different combinations used; per the WHO for *P. falciparum* average cost per treatment success for different combinations range from US$2.95 to US$6.97. For P. *vivax* the most effective treatment had an average cost per treatment success of $0.18 and greater than a 99% probability of the cost per success being less than $1.00.[86]

[86] Davis et al., *Cost–effectiveness of artemisinin combination therapy for uncomplicated malaria in children: data from Papua New Guinea*, (WHO, 2011).

III.b. Elimination and Modification

Of all the methods discussed in this essay, I find those suggesting total elimination (specicide) or genetic modification to be the most interesting and worthy of discussion. Let's start with the broad idea to kill all mosquitoes on the face of the earth using a gene drive. If it were feasible, and affordable, would it be worth undertaking? Or are the ethical and ecological implications too great for us to ever go through with it? From the effective altruist perspective, the argument would follow that it is worth eliminating mosquitoes once the benefits are greater than the costs and risks of making such a decision. If a gene drive was the most cost effective and risk-averse method to do so, then that would supersede other methods like insecticides. Of course, as we have discussed, it is impossible to actually know and quantify many of these factors. That is to say nothing of the ethical argument surrounding 'specicide.' Since eliminating mosquitoes would have a global impact, it presents a number of political and legal barriers such as: Which country or countries get to make the decision? Would it need to be unanimous among all countries? What if there are large dissenting groups of people within those countries? The list of political and legal challenges in undertaking such an initiative is overwhelming, to say the least, and, from a cynical perspective, impossible to answer properly.

From an interview about genetic elimination, Andrea Crisanti, an Italian parasitologist and genetic engineer, was asked the question, "is it ethical to eliminate any part of nature on purpose?" He responded (including interviewer paraphrasing):

> "Are you asking in a Darwinian way or a theological way? I think it's a species competition between us and the mosquito. And I don't think a species has the right to exist or not to exist." He says what species do have is "fitness"—they have adapted to flourish in their environmental niche. For species we hope to save, we might use gene

> drives to add beneficial genes, like ones for disease resistance. For species we despise, we can add ones that make them unfit for survival.[87]

There is no correct answer to this question, as it relies on underlying moral philosophy -- and we have yet to find a unified and undeniable theory of moral philosophy.

From a broad perspective, using gene drives on a global scale forces us to consider a number of risks and potential issues. First, it's very possible that a gene drive, such as CRISPR, would fail due to the resistance of some mosquitoes. In their review, "The gene drive bubble: new realities," James Bull and Harmit Malik assess:

> Simple changes in the target sequence can block the CRISPR nuclease, in turn reducing the rate at which heterozygotes become homozygotes. Champer et al. show that resistance to a CRISPR gene drive engineered in the germ line of D. *melanogaster* can arise rapidly, often in a single generation.[88]

Concerns of resistance may be abated, though, by the possibility of combining drives. For example, the Target Malaria project has suggested combining several drives targeting three different DNA sites simultaneously. Tony Nolan, who works on the Target Malaria project, speculated that, "Mosquitoes might eventually evolve resistance to all three, but maybe not before they're all dead."[89]

The ability to make these gene drives so powerful, though, is concerning in itself because it may make them harder to stop if and when we want to. Also, their potential to be applied to other species is no laughing matter; the FBI has been investigating the potential for gene drives to be

---

[87] Regaldo, *The Extinction Intervention*
[88] Bull and Malik, (2017).
[89] Regaldo, *The Extinction Intervention*

used as a bio-weapon against humans.[90] There is concern that a genetic eradication campaign on this level could become a slippery slope. If we eradicate mosquitoes, can we eradicate any species we deem to have a negative impact? Does that negative impact have to be on the collective environment or just on humans? It wouldn't take an overly critical eye to realize that as a species we haven't always acted in the best interest of the environment up to now. Because of the impact these drives can have, the risk of accidental release has become a concern. In August of 2016, 27 scientists penned a letter to *Science*, a peer-review journal, urging for "stringent confinement strategies" and calling on scientists to refuse requests to share their genetically-modified organisms until some kind of rules had be figured out.[91] The ability for small groups of people, even a team of two, to change an entire species using CRISPR and other gene drives should still be troubling. What happens when a team decides to bypass regulatory and legal roadblocks and decides to start releasing GM organisms from its lab into nature? Do we need to be policing scientists that in any way deal with gene drives?

We must also consider the effect that adopting gene drive methods may have on donation-based funding. Would people be less likely to support methods that they may disagree with ethically? Certainly some would. Could this be avoided by choosing methods that make mosquitoes sterile rather than methods that flip a "genetic kill switch?" Perhaps preventing mosquitoes from reproducing would be viewed as more acceptable than taking an active effort to kill the mosquito itself. If mosquitoes can feel pain (though there's not much evidence that they do), then reproductive prevention seems like the more humane option of the two gene drives. This illustrates a strength and weakness of centralized donation foundations such as the Gates

[90] Regaldo, *The Extinction Intervention*
[91] Regaldo, *The Extinction Intervention*

Foundation: they can make investments that not all individuals would be comfortable funding. Of course, for those who support the methods this is a plus -- we can bypass the less-informed public opinion to utilize the most effective methods. On the other hand, centralized donation could also make unpopular choices that turn out to be wrong and have negative impacts on the very people who opposed them in the first place.

All the issues mentioned up until now focus on the risks associated with gene drives even if we get them right. But what about the ecological impact of eliminating mosquitoes, specifically? This is an area of debate, and we cannot know for sure what the impact would be; what we should look to consider is whether the potential risks and uncertainty is significant enough to outweigh the current impact of mosquitoes. Not all mosquitoes bite humans; in fact, of the roughly 3500 named mosquito species only females from 6% of species bite humans.[92]

Because males do not produce eggs and therefore need less energy, they are able to obtain sufficient energy from feeding on plant pollens. This makes them important pollinators across the globe. That being said, Janet McAllister, a medical entomologist for the CDC, says they aren't crucial pollinators for any crops on which humans depend.[93] Mosquito larvae are also an important food source for some animals such as frogs and fish. Per Richard Merritt, an aquatic entomologist cited in "Ecology: a world without mosquitoes,":

> "Mosquitoes are delectable things to eat and they're easy to catch." In the absence of their larvae, hundreds of species of fish would have to change their diet to survive. "This may sound simple, but traits such as feeding behaviour are deeply imprinted, genetically, in those fish," says Harrison. The mosquitofish (*Gambusia affinis*), for example, is a

[92] Claire Bates, *Would it be wrong to eradicate mosquitoes?*, BBC News, 2016.
[93] Janet Fang, *Ecology: A world without mosquitoes*, (Nature, 2010).

> specialized predator — so effective at killing mosquitoes that it is stocked in rice fields and swimming pools as pest control — that could go extinct. And the loss of these or other fish could have major effects up and down the food chain.[94]

Mosquitoes serve as a food source for bats and birds too, though there is much greater skepticism that they are significant for those species. This is because outside of larvae, mosquitoes only provide a very small amount of nourishment. For example, only 2% of the gut content of bats is comprised of mosquitoes; ecologists believe that bats and birds would quickly substitute mosquitoes with other insects -- particularly if other insects filled the ecological space left by mosquitoes.[95] Perhaps one of the greatest, and least quantifiable, arguments is the possibility of gene drives "jumping" species to affect non-intended targets such as other insects or even humans. Species jumping hasn't been observed yet but remains a commonly cited concern considering the theoretical possibility of these bacteria nuclease systems moving to other hosts. There is also evidence that eliminating one species of mosquito can lead to another species replacing their role as vector-carriers. For example, researchers in Jacksonville, Florida found that Asian tiger mosquitoes had inseminated yellow-fever mosquitoes -- effectively sterilizing them -- which allowed them to overtake their niche. Both mosquitoes are vector carriers.[96]

An interesting consideration is where mosquitoes currently serve as active barriers to negative human behavior. Science writer David Quammen has argued that mosquitoes have made tropical rainforests virtually uninhabitable to humans and delayed global deforestation. "Nothing has done more to delay this catastrophe over the past 10,000 years, than the mosquito," Quammen says. Evaluating this claim brings to light why it is important to lay out what theory of

---

[94] Fang, *Ecology: A world without mosquitoes*
[95] Fang, *Ecology: A world without mosquitoes*
[96] Fang, *Ecology: A world without mosquitoes*

good we must use. From the *good-for-the-universe* point of view theory expressed in Schroeder's Value Theory, which has also been deeply covered by a number of moral philosophers including Peter Singer, mosquitoes may actually be viewed positively based on this role of safeguarding the environment from humans. In my opinion, this actually exemplifies why we can't use the *good-for-the-universe* point of view of good in order to guide antimalarial efforts. Based on the environmentally destructive impact of humans historically, this understanding of *good* would be further extended to recommend the end of all efforts to protect and extend human life. Any natural phenomenon, such as viruses and disease, would limit the destructive impact we would have. Such a recommendation is one that I am not willing to make, nor should we. For this reason, we will err on the side of a human-focused agglomerative theory while maintaining a consideration for other species and the environment at large. While there may also be intuition to think that the hundreds of thousands of lives saved may place a greater burden on already-struggling communities, eliminating mosquitoes and malaria would likely reduce stress on healthcare systems and return some of the GDP loss associated with malaria (estimates are generally between 1% and 2% of GDP lost due to malaria) per Jeffrey Hii, a malaria scientist for the WHO.[97]

Many arguments against eliminating the mosquito are focused on the global ecological impact of eliminating all mosquitoes. Most eradication campaigns currently considered, though, are focused on only a few species which largely impact humans. As mentioned, only females from 6% of mosquito species bite humans, and only half of these carry parasites that cause human disease. Of course, as we have seen, the impact of these 100 or so mosquito species is massive. We can abate some of the aforementioned criticisms against eradicating mosquitoes by

[97] Fang, *Ecology: A world without mosquitoes*

limiting our efforts to these particular species. By leaving most mosquito species in place, we can limit our effect on the ecology of many environments that are of concern. There is also reason to believe in the resilience of these environments. "We're not left with a wasteland every time a species vanishes" says biologist Olivia Judson, who has called for targeted extinction of some mosquitoes.[98]

Before diving into the specific campaigns currently under consideration or being employed, let's lay out the reasons that support gene drive methods. First, because gene drives can be so narrowly targeted, this actually gives us a greater ability to only get rid of mosquitoes that harm humans while leaving others untouched. In particular this stands in comparison to widespread insecticidal campaigns because insecticides can be harmful to all mosquitoes as well as other species of insect. Of course, many mosquitoes are vectors for diseases other than malaria such as yellow fever and dengue, which are responsible for 703,902 and 75,572 DALYs, respectively, in Africa. This could be viewed as an advantage to both insecticidal and gene drive methods which can target mosquitoes which carry human-affecting diseases, but serves as a limitation to malarial medication methods such as SMC and general treatment.

Another benefit to gene drives is that they don't require a change in human behavior the way nets, sprays, and medication do. The ability of gene drives to be implemented and let to run their course helps highlight why other methods haven't yet eliminated malaria. "Malaria is a problem of poverty, of instability and lack of political will. We are asking the drive to do what we can't do politically or economically," says Crisanti regarding the Target Malaria drive.[99] Per Antonio Regaldo, a writer for the MIT Technology Review, in "The Extinction Invention":

[98] Bates, *Would it be wrong to eradicate mosquitoes?*
[99] Regaldo, *The Extinction Intervention*

> "If it works, it will be incredibly cheap, easy to distribute, and egalitarian, benefiting everyone, rich or poor. It will also keep working once released, avoiding a common problem: often, the most difficult part of eradicating a disease is the endgame, when attention wanders elsewhere and spending per case skyrockets."[100]

We'll delve later into evaluating the specifics of the Target Malaria project, but staying on gene drives broadly, Crisanti and Regaldo make good points. One potential issue with effective altruism is how to decide who receives the benefits of intervention. Say, for example, there are 200 people at equal risk of malaria and they are all the same age. In other words, pretend the quantifiable DALYs are the same for all. Now assume that we can only distribute 100 nets, how can we decide who to help? Do we flip a coin? Gene drives can partially help overcome this issue. While we still may be faced with tough choices of which communities to start implementation in first, it leaves the chance of contraction to the natural probability of nature. In other words, we don't have to make potentially discriminatory choices over who sees the greatest benefit.

Gene drives do not necessarily mean eliminating the mosquito, either. Modification without eradication may be a viable pathway with potentially fewer, albeit different, ethical concerns. On one hand, modification means we can avoid "specicide." This helps overcome some of the ecological fears of eliminating mosquitoes. Yet, in a sense, modification can seem like a more slippery slope than elimination. Because the action is less extreme, it may be easier for us to justify modifying other species -- whereas we are much more hesitant to suggest annihilating another species (besides the mosquito) altogether. This could be because we have

[100] Regaldo, *The Extinction Intervention*

eliminated, or tried to eliminate, species before -- both unintentionally and intentionally -- but have never attempted genetic modification of animals on a large scale. For example, humans have launched campaigns to eradicate mosquitoes when building the Panama Canal in the early 1900s, larvacides in Brazil in the 1940s to fight malaria, DDT in the U.S. to rid the U.S. of malaria by 1949.[101] Some possible methods for modification, as we introduced earlier, are modifying the ability of mosquitoes to "find" humans and making *Aedes aegypti* carriers of the Wolbachia bacteria which prevents them from carrying yellow fever, dengue, and zika. For some modification methods, the goal still seems to be to make it difficult for mosquitoes to bite humans. Preventing them from feeding would still have the same long-term effect that either the mosquito populations will die out, or they will adapt to overcome the modifications. We must therefore be wary of expensive methods that seek to promote the same consequence as more cost-effective alternatives.

Having laid out the broad reasons in support of and against elimination and modification of mosquitoes through genetic means, there are still a number of specific interventions to evaluate. Oxitec's GM strategies in Brazil can serve as a good touchpoint to get an idea for cost-effectiveness of gene drives that attempt to eliminate mosquito populations. In a two-neighborhood trial, Oxitec says that to protect 5,600 people, it has been releasing three to four million mosquitoes a month. The company has not said how much the mosquitoes cost, but the government officials in Piracicaba, Brazil have said they expect to pay about 30 Brazilian reals, or $7.50 a year per person protected. That figure is approximately what the health department currently spends on sprays, larvicides, and associated costs like sick leave.[102]

---

[101] Janet Fang, *War against the winged*, (Nature, 2010).
[102] Flavio Devienne Ferriera, *Inside the Mosquito Factory That Could Stop Dengue and Zika*, (MIT Technology Review, 2016).

This is more expensive per person than protection via a net, but it has unique advantages. First, the protection is not limited, theoretically, by location. I.e. while you have to be inside a net for it to be effective, this method extends coverage to every nook and cranny of the city. Additionally, the cost of this method should actually scale with the physical size of the city (square mileage) rather than the number of people therein. In other words, while a highly dense city would need more nets, high population density would actually increase the cost-effectiveness of this method. If the method could reach a point that no longer needs consistent application, all the future lives in that area would also see the benefit of implementation.

We must consider, though, that Oxitec is a for-profit company. Funding businesses rather than charities or governments may mean that certain areas that are viewed as less profitable for implementation are ignored. This could include areas without infrastructure like roads, areas where people are not clustered together, or areas where the governments cannot afford to pay for the service. Conversely, foreign aid could provide funding to target areas that can't afford it while the business itself targets areas that can afford it (Brazil, Cayman islands, Florida, etc.). Additionally, we must consider whether the $7.50 figure is the actual cost of Oxitec's strategy, or whether that is the markup for the service. The figure could even be below-cost in order to gather important trial data and showcase the value of the technology.

While Oxitec is targeting reduction of Zika, yellow fever, dengue, and chikungunya, Target Malaria offers a vision to a malaria-free world using gene drives. While it may not be ready until 2029,[103] a representative has said that the Gates Foundation now views gene drives as "necessary" to end malaria and that it will be ready years before an effective vaccine could.[104]

[103] Per a Gates Foundation business plan, see Regaldo, *The Extinction Intervention*
[104] Per OpenPhil's report on Target Malaria, see https://www.openphilanthropy.org/focus/scientific-research/miscellaneous/target-malaria-general-support

Per the Open Philanthropy Project (OpenPhil), Target Malaria is the most likely project to develop a "fit-for-field construct first."[105] The project currently receives about $10 million per year in funding from the Gates Foundation to develop the gene drive technology itself for the Fertility and Sex Biasing strategies we outlined earlier. Despite being the largest research group working on a gene drive for malaria control, there is still room for effective funding to aid other areas of the project.

Using a basic analysis model built out by the Open Philanthropy Project, which I have updated to reflect recent figures, we can see the potential impact of donations to Target Malaria.[106] The following is taken from my updated version of OpenPhil's model:

Quantity	**Estimate**
Size of Open Phil grant to support Target Malaria	$18,700,000.00
Rough cost per life saved by donating to AMF	$3,280.00
Projected number of lives saved by donating this amount to AMF	5701
Malaria deaths per day in Sub-Saharan Africa	1717
Discount for several factors**	10
Estimated deaths averted per day of timeline acceleration	171.7
Number of days of timeline acceleration required to equal lives saved via AMF	**33.20**

Essentially what we see here is a series of calculations to see how many days OpenPhil's grant of $18.7 million would need to speed up the timeline for Target Malaria in order to avert the same amount of deaths that the funding could achieve with AMF bed nets. The most interesting piece of the model is the "Discount for several factors" in the left column. This is OpenPhil's effort to recognize that the current 1,717 daily deaths from IHME data would not be the same when

[105] Per OpenPhil's evaluation of Target Malaria, see https://www.openphilanthropy.org/focus/scientific-research/miscellaneous/target-malaria-general-support

[106] For OpenPhil's original model, see https://docs.google.com/spreadsheets/d/1E8qu474nUUvPjK21oBqGkdQ9FuXanigOgSGfumhL3_c/edit?usp=sharing

Target Malaria is finally deployed. Reaching the value of 10 came from discounting future funds by 3% annually. This means funds spent today for benefits in 10 years would be reduced by 25%. Additionally, the malaria burden in 10 years will likely be smaller given the other current efforts we have discussed. It is also possible that Target Malaria will not be the project to actually launch the antimalarial gene drive first. These uncertainties are also factored into the cumulative discount of 10. This leads us to estimate that at the time this technology is ready, it could save 171.7 lives per day rather than the 1,717 that are currently lost per day. This means that accelerating the timeline by roughly 33 days would avert as many deaths as a present-day donation to AMF which has a rough cost per life saved of $3,280.

This strikes at the core of what I discussed in the evaluation in Chapter I: a desire to make precise conclusions from imprecise data. This is not to say this model has no validity; to the contrary, I think it still effectively demonstrates the potential for funding to have a huge impact in saving lives through Target Malaria. The conclusion I would not be comfortable drawing, though, is that this demonstrates that we should abandon funding for other efficient causes.

Thirty-three days doesn't seem like a particularly absurd figure for $18.7 million to accelerate a timeline by in the context of 10 years. Let's look where that money would go and see if that intuition holds true. The Gates Foundation primarily funds the gene drive research, which leaves funding needs in the following areas, per OpenPhil:

1. Training graduate students and postdoctoral researchers in Burkina Faso, Mali, and Uganda
2. Designing materials and funding outreach to key regulators
3. Recruiting additional support to focus on shareholder engagement on an international level

4. Funding metagenomic analysis to better identify the ecological importance of the *Anepholes gambiae*
5. Developing methods to facilitate easier rearing and release of modified mosquitoes (such as better sex sorting)
6. Developing protocols that aim for maximally efficient release (conducting field trials)

OpenPhil believes that some of these areas may be bottlenecking the project. While we can't say for sure, it seems reasonable to think that $18.7 million could help accelerate processes such as the outreach to key regulators who could speed up regulation processes or be quicker to allow implementation of Target Malaria's strategies. Even a few months, as we can see, could have substantial effects.

Chapter IV: Recommendation and Conclusion

Recommendation

In my research thus far, the only thing I can say with 100% certainty is that there is no one-size-fits-all solution to malaria. We face significant challenges including, to name a few: (1) the geographic differences as certain regions host different species of malaria vectors and parasites with varying resistance to different methods, (2) the difficulty of deriving factors to balance future impact vs present impact, and (3) the issues with using quantifiable comparison and metrics which can miss or ignore other important interventions such as education. That being said, my goal here has been to call attention to certain methods that generally appear to carry higher potential for effective intervention. In spite of these challenges, I think we can still do that.

On the topic of education as a blind spot, I believe reducing health issues like malaria creates a foundation on which education can be built. When students are sick and schools aren't safe havens from disease, the willingness to embrace education can be significantly hurt.

Before all else, we must understand that a recommendation to move all funding to one method would be futile. Funding has marginal utility which decreases at a certain point for every additional dollar that can't be utilized to its maximum potential. At some point, the funds are wasted and the alternatives become more effective. Additionally, there is value to taking a diverse approach. Much like how the Target Malaria project will likely implement multiple gene drives simultaneously to reduce the probability of mosquito resistance to all drives, we can say that utilizing a diverse approach in antimalarial efforts protects us from the downside risk that one method may fail.

While the value of research and development (R&D) is difficult to quantify, we can certainly see with growing resistance to treatments and insecticidal methods that research to counter this resistance and to create substantial gene drives could have a significant impact beyond the cost of conducting said research. At minimum, I would recommend that a steady percentage of total funding going towards malaria is applied to R&D which is focused on gene drives and countering mosquito and plasmodium resistance to current prevention and treatment methods.

Outside of R&D, I would recommend that we continue to fund efforts that fall within the top cost-effective tier of evaluation. Particularly, LLIN distribution and SMC programs by AMF and the Malaria Consortium, respectively, are two charities within this top tier. With only 54% of at-risk individuals in sub-Saharan Africa sleeping under ITNs, there is still significant room for deployment and distribution of nets. I would recommend increasing funding to ramp up SMC programs. As we discussed, additional funding to the Malaria Consortium's SMC programs could extend coverage to 13 million more children, particularly in Nigeria, Chad, Burkina Faso, and Guinea Bissau.

As we saw from OpenPhil's model, funding to gene drive projects like Target Malaria can have significant returns on par with other prevention methods even when only accelerating implementation by short amounts of time. I would recommend that early funding to Target Malaria go towards funding metagenomic analysis of the *Anopheles gambiae* in order to identify any potentially negative ecological impacts of the gene drive. This would give us enough time to potentially correct the method or else to implement counter-measures in order to sustain ecosystems. For example, if there was concern of reducing the mosquito as a food source to other animals, we could look at another gene drive which will expand another substitute insect

population to take over the space left by the mosquito. Hopefully this would occur naturally, but funding these studies would help us obtain a higher degree of certainty. While Target Malaria's goal is to be implemented on a massive scale, late-stage trials (after safety is not a concern) should focus on high population density areas in order to stretch the cost-effectiveness and return from implementation.

Considerations

Of course, there are risks we must consider. While Target Malaria has pledged to work with local and national governments, and that the decision to implement the method will ultimately be up to them, is the project really prepared for countries to potentially reject the gene drive strategies they have developed? We have discussed the difficulty of containing such methods to borders; what could the results be from implementing in neighboring countries who disagree on its usage? What if, half-way through the expected 10 year timeline to launch Target Malaria, a better solution comes along? This, of course, should be factored into the discounts applied to valuing this method, but doing so accurately would be still rely on subjective factors.

Conclusion

Malaria continues to impact the lives of millions, and disproportionately takes the lives of young children in sub-Saharan Africa. While there are a number of valid causes that fit within the top tier of worthiness, this is undoubtedly one of them. I am excited by the prospects ahead for solving and combatting the malaria crisis as we continue to better leverage available data and technology. In a time when we continue to solve some of the world's most complex problems, I can't help but feel that there is a group that has been left behind to suffer from one of the most preventable diseases in the world. While there are a number of ethical and practical hurdles we

must overcome in doing so, I am confident that a better future awaits us when we choose to do more good than less.

Appendices:

Appendix A: The Malaria Cycle

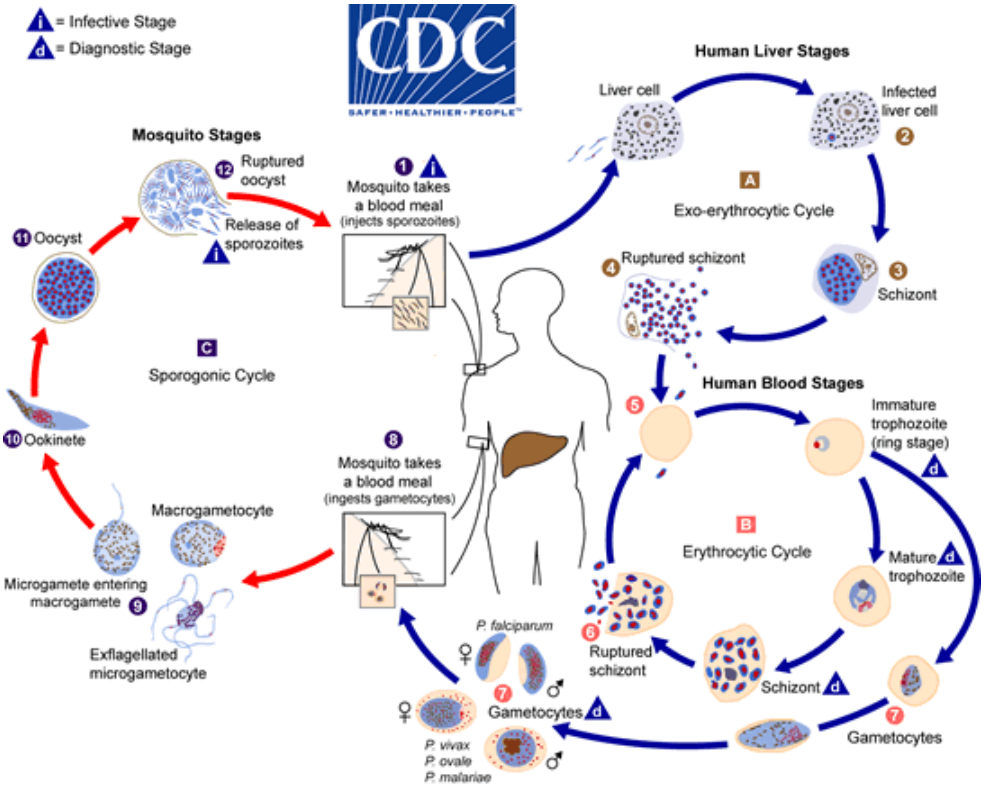

Appendix B: Global targets for 2030 and 5 year milestones

TABLE 1.1.

GTS: Global targets for 2030 and milestones for 2020 and 2025 (*1*)

Vision – A world free of malaria

Pillars	
Pillar 1	Ensure universal access to malaria prevention, diagnosis and treatment
Pillar 2	Accelerate efforts towards elimination and attainment of malaria free status
Pillar 3	Transform malaria surveillance into a core intervention

Goals	Milestones		Targets
	2020	2025	2030
1. Reduce malaria mortality rates globally compared with 2015	At least 40%	At least 75%	At least 90%
2. Reduce malaria case incidence globally compared with 2015	At least 40%	At least 75%	At least 90%
3. Eliminate malaria from countries in which malaria was transmitted in 2015	At least 10 countries	At least 20 countries	At least 35 countries
4. Prevent re-establishment of malaria in all countries that are malaria free	Re-establishment prevented	Re-establishment prevented	Re-establishment prevented

GTS, *Global technical strategy for malaria 2016–2030*

Appendix C: Sources of R&D Funding and Research Areas

FIG. 2.4.

Investments in malaria research and development by source and by research area[a], 2010–2015 (in US$ million) *Source: G-FINDER Public Search Tool Policy Cures.* https://gfinder.policycuresresearch.org/PublicSearchTool

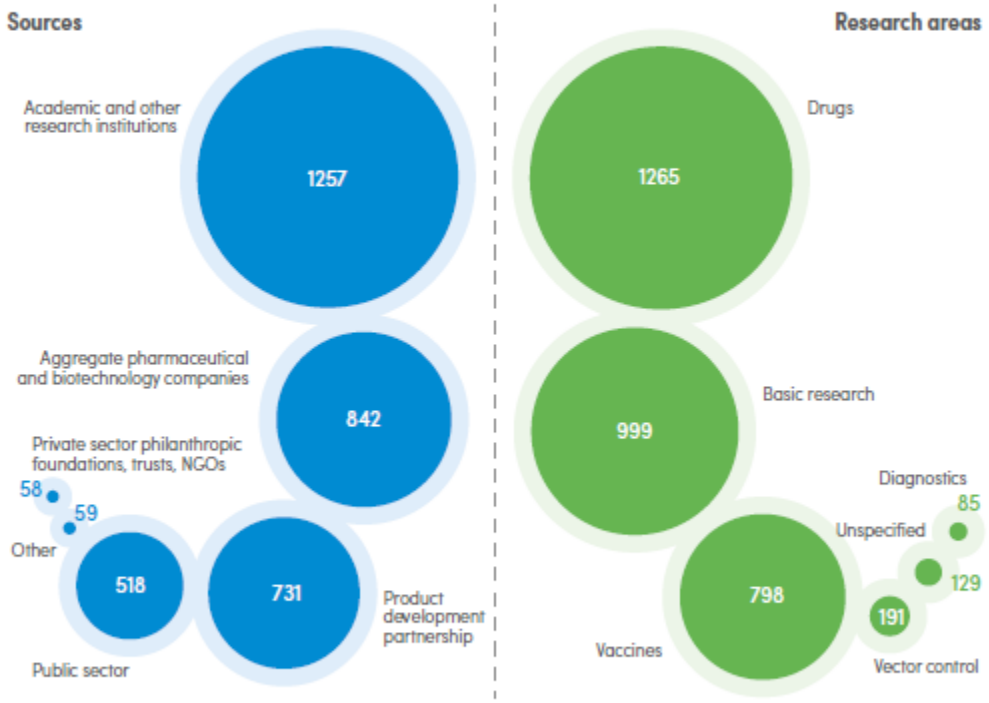

NGO, nongovernmental organization

[a] Public sector category includes governments, government agencies and government-affiliated research institutions.

Appendix D: ITN coverage and IRS coverage

FIG. 3.1.

Proportion of population at risk with access to an ITN and sleeping under an ITN, and proportion of households with at least one ITN and enough ITNs for all occupants, sub-Saharan Africa, 2010–2016
Source: Insecticide-treated mosquito net coverage model from Malaria Atlas Project[1]

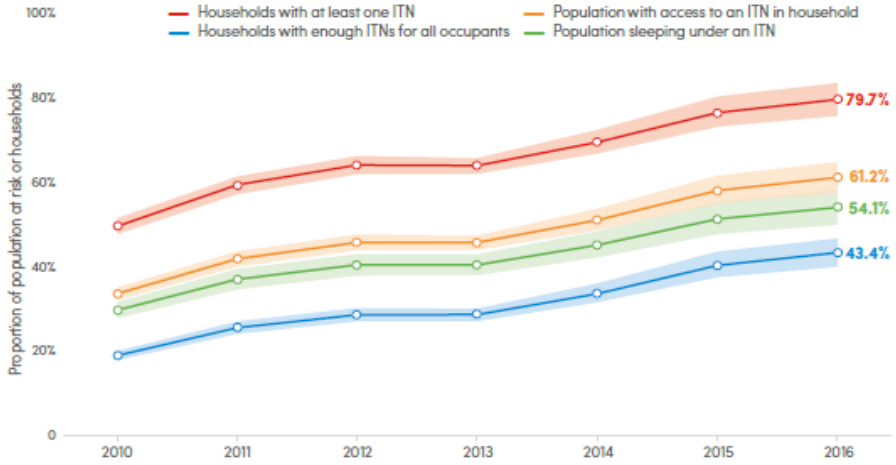

ITN, insecticide-treated mosquito net

FIG. 3.2.

Proportion of population at risk with access to an ITN, sub-Saharan Africa, 2010–2016 *Source: Insecticide-treated mosquito net coverage model from Malaria Atlas Project*[1]

FIG. 3.4.

Proportion of the population at risk protected by IRS by WHO region, 2010–2016 *Source: National malaria control programme reports*

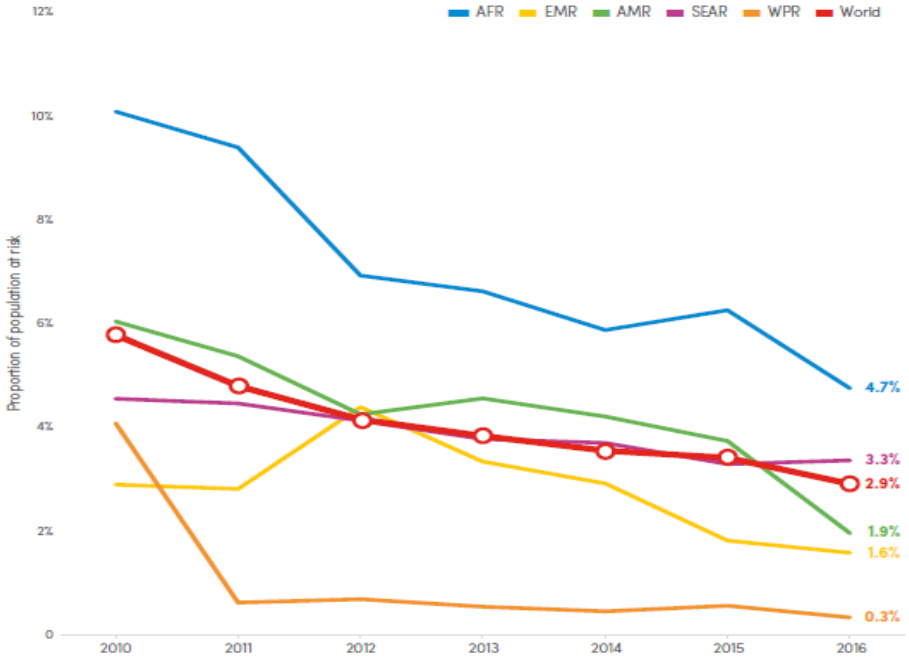

AFR, WHO African Region; AMR, WHO Region of the Americas; EMR, WHO Eastern Mediterranean Region; IRS, indoor residual spraying; SEAR, WHO South-East Asia Region; WPR, WHO Western Pacific Region

Appendix E: self-limiting gene technology

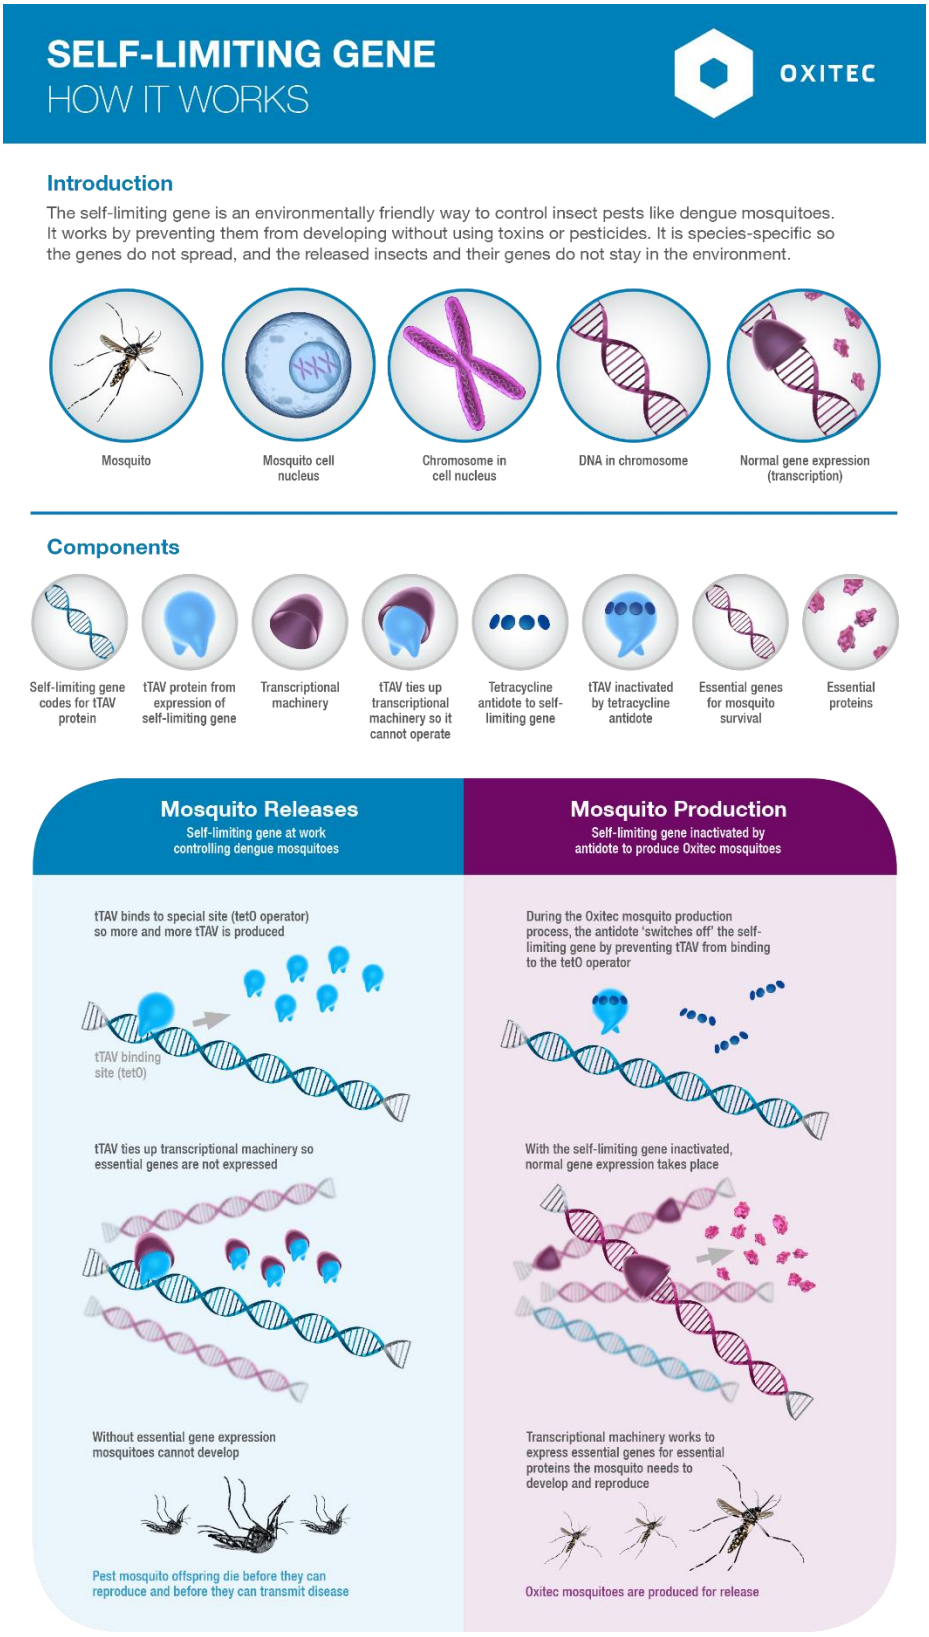

Works Cited

A.B. "Why Gene Drives May Never Eradicate Diseases or Pests." *The Economist*, 11 Oct. 2017, https://www.economist.com/blogs/economist-explains/2017/10/economist-explains-9.

Barnes, Karen I., and Nicholas J. White. "Population Biology and Antimalarial Resistance: The Transmission of Antimalarial Drug Resistance in Plasmodium Falciparum." *Acta Tropica*, vol. 94, no. 3, June 2005, pp. 230–40. *ScienceDirect*, doi:10.1016/j.actatropica.2005.04.014.

Bates, Claire. "Would It Be Wrong to Eradicate Mosquitoes?" *BBC News*, 28 Jan. 2016. *www.bbc.com*, http://www.bbc.com/news/magazine-35408835.

Berger, Alexander. "Public Priorities for Open Philanthropy Project - US Policy." *Google Docs*, https://docs.google.com/spreadsheets/d/1NpcmVO3DnOmtKRbYGkNbR0NTitclT5mwGBwBEwYlKHg/edit?usp=sharing&usp=embed_facebook. Accessed 31 Oct. 2017.

---. "Rough Target Malaria Cost-Effectiveness Calculation." *Google Docs*, https://docs.google.com/spreadsheets/u/1/d/1E8qu474nUUvPjK21oBqGkdQ9FuXanigOgSGfumhL3_c/edit?usp=drive_web&ouid=109160729313225908173&usp=embed_facebook. Accessed 18 Apr. 2018.

Berger, Ken, and Robert M. Penna. "The Elitist Philanthropy of So-Called Effective Altruism." *Stanford Social Innovation Review*, 25 Nov. 2013, https://ssir.org/articles/entry/the_elitist_philanthropy_of_so_called_effective_altruism.

Bill & Melinda Gates Foundation. "What We Do - Malaria Strategy Overview." *Bill & Melinda Gates Foundation*, https://www.gatesfoundation.org/What-We-Do/Global-Health/Malaria. Accessed 31 Oct. 2017.

Bowater, Donna. “Brazil Tackles Zika Virus with Genetically Modified Mosquitoes - Telegraph.” *The Telegraph*, 19 Jan. 2016, https://www.telegraph.co.uk/news/worldnews/southamerica/brazil/12107565/Brazil-tackles-zika-virus-with-genetically-modified-mosquitoes.html.

Brown, Steven. “Supporting the Best Charities Is Harder than It Seems.” *Journal of Global Ethics*, vol. 12, no. 2, Aug. 2016, pp. 240–44.

Bull, James J., and Harmit S. Malik. “The Gene Drive Bubble: New Realities.” *PLOS Genetics*, vol. 13, no. 7, July 2017, p. e1006850. *PLoS Journals*, doi:10.1371/journal.pgen.1006850.

Carrington, Damian. “Global Use of Mosquito Nets for Fishing ‘Endangering Humans and Wildlife.’” *The Guardian*, 31 Jan. 2018, http://www.theguardian.com/environment/2018/jan/31/global-use-of-mosquito-nets-for-fishing-endangering-humans-and-wildlife.

Centers for Disease Control and Prevention. *CDC - How Can Malaria Cases and Deaths Be Reduced? - Indoor Residual Spraying*. https://www.cdc.gov/malaria/malaria_worldwide/reduction/irs.html. Accessed 14 Dec. 2017.

---. *CDC - How Can Malaria Cases and Deaths Be Reduced? - Insecticide-Treated Bed Nets*. https://www.cdc.gov/malaria/malaria_worldwide/reduction/itn.html. Accessed 14 Dec. 2017.

---. *Malaria - About Malaria*. 13 July 2017, https://www.cdc.gov/malaria/about/index.html.

Clough, Emily. *Effective Altruism’s Political Blind Spot | Boston Review*. http://bostonreview.net/world/emily-clough-effective-altruism-ngos. Accessed 21 Feb. 2018.

“Cost Effectiveness.” *The Against Malaria Foundation*, http://www.againstmalaria.com/CostEffectiveness.aspx. Accessed 31 Oct. 2017.

Eckhouse, Seth, et al. “Trends in the Global Funding and Activity of Cancer Research.” *Molecular Oncology*, vol. 2, no. 1, June 2008, pp. 20–32. *ScienceDirect*, doi:10.1016/j.molonc.2008.03.007.

Fang, Janet. “Ecology: A World without Mosquitoes : Nature News.” *Nature*, https://www.nature.com/news/2010/100721/full/466432a.html. Accessed 29 Mar. 2018.

Friedman, Eric. “Reinventing Philanthropy: A Framework for More Effective Giving.” *Reinventing Philanthropy: A Framework for More Effective Giving*, Stanford Social Innovation Review, 2013, https://ssir.org/articles/entry/reinventing_philanthropy_a_framework_for_more_effective_giving.

Garrett, Aaron, “Joseph Butler’s Moral Philosophy”, The Stanford Encyclopedia of Philosophy (Spring 2018 Edition), Edwar.

GiveWell. “2018 GiveWell Cost-Effectiveness Analysis — Version 1.” *Google Docs*, 24 Jan. 2018, https://docs.google.com/spreadsheets/d/1FApXxsBziIo2bcCFBMp9G3F7sDk0VRjzCETbP9J84TE/edit?usp=drive_web&ouid=115586129363615452666&usp=embed_facebook.

---. “Against Malaria Foundation.” *GiveWell*, Nov. 2016, https://www.givewell.org/charities/against-malaria-foundation.

---. “Mass Distribution of Long-Lasting Insecticide-Treated Nets (LLINs).” *GiveWell*, Nov. 2015, https://www.givewell.org/international/technical/programs/insecticide-treated-nets.

Hammond, Andrew, et al. “A CRISPR-Cas9 Gene Drive System Targeting Female Reproduction in the Malaria Mosquito Vector *Anopheles Gambiae*.” *Nature Biotechnology*, vol. 34, no. 1, Jan. 2016, pp. 78–83. *www.nature.com*, doi:10.1038/nbt.3439.

“How Scientists Are Working To Eradicate Zika Mosquitoes.” NPR: Environment, 30 Jan. 2016. Science In Context, Http://Li.

Innis, Michelle. “In Australia, a New Tactic in Battle Against Zika Virus: Mosquito Breeding.” *The New York Times*, 4 Feb. 2016. *NYTimes.com*, https://www.nytimes.com/2016/02/05/world/australia/zika-virus-australia-mosquito-experiment.html.

Innovations for Poverty Action. “Free Malaria Bednets | Innovations for Poverty Action.” *Innovations for Poverty Action*, 1 Apr. 2011, https://www.poverty-action.org/impact/free-malaria-bednets.

Isaacs, Alison T., et al. “Insecticide-Induced Leg Loss Does Not Eliminate Biting and Reproduction in Anopheles Gambiae Mosquitoes.” *Scientific Reports*, vol. 7, Apr. 2017, p. 46674. *CrossRef*, doi:10.1038/srep46674.

Killeen, Gerry F., et al. “Going beyond Personal Protection against Mosquito Bites to Eliminate Malaria Transmission: Population Suppression of Malaria Vectors That Exploit Both Human and Animal Blood.” *BMJ Global Health*, vol. 2, no. 2, Apr. 2017, p. e000198. *CrossRef*, doi:10.1136/bmjgh-2016-000198.

Koekemoer, Lizette L., et al. “Characterization of Multiple Insecticide Resistance in Anopheles Gambiae (Diptera: Culicidae) from Pointe Noire, Republic of the Congo.” *Malaria Journal*, vol. 9, no. 2, Oct. 2010, p. P17. *BioMed Central*, doi:10.1186/1475-2875-9-S2-P17.

MacAskill, William. “Effective Altruism: Introduction.” *Essays in Philosophy*, vol. 18, no. 1, 2017. *CrossRef*, doi:10.7710/1526-0569.1580.

Matthews, Dylan. “You Have $8 Billion. You Want to Do as Much Good as Possible. What Do You Do?” *Vox*, 24 Apr. 2015, https://www.vox.com/2015/4/24/8457895/givewell-open-philanthropy-charity.

Murgia, Madhumita. “Genetically-Modified Mosquitos Released into the Wild to Wipe out Offspring.” *The Telegraph*, 14 Mar. 2016. *www.telegraph.co.uk*, http://www.telegraph.co.uk/technology/2016/03/14/genetically-modified-mosquitos-released-into-the-wild-to-wipe-ou/.

Muyskens, Kathryn. "The Other Half of Effective Altruism: Selective Asceticism." *Essays in Philosophy*, vol. 18, Jan. 2017, https://commons.pacificu.edu/cgi/viewcontent.cgi?referer=https://www.google.com/&httpsredir=1&article=1575&context=eip.

Nadia Whitehead. "What If A Drug Could Make Your Blood Deadly To Mosquitoes?" *NPR.Org*, 29 Mar. 2018, https://www.npr.org/sections/goatsandsoda/2018/03/29/597996321/what-if-a-drug-could-make-your-blood-deadly-to-mosquitoes.

Open Philanthropy Project. "Target Malaria — General Support." *Open Philanthropy Project*, 25 Aug. 2016, https://www.openphilanthropy.org/focus/scientific-research/miscellaneous/target-malaria-general-support.

"Our Work." *Target Malaria*, https://targetmalaria.org/our-work/. Accessed 9 Apr. 2018.

Parfit, Derek, "Future Generations: Further Problems," Philosophy and Public Affairs 11 (1982): 113-172, p. 131.

Pellegrino, Gianfranco. "Effective Altruism and the Altruistic Repugnant Conclusion." *Essays in Philosophy*, vol. 18, no. 1, 2017. *CrossRef*, doi:10.7710/1526-0569.1579.

Regalado, Antonio. "Bill Gates Sees CRISPR Gene Drive Eradicating Mosquitoes in Africa by 2029." *MIT Technology Review*, https://www.technologyreview.com/s/601213/the-extinction-invention/. Accessed 4 Apr. 2018.

Schroeder, Mark, "Value Theory", The Stanford Encyclopedia of Philosophy (Fall 2016 Edition), Edward N. Zalta (Ed.), URL.

Singer, Peter. *The Most Good You Can Do: How Effective Altruism Is Changing Ideas About Living Ethically*. Yale University Press, 2015.

Sohn, Emily. “Malaria Control: The Great Mosquito Hunt.” *Nature News*, vol. 511, no. 7508, July 2014, p. 144. *www.nature.com*, doi:10.1038/511144a.

Stern, Bastian. *Donating vs Investing*. 1 Dec. 2012, https://www.givingwhatwecan.org/sites/givingwhatwecan.org/files/attachments/donating-vs-investing.pdf.

Wolff, Jonathan. *Ethics and Public Policy: A Philisophical Inquiry*.

---. *The Human Right to Health*.

World Health Organization. *Global Report on Antimalarial Drug Efficacy and Drug Resistance: 2000-2010.* World Health Organization, 2010. *Open WorldCat*, http://whqlibdoc.who.int/publications/2010/9789241500470_eng.pdf.

---. *Indoor Residual Spraying: WHO Position Paper*. WHO - World Health Organization, http://apps.who.int/iris/bitstream/10665/69386/1/WHO_HTM_MAL_2006.1112_eng.pdf. Accessed 14 Dec. 2017.

---. “Malaria Fact Sheet.” *WHO*, http://www.who.int/mediacentre/factsheets/fs094/en/. Accessed 31 Oct. 2017.

---. *Malaria Vaccine: WHO Position Paper*. WHO - World Health Organization, Jan. 2016, http://www.who.int/wer/2016/wer9104.pdf?ua=1.

---. *Treatment Failure Plasmodium Falciparum by Drug*. Dec. 2017, http://www.who.int/malaria/areas/drug_resistance/treatment-failure-pf-by-drug.pdf?ua=1.

---. *Treatment Failure Plasmodium Vivax by Drug*. Dec. 2017, http://www.who.int/malaria/areas/drug_resistance/treatment-failure-pv-by-drug.pdf?ua=1.

---. “WHO | Antimalarial Drug Efficacy and Drug Resistance.” *WHO*, http://www.who.int/malaria/areas/treatment/drug_efficacy/en/. Accessed 31 Mar. 2018.

---. “WHO | Cost–effectiveness of Artemisinin Combination Therapy for Uncomplicated Malaria in Children: Data from Papua New Guinea.” *WHO*, http://www.who.int/bulletin/volumes/89/3/10-084103/en/. Accessed 5 Apr. 2018.

---. “WHO | How Malaria RDTs Work.” *WHO*, http://www.who.int/malaria/areas/diagnosis/rapid-diagnostic-tests/about-rdt/en/. Accessed 14 Dec. 2017.

---. *WHO World Malaria Report 2017*. WHO - World Health Organization, http://apps.who.int/iris/bitstream/10665/259492/1/9789241565523-eng.pdf?ua=1. Accessed 13 Dec. 2017.

Biography

Sam Williams was born in Austin, Texas in 1996 and enrolled in the University of Texas at Austin in the Fall of 2014. From UT, he earned a bachelors in business administration in finance and a bachelors in the arts from the Plan II Honors Program. In addition to his writing on effective altruism and antimalarial efforts, Sam has also written on the prevalence of high-frequency trading and dark pools in financial markets. While in college, Sam covered sports for the *Daily Texan* and undertook several internships in the financial industry. Upon graduation, he will be moving to New York City to work in management consulting with a specification in mergers and acquisitions.